

Multi-threaded mesh generation

J.-F. Remacle¹, Célestin Marot¹

¹ Université catholique de Louvain, Institute of Mechanics, Materials and Civil Engineering, [jean-francois.remacle, celestin.marot]@uclouvain.be

Résumé — This paper presents a fine grain parallel version of the 3D Mesh generation procedure using the OpenMP (Open Multi-Processing) API on standard NUMA multicore processors and on the new Intel Knight's Landing architecture. The most significant piece of algorithm of the meshing procedure is the Delaunay kernel. A set $S = \{p_1, \dots, p_n\}$ of n points is taken as input. S is initially sorted along a space-filling curve so that two points that are close in the insertion order are also close geometrically. A parallel radix sort is implemented on massively multithreaded architectures that allow to sort about half a billion points in about a tenth of a second. The sorted set of point is then divided into M subsets S_i , $1 \leq i \leq M$ of equal size n/M . The multithreaded version of the Delaunay kernel inserts M points at a time in the triangulation. OpenMP barriers provide the required synchronization that is needed after each multiple insertion in order to avoid data races. This simple approach exhibits two standard problems of parallel computing : load imbalance and parallel overheads. Those two issues are addressed using a two-level version of the multithreaded Delaunay kernel. Tests show that triangulations of about a billion tetrahedra can be generated on a 32 core machine (Intel Xeon E5-4610 v2 @ 2.30GHz with with 128 GB of memory) in less that 3 minutes of wall clock time, with a speedup of 18 compared to the single-threaded implementation. Then, a parallel point insertion algorithm is presented that allows to generate very large meshes one order of magnitude faster than state-of-the art implementations.

Mots clés — Mesh Generation, Multi-threading, KNL.

1 Introduction

In the last decades, the size of the finite element meshes that are used in industry for scientific computing has grown considerably due to the availability of massively parallel computers. It is nowadays not uncommon to generate meshes that have over 100 million tetrahedra. From a user's perspective, however, generating a mesh of a complex domain usually involves the generation of several intermediary meshes that are progressively enhanced in order to fulfill some adequate design requirements. Today's best 3D meshing algorithms are able to generate about five million tetrahedra per minute on one single core [1]. Each iteration in the meshing process thus takes long minutes and users eventually spend a significant portion of their time waiting for the mesh generator to provide outputs.

Computation platforms are increasingly based on multicore architectures with a large common shared memory but relatively low performance individual computing cores. The performance of today's mesh generation procedures, based on serial kernels, is thus largely stalling, and may actually degrade in the future.

Parallel mesh generation procedures have been developed for several years. Up to ten years ago, the great majority of parallel meshing algorithms worked at a coarse grain level (see [2] for an exhaustive survey). In such coarse grain algorithms, the 3D domain is first partitioned into subdomains ; the internal 2D boundaries between the subdomains are then meshed in order to ensure the compatibility of the 3D meshes in adjacent subdomains ; and serial meshing procedures are finally applied at the subdomain level [3, 4, 5]. The focus of these parallel algorithms is mostly about mesh size and not so much about fast meshing : indeed, generating very large meshes cannot be done on one single CPU and it is thus very important to distribute the meshing generation process in order to have access to sufficient memory. Multicore platforms offer the advantage of low latency communications through shared memory, leading to the possibility of fine-grain parallelism in the meshing kernels. Recently, CGAL [6] developers have proposed a fine-grained multithreaded version of the Delaunay triangulation in 3D based on locks [6]. Different locking strategies have been analyzed, with typical speedups of 5 with 8 computing cores, but

at the cost of a fairly complex algorithmic implementation.

This paper is essentially a preliminary work in multithreading the meshing process of Gmsh [7]. The overall aim is similar to CGAL's, i.e., to increase the speed of the central piece of the mesh generator through a fine-grained parallelization of the Delaunay kernel. The main difference is that we focus here on the simplicity of the implementation, using the OpenMP (Open Multi-Processing) Application Programming Interface (API).

2 The Delaunay Kernel

In what follows, a triangle is the generic term for a triangle in 2D or a tetrahedron in 3D. A triangulation $T(S)$ of S is a set of non overlapping triangles that exactly covers the convex hull $\Omega(S)$ with all points of S being among the vertices of the triangulation. The Delaunay triangulation $DT(S)$ is such that the empty circumcircle of any triangle in $DT(S)$ is empty, i.e., it contains no point of S . Delaunay triangulations are popular in the meshing community because fast algorithms exist that allow to generate $DT(S)$ in $O(n \log n)$ complexity.

