



The Voronoi Diagram of Weakly Smooth Planar Point Sets in $O(\log n)$ Deterministic Rounds on the Congested Clique

Jesper Jansson¹, Christos Levcopoulos^{2(✉)}, Andrzej Lingas², and Quan Xue³

¹ Graduate School of Informatics, Kyoto University, Kyoto, Japan
jj@i.kyoto-u.ac.jp

² Department of Computer Science, Lund University, Lund, Sweden
{Christos.Levcopoulos, Andrzej.Lingas}@cs.lth.se

³ The University of Hong Kong, Hong Kong, China
csxuequan@connect.hku.hk

Abstract. We study the problem of computing the Voronoi diagram of a set of n^2 points with $O(\log n)$ -bit coordinates in the Euclidean plane in a substantially sublinear in n number of rounds in the congested clique model with n nodes. Recently, Jansson et al. have shown that if the points are uniformly at random distributed in a unit square then their Voronoi diagram within the square can be computed in $O(1)$ rounds with high probability (w.h.p.). We show that if a very weak smoothness condition is satisfied by an input set of n^2 points with $O(\log n)$ -bit coordinates in the unit square then the Voronoi diagram of the point set within the unit square can be computed in $O(\log n)$ rounds in this model.

Keywords: Voronoi diagram · Delaunay triangulation · convex hull · distributed algorithm · the congested clique model

1 Introduction

The *congested clique* is a relatively new model of communication and computation introduced by Lotker *et al.* in 2005 [9]. It focuses on the cost of communication between the nodes in a network, ignoring the cost of local computation within each node. Hence, it can be seen as opposite to the Parallel Random Access Machine (PRAM) model, studied extensively in the 80s and 90s. The PRAM model focuses on the computation cost and ignores the communication cost [1].

Originally, the complexity of dense graph problems has been studied in the congested clique model under the following assumptions. Each node of the congested clique represents a distinct vertex of the input graph and knows its neighborhood in the graph. Every node also knows the unique ID numbers (between 1 and n) of itself and all the other nodes at the start of the computation. The computation proceeds in rounds. In each round, each of the n nodes can send a distinct message of $O(\log n)$ bits to each other node and can perform unlimited local computation. The primary complexity objective is to minimize the number of rounds necessary to solve a given problem on the input graph in this model.

For several basic graph problems such as the minimum spanning tree problem, one has even succeeded to design $O(1)$ -round protocols in the congested clique model [11, 14]. Observe that when the input graph is of bounded degree and the edge weights have $O(\log n)$ -bit representations, each node can send the ID numbers of all nodes in its neighborhood and the weights of its incident edges, e.g., to the first node in $O(1)$ rounds. After that, the first node can solve the whole problem locally. However, such a trivial solution would require $\Omega(n)$ rounds when the input graph is dense.

Matrix problems [3], sorting and routing [7], and geometric problems [6] have also been studied in the congested clique model. In all cases, the basic input items, i.e., matrix entries or keys, or points in the plane, respectively, are assumed to have $O(\log n)$ -bit representations and each node initially has a batch of n such items. *Note that the bound on bit representation of an input item is a natural consequence of the $O(\log n)$ -bit bound on the size of a single message which makes input items of unbounded bit representation incompatible with the assumed model.* As in the graph case, in every round, each node can send a distinct $O(\log n)$ -bit message to each other node and perform unlimited local computation. Significantly, it has been shown that matrix multiplication can be performed in a number of rounds substantially sublinear in n [3] while sorting and routing can be implemented in $O(1)$ rounds (Theorems 4.5 and 3.7 in [7]).

