

An Improved Parallel Algorithm for Delaunay Triangulation on Distributed Memory Parallel Computers

Sangyoon Lee Chan-Ik Park Chan-Mo Park
Dept. of Computer Science and Engineering/PIRL
POSTECH
San 31 HyoJa Dong
Pohang, 790-784 Korea
Email: cipark@vision.postech.ac.kr
Fax: +82-562-279-5699

Abstract

Delaunay triangulation has been much used in such applications as volume rendering, shape representation, terrain modeling and so on. The main disadvantage of Delaunay triangulation is large computation time required to obtain the triangulation on an input points set. This time can be reduced by using more than one processors, and several parallel algorithms for Delaunay triangulation have been proposed. In this paper, we propose an improved parallel algorithm for Delaunay triangulation, which partitions the bounding convex region of the input points set into a number of regions by using Delaunay edges and generates Delaunay triangles in each region by applying an incremental construction approach. Partitioning by Delaunay edges makes it possible to eliminate merging step required for integrating subresults. It is shown from the experiments that the proposed algorithm has good load balance and is more efficient than Cignoni et al.'s algorithm[8] and our previous algorithm[9].

1. Introduction

A triangulation for a given set of points is a well known topic of computational geometry [11, 12] and has many applications such as finite element analysis, solid modeling, shape representation, terrain modeling, volume rendering and computer vision [4, 5, 6]. Delaunay triangulation (DT) has been most attractive for triangulation due to its special feature that the circumscribed circle of a triangle must not contain other points than points of the triangle.

There have been much research in DT. The time complexity of constructing Delaunay triangles for a set of n points is $\Omega(n \log n)$ under uniprocessor environment[12].

Reducing the time complexity is made possible by using more than one processors [1, 3, 8, 9, 13]. Saxena et al. [13] have proposed two parallel algorithms on the orthogonal tree network of $n \times n$ processors; one is for 2-dimensional DT with the time complexity of $O(\log^2 n)$ and the other is for 3-dimensional DT with the time complexity of $O(m^{1/2} \log n)$, where n represents the number of points and m represents the number of tetrahedra constructed in case of 3-dimensional DT. Even though their algorithms are quite efficient, they require special architecture.

In general distributed memory parallel computers, Cignoni et al. have proposed two parallel algorithms in [8]; one (called 'ParDeWall') from a sequential algorithm (called 'Delaunay Wall[7]') based on divide & conquer paradigm and the other (called 'ParInCoDe') from a sequential algorithm (called 'Incremental Construction of Delaunay Triangulation[7]') based on an incremental construction method. ParDeWall algorithm exploits parallelism found when a set of points is recursively partitioned by some cutting planes and each subregion containing only one cutting plane is assigned to one processor. Each processor generates a set of triangles called simplex wall whose edges are crossing its corresponding cutting plane and then partitions its subregion by a newly generated cutting plane. Though this algorithm is simple, the degree of parallelism is very limited since subproblems are generated in the form of binary tree and parallelism comes from subproblems of the same level. We cannot obtain good scalability of it due to small degree of parallelism. The parallelism exploited in ParInCoDe algorithm comes from geometric partitioning of a set of points. Geometric partitioning means that the bounding box containing all points can be divided by sub-boxes (subregions) and each subregion is assigned to one processor. Each processor builds all Delaunay triangles possible for any points inside its assigned region. Then, we get

the final set of Delaunay triangles by simple merging partial results generated by each processor. However, this algorithm suffers from overlapping processing because each processor generates many identical triangles. And their method causes load imbalance among processors and, even worse, it does not work in a certain case. In addition, this type of algorithm like ParInCoDe requires good load balance among processors in order to get nontrivial speedup. But, the load balancing issue is completely ignored in [8].

