

## Parallel Mesh Generation

Nikos Chrisochoides

Computer Science Department, College of William and Mary, Williamsburg,  
VA 23185, USA

Division of Applied Mathematics, Brown University, 182 George Street, Providence,  
RI 02912, USA

nikos@cs.wm.edu

**Summary.** Parallel mesh generation is a relatively new research area between the boundaries of two scientific computing disciplines: computational geometry and parallel computing. In this chapter we present a survey of parallel unstructured mesh generation methods. Parallel mesh generation methods decompose the original mesh generation problem into smaller subproblems which are meshed in parallel. We organize the parallel mesh generation methods in terms of two basic attributes: (1) the sequential technique used for meshing the individual subproblems and (2) the degree of coupling between the subproblems. This survey shows that without compromising in the stability of parallel mesh generation methods it is possible to develop parallel meshing software using off-the-shelf sequential meshing codes. However, more research is required for the efficient use of the state-of-the-art codes which can scale from emerging chip multiprocessors (CMPs) to clusters built from CMPs.

### 7.1 Introduction

This chapter presents a survey of parallel unstructured mesh generation methods based on three widely used techniques: Delaunay [39], Advancing Front [66], and Edge Subdivision [59]. Parallel methods for quadrilateral [6] and hexahedral [54] mesh generation as well as block structured [93, 20, 90] and structured adaptive mesh refinement [1] methods are not reviewed in this chapter.

Parallel mesh generation procedures in general decompose the original 2-dimensional (2D) or 3-dimensional (3D) mesh generation problem into  $N_s$  smaller subproblems which are solved (i.e., meshed) concurrently using  $P$  processors. The subproblems can be formulated to be either tightly coupled [60, 56, 78], partially coupled [55, 31, 19] or even decoupled [38, 79, 52]. The coupling of the subproblems determines the intensity of the communication and the amount/type of synchronization required between the subproblems.

The challenges in parallel mesh generation methods are: to maintain *stability* of the parallel mesher (i.e., retain the quality of finite elements generated by state-of-the-art sequential codes) and at the same time achieve 100% *code re-use* (i.e., leverage the continuously evolving and fully functional off-the-shelf sequential mesh-

ers) without substantial deterioration of the *scalability* of the parallel mesher. In this chapter we review parallel mesh generation methods having in mind these three requirements.

We build on top of previous work [30, 39] where parallel mesh generation methods are classified in terms of the *way* and the *order* the artificial boundary surfaces (interfaces) of the subproblems are meshed. Specifically, in [31, 39] existing parallel methods are classified in three categories: (i) methods that first mesh (either in parallel [55] or sequentially [79]) the interfaces of the subproblems and then mesh in parallel the individual subproblems, (ii) methods that first solve the meshing problem in each of the subproblems in parallel and then mesh the interfaces so that the global mesh is conforming [36], and (iii) methods that simultaneously mesh and improve the interfaces as they mesh the individual subproblems [25, 19, 26].

In this chapter we organize the parallel mesh generation methods in terms of two basic attributes. First, the sequential technique used for meshing the individual subproblems: (1) *Delaunay*, (2) *Advancing Front*, and (3) *Edge Subdivision*. Second, the degree of coupling between the subproblems: (a) *tightly-coupled*, (b) *partially-coupled*, and (c) *decoupled methods*.

## 7.2 Domain Decomposition Approaches

Parallel mesh generation methods use a sequential pre-processing step for the data partitioning problem with the exception of [47, 48]. The data are partitioned using either the continuous domain which is decomposed into *subdomains* (see Figure 7.1, left) or a discrete approximation (i.e., an initial coarser mesh) of the domain which is decomposed into *submeshes* (see Figure 7.1, right). The internal boundaries between the subdomains or submeshes ( $S_i$ ) are called *interfaces* or *separators* ( $\partial S_i$ ). In both cases the number of generated subdomains or submeshes ( $N_s$ ) can be significantly greater than the number of processors  $P$  (*over-decomposition*). Over-decomposition was introduced in parallel computing in mid 80s. It is used to hide communication latency in message passing [50] and to mask information dissemination, decision making and data migration costs in dynamic load balancing [21].

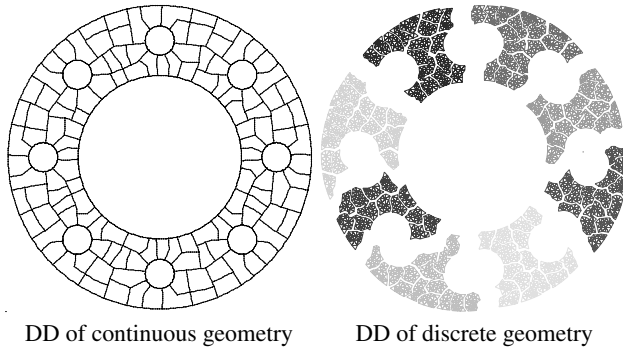
The domain decomposition (DD) problem in parallel mesh generation is defined as follows:

$$lfs(\Omega) \leq lfs(S_i) \quad i = 1, N_s \quad (7.1)$$

$$\min_{i=1, N_s} \frac{|\partial S_i|}{|S_i|} \quad (7.2)$$

$$\partial S_i \text{ form "good" angles between each other and the boundary } \partial \Omega. \quad (7.3)$$

where  $lfs(\Omega)$  and  $lfs(S_i)$  are the local feature size [85] of the original domain  $\Omega$  and the subdomains (or submeshes)  $S_i$ , respectively. The  $|\partial S_i|$  denotes the length (in 2D) and surface (in 3D) of the interfaces while the  $|S_i|$  denotes the area (in 2D) and volume (in 3D) of the subdomains (or submeshes)  $S_i$ .



**Fig. 7.1.** Domain decomposition of the continuous geometry [52] and the discrete geometry [17] of a cross section of a rocket pipe. (For the color version, see Figure A.17 on page 474).


### *Continuous Domain Decomposition*

The continuous domain decomposition methods partition the region  $\Omega$  into subdomains  $\Omega_i, i = 1, N_s$ . There are two types of continuous DD methods. The first and most popular approach is based on quadtree/octree methods [31, 56, 58]. The octree methods utilize an octree structure for the decomposition of  $\Omega$  into blocks (octants). The octants along with a description of the external boundary  $\partial\Omega$  define the subdomains. Another class of continuous DD methods [52] is based on auxiliary structures like the Medial Axis [9, 68, 91] so that the subdomains  $\Omega_i$  have no new features like small angles between the separators and the separators and external boundary [52].

Continuous DD approaches are attractive because they refine the individual subdomains by re-using existing well tested and fine-tuned sequential codes on each subdomain independently. However, independence in mesh refinement and high code re-use in some cases come at a price. The polyhedral surfaces which arise due to the decomposition of the initial mesh impose additional constraints on the execution of sequential meshing algorithms in each of the subdomains. Poorly generated interface surfaces can affect the termination of meshing algorithms and the quality of the elements. Moreover, the artificially imposed interfaces can affect the mesh gradation.

### *Discrete Domain Decomposition*

The Discrete DD methods partition an initial coarse mesh (usually a boundary conforming mesh),  $D$  into a number of simply-connected submeshes  $D_i, i = 1, N_s$  while they try to minimize the surface-to-volume ratio for each of the submeshes. Usually a coarse mesh is generated on a high-performance workstation using sequential mesh generators. The partitioning of a coarse mesh is performed either sequentially or in parallel using generic graph partitioning libraries like Metis/Parallel Metis [80] and Chaco [42]. Also, there are mesh partitioning libraries like Domain Decomposer [22, 23], Zoltan [33], Drama [3], Plum [64], and Jove [87] (to mention



**Fig. 7.2.** Bowyer-Watson kernel starts with a mesh  $M_i$  (left), computes the cavity (center left) of a newly inserted point, triangulates the cavity (center right) and updates the mesh into  $M_{i+1}$  (right).

a few) which extend and customize the generic data partitioning techniques for FEM calculations.

### 7.3 Parallel Mesh Generation Methods

In this section we review parallel mesh generation methods which are based on Delaunay triangulation in Section 7.3.1, Advancing Front Technique in Section 7.3.2, and Edge Subdivision methods in Section 7.3.3.


#### 7.3.1 Delaunay Based Methods

There are many approaches to generate Delaunay meshes [39], we focus on methods based on Bowyer-Watson [10, 96] kernel which can lead to: (1) more efficient parallel implementations due to easier optimizations for improving data locality and (2) simpler and more efficient data structures. The Bowyer-Watson (BW) kernel is described in Figure 7.2 and the loop below:

Algorithm 1 (BW( $M_o, p_1, \dots, p_n$ )).

1. **Input:** Mesh  $M_o$  an initial mesh and a set of  $n$  points
2. **for**  $i = 1, n$
3.     Compute the cavity of  $C_i$  of the point  $p_i$
4.     Compute the ball  $B_i$  of point  $p_i$
5.      $M_{i+1} = M_i - C_i + B_i$
6. **endfor**
7. **Output:** A new mesh  $M = M_{n+1}$

where the cavity  $C$  of a point  $p$  is defined as the set of all triangles whose circumcircle includes  $p$ ; the ball  $B$  of a point  $p$  is defined as set of new triangles defined by the point  $p$  and the vertices of the boundary of its cavity [39].