As for the geometric problems, Jansson et al. [6] recently provided low polylogarithmic, deterministic upper bounds on the number of rounds required to solve several basic geometric problems for a set of n^2 points in the plane with $O(\log n)$ -bit coordinates in the model of congested clique with n nodes. As for the construction of the Voronoi diagram and the dual Delaunay triangulation of the point set (see Fig. 1 for an illustration and Sect. 2 for the formal definition), they have shown an $O(1)$ upper bound on the number of required rounds under the assumption that the points are drawn uniformly at random from a unit square. On the other hand, already at the end of the 90s, Goodrich presented $O(1)$ -round randomized protocols for the construction of the three-dimensional convex hull of a set of points in three-dimensional Euclidean space in $O(1)$ communication rounds in the so-called Bulk Synchronous Processing model (BSP) [5]. His result also implies an $O(1)$ -round bound on the randomized construction of the Voronoi diagram and the dual Delaunay triangulation of a planar point set in the BSP model. By using the $O(1)$ -round routing protocol of Lenzen [7], Goodrich's $O(1)$ bound on the number of rounds necessary for the construction of the Voronoi diagram and Delaunay triangulation can be carried over from the BSP model to ours.

In this context, the major open problem is to derive a non-trivial upper bound on the number of rounds sufficient to deterministically construct the Voronoi diagram when the points are not necessarily drawn uniformly at random. The bottleneck in the design of efficient parallel or distributed algorithms for the Voronoi diagram of a planar point set using a direct divide-and-conquer approach is an efficient parallel or distributed merging of Voronoi diagrams. Aggarwal et al. [1] presented an $O(\log^2 n)$ -time CREW PRAM algorithm for the Voronoi diagram based on an involved $O(\log n)$ -time PRAM method for the parallel merging. Subsequently, Amato and Preparata [2] demonstrated an $O(\log n)$ -time CREW PRAM algorithm for the three-dimensional convex hull and consequently also for the two-dimensional Voronoi diagram of a point set.

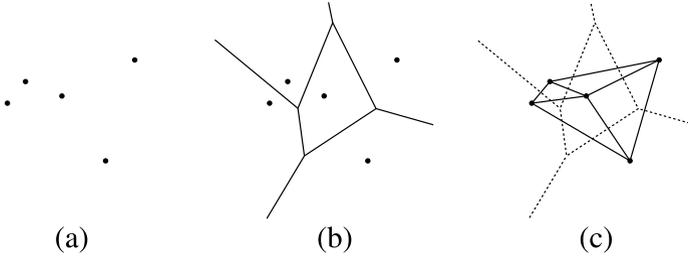


Fig. 1. An example of a planar point set, its Voronoi diagram, and the dual Delaunay triangulation.

We substantially extend the local approach to the construction of the Voronoi diagram used in the design of parallel and distributed algorithms for the Voronoi diagram of points drawn uniformly at random, e.g., from a unit square [6, 8, 15]. We show that already a very weak smoothness condition on the input set of n^2 points with $O(\log n)$ -bit coordinates within a unit square is sufficient to obtain an $O(\log n)$ upper bound on the number of rounds required to construct the Voronoi diagram of the set within the unit square on the congested n -clique. Roughly, our weak smoothness condition says that if a square Q of side length ℓ within the unit square contains at least n out of the n^2 input points then any square of the same size at distance at most $4\sqrt{2}\ell$ from Q and within the unit square has to contain at least one input point.

We obtain our result by combining a quadtree partition of the unit square with a local construction of the Voronoi diagram within the squares corresponding to the leaves of quadtree. The local construction is possible due to the fulfillment of the weak smoothness condition.

In order to simplify the presentation, we assume throughout the paper that the points in the input point sets are in general position (i.e., neither any three input points are co-linear nor any four input points are co-circular).

Our paper is structured as follows. The next section contains basic mathematical/geometric definitions, a lemma, and facts on routing and sorting in the congested clique model. Section 3 presents our protocol for the Voronoi diagram and Delaunay triangulation of a weakly smooth planar point set within a square. In Sect. 4, we conclude with final remarks.

2 Preliminaries

The cardinality of a set S is denoted by $|S|$.

For a positive integer r , $[r]$ stands for the set of positive integers not exceeding r .

For a finite set S of points in the Euclidean plane, the *Voronoi diagram* of S is the partition of the plane into $|S|$ regions such that each region consists of all points in the plane having the same closest point in S ; see Fig. 1.

