

The Inverse Voronoi Problem in Graphs I: Hardness

Édouard Bonnet* Sergio Cabello† Bojan Mohar‡ Hebert Pérez-Rosés§

May 1, 2020

Abstract

We introduce the inverse Voronoi diagram problem in graphs: given a graph G with positive edge-lengths and a collection \mathbb{U} of subsets of vertices of $V(G)$, decide whether \mathbb{U} is a Voronoi diagram in G with respect to the shortest-path metric. We show that the problem is NP-hard, even for planar graphs where all the edges have unit length. We also study the parameterized complexity of the problem and show that the problem is $W[1]$ -hard when parameterized by the number of Voronoi cells or by the pathwidth of the graph.

Keywords: distances in graphs, Voronoi diagram, inverse Voronoi problem, NP-complete, parameterized complexity.

1 Introduction

Let (X, d) be a metric space, where $d: X \times X \rightarrow \mathbb{R}_{\geq 0}$. Let Σ be a subset of X . We refer to each element of Σ as a *site*, to distinguish it from an arbitrary point of X . The *Voronoi cell* of each site $s \in \Sigma$ is then defined by

$$\text{cell}_{(X,d)}(s, \Sigma) = \{x \in X \mid \forall s' \in \Sigma : d(s, x) \leq d(s', x)\}.$$

The *Voronoi diagram* of Σ in (X, d) is

$$\mathbb{V}_{(X,d)}(\Sigma) = \{\text{cell}_{(X,d)}(s, \Sigma) \mid s \in \Sigma\}.$$

It is easy to see that, for each set Σ of sites, each element of X belongs to some Voronoi cell $\text{cell}_{(X,d)}(s, \Sigma)$. Therefore, the sets in $\mathbb{V}_{(X,d)}(\Sigma)$ cover X . On the other hand, the Voronoi cells do not need to be pairwise disjoint. In particular, when some point $x \in X$ is closest to two sites, then it is in both Voronoi cells.

In the *inverse Voronoi problem*, we are given a metric space (X, d) and a sequence X_1, \dots, X_k of subsets of X that cover X . The task is to decide whether $\{X_1, \dots, X_k\}$ is a Voronoi diagram in

*Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France. Email address: edouard.bonnet@ens-lyon.fr. Supported by the LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon, within the program "Investissements d'Avenir" (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR).

†**Corresponding author.** Faculty of Mathematics and Physics, University of Ljubljana, and IMFM, Slovenia. Supported by the Slovenian Research Agency, program P1-0297 and projects J1-8130, J1-8155, J1-9109. Email address: sergio.cabello@fmf.uni-lj.si.

‡Department of Mathematics, Simon Fraser University, Burnaby, BC, Canada. Email address: mohar@sfu.ca. On leave from IMFM & FME, Department of Mathematics, University of Ljubljana. Supported in part by the NSERC Discovery Grant R611450 (Canada), by the Canada Research Chairs program, and by the Research Project J1-8130 of ARRS (Slovenia).

§Departament d'Enginyeria Informàtica i Matemàtiques, Universitat Rovira i Virgili, Tarragona, Spain. Partially supported by Grant MTM2017-86767-R from the Spanish Ministry of Economy, Industry and Competitiveness.

(X, d) . This means that we have to decide whether there exist some sites s_1, \dots, s_k such that, for each index i , we have $X_i = \text{cell}_{(X,d)}(s_i, \{s_1, \dots, s_k\})$.

The inverse Voronoi problem is closely related to problems in classification and clustering. In pattern recognition, a classic paradigm to classify is to use the nearest neighbor rule: given a learning set \mathcal{L} of objects that are already classified, each new object is classified into the same class as its closest object from \mathcal{L} . To reduce the size of the learning set, Hart [17] introduced the concept of *consistent subsets*. A subset \mathcal{L}' of the learning set \mathcal{L} is a consistent subset if, for each object ℓ from \mathcal{L} , the object ℓ and its closest neighbor in \mathcal{L}' are in the same class. An equivalent, alternative perspective of this is given by Voronoi diagrams: in the Voronoi diagram of a consistent subset \mathcal{L}' , each object ℓ of \mathcal{L} belongs to a Voronoi cell defined by a site $s \in \mathcal{L}'$ if and only if ℓ and s belong to the same class. Ritter et al. [28] introduced the problem of finding consistent subsets of minimum size. Surveying the research in this applied area is beyond the scope of our research. We refer to Biniaz et al. [5] and Gottlieb et al. [16] for some of the latest algorithmic results on this topic. Considering each class as a Voronoi cell, the inverse Voronoi problem is asking precisely whether there exists a consistent subset with one element per class. Such consistent subset has of course to be of optimal size.

Graphic version. Let G be an undirected graph with n vertices and abstract, positive edge-lengths $\lambda: E(G) \rightarrow \mathbb{R}_{>0}$. The length of a path in G is the sum of the edge-lengths along the path. We define the (shortest-path) **distance** between two vertices x and y of G , denoted by $d_G(x, y)$, as the minimum length over all paths in G from x to y .

Since $(V(G), d_G)$ is a metric space, we can consider the concepts of Voronoi cells and Voronoi diagrams for this space. We denote them by $\text{cell}_G(s, \Sigma)$ and $\mathbb{V}_G(\Sigma)$ respectively. Moreover, when the graph is clear from the context, we remove the subscript and thus just talk about $\text{cell}(s, \Sigma)$ and $\mathbb{V}(\Sigma)$.

In this paper we consider computational aspects of the inverse Voronoi problem when the metric space is the shortest-path metric in a graph. Thus, we consider the following problem.

GRAPHIC INVERSE VORONOI

Input: (G, \mathbb{U}) , where G is a graph with positive edge-lengths and $\mathbb{U} = (U_1, \dots, U_k)$ is a sequence of subsets of vertices of G that cover $V(G)$.