**Fig. 7.3.** (a) Intersection of two cavities and (b) two cavities share an edge; solid lines represent the edges of the initial triangulation, and dashed lines edges created by the insertion of  $p_8$ ,  $p_9$ , and  $p_{10}$ .

The challenge, for parallel mesh generation methods based on the BW kernel, is to maintain the following loop invariant:  $M_i$  is conformal and Delaunay, for  $i = 1, n$ . Figure 7.3 depicts two cases where the concurrent point insertion violates the loop invariant. First, the cavities intersect i.e., there is a triangle  $\triangle p_3 p_6 p_7 \in \mathcal{C}(p_8) \cap \mathcal{C}(p_9)$ , then concurrent insertion of  $p_8$  and  $p_9$  results in a non-conformal mesh. Second, the cavities share an edge in 2D (or a face in 3D), an edge  $p_3 p_6$  is shared by  $\mathcal{C}(p_8) = \{\triangle p_1 p_2 p_7, \triangle p_2 p_3 p_7, \triangle p_3 p_6 p_7\}$  and  $\mathcal{C}(p_{10}) = \{\triangle p_3 p_5 p_6, \triangle p_3 p_4 p_5\}$ , then the new triangle  $\triangle p_3 p_{10} p_6$  can have point  $p_8$  inside its circle, thus, violating the Delaunay property.

The focus of this section is on parallel mesh generation methods that address this challenge. There is a number of parallel Delaunay and triangulation methods like the MIMD method in [92] and the HPF implementation in [14] which target parallel programming paradigms no longer in use for practical purposes. Other methods [27, 62, 63, 41] also contributed in shaping up this author's directions and work in parallel mesh generation and implicitly contribute in this chapter.

In [7] the authors describe a divide-and-conquer projection-based algorithm for constructing in parallel 2D Delaunay triangulations of a set of given points. The method extends to 3D, but its implementation is quite complex. The goal in parallel mesh generation, though, is to refine an existing mesh by inserting new points i.e., the set of points in the final mesh is not known in advance.

In [46, 48] the authors extended [7] for parallel 2D mesh generation which further eliminates the sequential step for an initial mesh, but does not address the issue of code re-use. The method in [46, 48] is partially coupled.

In [35] the authors define the points  $x$  and  $y$  as independent if the closures of their *prestars* (or *cavities*) are disjoint. The approach in [35] does not provide a way to schedule the concurrent insertion of points whose cavity closures are disjoint.

In [88] the authors presented the first theoretical analysis of the complexity of parallel Delaunay refinement algorithms. However, the assumption in [88] is that the global mesh is completely retriangulated each time a set of independent points is inserted. In [89] the authors developed a more practical algorithm.

In the rest of this section we describe five different practical (i.e., they have been implemented) parallel Delaunay mesh generation methods. These methods formulate the subproblems to be: (1) tightly coupled, (2) decoupled, and (3) partially coupled.

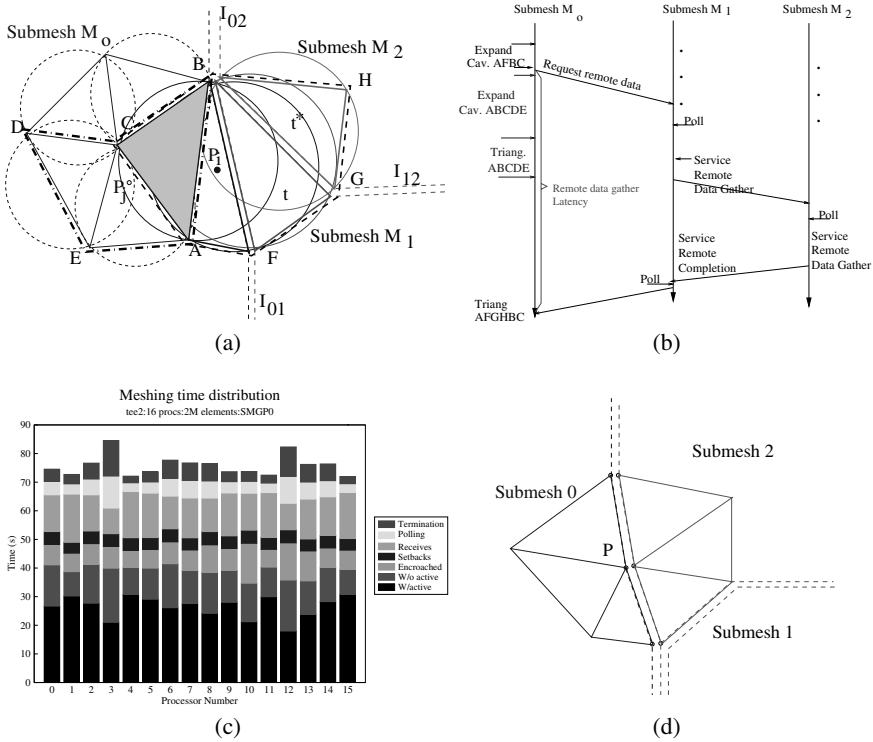
### Tightly Coupled Methods

A straight forward approach to parallel computing is based on identifying some partial order among the computations of well understood and successful sequential kernels and then in a brute-force fashion use message passing or threads to implement the computations on distributed and shared memory parallel machines, respectively. This approach leads to the tightly coupled method presented in [60] for parallel guaranteed quality Delaunay mesh generation.

#### *Parallel Optimistic Delaunay Meshing (PODM) Method*

In [60] the authors presented the first provable 3D parallel guaranteed quality Delaunay mesh generation method for polyhedral domains. PODM is based on discrete domain decomposition, but it is not constrained by the interfaces of the submeshes. The algorithm guarantees the stability by simultaneously re-partitioning and refining the interface surfaces and volume of the submeshes [26]—refinement due to a point insertion might extend across subproblem (or submesh) boundaries. The extension of a cavity beyond the interfaces is a source of intensive communication. However, PODM can tolerate most of the communication by concurrently refining other regions of the submeshes while it waits for remote data to arrive. Unfortunately, the concurrent refinement can create a number of inconsistencies in the mesh (see Figure 7.3). These inconsistencies are resolved at the cost of setbacks (or rollbacks [44]) and thus we call this method Parallel Optimistic Delaunay Meshing method. Setbacks is a source of major algorithm and code re-structuring (due to overlapping cavities) and they lead to zero code re-use. Unfortunately, the overlapping of the cavities becomes even more complex when they are near the external boundary, where a certain order of inserted points needs to be maintained due to encroachment rules that are used to maintain and prove the quality of the elements and thus satisfy the stability requirement.

Figure 7.4a depicts a cavity which extends beyond the submesh interfaces (because two of the cavity BHGFAC triangles  $t \in M_1$  and  $t^* \in M_2$  are non-local to submesh  $M_0$ ) in order to guarantee the quality of the mesh. The extension of the cavity beyond the interfaces is a source of intensive communication. However, as Figure 7.4b shows PODM can tolerate the communication by concurrently refining other regions (e.g. compute a new cavity ABCDE) of the submeshes while it waits for remote data (e.g. the partially completed cavity BCAF) to arrive (eg. rest of the cavity BFGH). Unfortunately, the concurrent refinement can lead the violation of the loop invariant by creating non-conforming meshes and/or the violation of the Delaunay property as is the case in Figure 7.4a where the point  $P_j$  is within the circumcenter of  $\triangle P_i C A$  which is a newly created triangle from the triangulation of the cavity (BHGFAC) that corresponds to the point  $P_1$ . These violations are resolved at




**Fig. 7.4.** a) cavity extension beyond submesh interfaces, b) time diagram with concurrent point insertion, c) a breakdown of execution time for PODM, and finally d) the refinement of a cavity with simultaneous distribution of the newly created elements. (For the color version, see Figure A.18 on page 475).

the cost of setbacks and frequent message polling shown in the performance graph of Figure 7.4c. With some additional communication cost the PODM becomes domain decomposition independent and moreover re-distributes new elements as they are generated (Figure 7.4d).

In summary, PODM does not depend on good domain decompositions before, during and after parallel meshing at the cost of being labor intensive approach. PODM is a stable and tightly coupled method, with zero code re-use.

### Decoupled Methods

In [38, 52] the authors present two approaches which achieve 100% code re-use and eliminate communication and synchronization. Both approaches rely on continuous domain decomposition and decouple the individual subdomains (subproblems) so



**Fig. 7.5.** The circumcenter of the face  $\triangle PQR$  lies on the plane  $\Pi$  which helps define a separator  $S_\Pi$ . Note that  $\triangle PQR \in S_\Pi$ .

that they can be meshed independently. Earlier, in [8] the authors presented similar approach for the parallel triangulation of a set of fixed points.


### *Parallel Projective Delaunay Meshing*

The Parallel Projective Delaunay Meshing ( $P^2DM$ ) method [38] starts by sequentially meshing the external surfaces of the geometry and by pre-computing domain separators whose facets are Delaunay-admissible (i.e., the precomputed interface faces of the separators will appear in the final Delaunay mesh). The separators decompose the continuous domain into subdomains which are meshed in parallel using a sequential Delaunay mesh generation method on each of the processors.