The fastest 3D algorithms that allow to build $DT(S)$ are based on the Delaunay kernel [8]. Let DT_k be the Delaunay triangulation of a point set $S_k = \{p_1, \dots, p_k\} \subset \mathbb{R}^d$. The *Delaunay kernel* is a procedure that allow the incremental insertion of a given point $p_{k+1} \in \Omega(S_k)$ into DT_k and to build the Delaunay triangulation DT_{k+1} of $S_{k+1} = \{p_1, \dots, p_k, p_{k+1}\}$. The *Delaunay kernel* can be written in the following abstract manner :

$$DT_{k+1} = DT_k - \mathcal{C}(DT_k, p_{k+1}) + \mathcal{B}(DT_k, p_{k+1}), \quad (1)$$

where the Delaunay cavity $\mathcal{C}(DT_k, p_{k+1})$ is the set of all triangles whose circumcircles contain the new point p_{k+1} (see Figure 1; the triangles of the cavity cannot belong to DT_{k+1}) and the Delaunay ball $\mathcal{B}(DT_k, p_{k+1})$ is a set of triangles that fill the polyhedral hole that has been left empty while removing the Delaunay cavity $\mathcal{C}(DT_k, p_{k+1})$ from DT_k .

The most important building block in any implementation of the Delaunay kernel is the computation of the simply connected Delaunay cavity $\mathcal{C}(DT_k, p_{k+1})$ [9]. One seed triangle t should be found that has its circumcircle containing p_{k+1} . Then, computing the Delaunay cavity can be done locally using a depth-first search technique. There are essentially two ways to compute the seed triangle t efficiently. The history of all triangles may be maintained and a search is made into the history dag to insert a new point [10]. This solution leads to logarithmic complexity for finding the seeding triangle t . It also requires some memory overhead because all triangles of all stages have to be maintained. A more straightforward algorithm consists in doing a “walk” into the triangulation : starting from any triangle τ , finding the next




FIGURE 1 – Delaunay triangulation DT_k (left), Delaunay cavity $\mathcal{C}_p(DT_k, p_{k+1})$ (center) and $DT_{k+1} = DT_k - \mathcal{C}(DT_k, p_{k+1}) + \mathcal{B}(DT_k, p_{k+1})$ (right).




FIGURE 2 – Delaunay cavity $\mathcal{C}(\text{DT}_k, p_{k+1})$ in red and Delaunay cavity $\mathcal{C}(\text{DT}_{k-1}, p_k)$ in yellow. A seed t for building $\mathcal{C}(\text{DT}_k, p_{k+1})$ is found by starting from an arbitrary triangle of $\mathcal{C}(\text{DT}_{k-1}, p_k)$ and walking through the triangulation. Here, 32 walking steps were necessary to find t .

n	10^3	10^4	10^5	10^6	10^3	10^4	10^5	10^6
	2D (random insertion)				3D (random insertion)			
N_{walk}	23	73	230	727	17	38	85	186
$t(\text{sec})$	$3.6 \cdot 10^{-3}$	$9.1 \cdot 10^{-2}$	3.98	187	$1.2 \cdot 10^{-2}$	$1.8 \cdot 10^{-1}$	3.42	73
	2D (Hilbert curve)				3D (Hilbert curve)			
N_{walk}	2.3	2.4	2.5	2.5	2.9	3.0	3.1	3.1
$t(\text{sec})$	$2 \cdot 10^{-3}$	$1.5 \cdot 10^{-2}$	$1.5 \cdot 10^{-1}$	1.47	$9.0 \cdot 10^{-3}$	$7.5 \cdot 10^{-2}$	$7.8 \cdot 10^{-1}$	7.81

TABLE 1 – Results of the Delaunay Triangulation algorithm applied to a set of random points uniformly distributed in $[0, 1]^d$. The table compares timings and average number of local searches in the case where points are inserted in a random fashion and in the case where points are inserted along a Hilbert curve.

triangle in the path to t consists in choosing one of the three neighbors τ_j of τ in such a way that p_{k+1} is on the other side of τ of the edge that is common to τ and τ_j . This walk can be shown to always terminate using Edelsbrunner’s acyclic theorem [11]. Figure 2 presents the construction of a Delaunay cavity using this approach. The walking path that allows to find the seed t is depicted with a thick black line. The Delaunay cavity of point p_{k+1} is shown in red. In our implementation, we choose τ as the first element of the Delaunay ball $\mathcal{B}(\text{DT}_{k-1}, p_k)$.