Question: Are there sites $s_1, \dots, s_k \in V(G)$ such that $\text{cell}_G(s_i, \{s_1, \dots, s_k\}) = U_i$ for each i ?

See Figure 1 for an example. As far as the existence of polynomial-time algorithms is concerned, it is equivalent to consider a graph or a finite metric space. Indeed, for each finite metric space we can build a graph that encodes those distances by using a complete graph with edge-lengths, and, inversely, given a graph, we can compute the matrix of distances between all pairs of vertices in polynomial time. However, considering special classes of graphs may be useful to get more efficient algorithms.

Our results. First we show that the problem GRAPHIC INVERSE VORONOI is NP-hard even for planar graphs where the candidate Voronoi cells are pairwise disjoint and each has at most 3 vertices. The reduction is from a variant of PLANAR 3-SAT. The bound on the number of vertices per cells is tight: when each candidate Voronoi cell has 2 vertices, the problem can be solved using 2-SAT.

Many graph decision and optimization problems admit fixed-parameter tractable (FPT) algorithms with respect to additional parameters that quantify how complex is the input; see for instance [9]. Using the framework of parameterized complexity, we provide stronger lower bounds when parameterized by the number k of sites and the pathwidth $p(G)$ of G . More precisely,




Figure 1: An instance with two solutions. The edges have unit length and the larger, filled dots represent the sites.

assuming the Exponential Time Hypothesis (ETH), we show that the problem cannot be solved in time $f(k)|V(G)|^{o(k/\log k)}$ nor in time $f(p(G))|V(G)|^{o(p(G))}$, for any computable function f . These hardness results hold for graphs where all the edges have unit length.

Related work. Voronoi diagrams on graphs were first investigated by Erwig [12], who showed that they can be efficiently computed. Subsequently, graph Voronoi diagrams have been used in a variety of applications. For instance, Okabe and Sugihara [27] describe several applications of graph Voronoi diagrams. More recent applications, many of them for planar graphs, can be found in [7, 8, 13, 14, 21, 25]. Voronoi diagrams in graphs have also been considered in the context of the so-called Voronoi game [3, 15] and in the context of topological data analysis [10].

On the other hand, the inverse Voronoi problem in the traditional, Euclidean setting has been studied since the mid 1980s, starting with the seminal paper by Ash and Bolker [2]. We are not aware of any previous work considering the graphic inverse Voronoi problem.

In an accompanying paper [6], we consider efficient algorithms for the problem GRAPHIC INVERSE VORONOI when the underlying graph is a tree.

2 Basics

Sets and Graphs. For a positive integer k we use the notation $[k] = \{1, \dots, k\}$. We use the standard graph-theoretic definitions and notations that can be found in Diestel's book [11]. In particular, we denote by $V(G)$, respectively by $E(G)$, the vertex-set, respectively the edge-set, of a graph G . If G is a graph and $S \subseteq V(G)$, we denote by $G[S]$ the subgraph induced by S , and $G - S$ is a short-hand for $G[V(G) \setminus S]$. A graph is said *planar* if its vertices can be drawn as distinct points of the real plane, its edges can be drawn as simple curves (in the plane) connecting the points that represent its vertices, and the interior of the curves representing the edges are pairwise non-crossing.

A *vertex-separator* in a graph G is a subset of its vertices $S \subseteq V(G)$ such that $G - S$ is a disconnected non-empty graph. A vertex-separator S is said *balanced* if all the connected components of $G - S$ have size at most $2|V(G)|/3$. Up to constant factors, we could define the notion of treewidth by means of repeated balanced vertex-separators. We choose not to do so, in order to follow the usual definition and to also introduce pathwidth. We will need the notion of balanced vertex-separators in Theorem 12 and treewidth/pathwidth in Section 5.

Treewidth and pathwidth. A *tree-decomposition* of a graph G is a tree T whose nodes are labeled by subsets of $V(G)$, called *bags*, such that for each vertex $v \in V(G)$, the bags containing v induce a subtree of T , and for each edge $e \in E(G)$, there is at least one bag containing both endpoints of e . The width of a tree-decomposition is the size of its largest bag minus one. The treewidth of a graph G is the minimum width of a tree-decomposition of G . The point of the “minus one” in the definition is that the tree-width of trees is 1, and not 2. The pathwidth is the same as treewidth except the tree T is now required to be a path, and hence is called a path-decomposition. In particular pathwidth is always at least as large as treewidth. We observe that if G has treewidth w , then it admits in particular a balanced vertex-separator of size at most $w + 1$. Indeed, any non-leaf bag is a vertex-separator. One can find one that is also balanced by starting from the root (any node) of the tree-decomposition labeled by say, S , and iteratively moving to a component of $G - S$ which is larger than $2/3$ of the whole graph. When this process is no longer possible, we have our balanced vertex-separator.

In Section 5, we will need to bound the pathwidth (and therefore the treewidth) of rather complicated graphs. Writing down the full description of a tree-decomposition or of a path-decomposition may be a bit tedious. Kirousis and Papadimitriou [22] showed the equality between the interval thickness number, which is equal to pathwidth plus one, and the *node searching number*. To give an upper bound to the pathwidth, we only need to prove that the number of cleaners required to win the following one-player game is bounded by a suitable function. We imagine the edges of a graph to be contaminated by a gas. We shall move around a team of cleaners, placed at the vertices, in order to clean all the edges. A move consists of removing a cleaner from the graph, adding a cleaner at an unoccupied vertex, or displacing a cleaner from a vertex to any other vertex (not necessarily adjacent). An edge is cleaned when both its endpoints are occupied by a cleaner. After each move, all the cleaned edges admitting a free-of-cleaners path from one of its endpoints to the endpoint of a contaminated edge are however recontaminated. The node searching number is the minimum number of cleaners required to win the game.