The basic idea for computing Delaunay-admissible separators can be explained easier in the context of the parallel triangulation of a convex hull for a set of points  $S \in \mathcal{R}^3$  [36, 30]. The convex hull of a set of points  $S$  is decomposed in two subdomains by computing a Delaunay admissible separator as follows: First, the position of a surface (in practice a plane  $\Pi$ ) is computed using an Inertia Axis Decomposition method [33]. The plane  $\Pi$  decomposes the convex hull of  $S$  into two almost equal pieces (in terms of points). Then the algorithm finds all faces  $(P,Q,R) \in \mathcal{R}^3$  (see Figure 7.5) for which there is an empty sphere whose center lies on the plane  $\Pi$  and passes through the points  $P,Q,R$ . These faces constitute a polyhedral separator  $S_\Pi$  which decomposes the domain into two subdomains assuming that the corresponding tetrahedra  $PQRX$  and  $PQRY$  contain the centers of their respective circumscribed spheres i.e., the quality of the initial mesh around the separators is very good which requires substantial refinement around the separators. In [38] it is shown that the faces of the polyhedral separator  $S_\Pi$  will appear in the final Delaunay triangulation of the convex hull. The generalization of the idea to complex geometries is possible, however it is much more difficult and it is explained in [38].

It is possible that the pre-constructed separators can not be Delaunay-admissible [38] and the whole process has to start from the beginning. This is a very difficult problem which for 2D has been solved in [52] using a different approach.





**Fig. 7.6.** The Medial Axis Transformation (a) which in turn is used to achieve high quality domain decomposition (b). For  $PD^3$  the interfaces of the subdomains are refined (c) in a pre-processing step in order to decouple the subdomains which are refined independently (d).

### *Parallel Delaunay Domain Decoupling $PD^3$ Method*

The  $PD^3$  method [52] like  $P^2DM$  is based on continuous domain decomposition.  $PD^3$ , for the domain decomposition of 2D geometries, uses medial axis of the domain and relies on the following simple geometric property [52]:

**Lemma 1.** *Let  $MA(\Omega)$  be the medial axis of  $\Omega$  and  $b$  a contact point of  $c \in MA(\Omega)$ . The angles formed by the segment  $cb$  and the tangent of the boundary  $\partial\Omega$  at  $b$  are at least  $\pi/2$ .*

The medial axis of a domain  $\Omega$  is approximated by Voronoi points of a discretization of the domain. Figure 7.6a depicts the medial axis approximation and a 8-way partition (b) for the same geometry. The level of the discretization of the boundary determines the quality of the approximation of the medial axis. However, the goal in [52] is not to approximate accurately the medial axis, but to obtain good angles from the separator. Therefore, the criteria for the discretization of the domain are determined from the quality of the angles formed between the separators and the external boundary of the domain [53].

After the decomposition of the domain (see Figure 7.6b),  $PD^3$  constructs a “zone” around the interfaces of the submeshes. The “zone” consist of the union of all diametral circles of the interface edges (see Figure 7.6c). The interfaces of the subdomains are refined using the  $lfs$  of the original domain. This leads into an overrefinement of the final distributed mesh. Experimental data from  $PD^3$  (see Table 7.1) suggest that the overrefinement is not as high as one could expect. However, the authors of [52] are working on a new approach which will use adaptive domain decomposition [53] and different  $lfs$  for different interfaces of the subdomains. This method is expected to reduce overrefinement of the interfaces and produce well graded meshes [51].

In [52] the authors prove that sequential Delaunay meshers will not insert any new points within a zone around the subdomain interfaces i.e., the sequential Delaunay meshing on the individual submeshes can terminate without inserting any new points on the interfaces and thus eliminate communication and modifications of the sequential codes. This way, the problem of parallel meshing is reduced into a

**Table 7.1.** Overrefinement data as we increase the number of subdomains for the decomposition of a cross section of a rocket pipe model.

Subs	1	16	32	64	128
Elms :	21,016,403	21,016,857	21,018,522	21,030,711	21,044,689
ORef.Elms/Sub	0	28	66	379	299

“proper” domain decomposition and a discretization of interfaces. However, the construction of decompositions that can decouple the mesh is a challenging problem, since its solution is based on medial axis which is very expensive and difficult to construct (even to approximate) for complex 3-dimensional geometries [40, 83, 29].

### Partially Coupled Methods

The parallel tightly-coupled and decoupled methods we have seen so far address some of the parallel mesh generation requirements we described in Section 7.1. For example, PODM is a 3D stable and domain decomposition independent, but it is zero code re-use with high communication method; while  $P^2DM$  and  $PD^3$  address the code re-use and communication issues, but their applicability in 3D is limited by the Delaunay-admissible and domain decomposition problem, respectively. In the rest of this section we present two partially coupled methods that make an attempt to balance trade-offs between all three requirements and the domain decomposition problem at the cost of some communication.

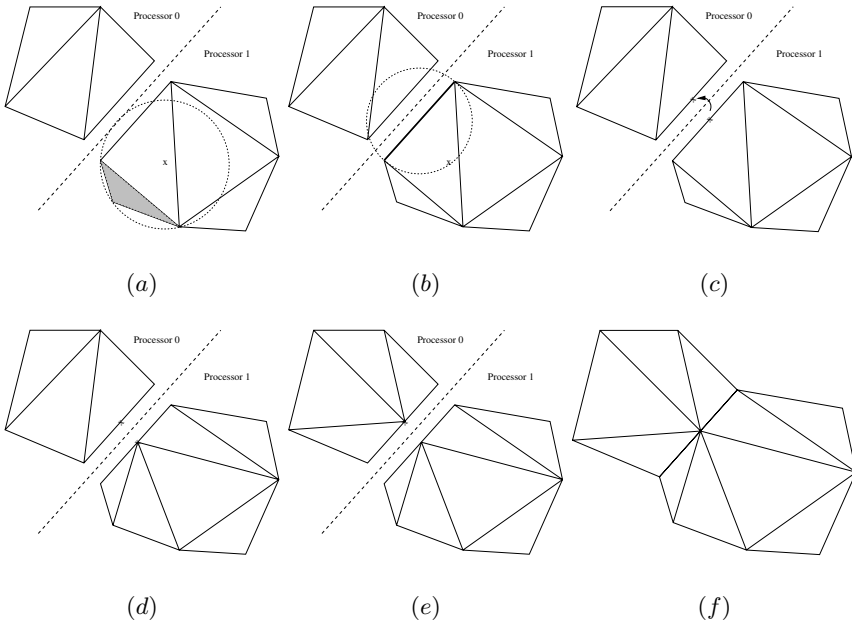
#### *Parallel Constrained Delaunay Meshing (PCDM) Method*

In order to address the communication and synchronization problem in [19], the authors developed the PCDM which is asynchronous and can reduce the variable and unpredictable communication patterns to irregular but bulk communication.

The PCDM [19] is based on the Constrained Delaunay Triangulation [18] and a discrete DD method. Each submesh is treated as an independent mesh defined by external boundary (if any) and/or constrained edges which are the edges of the interfaces between any pair of adjacent submeshes.

Intuitively, the constrained Delaunay triangulation is as close as one can get to the Delaunay triangulation given that one needs to preserve certain (constrained) edges and internal boundaries. It has been shown in [18] that the constrained internal edges do not affect the quality of the resulting mesh more than the edges and faces that define the external boundary. However, one might be able to identify such boundaries (interfaces for the PCDM) in the resulting mesh by noting the way in which triangle edges are aligned. Using the idea of a constrained Delaunay mesh generation one can introduce in the mesh artificial constrained edges (interfaces) which decompose the mesh into submeshes and can be meshed almost independently.

By the definition of the constrained Delaunay mesh, points inserted on one side of an interface have no effect on triangles on the other side; thus, no synchronization is required during the element creation process. In addition, communication between



**Fig. 7.7.** Processor P1 inserts a new point (a) which is encroaching upon an interface edge (b). Then P1 discards the new point and inserts the midpoint of the encroached edge (c) while at the same time it sends a request to split the same interface edge on processor P0. Processor P0 computes the cavity of the midpoint (d). The triangulation of the cavities (e) and (f) of the midpoint of the interface edge results in a new conforming and distributed Delaunay (in the CDT sense) triangulation which guarantees the quality of the elements.

submeshes is tremendously simplified: the only message between adjacent processes is of the form [19]: “Split this interface (i.e., constrained) edge” if a newly inserted point encroaches (see Fig. 7.7) upon an interface edge. Since interface edges are always split exactly in half, no additional information needs to be communicated.

The PCDM is an asynchronous with bulk communication and thus partially coupled method. Moreover, the number and size of messages can be reduced by message aggregation [17]. Although this optimization improves the performance of the PCDM it has its own problems when many “Split this interface edge” messages are delayed. This causes performance degradation due to: (1) the large number of accumulated messages which can consume memory, (2) redundant computation (by delaying messages), from neighboring processors which are unaware of each other’s interface splits. In [17] these problems are addressed by a mechanism which adaptively changes the number of messages allowed to be aggregated before a low-level message is send.