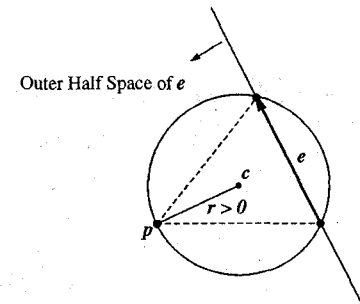
Another parallel algorithm has been proposed by Guy E. Blelloch et al. in [3]. The algorithm is based on projections and a divide & conquer paradigm. Each subproblem is determined by a region which is represented by a border consisting of Delaunay edges and the set of points inside or on the border. At each call the region is divided into two subregions by using a median line cut of the internal points and a corresponding path of Delaunay edges which is obtained through projections and a lower convex hull. Once the subproblem has no more internal points (all points of the region is on the border), Delaunay triangles in the region are generated by the special end-game strategy. This parallel algorithm is easy to implement and is of practical use. But the degree of parallelism is very limited for the same reason in ParDeWall algorithm. Moreover, the convex hull of points must be given in addition to points as inputs.

We have proposed a parallel algorithm in [9] which achieves nontrivial speedup by way of good load balance and elimination of overlapping processing in ParInCoDe algorithm. The main pitfall of our previous algorithm is that it requires a merging scheme on partial results and the merging scheme is executed sequentially on only one processor, resulting in limited speedup. In addition, it is required that the first Delaunay triangle generated by each processor must exist entirely in its own subregion.

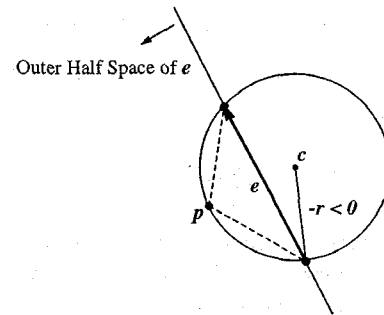
In this paper, we propose an improved parallel algorithm which eliminates all pitfalls of our previous algorithm proposed in [9]. The proposed algorithm uses the partitioning method which is based on projections, and generates Delaunay triangles in the subregion by applying an incremental construction approach.

2. Delaunay triangulation and the incremental construction

Given a finite set S of points in the plane, a triangulation[12] is defined as joining the points of S by nonintersecting straight line segments so that every region internal to the convex hull is a triangle. In addition, a triangulation is called Delaunay if it satisfies the empty circumcircle property: the circumcircle of a triangle in the triangulation does not contain any points of S in its interior[12]. The concept of a triangulation can be extended to E^d space. For instance, the triangulation in E^3 is to obtain tetrahedra



(a) $dd(e,p) = r$



(b) $dd(e,p) = -r$

Figure 1. Delaunay distance($dd(e,p)$) where e is a Delaunay edge and p is a point.

in the convex hull of points. In this paper, we consider Delaunay triangulation in E^2 space (a plane).

The duality between Delaunay triangulations and Voronoi diagrams is well known[11, 12], and thus algorithms are given for the construction of the former from the latter. However, it is generally more efficient to directly construct the triangulation, and in fact the construction time is $O(n)$ [12].

Incremental construction method was originally proposed by McLain[10] to construct Delaunay triangulation, which directly use the empty circumcircle property of the triangulation.

Our proposed algorithm uses a variant of McLain's original method, which was proposed by Cignoni et al. in [7]. Given a Delaunay edge e , the method constructs a new Delaunay triangle by selecting the point which minimizes Delaunay distance defined as $dd(e,p)$ for each point p :

$$dd(e,p) = \begin{cases} r & \text{if } c \in \text{the outer half space of } e \\ -r & \text{otherwise} \end{cases} \quad (1)$$

where r and c are the radius and the center of the circumcircle around e and p , respectively (see Figure 1).

The efficiency of the incremental construction method is low ($O(n^2)$ in the worst case), but the method is practically

good by using some speedup techniques such as hashing and the uniform grid[2].

3. Partitioning methods

The proposed algorithm partitions a convex bounding region of given points into subregions according to paths consisting of Delaunay edges. A path which partitions subregions is constructed by a projection-based method on Delaunay edges[3].