ETH and Sparsification Lemma. The Exponential-Time Hypothesis (ETH, for short) of Impagliazzo and Paturi [18] asserts that there is no subexponential-time algorithm solving 3-SAT. More precisely, there is a positive real number $\delta > 0$ such that 3-SAT cannot be solved in time $2^{\delta n}$ on n -variable instances. Impagliazzo et al. [19] present a subexponential-time Turing-reduction parameterized by a positive real number $\varepsilon > 0$ which, given a 3-SAT-instance ϕ with n variables and m clauses, produces at most $2^{\varepsilon n}$ 3-SAT-instances ϕ_1, \dots, ϕ_t such that $\phi \Leftrightarrow \bigvee_{i \in [t]} \phi_i$, each ϕ_i having no more than n variables and $C_\varepsilon n$ clauses for some constant C_ε (depending solely on ε , and *not* on n and m). This important reduction is known as the Sparsification Lemma. Due to the Sparsification Lemma, there exists a positive real number $\delta' > 0$ such that no algorithm can solve 3-SAT in time $2^{\delta'(n+m)}$ on n -variable m -clause instances, assuming that the ETH holds.

Inverse Voronoi and compatibility. Consider an instance $(G, (U_1, \dots, U_k))$ to the GRAPHIC INVERSE VORONOI and a candidate solution $s_1, \dots, s_k \in V(G)$. We say that s_i and s_j ($i \neq j$) are **compatible** if we have $d(s_i, u) = d(s_j, u)$ for each $u \in U_i \cap U_j$, $d(s_i, u) < d(s_j, u)$ for each $u \in U_i \setminus U_j$, and $d(s_j, u) < d(s_i, u)$ for each $u \in U_j \setminus U_i$. Consider a fixed index $i \in [k]$. It is straightforward from the definition that $\text{cell}_G(s_i, \{s_1, \dots, s_k\}) = U_i$ if and only if s_i and s_j are compatible for all $j \neq i$. (Here the relevant assumption is that $U_1 \cup \dots \cup U_k$ is $V(G)$.)

In all cases we use G as the ground graph that defines the metric. Note that in the following claims it is important that G has positive edge-lengths.

We have remarked before that Voronoi cells need not be disjoint. A vertex belongs to various Voronoi cells if it is equidistant to different sites. An alternative is to define cells using strict inequalities. More precisely, for a set Σ of sites, the **open Voronoi cell** of each site $s \in \Sigma$ is then




Figure 2: Left: incidence graph for the POSITIVE 1-IN-3-SAT instance with $\mathcal{V} = \{x_1, x_2, \dots, x_5\}$ and $\mathcal{C} = \{C_1 = \{x_1, x_3, x_4\}, C_2 = \{x_2, x_3, x_4\}, \dots, C_5 = \{x_1, x_2, x_3\}\}$. Right: resulting instance for GRAPHIC INVERSE VORONOI. Each connected shaded region corresponds to one set of \mathbb{U} .

defined by

$$\text{cell}^{\leftarrow}(s, \Sigma) = \{x \in X \mid \forall s' \in \Sigma \setminus \{s\} : d(s, x) < d(s', x)\}.$$

In this case, the cells are disjoint but they do not necessarily form a partition of X . The following lemma is straightforward and we omit its proof.

Lemma 1. *For each set Σ of sites and each site $s \in \Sigma$ we have $s \in \text{cell}^{\leftarrow}(s, \Sigma)$ and*

$$\text{cell}^{\leftarrow}(s, \Sigma) = \text{cell}(s, \Sigma) \setminus \left(\bigcup_{s' \neq s} \text{cell}(s', \Sigma) \right).$$

3 Hardness of the Graphic Inverse Voronoi

In this Section we show that the problem GRAPHIC INVERSE VORONOI is NP-hard, even for planar graphs. Stronger lower bounds are derived assuming the Exponential Time Hypothesis (ETH). We will make a reduction from a variant of the satisfiability (SAT) where each clause has 3 literals, all the literals are positive, and we want that each clause is satisfied at exactly one literal. The problem can be stated combinatorially as follows.

POSITIVE 1-IN-3-SAT

Input: $(\mathcal{V}, \mathcal{C})$, where \mathcal{V} is a ground set and \mathcal{C} is a family of subsets of \mathcal{V} of size 3.

Question: Is there a subset $T \subseteq \mathcal{V}$ such that $|C \cap T| = 1$ for each $C \in \mathcal{C}$?

In this combinatorial setting, \mathcal{V} represents the variables, \mathcal{C} represents the clauses with 3 positive literals each, and T represents the variables that are set to true.

The incidence graph $I(\mathcal{V}, \mathcal{C})$ of an instance $(\mathcal{V}, \mathcal{C})$ has vertex set $\mathcal{V} \cup \mathcal{C}$ and an edge between $v \in \mathcal{V}$ and $C \in \mathcal{C}$ precisely when $v \in C$. The graph is bipartite.

As shown by Mulzer and Rote [26], the problem POSITIVE 1-IN-3-SAT is NP-complete even when the incidence graph is planar.

Theorem 2. *The GRAPHIC INVERSE VORONOI problem is NP-hard on planar graphs with unit edge-lengths, even when the candidate Voronoi cells are disjoint sets of size at most 3.*

Proof. We make a reduction from POSITIVE 1-IN-3-SAT with planar incidence graphs. Let $(\mathcal{V} = \{x_1, \dots, x_n\}, \mathcal{C} = \{C_1, \dots, C_m\})$ be an instance of POSITIVE 1-IN-3-SAT with planar incidence graph. We produce an equivalent instance (G, \mathbb{U}) of GRAPHIC INVERSE VORONOI as follows. See Figure 2.