However, code re-use remains a problem due to “Split this interface edge” message and optimizations required for reducing the fine-grain communication to a bulk and asynchronous message passing.

### *Parallel Delaunay Refinement (PDR) Method*

The above tightly coupled (PODM) and partially coupled (PCDM) methods [61, 24, 19] require algorithm re-structuring and thus lead to completely new implementations for parallel Delaunay mesh generation. The implementation of sequential mesh generation codes is labor intensive and requires multi-disciplinary effort; it takes about ten to fifteen years to develop the algorithmic and software infrastructure for sequential industrial strength mesh generation libraries. Moreover, improvements in terms of quality, speed, and functionality are open ended and permanent which makes the task of delivering state-of-the-art parallel mesh generation codes much more difficult.

This problem is addressed by  $P^2DM$  and  $PD^3$  in [38, 52], where two decoupling methods are presented in order to use (without modifications) optimized and fully functional sequential codes on each of the subproblems and eliminate communication and synchronization. However,  $P^2DM$  can suffer setbacks due to difficulty of constructing Delaunay-admissible separators and  $PD^3$ , for 3D geometries, is expected to be suffer high pre-processing overhead due the construction (or approximation) of the medial axis.

With PDR in [16, 15] the authors try to balance trade-offs between the data decomposition, communication and code re-use i.e., maintain stability and achieve high code re-use using a simple domain decomposition method at the cost of some communication. The key idea of the PDR method is based on the concurrent point insertion of more than two points without calculating their corresponding cavities ahead of time in order to decide whether they violate the conformity and Delaunay properties of the mesh. PDR accomplishes this objective by introducing for the first time a practical Delaunay-independence criterion for concurrent point insertion [16]:

**Theorem 1** *Let  $\bar{r}$  be the upper bound on triangle circumradius in the mesh and  $p_i, p_j \in \Omega \subset \mathbb{R}^2$ . Then if  $\|p_i - p_j\| \geq 4\bar{r}$ , then independent insertion of  $p_i$  and  $p_j$  will result in a mesh which is both conformal and Delaunay.*

Theorem 1 is applicable throughout the run of the algorithm, since the execution of the Bowyer-Watson kernel, either sequentially [10, 96] or in parallel [61], does not violate the condition that  $\bar{r}$  is the upper bound on triangle circumradius in the entire mesh [16]. However, checking the inequality of the theorem, for every pair of candidate points, would be quite expensive task. In [16] the authors present a simple block domain decomposition scheme<sup>1</sup> which guarantees that any pair of points in non-adjacent cells are far apart no less than  $4\bar{r}$ . To enforce the  $\bar{r}$  circumradius bound in the mesh they derive the following relation which allows the use of

---

<sup>1</sup>This scheme is based on a simple block decomposition for uniform mesh refinement [16] and octree decomposition for graded mesh refinement [15].

a standard sequential Delaunay refinement algorithm/software like Triangle [84] for preprocessing [16]:

**Theorem 2** *If  $\bar{\rho}$  and  $\bar{\Delta}$  are upper bounds on triangle circumradius-to-shortest edge ratio and area, respectively, then  $\bar{r} = 2(\bar{\rho})^{3/2}\sqrt{\bar{\Delta}}$  is an upper bound on triangle circumradius.*

### 7.3.2 Advancing Front Based Methods

All five parallel Delaunay methods we present in Section 7.3.1 maintain the stability of the parallel mesher. However, parallel finite element codes require only “good” quality of elements and the definition of quality depends on the field solver and varies from code to code. For example, in [79] although the stability is not guaranteed, it appears that the generated meshes are practical and of “good” quality. This raises the following two questions: Is the stability of parallel mesher important? Does the parallel mesh generation without the stability requirement become easier?

The answer to the first question depends on the upstream solver. Regarding the second question, our experience<sup>2</sup> suggests that even if we relax the stability criterion the problem of parallel mesh generation does not become easier. In fact, the termination problem (which is even more fundamental than the stability) becomes, for some cases, a very important issue. In some cases, subdomains or submeshes obtained from state-of-the-art partitioning libraries can not be meshed even by industrial strength advancing front sequential meshers. Parallel mesh smoothing techniques [57] are helpful, but do not work always.

There is a trade-off between the domain decomposition and the capability of the sequential mesher required to mesh the individual subdomains. A balance between the two is important not only for stability but even for termination. Two successful Parallel Advancing Front Techniques [56, 28] address this issue by what we refer to as *guided re-partitioning or shifting of the separators*. In [56] the authors present a tightly coupled method for shared memory machines and in [28] the authors present a partially coupled method for distributed memory machines. We review both methods in the rest of this section.

#### Tightly Coupled Methods

Löhner et al. in [56] revisit a partially coupled Parallel Advancing Front Technique (PAFT) they developed in [55] (see below) in order to address the termination, stability, and code re-use requirements. In [56] they address these issues by developing a PAFT for shared memory computers (PAFT<sub>SM</sub>). However, instead of generating and partitioning a very fine-grain octree as in [58] on a single processor, for the whole geometry, they use an octree to identify the zones where elements can be introduced concurrently. They set the edge length of the smallest octree box to be an order of magnitude larger than the specified size of elements and they use the “shift

---

<sup>2</sup>From the implementation of a method similar to one appeared in [79].

and regrid” technique, but in a completely different way from the method in [32]. The PAFT<sub>SM</sub> is broken into two phases: (1) the AFT phases and (2) “shift” or as we call it here guided re-partitioning phase. At each AFT phase the active front expands and a new one is created. The process continues until the whole domain is meshed. The PAFT<sub>SM</sub> method synchronizes at the beginning of each AFT phase in order to sequentially refine and re-partition the global octree, for the new active front, whose leaves will be refined in parallel. The method is suitable for shared memory machines but can not be used in large-scale distributed memory parallel platforms, because of the global synchronization required between the mesh generation and re-partitioning phases.

The PAFT<sub>SM</sub> is stable and code re-use is achieved at the cost of global synchronization which is not expensive on shared memory machines.

### Partially Coupled Methods

In [55] Löhner et al. introduced the first 2D PAFT. The initial mesh is subdivided into submeshes using a discrete domain decomposition approach. Each submesh is further separated into an interior region and interface regions, where interface regions of a submesh are defined to be the set of elements that are adjacent to elements that belong to different submeshes. The interior regions of each submesh are refined independently. The interface regions and then the corners are refined once all the interior and interface regions are meshed, respectively (a posteriori approach). The order of meshing interface and interior regions can change i.e., interfaces can be refined first (a priori approach) and the interior regions refined last [55]. The submeshes synchronize locally, because no new elements can be inserted in the interfaces and corner regions before the meshing of adjacent interior and interface regions, respectively. The pre-computed interface regions work well for AFT because they create buffer zones which fully decouple the interior regions of the submeshes.

#### *Parallel Octree AFT (POAFT) Method*

The 3D POAFT in [28], contrary to the PAFT in [55], is based on continuous domain decomposition method. The POAFT method generates a distributed coarse-grain octree using a divide-and-conquer algorithm. The terminal octants and the geometric model of a domain define the subdomains. The terminal octants of the octree are classified into: interior, interface, boundary, and complete. Interface octants have at least one adjacent octant which is not local. Boundary octants include mesh entities from the input surface mesh. Complete octants have no front faces in their volumes. The subdomains represented as subtrees (on each processor) which are refined further until their leaves reach to a predefined size to use tetrahedral meshing templates. The new octrees are repartitioned (using stop-and-repartition methods) in order to guarantee load balancing during the execution of meshing templates. After the redistribution of interior octants, mesh templates are applied so that the triangulations conform on both sides of interface-octant faces. The interior octants of one processor are independent of the interior octants in other subdomains and thus can be meshed

in parallel. At this step code re-use is high, since the meshing templates of sequential octree meshing code [82] can be used on each processor. Existing scalar meshing templates for the interface octants can be used, but some communication will be required during the meshing process. Instead in [28] meshing templates were re-designed in order to guarantee conformity without compromising stability and eliminating communication. The potential for ambiguous splits of faces is addressed and resolved in [69].

Before the boundary octants are refined and meshed a re-partitioning might take place if it is necessary. Any parallel partitioning algorithm can be used; in [28] the parallel recursive inertia bisection method is applied. The meshing of boundary octants is a challenging task. Every processor applies a tree-based face removal procedure [28] in order to connect the input surface mesh with the mesh of the interior octants. The face removal (from the active front) is a basic operation in AFT and it consists of connecting a front mesh face to a target mesh vertex which is drawn from a “neighborhood” of the face [28]. In the parallel face removal, portion of the “neighborhood” might be on a remote processor and a target vertex can not be found locally; in this case the face removal is postponed. This will create unmeshed regions between the terminal interface boundary octants and input surface mesh. In [28] active terminal and boundary interface octants are repartitioned so that the remaining unmeshed “neighborhoods” become local and thus the face removal becomes a local operation. This permits code re-use. This process is repeated until there are no unmeshed regions. The “guided” repartitioning is a very challenging problem.