A *Delaunay triangulation* of S is a maximal set of non-crossing edges between pairs of points from S such that no point from S is placed inside any of the formed triangles' circumcircles. It is well known that if no four points in S are co-circular then

the Delaunay triangulation of S is a dual of the Voronoi diagram of S in the following sense [13]: for each edge e of each region in the Voronoi diagram of S , if e is a part of the bisector of the points u, v in S then (u, v) is an edge of the Delaunay triangulation of S ; again, see Fig. 1.

Our concept of weak smoothness is formally defined in terms of two parameters as follows.

Definition 1. Let ε, d be two positive real constants. A set of N points in a unit square is (ε, d) -smooth if for any two equal size squares Q, R within the unit square the following implication holds:

if Q contains at least N^ε points of S and R is at distance at most $d \cdot \ell$ from Q , where ℓ is the length of each edge of Q and R , then R contains at least one point of S .

We also need to define a sequence of grids within a unit square and related notions.

Definition 2. For a nonnegative integer i , we shall denote by $G_i(U)$ the orthogonal grid within the unit orthogonal square U that includes the edges of U such that the distance between two neighboring vertical or horizontal grid line segments is $\frac{1}{2^i}$. A basic square of $G_i(U)$ is a square within U such that the endpoints of each its edge is a pair of neighboring grid points. For a basic square R in $G_i(U)$, we shall denote the orthogonal region consisting of R and the two layers of basic squares around R by $TL_i(R)$ (if between R and an edge of the unit square there is a place only for one or zero layers then $TL_i(R)$ includes only one or zero layers on this side, respectively).

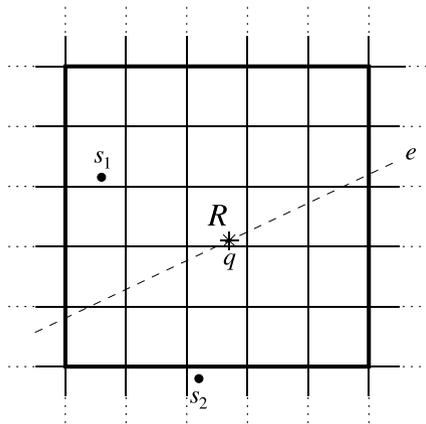


Fig. 2. An example of the configuration in the proof of Lemma 1.

The proof of the following lemma corresponds to the second paragraph of the proof of Theorem 4 in [6].

Lemma 1. *Let R be a basic square in a grid $G_i(U)$ within the unit square U . Consider a finite set S of points within the unit square. If R contains a point in S then the Voronoi diagram of S within R can be computed by taking into account only the points of S within $TL_i(R)$. Hence, in particular all edges (u, v) of the Delaunay triangulation of S such that a part of the bisector of u and v borders some region of the Voronoi diagram of S within R can be determined.*

Proof. Let e be an edge of the Voronoi diagram of S within R . The edge e has to be a part of the bisector of some couple of points s_1 and s_2 in S . Consider an arbitrary point q on e . Suppose that s_1 or s_2 is placed outside $TL_i(R)$, i.e., the orthogonal area consisting of at most $1 + 8 + 16 = 25$ squares including R . See Fig. 2. Without loss of generality, let s_2 be such a point. Then the distance between q and s_2 is at least $2 \cdot \frac{1}{2^i}$, while the distance between q and every point inside R is at most $\sqrt{2} \cdot \frac{1}{2^i}$. We obtain a contradiction because R contains at least one point from S and q is closer to such a point than to s_2 . \square

Lenzen gave an efficient solution to the following fundamental routing problem in the congested clique model, known as the *Information Distribution Task* (IDT) [7]: Each node of the congested n -clique holds a set of exactly n $O(\log n)$ -bit messages with their destinations, with multiple messages from the same source node to the same destination node allowed. Initially, the destination of each message is known only to its source node. Each node is the destination of exactly n of the aforementioned messages. The messages are globally lexicographically ordered by their source node, their destination, and their number within the source node. For simplicity, each such message explicitly contains these values, in particular making them distinguishable. The goal is to deliver all messages to their destinations, minimizing the total number of rounds.