- For each element $x_i \in \mathcal{V}$, we add two vertices $v(x_i)$ and $\bar{v}(x_i)$ to the vertex set of G , and we connect them by an edge. We add the set $\{v(x_i), \bar{v}(x_i)\}$ to the candidate Voronoi cells \mathbb{U} .
- For each subset $C_j = \{x_a, x_b, x_c\}$, we add three vertices $v(j, a)$, $v(j, b)$, and $v(j, c)$ to $V(G)$, and we connect the three pairs by edges, forming a triangle. We add the set $\Delta(j) = \{v(j, a), v(j, b), v(j, c)\}$ to \mathbb{U} .
- Finally, for each $x_a \in \mathcal{V}$ and each $C_j \in \mathcal{C}$ with $x_a \in C_j$, we link $v(j, a)$ to $v(x_a)$ by an edge.

This finishes the construction of (G, \mathbb{U}) . We observe that the sets of \mathbb{U} are indeed pairwise disjoint and of size 2 or 3. The graph G is planar since it is obtained from the planar incidence graph $I(\mathcal{V}, \mathcal{C})$ by adding pendant vertices and splitting each vertex representing a subset (with three neighbors) into a triangle in which each vertex is linked to one distinct neighbor.

If there is a solution T to the instance $(\mathcal{V}, \mathcal{C})$, we position the sites in the following way. For each $x_i \in \mathcal{V}$, we place the site of $\{v(x_i), \bar{v}(x_i)\}$ in $v(x_i)$ if $x_i \in T$, and in $\bar{v}(x_i)$, otherwise. For each $C_j = \{x_a, x_b, x_c\} \in \mathcal{C}$, we place the site of $\Delta(j)$ in $v(j, z)$, where x_z is the unique element of $C_j \cap T$. We denote by Σ the obtained set of sites. We check that this placement defines the same Voronoi cells as specified by \mathbb{U} .

- For each $\bar{v}(x_i) \in \Sigma$, we have $\text{cell}_G(\bar{v}(x_i), \Sigma) \supseteq \{v(x_i), \bar{v}(x_i)\}$, since by construction there is no site in $v(x_i)$. The only neighbors of $\{v(x_i), \bar{v}(x_i)\}$ are vertices $v(j, i)$ for some values of $j \in [m]$. However, those neighbors do not contain a site of Σ by construction. On the other, there is always a site of Σ at distance at most 1 of $v(j, i)$, whereas $\bar{v}(x_i)$ is at distance 2 of $v(j, i)$. Hence, $\text{cell}_G(\bar{v}(x_i), \Sigma) = \{v(x_i), \bar{v}(x_i)\}$.
- Similarly, for each $v(x_i) \in \Sigma$, we have $\text{cell}_G(v(x_i), \Sigma) \supseteq \{v(x_i), \bar{v}(x_i)\}$, since by construction there is no site in $\bar{v}(x_i)$. The only neighbors of $\{v(x_i), \bar{v}(x_i)\}$ are vertices $v(j, i)$ for some values of $j \in [m]$, but since $v(x_i) \in \Sigma$, by construction, $v(j, i)$ also belongs to Σ . Therefore, $\text{cell}_G(v(x_i), \Sigma) = \{v(x_i), \bar{v}(x_i)\}$.
- Finally, consider some $v(j, z) \in \Sigma$, where $C_j = \{x_a, x_b, x_c\}$. We have $\text{cell}_G(v(j, z), \Sigma) \supseteq \Delta(j)$ because $v(j, z)$ is the only site in $\Delta(j)$. The only other neighbor of $v(j, z)$ is $v(x_z)$, which is in Σ . The only neighbor of $v(j, z')$ with $z' \in \{a, b, c\} \setminus \{z\}$ is $v(x_{z'})$ which is at distance 2 of $v(j, z)$ and at distance 1 of the site $\bar{v}(x_{z'}) \in \Sigma$. Thus, $\text{cell}_G(v(j, z), \Sigma) = \Delta(j)$.

If there is no solution to the instance $(\mathcal{V}, \mathcal{C})$, we show that there is no solution to the GRAPHIC INVERSE VORONOI instance (G, \mathbb{U}) . Fix a position of the sites. The set of sites Σ has to intersect each $\{v(x_i), \bar{v}(x_i)\}$ exactly once. Define the set

$$T = \{x_i \in \mathcal{V} \mid \text{the site chosen for } \{v(x_i), \bar{v}(x_i)\} \text{ is } v(x_i)\}.$$

As T is not a solution for the POSITIVE 1-IN-3-SAT instance, there is a set $C_j = \{x_a, x_b, x_c\} \in \mathcal{C}$ such that $|C_j \cap T| \neq 1$. We now turn our attention to the site chosen for $\Delta(j)$. We distinguish two cases: $|C_j \cap T| = 0$ and $|C_j \cap T| \geq 2$. If $|C_j \cap T| = 0$, then, for every position of the site, say in $v(j, z)$ (with $z \in \{a, b, c\}$), $\text{cell}_G(v(j, z), \Sigma)$ contains $v(x_z)$, and therefore cannot be equal to $\Delta(j)$. Now if $|C_j \cap T| \geq 2$, let $v(x_z)$ and $v(x_{z'})$ be two sites of Σ with $z \neq z' \in \{a, b, c\}$. Since $\Delta(j)$ contains precisely one site, we have $v(j, z) \notin \Sigma$ or $v(j, z') \notin \Sigma$. If $v(j, z) \notin \Sigma$, then $\text{cell}_G(v(x_z), \Sigma)$ contains $v(j, z)$, and therefore cannot be equal to $\{v(x_z), \bar{v}(x_z)\}$. Similarly, if $v(j, z') \notin \Sigma$, then $\text{cell}_G(v(x_{z'}), \Sigma) \neq \{v(x_{z'}), \bar{v}(x_{z'})\}$. In both cases, we reach the conclusion that there cannot be a solution for the instance (G, \mathbb{U}) .