### 7.3.3 Edge Subdivision Based Methods

Parallel Edge Subdivision (PES) methods have been used successfully for both 2D domains [97, 45] and 3D geometries [13, 28, 65, 78]. PES methods use discrete DD for data decomposition and their termination and stability does not depend on the geometric properties of the submeshes. Once a coarse mesh is partitioned into submeshes, the individual submeshes are refined in parallel by splitting tetrahedra using sequential subdivision techniques. The longest-edge bisection method [72, 76] is the most commonly used for parallel refinement/derefinement [97, 45, 13, 78]. In 2D an element is refined into two triangles by adding an edge defined by the longest-edge midpoint and its opposite vertex, while in 3D an element is refined into two tetrahedra by adding a triangle defined by the longest-edge midpoint and its two opposite vertices. The longest-edge bisection technique is attractive because it simplifies the management of intermediate non-conforming points throughout the process. With the introduction of terminal-edges in [78] this management is localized in a similar way the cavity localizes the computation of Delaunay based methods.

Like all parallel mesh generation methods PES refinement methods should satisfy all three requirements we listed in the introduction of this chapter. Existing PES methods address some of these requirements successfully and have the potential to meet all the requirements in the future. In [12] the authors present a termination proof, for parallel longest-edge bisection algorithms, using Dijkstra’s general termi-

nation algorithm [34, 5]. Moreover, they prove the stability and even show that the mesh refined in parallel is identical to a sequentially generated mesh.

The scalability of PES methods depends on the way they address the *refinement collision: more than one processor split concurrently two different copies of the same interface edge*. Other factors that affect the scalability is the choice of dynamic load balancing methods and the degree of code re-use. For example, frequent use of stop-and-repartition methods due to global barrier operations can deteriorate the scalability of computationally inexpensive parallel mesh generation methods [2]. In general parallel mesh generation methods that do not take advantage of highly optimized sequential codes have difficulty to demonstrate good scalability against the best sequential codes.

In [65] it has been shown that 100% code re-use is possible at the cost of 10% overhead by putting a wrapper around the sequential data structure in order to handle data distribution and remote memory accesses. Communication is another aspect of parallel codes that affects scalability. In [28] the authors present a number of subdivision templates that can be used to decouple the refinement on different processors and thus eliminate communication completely.

The main challenge in PES methods is the *collision refinement* problem. In order to achieve mesh conformity and correctness the interface faces between the submeshes should be subdivided the same way from all submeshes that share them. Thus interface edges that are subdivided in one submesh are marked to be subdivided from all other submeshes that share them. This causes communication which is handled by sending, at the end of the refinement of the interface faces, a message to submeshes that share refined faces and edges. Based on the communication and synchronization requirements for handling the refinement collision problem, the PES methods are classified into three categories: tightly coupled methods [78], partially coupled methods [45, 28, 13, 65] and decoupled methods [74].

### Tightly Coupled Methods

The 3D Parallel Terminal-Edge (PTE) method described in [78] is an inherently decoupled method. However, the PTE method in [78] is implemented as a coupled method. In [74] a new design and implementation is presented so that the stability and code re-use requirements are satisfied while the global synchronization for maintaining the global name of all bisected edges is eliminated.

#### *Parallel Coupled Terminal-Edge (PCTE) Method*

In this paragraph we will refer to the tightly coupled implementation of the PTE method as PCTE. The PCTE method is based on longest-edge bisection approach introduced by Rivara [71, 72, 76]. Triangles/ tetrahedra are refined by bisecting their longest-edge. The longest-edge bisection algorithm requires the management of sequences of intermediate non-conforming mesh points throughout the refinement process. This complicates its parallel implementation because it requires some synchronization in order to handle the collision refinement and global name of newly inserted vertices, both are required to maintain the conformity of the distributed mesh.



The PCTE method [78], although it requires zero communication between processors, relies on a central processor for global name-assignment of new mesh points. The use of the central processor limits the scalability of the method for more than 60 processors and reduces the speed (tetrahedra per second) of the method by an order of magnitude. However, in [74] the authors present a decoupled method and implementation which takes full advantage of the terminal-edge algorithm introduced in [78]. The terminal-edge of a longest edge propagation path of  $t$ ,  $\text{Lepp}(t)$ , is the longest-edge between all the edges involved in  $\text{Lepp}(t)$  including the boundary of the Lepp polygon [73, 75, 77]. We review this method at the end of this Section.

### Partially Coupled Methods

Partially coupled methods resolve inconsistencies during the collision refinement by processing interface edges in 2D (or faces in 3D) using independent sets of elements [45] and by breaking the mesh refinement process into two phases [28, 13, 65]: computation (actual refinement of elements) and communication (exchange of information about the newly created points and elements due to refinement of interfaces).

#### *Parallel Independent Set Method*

In [45] the refinement of a 2D mesh takes place in phases (refinement of one independent set at a time). This guarantees the conformity of the mesh and the elimination of the collision refinement problem, since non-local adjacent elements never refine interface edges concurrently and the processors are always aware of bisections of their interface edges. Specifically, the authors in [45] use a vertex-based partition of a 2D mesh to generate  $P$  submeshes, where  $P$  is the number of processors. Then all non-local adjacent elements (i.e., elements that share an edge) and adjacent vertices to the elements and vertices of submeshes are computed to create a layer of “ghost” mesh entities which are used to minimize communication in the independent set (IS) phase. The distributed memory implementation of the IS phase in [45] computes a distributed independent set  $I = \cup_{i=1}^P I_{M_i}$ , where  $I_{M_i} = I \cap M_i$  and  $M_i$  is a submesh, as follows: a triangle  $t_a \in I_{M_i}$  if:  $\forall t_b \in \text{adj}(t_a)$  and one of the following three holds (1)  $t_a, t_b \in M_i$ , (2)  $\rho(t_a) > \rho(t_b)$ , and (3)  $t_b$  is not a marked triangle for refinement, where  $\text{adj}(t)$  is the set of adjacent triangles of  $t$ , and  $\rho(t)$  is a unique random number assigned to each element in the mesh  $M_i$ ,  $i = 1, P$ . Note that due to the ghost elements, the communication for checking the above conditions is eliminated. The algorithm requires communication only for: (a) the update of the bisections of ghost elements, and (b) a global reduction operation for termination. Both take place at the end of the refinement of an independent set. These two types of communication make the algorithm partially coupled, since experimental data in [45] indicate that the number of refinement phases (or loop iterations) is small (10 to 20) as the number of processors and the size of the mesh increase to 200 processors and a million elements, respectively.

### *Parallel Alternate Bisection Method*

DeCougny et. al [28] addresses the collision refinement problem by using, first, alternate bisection on the interface faces then by applying region subdivision templates on the rest of the tetrahedra. After the mesh faces are subdivided, it is possible to create non-conforming interface edges on the interfaces. The non-conforming interface edges are sent to the corresponding adjacent submeshes that are refined by different processors. This will start a new mesh face subdivision followed by a communication phase, until no mesh faces need to be subdivided. Upon termination of face subdivisions, the mesh is conforming across the interfaces and then a region subdivision using sequential templates is applied in parallel to the rest of the interior tetrahedra.

### *Parallel Nested Elements Method*

Castaños and Savage [13] have parallelized the non-conforming longest edge bisection algorithm both in 2D and 3D. In this case the refinement propagation implies the creation of sequences of non-conforming edges that can cross several submeshes involving several processors. This also means the creation of non-conforming interface edges which is particularly complex to deal with in 3D. To perform this task each processor  $P_i$  iterates between a no-communication phase (where refinement propagation between processors is delayed) and an interprocessor communication phase. Different processors can be in different phases during the refinement process, their termination is coordinated by a central processor  $P_0$ . The subdivision of an interface edge might leads to either a non-conforming edge or to a conforming edge, but the creation of different copies (one per subdomain) of its midpoint. However, after the communication phase a remote cross reference for each newly created interface edge midpoint along with *nested elements* information guarantee a unique logical name for these newly created vertices [11].

## **Decoupled Methods**

The PTE method [74] in addition to the terminal-edge of a Lepp(t) takes full advantage of the terminal-star, which is the set of tetrahedra that share a terminal-edge. The terminal-star can play the same role in PTE the cavity plays in PCDM. Contrary to the method in [13] the terminal-star refinement algorithm completely avoids the management of non-conforming edges both in the interior of the submeshes and in the inter-subdomain interface. This eliminates the communication among subdomains and thus processors. Similarly to Castaños et al. the terminal-star method can terminate using a single processor as coordinator for adaptive mesh refinement i.e., when a global stopping criterion like the minimum-edge length of terminal-edges is not used.

The decoupled PTE algorithm and its implementation lead to an order of magnitude performance improvements compared to a previous tightly coupled implementation [74] of the same algorithm. Although the algorithm is theoretically scalable, our performance data indicate the contrary; the reason is the work-load imbalances and heterogeneity of the clusters we use. We will address these two issues in Section 7.5.

	Mesh Generation Technique		
Coupling	Deluanay	Advancing Front	Edge Bisection
Tight	PODM	PAFT <sub>SM</sub>	PCTE
Partial	PCDT, PDR	POAFT	PIS, PNE
None	P <sup>2</sup> DM, PD <sup>3</sup>	PAFT	PTE

Fig. 7.8. Taxonomy of Parallel Mesh Generation Methods.