Lenzen proved that IDT can be solved in $O(1)$ rounds (Theorem 3.7 in [7]). He also noted that the relaxed IDT, where each node is required to send and receive at most n messages, reduces to IDT in $O(1)$ rounds. From here on, we shall refer to this important result as:

Fact 1 [7]. *The relaxed Information Distribution Task can be solved deterministically within $O(1)$ rounds.*

The *Sorting Problem* (SP) is defined as follows:

Each node i of the congested n -clique holds a set of n $O(\log n)$ -bit keys. All the keys are different w.l.o.g. Each node i needs to learn all the keys of indices in $[n(i-1)+1, ni]$ (if any) in the total order of all keys.

Lenzen showed that SP can be solved in $O(1)$ rounds if each node holds a set of exactly n keys (Theorem 4.5 in [7]). In order to relax the requirement that each node holds exactly n keys to that of with most n keys, we can determine the maximum key and add appropriate different dummy keys in $O(1)$ rounds. We summarize this result as:

Fact 2 [7]. *The relaxed Sorting Problem can be solved in $O(1)$ rounds.*

3 The Local Approach

Consider a $(\frac{1}{2}, 4\sqrt{2})$ -smooth set of n^2 points with $O(\log n)$ -bit coordinates in a unit orthogonal square. We shall first describe a protocol for listing the edges of the Delaunay triangulation of the set that are dual to the edges of the Voronoi diagram of the set within the unit square. Roughly, it implicitly grows a quadtree of squares rooted at the unit square in phases corresponding to the levels of the quadtree. If a square R currently at a leaf of the quadtree jointly with the two layers of equal size squares around it include at most cn input points, for an appropriate positive constant c , then the intersection of the Voronoi diagram of the input point set with R and the dual edges of the Delaunay triangulation of the input point set can be computed locally. This follows from the smoothness condition and Lemma 1 combined with the fact that the parent square of R jointly with the two layers of equal size squares around it include more than cn input points. Otherwise, four child squares whose union forms R are created on the next level of the quadtree. In particular, checking the aforementioned condition in parallel for the squares at the current front level of the quadtree and delivering the necessary points to the nodes representing respective frontier squares in $O(1)$ rounds on the congested n -clique are highly non-trivial.

protocol $DT - SQUARE(S, U)$

Input: A $(\frac{1}{2}, 4\sqrt{2})$ -smooth set of n^2 points with $O(\log n)$ -bit coordinates in a unit orthogonal square U held in n -point batches at the n nodes of the congested clique.

Output: The set of the edges of the Delaunay triangulation of S dual to the edges of the Voronoi diagram of S within U held in $O(n)$ -edge batches at clique nodes.

1. Initialize an empty list L of edges of the Delaunay triangulation of S .
2. Activate the basic square U in $G_0(U)$ and assign it to the first node.
3. For $i = 0, 1, \dots$ do
 - (a) Each node for each point p in its batch determines the number $num(p)$ of the basic square of $G_i(U)$ containing p in a common fixed numbering of the basic squares in $G_i(U)$ (e.g., column-wise and top-down-wise in each column). Next, a *prefixed representation* of p is formed by the concatenation of the fixed size bit representation of $num(p)$ with the fixed size bit representations of the coordinates of p .
 - (b) The points in S are sorted by their prefixed representation. After that each node informs all other nodes about the range of numbers of the basic squares in $G_i(U)$ holding the prefixed representations of points in S that landed at the node after the sorting of the prefixed representations of all the points.
 - (c) For each basic square W in $G_i(U)$ such that the prefixed representations of points belonging to W that are held in a sequence C of at least two consecutive nodes, the nodes in C inform additionally the other nodes in C about the number of the prefixed representations of the points in W they got so in particular the node in C with the smallest index can compute the total number of the points in W .
 - (d) For each active basic square R in $G_i(U)$, the node representing R sends queries to the nodes holding the prefixed point representations of the points in the basic

squares in $TL_i(R)$ (i.e., in R and the two layers of basic squares around R in $G_i(U)$) about the number of points in these squares. If several nodes hold the prefixed point representation of points in a basic square in $TL_i(R)$, the query is sent only to that node with the smallest index.