We consider two partitioning methods using the paths. One divides a region into two subregions according to a vertical path, and then each subregion divides according to a horizontal path. This is recursively continued until the number of subregions equals to that of processors. The other divides a region into m subregions according to just vertical (or, horizontal) paths where m is the number of processors.

3.1. The projection-based construction of a path consisting of Delaunay edges

Given a set P of points in E^2 space, a vertical (or horizontal) path which consists only of Delaunay edges and passes through a point $p \in P$ and divides the convex hull of P into two regions, can be constructed by the following method[3]:

1. Let L be the line $x = p_x$ (or $y = p_y$) where $p = (p_x, p_y)$ (see Figure 2(a)).
2. Project all points of P on a 3-D paraboloid centered at p (see Figure 2(b)), and then project them on the vertical plane (or the horizontal plane) along the line L (see Figure 2(c)). The set P' of projected points is:

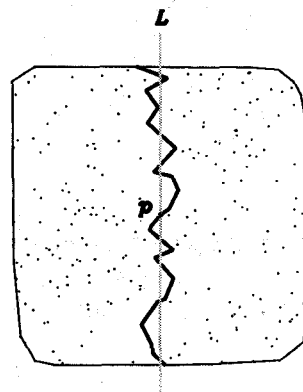
$$\{(q_y - p_y, \|q - p\|^2) \mid q \in P\}$$
 (or $\{(q_x - p_x, \|q - p\|^2) \mid q \in P\}$).
3. Construct the lower convex hull of P' . The set of edges in P corresponding to this lower convex hull is a path consisting of Delaunay edges.

From now on, we assume that all paths are constructed by this method, thereby satisfying all these properties.

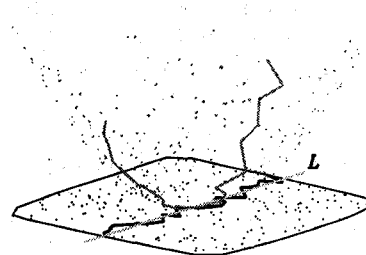
3.2. The partitioning method 1

Given a set P of points, the convex bounding region of P is divided into 2^k subregions, where the number of processors is 2^k for an integer k , by the following method:

1. Find the point q that is the median along the x (or y) axis of points of the region.
2. Construct a vertical (or horizontal) path H passing q . Since a region is specified by a set of directed paths,



(a) A path of Delaunay edges through a point $p \in P$



(b) Points projected on the paraboloid



(c) The lower convex hull of points ($\in P'$) projected on the vertical plane

Figure 2. A path of Delaunay edges constructed by projections and a lower convex hull

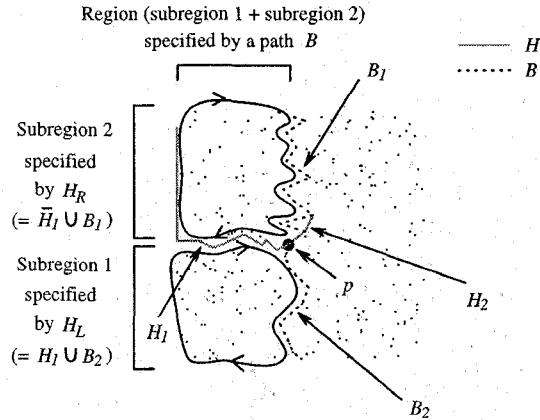


Figure 3. An example of the merged paths at the 2nd level. $H_1 \cup B_2$ and $\bar{H}_1 \cup B_1$ are the merged paths H_L and H_R for the subregion 1 and 2, respectively.

we assume that H has its own direction. Without loss of generality, assume that a vertical path has a direction from top to bottom and a horizontal path has a direction from left to right. Let \bar{H} be the path traversed in the opposite direction of H .

3. Divide the region into two subregions by the path H :

- (a) One subregion includes all the points on the left (or down) side of H as well as the points on H . We can specify this subregion by the path H_L obtained from H and the path representing the given region.
- (b) The other subregion includes all the points on the right (or up) side of H as well as the points on H . This subregion is specified by the path H_R obtained from \bar{H} and the path representing the given region.