## 7.4 Taxonomy


The taxonomy in Figure 7.8 helps to clarify basic similarities and differences between parallel tetrahedral meshing methods. The taxonomy is based on the two attributes we used to classify the methods reviewed in this chapter: (i) the basic sequential meshing technique used for each subproblem and (ii) the degree of coupling between the subproblems. The coupling (i.e., the amount of communication and synchronization between the subproblems) is determined by the degree of dependency between the subproblems.

## 7.5 Implementation

The complexity of implementing efficient parallel mesh generation codes arises from the dynamic, data-dependent, and irregular computation and communication patterns of the algorithms. This inherent complexity, when combined with challenges from using primitive tools for communication like message-passing libraries [86, 4], makes the development of parallel mesh generation codes even more time-consuming and error-prone.

In the rest of the section we focus on dynamic load balancing issue. The scientific computing community has developed application-specific runtime libraries and software systems [3, 33, 49, 64, 67, 98] for dynamic load balancing. These systems are designed to support the development of parallel multi-phase applications which are computationally-intensive and consist of phases that are separated by computations such as the global error estimation. In these cases the load-balancing is accomplished by dynamically repartitioning the data after a global synchronization [95]. Throughout this chapter we call this approach to load balancing the *stop-and-repartition* method.

The stop-and-repartition approaches are good for loosely-synchronous applications like iterative PDE solvers, however they are not well-suited for applications such as adaptive mesh generation and refinement. Because for asynchronous and not



**Fig. 7.9.** Execution time of PAFT on Whirlwind subcluster with 64 homogeneous processors for the simplified human brain model without load balancing (left) and with load balancing using PREMA (right). The final mesh in both cases is 1.2 billion tetrahedrons.


computation-intensive applications the global synchronization overhead can overwhelm the benefits from load balancing. This problem is exacerbated as the number of processors in the parallel system grows. In order to address this issue, the authors in [2] developed a *Parallel Runtime Environment for Multi-computer Applications* (PREMA).

PREMA is a software library which provides a set of tools to application developers via a concise and intuitive interface. It supports single-sided communication primitives which conform to the *active messages* paradigm [94], a global namespace, message forwarding mechanisms to cope with object/data migration and a preemptive dynamic load balancing [2].


### *Performance Evaluation*

In the rest of this section we present some performance data that show the effects of two sources of imbalance: (1) work-load due to geometric complexity of the sub-domains/submeshes, and (2) processor heterogeneity. The experimental study was performed on Sciclone [81] cluster at the College of William and Mary which consists of many different heterogeneous subclusters. We have used three subclusters: (1) Whirlwind subcluster which consists of 64 single-cpu Sun Fire V120 nodes (650 MHz, 1 GB RAM), (2) Tornado which consists of 32 dual-cpu Sun Ultra 60 nodes (360 MHz, 512 MB RAM) and (3) Typhoon which consists of 64 single-cpu Sun Ultra 5 nodes (333 MHz, 256 MB RAM). The models we used are: (i) a cross-section of the rocket pipe (see Figure 7.1) and (ii) a simplified model of a human brain (see Figure 7.11, left).

Figure 7.9 shows the impact of dynamic load balancing on the performance of PAFT on the human brain model. The work-load imbalances are due to differences in the geometric complexity of the submeshes. The PAFT with dynamic load balancing (using PREMA) took 1.7 hours to generate the 1.2 billion elements while without dynamic load balancing it took 2.7 hours. The dynamic load balancing improved the



**Fig. 7.10.** The execution time of PCDM for the cross section of the rocket pipe whose data are equidistributed on 128 heterogenous processors; without load balancing (left) and with load balancing using PREMA (right).



**Fig. 7.11.** Surface of the tetrahedral mesh for a simplified model of a human brain generated from an advancing front method [43].

performance of PAFT by more than 30%. Sequentially using Solidmesh [37], it takes three days one hour and 27 minutes, by executing the subdomains one at a time.

Figure 7.10 shows the impact of dynamic load balancing (using PREMA) on the performance of the PCDM for the cross section of rocket pipe. Although we used state-of-the-art ab-initio data partition methods for equidistributing the data and thus the computation among all 128 processors, the imbalances are due to heterogeneity of the three different clusters; the first 64 processors are from Typhoon (slowest cluster), the next 32 processors are from Tornado and the last 32 processors are from Whirlwind (the fastest cluster). Again, the dynamic load balancing (using PREMA) improved the performance of parallel mesh generation by 23%.

Finally, the data from Figure 7.11 and Table 7.2 indicate the impact of work load imbalances due to: (1) the differences in the work-load of submeshes and (2) heterogeneity of processors using the PTE method. Figure 7.11(right), shows that the speed of the PTE method is substantially lower, for the brain model (see Figure 7.11, left), due to work-load imbalances; while for a more regular geometry (the semiconductor test case [74]), the PTE speed is almost twice higher, because of better load balancing

**Table 7.2.** PTE speed (in tetrahedra per second) for the simplified human brain model using min-edge = 2.0. The final mesh is about 2.4 million tetrahedra.

Processors	8	16	32	48	64	96
Whirlwind	5427	9920	16195	21890	29035	23571
Tornado	3852	7025	11408	15526	20312	23571

due to more uniform point distribution. Also, Table 7.2 indicates a 19% slowdown in the PTE's speed once we increase the number of processors from 64 to 96 using additional 32 slower processors, despite the fact the PTE is a scalable method. Finally, a comparison between the speed data from the Figure 7.11 (right) and Table 7.2, for the brain model, indicate that the coupling (i.e., global synchronization) in the PCTE method slows down the speed of the code by an order of magnitude.

These data (and some more from [24, 2, 17, 74]) suggest that the tightly coupling methods should be used as a last resort. In addition, these data suggest that work-load imbalances are no longer a problem and it should not limit our creativity in the second round of our search for practical and effective parallel mesh generation methods. Runtime software systems like PREMA [2] can handle work-load imbalances quite successfully.

## 7.6 Future Directions

It takes about ten to fifteen years to develop the algorithmic and software infrastructure for *sequential industrial strength mesh generation libraries*. Moreover, improvements in terms of quality, speed, and functionality are open ended and permanent which makes the task of delivering state-of-the-art parallel mesh generation codes even more difficult.

This survey demonstrates that without compromising in the stability of parallel mesh generation methods it is possible for all three mesh generation classes of techniques to develop parallel meshing software using off-the-shelf sequential meshing codes.

An area with immediate high benefits to parallel mesh generation is domain decomposition. The DD problem as it is posed in Section 7.2 is still open for 3D geometries and its solution will help to deliver stable and scalable methods that rely on off-the-shelf mesh generation codes for Delaunay and Advancing Front Techniques. The edge subdivision methods are independent off the domain decomposition.

A longer term goal should be the development of both theoretical and software frameworks like PDR to implement new mesh generation methods which can: (1) take advantage of multicore architectures with more than two hardware contexts for the next generation of high-end workstations and (2) scale without any substantial implementation costs for clusters of high-end workstations.

Finally, a long term investment to parallel mesh generation is to attract the attention of mathematicians with open problems in mesh generation and broader impact in

mathematics. For example, develop theoretical frameworks able to prove the correctness of single threaded guaranteed quality Delaunay theory in the context of partial order [70].

## Acknowledgment

This work is supported by the IBM professorship at Brown University, the Alumni Memorial Endowment at the College of William and Mary, and National Science Foundation (NSF) grants: CNS-0312980, ANI-0203974, ACI-0085969, Career Award CCR-0049086, EIA-9972853, EIA-9726388. The experimental work was performed using computational facilities at the College of William and Mary which were enabled by grants from Sun Microsystems, the NSF, and Virginia's Commonwealth Technology Research Fund. Andrey Chernikov made many insightful recommendations and Leonidas Linardakis made useful comments for the Delaunay methods. Most of the figures and tables are the result of my collaboration with my students from the College of William and Mary Kevin Barker, Andrey Chernikov, Andriy Fedorov, Brian Holinka, Leonidas Linardakis, from Notre Dame Demian Nave, and from Universidad de Chile Carlo Calderon and Daniel Pizarro. The anonymous referee for his many useful comments.