- (e) After getting answers to the queries, each node proceeds as follows for each active basic square in R in $G_i(U)$ it represents. If the total number of points of S in $TL_i(R)$ does not exceed $100n$ then the node asks the nodes holding the prefixed representations of the points in the basic squares in $TL_i(R)$ for sending the points to the node. After that the node computes the Voronoi diagram of all these points and then the intersection of the diagram with R locally. Next, the node appends to L all edges (u, v) where a part of the bisector of u and v borders some region of the Voronoi diagram in the computed intersection. Otherwise, the node activates the four squares in $G_{i+1}(U)$ whose union forms R and assigns them temporarily to itself.
- (f) The nodes balance the assignment of active basic squares in $G_{i+1}(U)$ by informing all other nodes about the number of active basic squares in $G_{i+1}(U)$ they are assigned and following the results of the same assignment balancing algorithm run by each of them separately locally.
- (g) The list L is sorted in order to remove multiple copies of the same edge.

Lemma 2. $DT - SQUARE(S, U)$ activates basic squares solely in the grids $G_i(U)$, where $i = O(\log n)$.

Proof. Simply, the points in S have $O(\log n)$ -bit coordinates so at depth at most $O(\log n)$ the condition in Step 3(e) of $DT - SQUARE(S, U)$ has to be satisfied. □

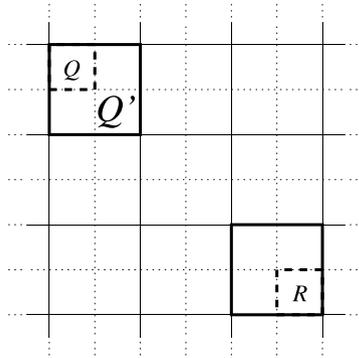


Fig. 3. An example of the configuration in the proof of Lemma 3.

Lemma 3. The protocol $DT - SQUARE(S, U)$ is correct.

Proof. When the Voronoi diagram of the points of S in $TL_i(R)$ for a basic square R of the grid $G_i(U)$ is computed then there must be square Q' in the grid $G_{i-1}(U)$ that contains at least $100n/25$ points in S and is at distance at most $\frac{\sqrt{2}}{2^{i-1}}$ from the basic square in $G_{i-1}(U)$ that is the parent of R . Hence, there is a basic square Q in $G_i(U)$ that is part of Q' and contains at least $100n/100$ points in S ; see Fig. 3. By straightforward calculations, the distance between Q and R is at most $4\sqrt{2}\frac{1}{2^i}$. Thus, by the assumed $(\frac{1}{2}, 4\sqrt{2})$ -smoothness property, the square R contains at least one point in S . It follows from Lemma 1 that the intersection of the Voronoi diagram of the points of S in $TL_i(R)$ with R yields the Voronoi diagram of S within R . Hence, the edges appended to the list L are the edges of the Delaunay triangulation of S dual to the edges of the Voronoi diagram of S within U . It easily follows by induction on i during forming the quadtree of active basic squares that the leaf active basic squares form a partition of the unit square U . Therefore, for each edge (u, v) of the Delaunay triangulation of S dual to an edge of the Voronoi diagram of S within U there must exist a positive integer i and an active basic square R in $G_i(U)$ such that R does not have any child active basic squares in $G_{i+1}(U)$ and a part of the bisector of u and v borders some region in the Voronoi diagram of S within R . Hence, the list L is complete. \square

Lemma 4. *For $i = 0, 1, \dots, O(\log n)$, the number of active basic squares in the grid $G_i(U)$ is $O(n)$ during the execution of $DT - SQUARE(S, U)$.*

Proof. We argue similarly as at the beginning of the proof of Lemma 3. If R is an active basic square in $G_i(U)$ different from the unit square U then there must exist a basic square Q in $TL_{i-1}(R')$, where R' is the parent of R in $G_{i-1}(U)$, such that Q contains at least $100n/25$ points in S . Now it is sufficient to note that: (i) there are at most $O(n)$ basic squares in $G_{i-1}(U)$ that contain at least $100n/25$ points in S ; (ii) there are at most $O(1)$ basic squares Q' in $G_{i-1}(U)$ different from R' such that Q is included in $TL_{i-1}(Q')$; (iii) an active basic square in $G_{i-1}(U)$ can be a parent to at most four active basic squares in $G_i(U)$. \square

Lemma 5. *The protocol $DT - SQUARE(S, U)$ can be implemented in $O(\log n)$ rounds on the congested clique.*

Proof. Steps 1, 2 can be easily implemented in $O(1)$ rounds. By Lemma 2, the block under the for loop in Step 3 is iterated $O(\log n)$ times. It is sufficient to show that this block (a-g) can be implemented in $O(1)$ rounds.