Note that in the argument we did not use that $I(\mathcal{V}, \mathcal{C})$ or G are planar. \square

Using additional properties of the reduction from (PLANAR) 3-SAT to (PLANAR) POSITIVE 1-IN-3-SAT given by Mulzer and Rote [26] and the Sparsification Lemma, we derive the following conditional lower bound.

Corollary 3. *Unless the Exponential Time Hypothesis fails, the problem GRAPHIC INVERSE VORONOI cannot be solved in time $2^{o(n)}$ in general graphs and in time $2^{o(\sqrt{n})}$ in planar graphs, where n is the number of vertices, even when the potential Voronoi cells are disjoint and of size at most 3.*

Proof. Applying the reduction of Mulzer and Rote [26] to a 3-SAT instance with n variables and m clauses gives an instance to POSITIVE 1-IN-3-SAT with $O(n + m)$ variables and $O(m)$ clauses. This is so because in their reduction each clause is replaced locally using $O(1)$ new variables and clauses. The reduction and the proof used in Theorem 2 then give an instance with $O(n + m)$ vertices. (The reduction also works for non-planar instances, as mentioned at the end of the proof.) Therefore, if we could solve GRAPHIC INVERSE VORONOI in time $2^{o(|V(G)|)}$, we could solve any 3-SAT instance with n variables and m clauses in time $2^{o(|V(G)|)} = 2^{o(n+m)}$. However, the Sparsification Lemma [20] rules out, under the Exponential Time Hypothesis, a running time $2^{o(n+m)}$ for 3-SAT.

The reduction from 3-SAT to PLANAR 3-SAT given by Lichtenstein [23] increases quadratically the number of variables and clauses. Together with the reduction of Mulzer and Rote from (PLANAR) 3-SAT to (PLANAR) POSITIVE 1-IN-3-SAT and our reduction in the proof of Theorem 2, we conclude that each instance of 3-SAT with n variables and m clauses becomes an instance of GRAPHIC INVERSE VORONOI where the graph G is *planar* and has $O((n + m)^2)$ vertices. Again, solving the problem in planar graphs in time $2^{o(\sqrt{|V(G)|})}$ time for planar graphs would contradict the Sparsification Lemma. \square

This upper bound of 3 for the size of the potential Voronoi cells is sharp.

We show that the problem can be solved in polynomial time when each potential Voronoi cell has at most two points not contained in other potential cells. For this, one uses a reduction to 2-SAT. Inspired by Lemma 1, we say that each $U \in \mathbb{U}$ defines the *potential open Voronoi cell*

$$U \setminus \left(\bigcup_{U' \in \mathbb{U} \setminus \{U\}} U' \right).$$

Theorem 4. *The GRAPHIC INVERSE VORONOI problem can be solved in polynomial time when all the potential open Voronoi cells are of size at most 2.*

Proof. We present a polynomial reduction to 2-SAT. See Figure 3 for an example. Let $(G, \mathbb{U} = \{U_1, \dots, U_k\})$ be the GRAPHIC INVERSE VORONOI instance. We denote by U'_i the open potential Voronoi cell of the potential Voronoi cell U_i . By assumption, $|U'_i| \leq 2$. Because of Lemma 1, if the instance has a solution, then $s_i \in U'_i$. For each open cell U'_i , we introduce a variable x_i . We interpret putting the site on one fixed but arbitrary vertex of U'_i to setting x_i to true, and putting the site on the other vertex (if it exists) to setting x_i to false. Now, $\mathbb{V}_G(\Sigma) = \mathbb{U}$ if and only if for each pair of sites $s_i, s_j \in \Sigma$ with $s_i \in U'_i$ and $s_j \in U'_j$:

- every vertex of $U_i \setminus U_j$ is strictly closer to s_i than to s_j , and
- every vertex of $U_j \setminus U_i$ is strictly closer to s_j than to s_i , and
- every vertex of $U_i \cap U_j$ is equidistant to s_i and s_j .

Therefore, one just needs to check that each pair of sites of Σ is *compatible*, that is, satisfies those three conditions.

We define the following set of 2-SAT constraints. For each open cell U'_i of size 1, we add the clause x_i , which forces to set x_i to true. For each pair $s_i \in U'_i, s_j \in U'_j$ which is *not* compatible we add the clause $\ell_i \vee \ell_j$ where ℓ_i (resp. ℓ_j) is the opposite literal to the one chosen by placing a site in s_i (resp. s_j).

It is easy to check that the produced 2-SAT formula is satisfiable if and only if there is a pairwise compatible set of sites. This is in turn equivalent to the existence of a solution for the GRAPHIC INVERSE VORONOI instance. \square




Figure 3: An instance satisfying the assumption of Theorem 4. The vertices of each set $U \in \mathbb{U}$ are enclosed by dashed curve. The crosses indicate the position when the variables are true. Some of the compatibility clauses are $x_5 \vee \overline{x_4}$, $\overline{x_5} \vee x_4$, $x_5 \vee x_4$, $\overline{x_2} \vee \overline{x_3}$, etc., as well as the 1-clause x_3 .

4 Hardness parameterized by the number of Voronoi cells

In the previous section we showed that the problem GRAPHIC INVERSE VORONOI is NP-hard. Stronger lower bounds are derived under the assumption of the Exponential Time Hypothesis (ETH). We will prove the following result.

Theorem 5. *The GRAPHIC INVERSE VORONOI problem is $W[1]$ -hard parameterized by the number of candidate Voronoi cells. Furthermore, for n -vertex graphs and k subsets to be candidate Voronoi cells, for any computable function f , there is no algorithm to solve the GRAPHIC INVERSE VORONOI problem in $f(k)n^{o(k/\log k)}$ time, unless the Exponential Time Hypothesis fails. The claim holds even for graphs with unit edge-lengths.*

Note that it is trivial to solve the problem in $n^{O(k)}$ time: just try each n^k tuples of k vertices as candidate sites and check each of them. The remaining of this section is devoted to prove Theorem 5. We will make a reduction from the following problem.