## References

1. S. Baden, N. Chrisochoides, D. Gannon, and M. Norman, editors. *Structured Adaptive Mesh Refinement Grid Methods*. Springer-Verlag, 1999.
2. K. Barker, A. Chernikov, N. Chrisochoides, and K. Pingali. A load balancing framework for adaptive and asynchronous applications. *IEEE Transactions on Parallel and Distributed Systems*, 15(2):183–192, Feb. 2004.
3. A. Basermann, J. Clinckemaillie, T. Coupez, J. Fingberg, H. Dignonnet, R. Ducloux, J. Gratien, U. Hartmann, G. Lonsdale, B. Maerten, D. Roose, and C. Walshaw. Dynamic load-balancing of finite element applications with the drama library. *Applied Mathematical Modeling*, 25:83–98, 2000.
4. A. Belguelin, J. Dongarra, A. Geist, R. Manchek, S. Otto, and J. Walpore. PVM: Experiences, current status, and future direction. In *Supercomputing '93 Proceedings*, pp. 765–766, 1993.
5. D. Bertsekas and J. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice Hall, 1989.
6. T. B.H.V. and B. Cheng. Parallel adaptive quadrilateral mesh generation. *Computers and Structures*, 73:519–536, 1999.
7. G. E. Blelloch, J. Hardwick, G. L. Miller, and D. Talmor. Design and implementation of a practical parallel Delaunay algorithm. *Algorithmica*, 24:243–269, 1999.
8. G. E. Blelloch, G. L. Miller, and D. Talmor. Developing a practical projection-based parallel Delaunay algorithm. In *12th Annual Symposium on Computational Geometry*, pp. 186–195, 1996.
9. H. Blum. A transformation for extracting new descriptors of shape. In *Models for the Perception of speech and Visual Form*, pp. 362–380. MIT Press, 1967.

10. A. Bowyer. Computing Dirichlet tessellations. *Computer Journal*, 24:162–166, 1981.
11. J. G. Castaños and J. E. Savage. The dynamic adaptation of parallel mesh-based computation. In *SIAM 7th Symposium on Parallel and Scientific Computation*, 1997.
12. J. G. Castaños and J. E. Savage. Parallel refinement of unstructured meshes. In *Proceedings of the IASTED, International Conference of Parallel and Distributed Computing and Systems*, 1999.
13. J. G. Castaños and J. E. Savage. PARED: a framework for the adaptive solution of PDEs. In *8th IEEE Symposium on High Performance Distributed Computing*, 1999.
14. M.-B. Chen, T. R. Chuang, and J.-J. Wu. Efficient parallel implementations of near Delaunay triangulation with high performance Fortran. *Concurrency: Practice and Experience*, 16(12), 2004.
15. A. N. Chernikov and N. P. Chrisochoides. Parallel guaranteed quality planar Delaunay mesh generation by concurrent point insertion. In *14th Annual Fall Workshop on Computational Geometry*, pp. 55–56. MIT, Nov. 2004.
16. A. N. Chernikov and N. P. Chrisochoides. Practical and efficient point insertion scheduling method for parallel guaranteed quality Delaunay refinement. In *Proceedings of the 18th annual international conference on Supercomputing*, pp. 48–57. ACM Press, 2004.
17. A. N. Chernikov, N. P. Chrisochoides, and L. P. Chew. Design of a parallel constrained Delaunay meshing algorithm, 2005.
18. L. P. Chew. Constrained Delaunay triangulations. *Algorithmica*, 4(1):97–108, 1989.
19. L. P. Chew, N. Chrisochoides, and F. Sukup. Parallel constrained Delaunay meshing. In *ASME/ASCE/SES Summer Meeting, Special Symposium on Trends in Unstructured Mesh Generation*, pp. 89–96, Northwestern University, Evanston, IL, 1997.
20. N. Chrisochoides. An alternative to data-mapping for parallel iterative PDE solvers: Parallel grid generation. In *Scalable Parallel Libraries Conference*, pp. 36–44. IEEE, 1993.
21. N. Chrisochoides. Multithreaded model for load balancing parallel adaptive computations. *Applied Numerical Mathematics*, 6:1–17, 1996.
22. N. Chrisochoides, C. Houstis, E.N.Houstis, P. Papachiou, S. Kortesis, and J. Rice. Domain decomposer: A software tool for partitioning and allocation of PDE computations based on geometry decomposition strategies. In *4th International Symposium on Domain Decomposition Methods*, pp. 341–357. SIAM, 1991.
23. N. Chrisochoides, E. Houstis, and J. Rice. Mapping algorithms and software environment for data parallel PDE iterative solvers. *Special issue of the Journal of Parallel and Distributed Computing on Data-Parallel Algorithms and Programming*, 21(1):75–95, 1994.
24. N. Chrisochoides and D. Nave. Parallel Delaunay mesh generation kernel. *Int. J. Numer. Meth. Engng.*, 58:161–176, 2003.
25. N. Chrisochoides and F. Sukup. Task parallel implementation of the Bowyer-Watson algorithm. In *Proceedings of Fifth International Conference on Numerical Grid Generation in Computational Fluid Dynamics and Related Fields*, 1996.
26. N. P. Chrisochoides. A new approach to parallel mesh generation and partitioning problems. *Computational Science, Mathematics and Software*, pp. 335–359, 2002.
27. P. Cignoni, D. Laforenza, C. Montani, R. Perego, and R. Scopigno. Evaluation of parallelization strategies for an incremental Delaunay triangulator in  $E^3$ . *Concurrency: Practice and Experience*, 7(1):61–80, 1995.
28. H. D. Cougny and M. Shephard. Parallel refinement and coarsening of tetrahedral meshes. *Int. J. Meth. Eng.*, 46(1101-1125), 1999.
29. T. Culver. *Computing the Medial Axis of a Polyhedron Reliably and Efficiently*. PhD thesis, The University of North Carolina at Chapel Hill, 2000.
30. H. de Cougny and M. Shephard. *CRC Handbook of Grid Generation*, chapter Parallel unstructured grid generation, pp. 24.1–24.18. CRC Press, Inc., 1999.



31. H. de Cougny and M. Shephard. Parallel volume meshing using face removals and hierarchical repartitioning. *Comp. Meth. Appl. Mech. Engng.*, 174(3-4):275–298, 1999.
32. H. L. de Cougny, M. S. Shephard, and C. Ozturan. Parallel three-dimensional mesh generation on distributed memory mimd computers. *Engineering with Computers*, 12:94–106, 1995.
33. K. Devine, B. Hendrickson, E. Boman, M. S. John, and C. Vaughan. Design of dynamic load-balancing tools for parallel applications. In *Proc. of the Int. Conf. on Supercomputing*, Santa Fe, May 2000.
34. E. W. Dijkstra and C. Sholten. Termination detection for diffusing computations. *Inf. Proc. Lettres*, 11, 1980.
35. H. Edelsbrunner and D. Guoy. Sink-insertion for mesh improvement. In *Proceedings of the Seventeenth Annual Symposium on Computational Geometry*, pp. 115–123. ACM Press, 2001.
36. P. J. Frey and P. L. George. *Mesh Generation: Applications to Finite Element*. Hermis; Paris, 2000.
37. J. Gaither, D. Marcum, and B. Mitchell. Solidmesh: A solid modeling approach to unstructured grid generation. In *7th International Conference on Numerical Grid Generation in Computational Field Simulations*, 2000.
38. J. Galtier and P. L. George. Prepartitioning as a way to mesh subdomains in parallel. In *Special Symposium on Trends in Unstructured Mesh Generation*, pp. 107–122. ASME/ASCE/SES, 1997.
39. P. L. George and H. Borouchaki. *Delaunay Triangulation and Meshing: Applications to Finite Element*. Hermis; Paris, 1998.
40. H. N. Gürsoy. *Shape interrogation by medial axis transform for automated analysis*. PhD thesis, Massachusetts Institute of Technology, 1989.
41. J. C. Hardwick. Implementation and evaluation of an efficient 2D parallel Delaunay triangulation algorithm. In *Proceedings of the 9th Annual ACM Symposium on Parallel Algorithms and Architectures*, 1997.
42. B. Hendrickson and R. Leland. The Chaco user's guide, version 2.0. Technical Report SAND94-2692, Sandia National Laboratories., 1994.
43. A. M. S. Ito, Yasushi and B. K. Soni. Reliable isotropic tetrahedral mesh generation based on an advancing front method. In *Proceedings 13th International Meshing Roundtable, Williamsburg, VA, Sandia National Laboratories*, pp. 95–106, 2004.
44. D. A. Jefferson. Virtual time. In *ACM Transactions on Programming Languages and Systems*, volume 7, pp. 404–425, July 1985.
45. M. T. Jones and P. E. Plassmann. Parallel algorithms for the adaptive refinement and partitioning of unstructured meshes. In *Proceedings of the Scalable High-Performance Computing Conference*, 1994.
46. C. Kadow. Adaptive dynamic projection-based partitioning for parallel Delaunay mesh generation algorithms. In *SIAM Workshop on Combinatorial Scientific Computing*, Feb. 2004.
47. C. Kadow and N. Walkington. Design of a projection-based parallel Delaunay mesh generation and refinement algorithm. In *Fourth Symposium on Trends in Unstructured Mesh Generation*, July 2003. <http://www.andrew.cmu.edu/user/sowen/usnccm03/agenda.html>.
48. C. M. Kadow. *Parallel Delaunay Refinement Mesh Generation*. PhD thesis, Carnegie Mellon University, May 2004.
49. S. Kohn and S. Baden. Parallel software abstractions for structured adaptive mesh methods. *Journal of Par. and Dist. Comp.*, 61(6):713–736, 2001.