Step 3(a) can be performed totally locally.

The sorting of the prefixed representations of points in S in Step 3(b) can be done in $O(1)$ rounds by Fact 2.

For each node, the range of the numbers of the basic squares in $G_i(U)$ holding the prefixed representations of points in S at the node after the sorting of the prefix representations of the points can be specified by two $O(\log n)$ -bit numbers. Hence, all nodes can inform all other nodes about their ranges in $O(1)$ rounds. Thus, Step 3(b) requires $O(1)$ rounds in total.

The situation described in Step 3(c) can happen for at most n basic squares W in $G_i(U)$. It requires sending by each node at most two different messages to at most n

nodes in total and also receiving at most n messages. Hence, Step 3(c) can be implemented in $O(1)$ rounds by using the routing protocol from Fact 1.

In Step 3(d), for each active basic square, a node representing the square has to send $O(1)$ $O(\log n)$ -bit queries to $O(1)$ other nodes. The total number of active basic squares in $G_i(U)$ is $O(n)$ by Lemma 4. Hence, by using the routing protocol from Fact 1 this task can be done in $O(1)$ rounds.

Consider Step 3(e). Answering the queries sent in Step 3(c) can be done by local computations and the routing reverse to that in Step 3(d) in $O(1)$ rounds. After that each node for each active square assigned to it determines locally if the criterion for computing the Voronoi diagram of S within R is satisfied. If so the node sends messages asking the nodes holding the prefixed representations of points in the squares of $TL_i(R)$ for sending the points. This requires sending $O(n)$ messages for each active basic square in $G_i(U)$. Since the total number of such squares is $O(n)$ by Lemma 4 and each node represents $O(1)$ active squares in $G_i(U)$, it can be accomplished in $O(1)$ rounds by Fact 1. Delivering the requested points to the nodes representing respective active basic squares can also be done in $O(1)$ rounds for the following reasons. For each active basic square the node representing it needs to receive $O(n)$ points. Furthermore, by Lemma 4 there are $O(n)$ active basic squares in $G_i(U)$. Hence, since the active squares are assigned to the n nodes in a balanced way, each node needs to receive $O(n)$ points. Also, the points contained in a given basic square in $G_i(U)$ can be requested by at most $O(1)$ nodes since there are at most $O(1)$ active basic squares behind these requests to the given square. Since the sorted prefixed representations of the points in S are divided between the nodes in a balanced way, each node needs to send $O(n)$ points, each of them to $O(1)$ nodes. We conclude that this part of Step 3(d) can be implemented in $O(1)$ rounds by Fact 1. The remaining parts of Step 3(d) are done locally.

Step 3(f) requires sending and receiving by each node $O(1)$ messages so it can be done in $O(1)$ rounds.

Consider an edge (u, v) dual to some edge of the Voronoi diagram of the points of S included in $TL_i(R)$ within an active basic square R in $G_i(U)$. The edge can be appended to L at most for $O(1)$ different squares R as u, v are in $TL_i(R)$. Therefore, the list L may contain at most $O(1)$ copies of an edge of the Delaunay triangulation of S so Step 3(g) can be implemented in $O(1)$ rounds by using the sorting protocol from Fact 2. □

Lemmas 3, 5 yield our first main result.