MULTICOLORED SUBGRAPH ISOMORPHISM

Input: (H, P) , where H is a graph whose vertex set $V(H)$ is partitioned into ℓ pairwise disjoint sets $V_1 \uplus \dots \uplus V_\ell$, and a pattern graph P with vertex set $V(P) = [\ell]$.

Question: Can we select vertices $v_i \in V_i$ for every $i \in [\ell]$ such that we have $v_i v_j \in E(H)$ for each $ij \in E(P)$?

When the answer is positive, we say that P is isomorphic to a multicolored subgraph of H . It follows from the work of Marx [24] that, assuming the Exponential Time Hypothesis, the MULTICOLORED SUBGRAPH ISOMORPHISM cannot be solved in time $f(\ell)n^{o(\ell/\log \ell)}$ for any computable function f , even when the pattern P has $\Theta(\ell)$ edges. This lower bound is made explicit for example in [25, Corollary 5.5], where P is assumed to be 3-regular.

Consider an instance (H, P) to the MULTICOLORED SUBGRAPH ISOMORPHISM problem, where V_1, \dots, V_k are the partite classes of $V(H)$ and P has $\Theta(\ell)$ edges. We assume for simplicity that each vertex of P has degree at least 2. For each $i, j \in [\ell]$, let $E_H(V_i, V_j)$ denote the edges of H with one endpoint in V_i and the other endpoint in V_j . We shall assume that $E_H(V_i, V_j)$ is empty




Figure 4: A graph H (left) and a pattern P (right) for the MULTICOLORED SUBGRAPH ISOMORPHISM.

whenever $ij \notin E(P)$ because those edges can be removed without affecting the instance. We build an instance $(G, \mathbb{U}) = (G(H, P), \mathbb{U}(H, P))$ for the GRAPHIC INVERSE VORONOI problem as follows.

Figures 4 and 5 may be helpful to follow the construction.

- We start with $V(G) = V(H)$ and $E(G) = E(H)$.
- We subdivide each edge e of G with a new vertex, which we call $w(e)$.
- For each $i \in [\ell]$, we add all edges between all the vertices in V_i .
- For each $ij \in E(P)$, let W_{ij} be the vertices $w(e)$ used to subdivide $E_H(V_i, V_j)$. We add all edges between all the vertices in $W_{i,j}$.
- For each $ij \in E(P)$, we add $U_{ij} = W_{ij} \cup V_i \cup V_j$ to \mathbb{U} .

All the edges have unit length. This completes the construction of $G = G(H, P)$ and $\mathbb{U} = \mathbb{U}(H, P)$. Note that \mathbb{U} has $|E(P)| = \Theta(\ell)$ candidate Voronoi regions, while G has $|V(H)| + |E(H)| = \Theta(|V(H)|^2)$ vertices and

$$\sum_{i \in [\ell]} \binom{|V_i|}{2} + \sum_{ij \in E(P)} \binom{|E_H(V_i, V_j)|}{2} + 2|E(H)|$$

edges.

The next two lemmas show that the pair (G, \mathbb{U}) is a correct reduction from MULTICOLORED SUBGRAPH ISOMORPHISM to GRAPHIC INVERSE VORONOI. The intuition of the reduction is that selecting the site of each Voronoi cell corresponds to selecting an edge of $E_H(V_i, V_j)$ for each $ij \in E(P)$. Moreover, the selection of the edges we make needs to have compatible endpoints in each partite set V_i , as otherwise we do not get the correct Voronoi cells.

Lemma 6. *If P is isomorphic to a multicolored subgraph of H , then G has a set Σ of sites such that $\mathbb{V}(\Sigma) = \mathbb{U}$.*

Proof. Assume that P is isomorphic to a multicolored subgraph of H . This means that there are vertices $v_i \in V_i$, for all $i \in [\ell]$, such that $v_i v_j \in E(H)$ for every $ij \in E(P)$. This means that, for




Figure 5: Example showing the construction for (H, P) of the Figure 4. Left: the graph G . The vertices inside a connected shaded region form a clique whose edges are not shown in the drawing. Right: three (of the five) candidate Voronoi cells of \mathbb{U} are indicated by shaded regions of different colors. The cliques induced by V_i and W_{ij} are not shown in this figure.

every $ij \in E(P)$, the vertex $w(v_i v_j)$ obtained when subdividing $v_i v_j$ belongs to G . We define $s_{ij} = w(v_i v_j)$ for every $ij \in E(P)$ and $\Sigma = \{w(v_i v_j) \mid ij \in E(P)\}$.

We claim that Σ is a set of sites in G such that $\text{cell}(s_{ij}, \Sigma) = U_{ij}$, for every $ij \in E(P)$. This claim implies the lemma.

For each $s_{ij} \in \Sigma$ and for each vertex u of G we have the following distances

$$d_G(s_{ij}, u) = \begin{cases} 0 & \text{if } u = s_{ij}, \\ 1 & \text{if } u \in W_{ij} \setminus \{s_{ij}\}, \\ 1 & \text{if } u = v_i \text{ or } v = u_j, \\ 2 & \text{if } u \in (V_i \cup V_j) \setminus \{v_i, v_j\}, \\ \geq 2 & \text{if } u \in W_{i'j'} \text{ for some } i'j' \neq ij, \\ \geq 3 & \text{if } u \in V_t \text{ for some } t \neq i, j. \end{cases}$$

Therefore, each vertex in $V_1 \cup \dots \cup V_\ell = V(H)$ is at distance at most 2 from some vertex of Σ and each vertex in $\bigcup_{ij \in E(P)} W_{ij}$ is at distance at most 1 from some vertex of Σ .