4. Recursively repeat 1-3 for each subregion vertically or horizontally in turn until the number of subregions becomes 2^k .

Checking the geometric relationships between Delaunay edges of H and those of the current border is required when we construct merged paths H_L and H_R for a given H because a portion of the path H may be outside of the current region. For example of Figure 3, consider a region specified by a path B . The region is partitioned by a path H and we can see that a portion of the path H , H_2 , is outside of the region, where H and B intersect at a point p and are divided by the point into H_1 , H_2 , B_1 , and B_2 , respectively. This portion of the path H is not needed to specify a subregion,

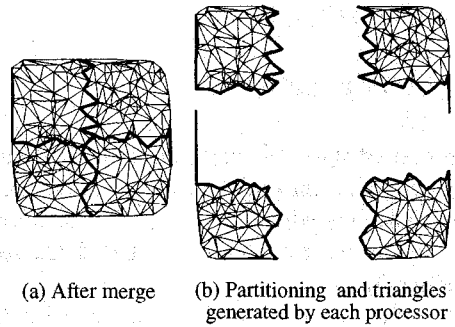


Figure 4. An example of the partitioning using the method 1 for 4 processors

i.e., subregion 1 can be specified by $H_L (= H_1 \cup B_2)$ where H_2 is not included.

Figure 4 shows an example of a partitioning result using this method when there are 4 ($= 2^2$) processors. Note that the construction of the merged paths H_L and H_R is not needed at the first level, since the convex bounding region of P is completely divided into two subregions by the path H .

3.3. The partitioning method 2

Given a set P of n points, the convex bounding region of P is divided into m subregions, where m is the number of processors, by the following method:

1. Find the $(m - 1)$ points q_i that are the $\lceil \frac{ni}{m} \rceil$ th largest points along the x (or y) axis of all points of P ($i = 1 \dots (m - 1)$).
2. Construct vertical (or horizontal) paths H_i passing q_i . Let \bar{H}_i be the path traversed in the opposite direction of H_i .
3. The 1st partition is the subregion on the left (or down) side of the path H_1 , which is represented by H_1 and the points on the left (or down) side of or on H_1 .
4. The m th partition is the subregion on the right (or up) side of the path H_{m-1} , which is represented by \bar{H}_{m-1} and the points on the right (or up) side of or on H_{m-1} .
5. The other partitions are the subregions between the paths \bar{H}_{i-1} and H_i , which are represented by the merged path H_m^i of the two paths and includes all the points between or on the paths.

The construction of H_m^i is more simpler than that of the merged path H_L or H_R in the partitioning method 1. H_m^i is

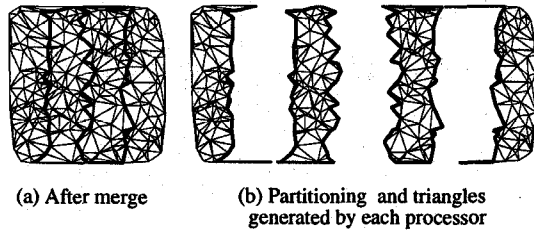


Figure 5. An example of the partitioning using the method 2 for 4 processors

constructed by just selecting the Delaunay edges that exist on either the paths \bar{H}_{i-1} or H_i , but not both.

Figure 5 shows an example of a partitioning result using this method when there are 4 ($= 2^2$) processors.

4. The proposed algorithm

Given a set P of n points and the number of processors 2^k for an integer k , the proposed algorithm consists of three steps:

Step 1: Partition the convex bounding region of P into 2^k subregions by either the partitioning method 1 or 2 described in section 3 and assign each subregion to disjoint processor.

Step 2: Each processor executes the procedure *ParDeTri* of Figure 6 for its assigned subregion.

Step 3: Simply merge *Partial_DT_list* generated by each processor into one set of triangles, called *DT_list*.