Theorem 1. *Let S be a $(\frac{1}{2}, 4\sqrt{2})$ -smooth set of n^2 points with $O(\log n)$ -bit coordinates in an orthogonal unit square, held in n -point batches at the n nodes of the congested clique. The set of edges of the Delaunay triangulation of S dual to the edges of the Voronoi diagram of S within the unit square can be constructed in $O(\log n)$ rounds on the congested clique.*

Lemma 6. *Let S be defined as in Theorem 1. Suppose that a list L of the edges of the Delaunay triangulation of S dual to the edges of the Voronoi diagram of S within the unit square is held in $O(n)$ -edge batches at the n nodes of the congested clique. The Voronoi diagram of S within the unit square can be constructed in $O(1)$ rounds on the congested clique.*

Proof. Double the list L by inserting for each $(u, v) \in L$ also (v, u) into L . For each edge (u, v) determine locally an $O(\log n)$ -bit representation of the angle $\beta(u, v)$ between (u, v) and the horizontal line passing through u . For instance, the representation can specify the tangent of the angle by $(v_y - u_y, v_x - u_x)$. Sort the edges (x, y) by $(x, \beta(u, v))$, letting the nodes compare the angle tangents locally, using the sorting protocol from Fact 2. In this way, for each point $u \in S$, a sub-list of all edges of the Delaunay triangulation incident to u in the angular order is created. Some of the sub-lists can stretch through several nodes of the clique network. Given the edges of the Delaunay triangulation incident to u in the angular order, the edges of the Voronoi region of u within the unit square can be easily produced. This is done by intersecting the bisectors of u and the other endpoints of consecutive edges incident to u in the angular order as long as the intersection of two consecutive bisectors is within the unit square. Otherwise, the border of the region of u has to be filled with the fragment of the perimeter of the unit square between the intersections of the two bisectors with the perimeter. \square

Theorem 1 combined with Lemma 6 yield our second main result.

Theorem 2. *Let S be a $(\frac{1}{2}, 4\sqrt{2})$ -smooth set of n^2 points with $O(\log n)$ -bit coordinates in an orthogonal unit square held in n -point batches at the n nodes of the congested clique. The Voronoi diagram of S within the unit square can be constructed in $O(\log n)$ rounds on the congested clique.*

4 Final Remarks

The *message complexity* of a protocol in the congested clique model is the maximum total number of $O(\log n)$ -bit messages exchanged by the n nodes of the congested clique during a run of the protocol (e.g., see [12]). In case of our protocols, it is easily seen to be the product of the maximum number of messages that can be exchanged in a single round, i.e., $\Theta(n^2)$, times the number of required rounds. Thus, the message complexity of our deterministic protocols for the Delaunay triangulation and the Voronoi diagram of n^2 point sets from Sect. 3 is $O(n^2 \log n)$.

The remaining major open problem is the derivation of a low polylogarithmic upper bound on the number of rounds sufficient to deterministically construct the Voronoi diagram of n^2 points with $O(\log n)$ -bit coordinates in the Euclidean plane (when the points are not necessarily randomly distributed) on the congested clique with n nodes. This seems feasible but it might require a substantial effort as in the PRAM case [1, 2].

Note here that the existence of an $O(\log n)$ -time (unit cost) PRAM algorithm for a geometric problem on a point set (e.g., [2]) does not guarantee the membership of the problem in the NC^1 class defined in terms of Boolean circuits [4, 10]. Simply, assuming that the input points have $O(\log n)$ -bit coordinates, the arithmetic operations of the PRAM implemented by Boolean circuits of bounded fan-in have a non-constant depth, at least $\Omega(\log \log n)$. Consequently, the Boolean circuit simulating the $O(\log n)$ -time PRAM for fixed input size can have a super-logarithmic depth. This is a subtle and important point in the context of relatively recent results of Frei and Wada providing simulations of the classes NC^k , $k > 0$, by MapReduce (see Theorems 9 and 10 in

[4]), and consequently, in the Massively Parallel Computation (MPC) and BSP models (see Theorem 1 in [10]). Due to the $O(1)$ -round routing protocol of Lenzen [7], the congested clique model in our setting can be roughly regarded as a special case of MPC, where the size of the input is approximately the square of the number of processors. For this reason, the NC simulation results from [4] are relevant to our model only when the parameter ϵ in the exponent of the space bounds in [4] equals $\frac{1}{2}$. This is possible in case of Theorem 9 in [4] on NC^1 simulation but not possible in case of Theorem 10 in [4] on NC^k , $k > 0$, simulation. However, the proof of the former theorem in [4] relies on a strict logarithmic upper bound on the depth of Boolean circuit of bounded fan-in required by Barrington's characterization of the NC^1 class in terms of bounded-width polynomial-size branching programs. Otherwise, one has to adhere to the direct circuit simulation method from [4] that does not work for $\epsilon = \frac{1}{2}$. In summary, Theorems 9 and 10 in [4] do not seem to have any direct consequences for geometric problems on point sets in our model setting.