The top half of Table 1 presents the average number of walking steps N_{walk} required to find an initial triangle of the Delaunay cavity in terms of the number of points to insert n , when the points are inserted randomly. In 2D N_{walk} is of the order $O(n^{1/2})$ as expected. It is interesting to see that the CPU time for computing the triangulation is asymptotically lower in 3D because N_{walk} is of the order $O(n^{1/3})$.

This behaviour can be improved dramatically by sorting the points in such a way that two successive points in the set are close to each other geometrically. On Figure 3, a set of 10^5 points are sorted using a Hilbert curve (two successive points in the sorted list are linked with a solid line). In the context of the Delaunay kernel, this kind of data locality can decrease the number of local searches N_{walk} that are required to find the next invalid triangle : one can use as a starting triangle t for point p_{k+1} one of the triangles of cavity $\mathcal{C}(\text{DT}_{k-1}, p_k)$. In practice, points are initially separated in groups of increasing sizes. Then, points are sorted within each group using a Hilbert space-filling curve. This has the advantage to produce smaller Delaunay cavities during the incremental insertion process than a simple Hilbert sort. Adding some randomness in the insertion process allows to reduce the average cavity size. Here, we use the Biased Randomized Insertion Order (BRIO) approach [12] that allows to recover enough randomness in the process.

Figure 4 shows Hilbert sorting time for a random set of $5 \cdot 10^7$ points in 3D using a radii sort (linear complexity). The implementation has been done both on a high end NUMA multicore machine and on a new Knight’s landing architecture. It is interesting to see that, even using one single core, the new KNL beats a high end core I7 if vectorization is used wisely.




FIGURE 3 – Hilbert sort of a sets of 10^5 random points.

The bottom half of Table 1 presents the average number of walking steps N_{walk} as well as timings in the case where points are inserted along a Hilbert curve. The number of walks N_{walk} is now almost independent of n and both 2D and 3D point insertion algorithms have an overall linear complexity. The difference between 2D and 3D timings is essentially due to the size of Delaunay cavities : 4 on average in 2D and above 20 in 3D. Assuming that geometrical predicates are slightly more expensive to compute in 3D, a factor of about 6 between the 2D and 3D timings is to be expected. Note that the Hilbert sorting has a $O(n \log n)$ complexity : the overall complexity of the Delaunay triangulation is $O(n \log n)$ as well. Yet, this is an asymptotic bound : with 1 million points, the time required to sort the points is still typically 100 times lower than the time for generating the mesh.

Before proceeding to the description of the proposed fine-grained parallelization of the Delaunay kernel, let us note that its serial performance is comparable to state of the art algorithms. For the same point set and for the same machine (same compiler and compiler options), Tetgen 1.5 [13] takes 7.7 seconds to tetrahedralize the set of points while ours takes 7.81 sec (a difference of less than 2%). The source code of both the 2D and the 3D version of the algorithm is available on Gmsh's website www.gmsh.info.

3 A Multithreaded Delaunay Kernel

Assume M computational threads that aim at inserting M points in the triangulation at the same time. At the end, each thread is going to insert $\frac{n}{M}$ points and our hope is of course to obtain a speedup close to M . The situation is of course not that simple : two points p_i and p_j can only be inserted at the same time in DT_k if their respective Delaunay cavities $\mathcal{C}(DT_k, p_i)$ and $\mathcal{C}(DT_k, p_j)$ do not overlap, i.e., if they do not have triangles in common :

$$\mathcal{C}(DT_k, p_i) \cap \mathcal{C}(DT_k, p_j) = \emptyset.$$

A non-overlapping situation is more likely to happen if points p_i and p_j are not close geometrically. For that purpose, we split the Hilbert curve into M equal parts and assign each part to one thread. Threads process their assigned points in order. A first chunk of points that correspond to about $20 \times M$ points is inserted in a serial fashion at first [6]. This allows to avoid inevitable cavities overlap in the first stages of the algorithm. Then, the rest of the points is inserted in parallel. Figure 5 shows different stages of the algorithm. Delaunay cavities are far apart most of the time thanks to the property of the Hilbert curve.

The *multithreaded Delaunay kernel* can be written in the following abstract manner :

$$DT_{k+1} = DT_k + \sum_{i=0}^{M-1} \left[-\mathcal{C}(DT_k, p_{k+i\frac{n}{M}}) + \mathcal{B}(DT_k, p_{k+i\frac{n}{M}}) \right]. \quad (2)$$




FIGURE 4 – Hilbert sort of a sets of $5 \cdot 10^7$ 3D random points.