In Step 1, the partitioning is performed on each processor, and each processor just constructs its own subregion which is represented by the points and a border of Delaunay edges. The border, called *DE_list*, is used for the input of the procedure *ParDeTri* in Step 2.

The function *MakeTriangle* in *ParDeTri* of Figure 6 generates a Delaunay triangle from a Delaunay edge by selecting the point which minimizes the function *dd* for each point in the subregion, described in Section 2.

The procedure *ParDeTri* differs from the routines of Cignoni et al.'s *ParInCoDe*[8] and our previous algorithm[9] in that it does not have to specially generate the first triangle for the incremental construction since it directly uses *DE_list* for the construction, and it does not check if or not the edge exists in its own subregion. Thus it is released from the condition that the first triangle must be entirely in its own subregion. Note that this condition is essential in Cignoni et al.'s *ParInCoDe* and our previous algorithm.

```

Procedure ParDeTri( $P$  : pointset,  $DE\_list$  : edge_list, var
 $Partial\_DT\_list$  : triangle_list);
var
   $e$ : edge;
   $t$ : triangle;
begin
  while notempty( $DE\_list$ ) do begin
     $e := Extract(DE\_list)$ ;
     $t := MakeTriangle(e,P)$ ;
    if  $t \neq null$  then begin
      Insert( $t, Partial\_DT\_list$ );
      for each  $e'$ :  $e' \in edges(t)$  AND  $e' \neq e$  do
        Update( $e', DE\_list$ );
    end;
  end;
end;

Procedure Update( $e$ :edge, var  $L$ :edge_list);
begin
  if  $e \in L$  then Delete( $e, L$ )
  else Insert( $e, L$ );
end;

```

Figure 6. The procedure ParDeTri executed on each processor.

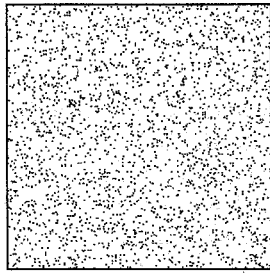
5. Experimental results

The algorithms are implemented on the INMOS TRAM networks which have 32 T800 processors. The input point sets for the experiment are generated by using various distributions: the uniform distribution, the normal distribution, the bubble distribution and the uniform distribution with a much narrower width, which are shown in Figure 7(a), 7(b), 7(c) and 7(d), respectively. In the experiment, the partitioning method 2 divides a region into subregions according to horizontal paths.

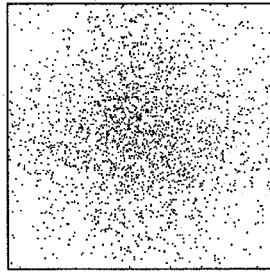
The speedups of the algorithms on the INMOS TRAM networks are shown in Figure 8. As compared with these, the proposed algorithm with the partitioning method 2 is always more efficient than the others except the last case. In the last case, the proposed algorithm with the partitioning method 1 is the best of the algorithms. This means that the proposed algorithm is more efficient than Cignoni et al.'s algorithm and our previous algorithm in the experiment.

Figure 8(d) shows a problem of our previous algorithm that the algorithm have poor speedup with a number of processors since it should execute the sequential merging scheme on one processor (in this case, the algorithm have the more speedup with 16 processors than 32 processors).

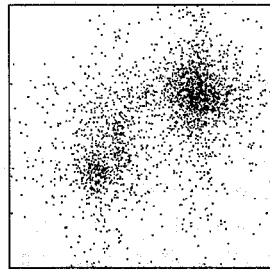
Figure 9 shows the breakdown of the execution time into