Now we note that, for each $ij \in E(P)$, each vertex of W_{ij} is strictly closer to s_{ij} than to any other site. Furthermore, for each $i \in [\ell]$, each vertex of V_i has the same distance to each site $s_{ij'}$ with $ij' \in E(P)$, and a larger distance to each $s_{i'j'}$ with $i'j' \in E(P - i)$. Therefore $\text{cell}(s_{ij}, \Sigma) = W_{ij} \cup V_i \cup V_j = U_{ij}$. The result follows. \square

Lemma 7. *If G has a set Σ of sites such that $\mathbb{V}(\Sigma) = \mathbb{U}$, then P is isomorphic to a multicolored subgraph of H .*

Proof. Let Σ be a set of sites in G such that $\mathbb{V}(\Sigma) = \mathbb{U}$. For each $ij \in E(P)$, let s_{ij} be the site of Σ with $\text{cell}(s_{ij}, \Sigma) = U_{ij}$.

Because of Lemma 1, each $s_{ij} \in \Sigma$ belongs to

$$\text{cell}^<(s_{ij}, \Sigma) = \text{cell}(s_{ij}, \Sigma) \setminus \left(\bigcup_{s \in \Sigma \setminus \{s_{ij}\}} \text{cell}(s, \Sigma) \right) = U_{ij} \setminus \left(\bigcup_{i'j' \in E(P) \setminus \{ij\}} U_{i'j'} \right) = W_{ij}.$$

In the last equality we have used that each vertex of P has degree at least 2, which means that each V_i is contained in at least 2 sets U_{ij} of \mathbb{U} . We conclude that, for each $ij \in E(P)$, the site s_{ij} must be in W_{ij} .

Since each site s_{ij} is in W_{ij} , for each $ij \in E(P)$, the construction of G implies that there are unique vertices $v(i, ij) \in V_i$ and $v(j, ij) \in V_j$ such that s_{ij} is the vertex obtained when subdividing the edge connecting $v(i, ij)$ and $v(j, ij)$. In particular, $v(i, ij)v(j, ij)$ is an edge of $E(H)$.

Fix the index $i \in [\ell]$ and consider two edges $ij, ij' \in E(P)$ incident to i . We must have $v(i, ij) = v(i, ij')$, as otherwise we would have $d_G(s_{ij}, v(i, ij)) = 1 < 2 = d_G(s_{ij'}, v(i, ij'))$, which would imply that $V_i \not\subseteq \text{cell}(s_{ij'}, \Sigma) = U_{ij'}$ and would contradict the definition of $U_{ij'} = V_i \cup V_{j'} \cup W_{ij'}$. Therefore, each of the (three) edges ij of $E(P)$ defines the same vertex $v(i, ij) \in V_i$. Henceforth we denote this vertex by v_i .

We have found ℓ vertices v_1, \dots, v_ℓ with the property that $v_i \in V_i$, for each $i \in [\ell]$, and such that the edge $v_i v_j = v(i, ij)v(j, ij)$ in $E(H)$, for each $ij \in E(P)$. This means that P is isomorphic to the multicolored subgraph of H defined by $\{v_1, \dots, v_\ell\}$. \square

Proof of Theorem 5. As shown in Lemmas 6 and 7, H has a multicolored subgraph isomorphic to P if and only if \mathbb{U} is a valid Voronoi diagram of G . Thus, the answers to MULTICOLORED SUBGRAPH ISOMORPHISM(H, P) and GRAPHIC INVERSE VORONOI(G, \mathbb{U}) are the same.

Recall that \mathbb{U} has $|E(P)| = \Theta(\ell)$ candidate Voronoi regions. If we could solve each instance of the GRAPHIC INVERSE VORONOI problem with n vertices and k sites in time $f(k)n^{o(k/\log k)}$, for some computable function f , then we could solve the instance (G, \mathbb{U}) in

$$f(|\mathbb{U}|) \cdot |V(G)|^{o(|\mathbb{U}|/\log |\mathbb{U}|)} \leq f(\Theta(\ell))(\Theta(|V(H)|^2))^{o(\Theta(\ell)/\log(\Theta(\ell)))} \leq g(\ell)|V(H)|^{o(\ell/\log \ell)}$$

time, for some computable function g . However, this also means that we could solve the MULTICOLORED SUBGRAPH ISOMORPHISM in H with pattern P in $g(\ell)|V(H)|^{o(\ell/\log \ell)}$ time, and this contradicts the Exponential Time Hypothesis. \square

5 Hardness parameterized by the pathwidth and the treewidth

In this section we show that the GRAPHIC INVERSE VORONOI problem is unlikely to be fixed parameter tractable with respect to the pathwidth of the graph. Since the pathwidth is always smaller than the treewidth, this implies the same result for the treewidth. More precisely, in this section we will prove the following.

Theorem 8. *The GRAPHIC INVERSE VORONOI problem is $W[1]$ -hard parameterized by the pathwidth of the input graph. Furthermore, for n -vertex graphs with pathwidth p , there is no algorithm to solve the GRAPHIC INVERSE VORONOI problem in time $f(p)n^{o(p)}$ for any computable function f , unless the Exponential Time Hypothesis fails. The claims hold even for graphs with unit edge-lengths and disjoint candidate Voronoi cells.*

In order to show that, we will reduce the following $W[1]$ -hard problem.

MULTICOLORED INDEPENDENT SET

Input: A graph $H = (V, E)$ whose vertex set V is partitioned into ℓ pairwise disjoint sets $V_1 \uplus \dots \uplus V_\ell$.

Question: Is there an independent set X of size ℓ in H such that $|X \cap V_i| = 1, \forall i \in [\ell]$?

An independent set of H is said *multicolored* if it satisfies $|X \cap V_i| = 1, \forall i \in [\ell]$. The MULTICOLORED INDEPENDENT SET problem is $W[1]$ -hard with respect to ℓ and cannot be solved in time $f(\ell)n^{o(\ell)}$ for any computable function f , assuming the Exponential Time Hypothesis [9, Corollary 14.23]. The lower bounds still hold if all the partite sets V_i have the same cardinality and there are no edges connecting any two vertices within a set V_i .