Acknowledgments. The authors are grateful to the anonymous reviewers for valuable comments. This research was in part funded by Swedish Research Council grant 2018-04001.

References

1. Aggarwal, A., Chazelle, B., Guibas, L.J., Ó'Dúnlaing, C., Yap, C.-K.: Parallel computational geometry. *Algorithmica* **3**, 293–327 (1988)
2. Amato, M., Preparata, F.: A time-optimal parallel algorithm for three-dimensional convex hull. *Algorithmica* **14**(2), 169–182 (1995)
3. Censor-Hillel, K., Kaski, P., Korhonen, J.H., Lenzen, C., Paz, A., Suomela, J.: Algebraic methods in the congested clique. *Distrib. Comput.* **32**(6), 461–478 (2019)
4. Frei, F., Wada, K.: Efficient deterministic MapReduce algorithms for parallelizable problems. *J. Parallel Distrib. Comput.* **177**, 28–38 (2023)
5. Goodrich, M.: Randomized fully-scalable BSP techniques for multi-searching and convex hull construction. In: *Proceedings of the Eighth Annual Symposium on Discrete Algorithms (SODA 1997)*, pp. 767–776. ACM-SIAM (1997)
6. Jansson, J., Levcopoulos, C., Lingas, A., Polishchuk, V.: Convex hulls, triangulations, and Voronoi diagrams of planar point sets on the congested clique. [arXiv:2305.09987](https://arxiv.org/abs/2305.09987) (2023). Preliminary version in *Proceedings of the Thirty-Fifth Canadian Conference on Computational Geometry (CCCG 2023)*, pp. 183–189 (2023)
7. Lenzen, C.: Optimal deterministic routing and sorting on the congested clique. In: *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing (PODC 2013)*, pp. 42–50. ACM (2013)
8. Levcopoulos, C., Katajainen, J., Lingas, A.: An optimal expected-time parallel algorithm for Voronoi diagrams. In: Karlsson, R., Lingas, A. (eds.) *SWAT 1988*. LNCS, vol. 318, pp. 190–198. Springer, Heidelberg (1988). https://doi.org/10.1007/3-540-19487-8_21
9. Lotker, Z., Patt-Shamir, B., Pavlov, E., Peleg, D.: Minimum-weight spanning tree construction in $O(\log \log n)$ communication rounds. *SIAM J. Comput.* **35**(1), 120–131 (2005)
10. Nanongkai, D., Scquizzato, M.: Equivalence classes and conditional hardness in massively parallel computing. *Distrib. Comput.* **35**, 165–183 (2022)
11. Nowicki, K.: A deterministic algorithm for the MST problem in constant rounds of congested clique. In: *Proceedings of the Fifty-Third Annual ACM SIGACT Symposium on Theory of Computing (STOC 2021)*, pp. 1154–1165. ACM (2021)

12. Pemmaraju, S., Sardeshmukh, V.: Super-fast MST algorithms in the congested clique using $o(m)$ messages. In: Proceedings of the 36th Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2016), pp. 47:1–47:15. LIPICS (2016)
13. Preparata, F., Shamos, M.: Computational Geometry: An Introduction, vol. 10. In: Texts and Monographs in Computer Science. Springer-Verlag (1985). <https://doi.org/10.1007/978-1-4612-1098-6>
14. Robinson, P.: Brief announcement: what can we compute in a single round of the congested clique? In: Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing (PODC 2023), pp. 168–171. ACM (2023)
15. Vemuri, B.C., Varadarajan, R., Mayya, N.: An efficient expected time parallel algorithm for Voronoi construction. In: Proceedings of the Fourth Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 1992), pp. 392–401. ACM (1992)