We have implemented the multithreaded Delaunay kernel using OpenMP [14]. Two OpenMP barriers were used at each iteration k . A first barrier is used after the computation of the M cavities : every thread i has to complete its cavity $\mathcal{C}(\text{DT}_k, p_{k+i\frac{n}{M}})$ at iteration k in order to be able to verify that the cavity does not overlap other cavities. When several cavities overlap, only the point corresponding to the smallest thread number is processed. The other points are delayed to the next iteration. A second barrier is used after the construction of $\mathcal{B}(\text{DT}_k, p_{k+i\frac{n}{M}})$: every thread has to finish computing the Delaunay kernel in order to start iteration $k + 1$ with a valid mesh.

Three potential threats can definitively harm the parallel speedup :

1. **Load balancing** : at stage k , the size of the M Delaunay cavities $\mathcal{C}(\text{DT}_k, p_{k+i\frac{n}{M}})$ may vary from one thread to another, leading to some load imbalance. Threads with small cavities will wait at the barrier for the thread that has the largest cavity.
2. **Overlaps** : it should be verified that cavity overlaps are rare events : each overlap adds one iteration in the insertion process.
3. **Overheads** : The OpenMP parallelization adds its own overhead : two barriers are used per iteration.

4 A Two-Level Multithreaded Delaunay Kernel

A two-level strategy is a partial solution to both issues of cavity size discrepancy and barrier overheads. At iteration k , we assume that each thread i inserts M_2 points at a time in a serial manner : M_2 Delaunay cavities are computed by thread i which leads to a Delaunay kernel that inserts $M \times M_2$ points at each iteration k . Figure 6 shows how M_2 affects the distribution of sizes of 3D Delaunay cavities for a set of on million points that are uniformly distributed on the unit cube. The distribution gets sharper around the average cavity size, which is definitively advantageous for scalability.

The *two-level multithreaded Delaunay kernel* can be written in the following abstract manner :

$$\text{DT}_{k+1} = \text{DT}_k + \sum_{i=0}^{M-1} \sum_{j=0}^{M_2-1} \left[-\mathcal{C}(\text{DT}_k, p_{k+i\frac{n}{MM_2}+j\frac{n}{M_2}}) + \mathcal{B}(\text{DT}_k, p_{k+i\frac{n}{MM_2}+j\frac{n}{M_2}}) \right]. \quad (3)$$

This procedure generates exactly the same insertion sequence as the multithreaded Delaunay kernel with $M \times M_2$ threads.




FIGURE 5 – Different stages of the multithreaded Delaunay insertion ($M = 4$). Delaunay cavities at the four different stages are colored w.r.t. their thread number.

Figure 7 reports strong scaling results for the two-scale multithreaded Delaunay kernel for $M_2 = 1, 2, 4, 8$. Maximal theoretical speedups S_{\max} are reported as well as true speedups S . The computer that has been used in is the same 4 socket node with a 8-cores processor on each socket for a total of 32 cores. The 8 first threads were executed within the scope of one single processor, allowing to maintain the nearness of threads and their data. Figure 7 shows that speedups increase with M_2 as predicted, and that for $M \leq 8$ the actual speedup of our procedure is close to the maximal possible one. A speedup of $S = 5.8$ is attained for $M = 8$ and $M_2 = 8$, which is slightly better than the ones observed in [6] and which is close to $S_{\max} = 6.7$. For a higher number of cores, speedups are moving away from their maximal values even though we made sure that thread affinity was maintained by allocating vertices and tetrahedra on the local node [15].

Yet, the good news is that speedups always increase with M : a triangulation of one million points in 3D was performed in 0.89 seconds of wall clock time on 32 cores for a total speedup of $S = 13.2$. This is one order of magnitude faster than the fastest procedure available.

Figure 8 gives scaling results for different sizes for the point set : it starts from $n = 10^4$ and increases that number to $n = 1.5 \cdot 10^8$ points, i.e., to about 10^9 tetrahedra. With $M > 8$ speedups are even more important for very large sets of points. This is essentially due to the fact that data locality is better on 32 cores than using one single node for such a large amount of points.

In our implementation, one tetrahedron requires 72 bytes of memory. Less than 100 GigaBytes of memory were necessary for 10^9 tetrahedra and it took 142.8 seconds of wall clock time for generating such a mesh on $M = 32$ cores with $M_2 = 8$. This corresponds to a rate of about 7 million of tetrahedra per second.




FIGURE 6 – Average size of 3D Delaunay cavities for $M_2 = 1$ to $M_2 = 16$.