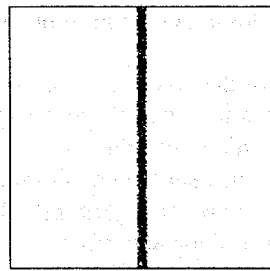
(a) The uniform distribution



(b) The normal distribution

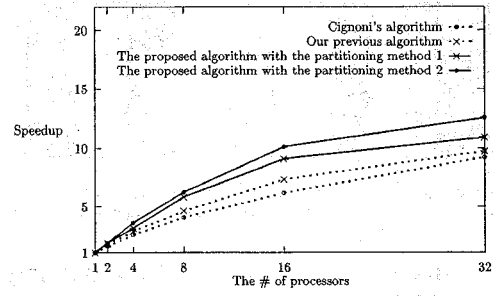


(c) The bubble distribution

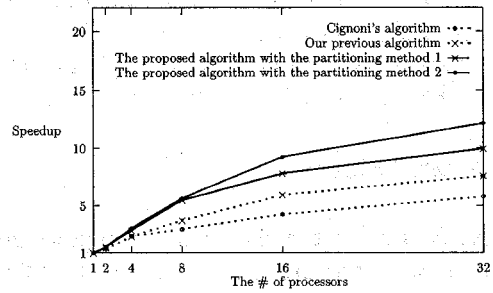


(d) The uniform distribution (width:height = 3:100)

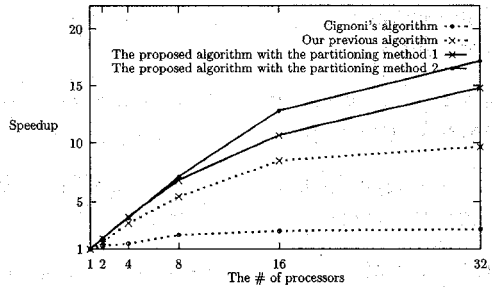
Figure 7. The distributions for the experiment. The number of points are 10000.



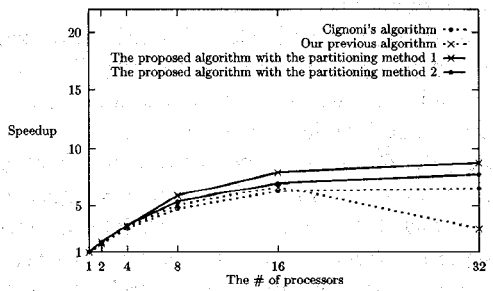
(a) The uniform distribution



(b) The normal distribution



(c) The bubble distribution



(d) The uniform distribution (width:height = 3:100)

Figure 8. The speedup graph for the distributions.

Distribution	Standard deviation (average)
Uniform	0.81 (8.29) secs
Normal	0.96 (9.86) secs
Bubble	1.32 (10.25) secs
Uniform (width:height=3:100)	1.62 (15.19) secs

(a) With the partitioning method 1

Distribution	Standard deviation (average)
Uniform	1.13 (7.51) secs
Normal	1.38 (7.92) secs
Bubble	1.54 (8.47) secs
Uniform (width:height=3:100)	1.69 (16.95) secs

(b) With the partitioning method 2

Table 1. The standard deviations of running times obtained in executing the proposed algorithm with 32 processors.

the two components of the proposed algorithm (the partitioning time and the execution time of *ParDeTri*) for the distributions. As the figure shows, the execution time of *ParDeTri* is approximately the same across the partitioning methods and the distributions. For the partitioning time, the algorithm using the partitioning method 2 requires more time than the other in the uniform (width:height = 3:100) distribution, and less time in the other distributions. It is because the partitioning method 2 uses just horizontal paths and thus the work of the partitioning method 2 to get lower convex hulls of the projected points is more than that of the partitioning method 1 for the distribution. This explains the reason why the algorithm using the method 1 is faster than the algorithm using the method 2 in Figure 8(d).

In order to show that the proposed algorithm has a good load balance, we calculate the standard deviation of execution times on processors for each case of the input sets. On executing the proposed algorithm on 32 processors, the standard deviations and averages of times for the distributions are shown in Table 1. Particularly, Figure 10 shows the load of each processor for the bubble distribution.