50. R. Konuru, J. Casas, R. Prouty, S. Oto, and J. Walpore. A user level process package for pvm. In *Proceedings of Scalable High-Performance Computing Conferene*, pp. 48–55. IEEE, 1997.
51. L. Linardakis and N. Chrisochoides. Parallel Delaunay domain decoupling method for non-uniform mesh generation. *SIAM Journal on Scientific Computing*, 2005.
52. L. Linardakis and N. Chrisochoides. Parallel domain decoupling Delaunay method. *SIAM Journal on Scientific Computing*, in print, accepted Nov. 2004.
53. L. Linardakis and N. Chrisochoides. Medial axis domain decomposition method. *ACM Trans. Math. Software*, To be submitted, 2005.
54. R. Lober, T. Tautges, and R. Cairncross. The parallelization of an advancing-front, all-quadrilateral meshing algorithm for adaptive analysis. In *4th International Meshing Roundtable*, pp. 59–70, October 1995.
55. R. Löhner, J. Camberos, and M. Marshal. *Unstructured Scientific Computation on Scalable Multiprocessors* (Eds. Piyush Mehrotra and Joel Saltz), chapter Parallel Unstructured Grid Generation, pp. 31–64. MIT Press, 1990.
56. R. Löhner and J. R. Cebral. Parallel advancing front grid generation. In *Proceedings of the Eighth International Meshing Roundtable*, pp. 67–74, 1999.
57. F. Lori, M. Jones, and P. Plassmann. An efficient parallel algorithm for mesh smoothing. In *Proceedings 4th International Meshing Roundtable*, pp. 47–58, 1995.
58. R. P. M. Saxena. Parallel FEM algorithm based on recursive spatial decomposition. *Computers and Structures*, 45(9-6):817–831, 1992.
59. S. N. Muthukrishnan, P. S. Shiakolos, R. V. Nambiar, and K. L. Lawrence. Simple algorithm for adaptative refinement of three-dimensional finite element tetrahedral meshes. *AIAA Journal*, 33:928–932, 1995.
60. D. Nave, N. Chrisochoides, and L. P. Chew. Guaranteed: quality parallel Delaunay refinement for restricted polyhedral domains. In *SCG '02: Proceedings of the eighteenth annual symposium on Computational geometry*, pp. 135–144. ACM Press, 2002.
61. D. Nave, N. Chrisochoides, and L. P. Chew. Guaranteed-quality parallel Delaunay refinement for restricted polyhedral domains. *Computational Geometry: Theory and Applications*, 28:191–215, 2004.
62. T. Okusanya and J. Peraire. Parallel unstructured mesh generation. In *5th International Conference on Numerical Grid Generation on Computational Field Simulations*, pp. 719–729, April 1996.
63. T. Okusanya and J. Peraire. 3D parallel unstructured mesh generation. In S. A. Canann and S. Saigal, editors, *Trends in Unstructured Mesh Generation*, pp. 109–116, 1997.
64. L. Oliker and R. Biswas. Plum: Parallel load balancing for adaptive unstructured meshes. *Journal of Par. and Dist. Comp.*, 52(2):150–177, 1998.
65. L. Oliker, R. Biswas, and H. Gabow. Parallel tetrahedral mesh adaptation with dynamic load balancing. *Parallel Computing Journal*, pp. 1583–1608, 2000.
66. S. Owen. A survey of unstructured mesh generation. Technical report, ANSYS Inc., 2000.
67. M. Parashar and J. Browne. Dagh: A data-management infrastructure for parallel adaptive mesh refinement techniques. Technical report, Dept. of Comp. Sci., Univ. of Texas at Austin, 1995.
68. N. Patrikalakis and H. Gürsoy. Shape interrogation by medial axis transform. In *Design Automation Conference (ASME)*, pp. 77–88, 1990.
69. P. Pebay and D. Thompson. Parallel mesh refinement without communication. In *Proceedings of International Meshing Roundtable*, pp. 437–443, 2004.
70. S. Prassidis and N. Chrisochoides. A categorical approach for parallel Delaunay mesh generation, July 2004.

71. M. C. Rivara. Algorithms for refining triangular grids suitable for adaptive and multigrid techniques. *International Journal for Numerical Methods in Engineering*, 20:745–756, 1984.
72. M. C. Rivara. Selective refinement/derefinement algorithms for sequences of nested triangulations. *International Journal for Numerical Methods in Engineering*, 28:2889–2906, 1989.
73. M. C. Rivara. New longest-edge algorithms for the refinement and/or improvement of unstructured triangulations. *International Journal for Numerical Methods in Engineering*, 40:3313–3324, 1997.
74. M.-C. Rivara, C. Calderon, D. Pizarro, A. Fedorov, and N. Chrisochoides. Parallel decoupled terminal-edge bisection algorithm for 3D meshes. *(Invited) Engineering with Computers*, 2005.
75. M. C. Rivara, N. Hitschfeld, and R. B. Simpson. Terminal edges Delaunay (small angle based) algorithm for the quality triangulation problem. *Computer-Aided Design*, 33:263–277, 2001.
76. M. C. Rivara and C. Levin. A 3D refinement algorithm for adaptive and multigrid techniques. *Communications in Applied Numerical Methods*, 8:281–290, 1992.
77. M. C. Rivara and M. Palma. New lepp algorithms for quality polygon and volume triangulation: Implementation issues and practical behavior. In *In Trends unstructured mesh generation Eds: S. A. Cannan . Saigal, AMD*, volume 220, pp. 1–8, 1997.
78. M.-C. Rivara, D. Pizarro, and N. Chrisochoides. Parallel refinement of tetrahedral meshes using terminal-edge bisection algorithm. In *13th International Meshing Roundtable*, Sept. 2004.
79. R. Said, N. Weatherill, K. Morgan, and N. Verhoeven. Distributed parallel Delaunay mesh generation. *Computer Methods in Applied Mechanics and Engineering*, (177):109–125, 1999.
80. K. Schloegel, G. Karypis, and V. Kumar. Parallel multilevel diffusion schemes for repartitioning of adaptive meshes. Technical Report 97-014, Univ. of Minnesota, 1997.
81. Sciclone cluster project. Last accessed, March 2005. <http://www.compsci.wm.edu/SciClone/>.
82. M. Shephard and M. Georges. Automatic three-dimensional mesh generation by the finite octree technique. *International Journal for Numerical Methods in Engineering*, 32:709–749, 1991.
83. E. C. Sherbrooke. *3-D shape interrogation by medial axial transform*. PhD thesis, Massachusetts Institute of Technology, 1995.
84. J. Shewchuk. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. In *Proceedings of the First workshop on Applied Computational Geometry*, pp. 123–133, Philadelphia, PA, 1996.
85. J. R. Shewchuk. *Delaunay Refinement Mesh Generation*. PhD thesis, Carnegie Mellon University, 1997.
86. M. Snir, S. Otto, S. Huss-Lederman, and D. Walker. MPI the complete reference. MIT Press, 1996.
87. A. Sohn and H. Simon. Jove: A dynamic load balancing framework for adaptive computations on an SP-2 distributed memory multiprocessor, 1994. Technical Report 94-60, Dept. of Comp. and Inf. Sci., New Jersey Institute of Technology, 1994.
88. D. A. Spielman, S.-H. Teng, and A. Üngör. Parallel Delaunay refinement: Algorithms and analyses. In *Proceedings of the Eleventh International Meshing Roundtable*, pp. 205–217, 2001.

89. D. A. Spielman, S.-H. Teng, and A. Üngör. Time complexity of practical parallel Steiner point insertion algorithms. In *Proceedings of the sixteenth annual ACM symposium on Parallelism in algorithms and architectures*, pp. 267–268. ACM Press, 2004.
90. M. Stuti and A. Moitra. Considerations of computational optimality in parallel algorithms for grid generation. In *5th International Conference on Numerical Grid Generation in Computational Field Simulations*, pp. 753–762, 1996.
91. T. Tam, M. Price, C. Armstrong, and R. McKeag. Computing the critical points on the medial axis of a planar object using a Delaunay point triangulation algorithm.
92. Y. A. Teng, F. Sullivan, I. Beichl, and E. Puppo. A data-parallel algorithm for three-dimensional Delaunay triangulation and its implementation. In *SuperComputing*, pp. 112–121. ACM, 1993.
93. J. F. Thompson, B. K. Soni, and N. P. Weatherill. *Handbook of Grid Generation*. CRC Press, 1999.
94. T. von Eicken, D. Culler, S. Goldstein, and K. Schauer. Active messages: A mechanism for integrated communication and computation. In *Proceedings of the 19th Int. Symp. on Comp. Arch.*, pp. 256–266. ACM Press, May 1992.
95. C. Walshaw, M. Cross, and M. Everett. Parallel dynamic graph partitioning for adaptive unstructured meshes. *Journal of Par. and Dist. Comp.*, 47:102–108, 1997.
96. D. F. Watson. Computing the n-dimensional Delaunay tessellation with application to Voronoi polytopes. *Computer Journal*, 24:167–172, 1981.
97. R. Williams. *Adaptive parallel meshes with complex geometry*. Numerical Grid Generation in Computational Fluid Dynamics and Related Fields, 1991.
98. X. Yuan, C. Salisbury, D. Balsara, and R. Melhem. Load balancing package on distributed memory systems and its application particle-particle and particle-mesh (P3M) methods. *Parallel Computing*, 23(10):1525–1544, 1997.