Références

- [1] Paul Louis George, Frédéric Hecht, and E Saltel. Automatic mesh generator with specified boundary. *Computer methods in applied mechanics and engineering*, 92(3) :269–288, 1991.
- [2] Nikos P Chrisochoides et al. A survey of parallel mesh generation methods. *Brown University, Providence RI-2005*, 2005.
- [3] Nikos Chrisochoides and Démian Nave. Parallel delaunay mesh generation kernel. *International Journal for Numerical Methods in Engineering*, 58(2) :161–176, 2003.
- [4] Rainald Löhner, Jose Camberos, and Marshal Merriam. Parallel unstructured grid generation. *Computer Methods in Applied Mechanics and Engineering*, 95(3) :343–357, 1992.
- [5] Hugues L De Cougny, Mark S Shephard, and Can Ozturan. Parallel three-dimensional mesh generation. *Computing Systems in Engineering*, 5(4) :311–323, 1994.
- [6] Vicente HF Batista, David L Millman, Sylvain Pion, and Johannes Singler. Parallel geometric algorithms for multi-core computers. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 217–226. ACM, 2009.
- [7] C. Geuzaine and J.-F. Remacle. Gmsh : a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11) :1309–1331, 2009.
- [8] Pascal Jean Frey and Paul L George. *Mesh generation : application to finite elements*. ISTE London, 2008.
- [9] Martin Isenburg, Yuanxin Liu, Jonathan Shewchuk, and Jack Snoeyink. Streaming computation of delaunay triangulations. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 1049–1056. ACM, 2006.
- [10] C Bradford Barber and Hannu Huhdanpaa. Qhull, softwarepackage, 1995.
- [11] Herbert Edelsbrunner. *Geometry and topology for mesh generation*. Cambridge University Press, 2001.
- [12] Nina Amenta, Sunghye Choi, and Günter Rote. Incremental constructions con brio. In *Proceedings of the nineteenth annual symposium on Computational geometry*, pages 211–219. ACM, 2003.
- [13] Hang Si. Tetgen, a delaunay-based quality tetrahedral mesh generator. *ACM Transactions on Mathematical Software (TOMS)*, 41(2) :11, 2015.
- [14] Leonardo Dagum and Ramesh Menon. Openmp : an industry standard api for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1) :46–55, 1998.
- [15] Andi Kleen. A numa api for linux. *Novel Inc*, 2005.

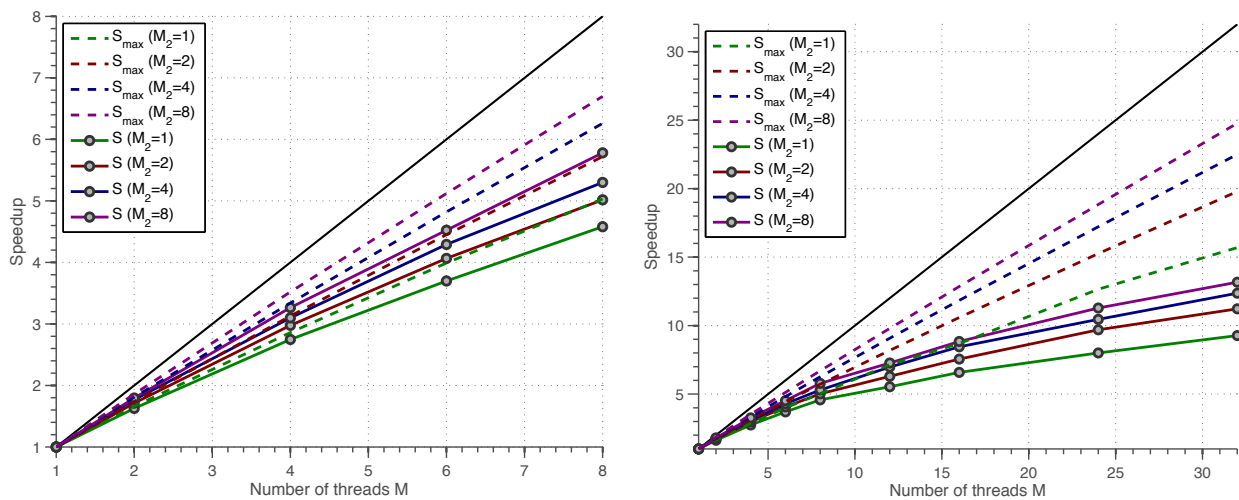


FIGURE 7 – Strong scaling of the two-level multithreaded 3D Delaunay kernel. Left figure is a zoom ($1 \leq M \leq 8$) of the right Figure ($1 \leq M \leq 32$).




FIGURE 8 – Scaling of the two-level multithreaded Delaunay kernel with respect to N .