6. Conclusions

In this paper, we described an improved parallel algorithm for Delaunay triangulation that eliminates the sequential merging step executed on one processor. The sequen-

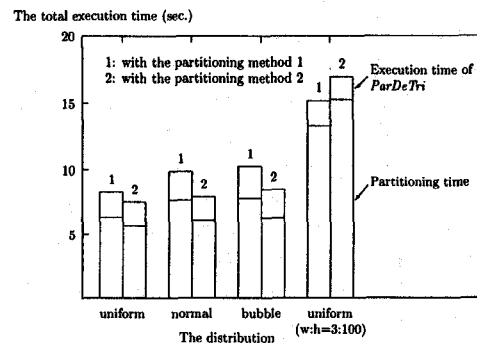


Figure 9. The breakdown of the execution time into the two components of the proposed algorithm for the distributions.

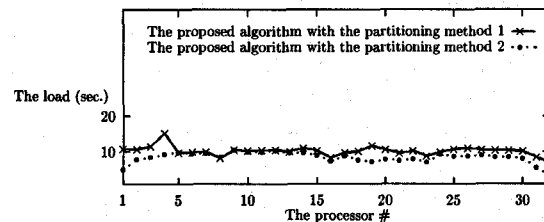


Figure 10. The load of each processor for the bubble distribution.

tial merging step significantly limits the speed up. The algorithm is also released from the restriction that the first Delaunay triangle generated by each processor must exist entirely in its own subregion. All these improvements are possible by using projection-based partitioning methods.

We consider Delaunay triangulation of points in just E^2 space. However, it is remained as our future work to extend the proposed algorithm to E^d spaces for $d > 2$.

References

- [1] S. G. Akl and K. A. Lyons, *Parallel Computational Geometry*, Prentice-Hall, 1993.
- [2] V. Akman, W. R. Franklin, M. Kankanhalli and C. Narayanaswami, Geometric Computing and Uniform Grid Technique, *Computer-Aided Design*, Vol. 21, No. 7, pp. 410-420, September 1989.
- [3] Guy E. Blelloch, Gary L. Miller and Dafna Talmor, Developing a Practical Projection-Bases Parallel Delaunay Algorithm, *ACM Symposium on Computational Geometry*, May 1996.

- [4] J. D. Boissonnat, Geometric Structures for Three Dimensional Shape Representation, *ACM Transactions on Graphics*, Vol. 3, pp. 266-286, Oct. 1984.
- [5] J. C. Cavendish, Automatic Triangulation of Arbitrary Planar Domains for Finite Element Method, *International Journal for Numerical Methods in Engineering*, Vol. 8, pp. 679-696, 1974.
- [6] H. J. Choi, An Implementation of a Flight Path Visualization System Using Optimization Algorithms of the Triangulated Irregular Network, *MS. Thesis*, POSTECH, 1996.
- [7] P. Cignoni, C. Montani and R. Scopigno, A Merge-First Divide & Conquer Algorithm for E^d triangulations, Technical Report 92-16, Istituto CNUCE-C.N.R., Pisa, ITALY, October 1992.
- [8] P. Cignoni, C. Montani, R. Perego and R. Scopigno, Parallel 3D Delaunay Triangulation, *Eurographic 93*, 1993.
- [9] Sangyoon Lee, Chan-Ik Park and Chan-Mo Park, An Efficient Parallel Algorithm for Delaunay Triangulation on Distributed Memory Parallel Computers, *The 1996 International Conference on Parallel and Distributed Processing Techniques and Applications*, August 1996.
- [10] D. H. McLain, Two Dimensional Interpolation from Random Data, *The Computer Journal*, Vol. 19, No. 2, pp. 178-181, 1976.
- [11] J. O'Rourke, *Computational Geometry in C*, Cambridge University Press, 1994.
- [12] F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*, Springer-Verlag, New York, 1985.
- [13] S. Saxena, P. C. P. Bhatt and V. C. Prasad, Efficient VLSI Parallel Algorithm for Delaunay Triangulation on Orthogonal Tree Network in Two and Three Dimensions, *IEEE Transactions on Computers*, Vol. C-39, No. 3, pp. 400-404, March 1990.