Figure 6: A graph H whose vertex set is partitioned into $\ell = 3$ partite sets, each with $t = 4$ vertices.

Let $H = (V_1 \uplus \dots \uplus V_\ell, E)$ be an instance of MULTICOLORED INDEPENDENT SET such that $|V_1| = |V_2| = \dots = |V_\ell| = t$. Let m be the number of edges in H . We build an equivalent GRAPHIC INVERSE VORONOI instance (G, \mathbb{U}) where the pathwidth of G is $\Theta(\ell)$. This instance will have unit edge-length edges and the sets in \mathbb{U} will be pairwise disjoint.

Our global strategy for the reduction is to propagate a vertex choice in each V_i with a *path-like* structure with ℓ rows and $m = |E|$ columns. In each column, we introduce a single distinct edge of E so that the pathwidth of the built graph stays in $\Theta(\ell)$. Figure 7 shows the whole reduction in the graph H of Figure 6. (Seeing the details requires zooming in.) In Figure 8 we show a part of the construction in detail showing also the notation we employ. The detailed construction is as follows.

- For each $i \in [\ell]$ and $j \in [m]$, we add to $V(G)$ an independent set $I(i, j)$ of size $|V_i| = t$. The vertices of the independent set $I(i, j)$ are denoted by $v(i, j, 1)$ to $v(i, j, t)$, the third index being in one-to-one correspondence with the vertices of V_i .
- For each $i \in [\ell]$ and $j \in [m]$, we add two vertices $a(i, j)$ and $b(i, j)$. Furthermore, for each $h \in [t]$, we connect $a(i, j)$ to $v(i, j, h)$ by a private path $P_a(i, j, h)$ of length $t + h$, and we connect $b(i, j)$ to $v(i, j, h)$ by a private path $P_b(i, j, h)$ of length $t + h$. For each $i \in [\ell]$ and $j \in [m - 1]$, we connect $b(i, j)$ and $a(i, j + 1)$ by an edge.
- For each $i \in [\ell]$ and $j \in [m]$, we add three new vertices $c(i, j)$, $e(i, j)$ and $z(i, j)$. For each $h \in [t]$, we add a private path $P_c(i, j, h)$ of length t between $v(i, j, h)$ and $c(i, j)$. Furthermore, we connect $e(i, j)$ and $z(i, j)$ with an edge and add a path $P_e(i, j)$ of length t with one extreme on $e(i, j)$ and the other extreme connected through an edge to $c(i, j)$. (Thus $e(i, j)$ and $c(i, j)$ are connected with a path of length $t + 1$.)
- For each $i \in [\ell]$ and $j \in [m]$, we denote by $U(i, j)$ the set of vertices comprising $I(i, j)$ and all the paths going from this independent set to $a(i, j)$, $b(i, j)$, and $c(i, j)$, including those three vertices. We add $U(i, j)$, $V(P_e(i, j))$, and $Z(i, j) = \{z(i, j)\}$ to the candidate Voronoi cells \mathbb{U} .
- We call the set $\bigcup_{i=1}^{\ell} U(i, j) \cup P_e(i, j) \cup Z(i, j)$ the *j -th column*, for a fixed $j \in [m]$. We introduce exactly one distinct edge of E per column. Let e_1, \dots, e_m be any ordering of the edges of E . We put an edge gadget encoding e_j in the j -th column, for every $j \in [m]$.




Figure 7: Whole graph showing the reduction used to prove Theorem 8 for the graph H in Figure 6. Each connected gray area corresponds to one candidate Voronoi region. Figure 8 shows details for a part of the construction.




Figure 8: Zoom into a part of the reduction shown in Figure 7 with some notation. Each connected gray area corresponds to one candidate Voronoi region. Some selection of sites marked with crosses that is *locally* correct (but globally would have a problem). This selection corresponds to selecting vertex 3 of V_1 , vertex of 2 of V_2 and vertex 1 of V_3 .

Assume that e_j is an edge between the h -th vertex of V_i and the h' -th vertex of $V_{i'}$ where $i \neq i'$. We add a path $P(e_j)$ of length $2t + 2$ between $v(i, j, h)$ and $v(i', j, h')$. We add a path $Q(e_j)$ of length t between the middle vertex of $P(e_j)$ and a new vertex, denoted $f(j)$. (The vertex $f(j)$ has degree 1 and it is at distance $2t + 1$ from $v(i, j, h)$ and from $v(i', j, h')$.) We add $R(j) = V(Q(e_j)) \cup V(P(e_j)) \setminus \{v(i, j, h), v(i', j, h')\}$ as a candidate Voronoi region to \mathbb{U} . The subgraph induced by $R(j)$ is the **edge gadget** of e_j .

That finishes the construction of $G = G(H)$ and of $\mathbb{U} = \mathbb{U}(H)$. All the edges of G have unit length. One can observe that \mathbb{U} is made of pairwise disjoint sets and it contains $(3\ell + 1)m$ candidate Voronoi cells. We first show that the pathwidth (and thus also the treewidth) of G is at most $2\ell + 5$. For that, we use the cleaning game characterization of pathwidth presented in Section 2.

Lemma 9. *The pathwidth of G is at most $2\ell + 5$.*

Proof. We present a winning strategy for cleaning G using $2\ell + 6$ cleaners. We make m rounds where in the j -th round, $j = 1, \dots, m$, we scan completely the j -th column and the gadget for e_j .

At the start of the j -th round we have 2ℓ cleaners placed at the vertices $a(i, j)$ and $b(i, j)$ for all $i \in [\ell]$. Assume that the edge $e_j \in E$ is between the h -th vertex of V_i and the h' -th vertex of $V_{i'}$. We place two cleaners at $v(i, j, h)$ and $v(i', j, h')$. Let X_j be the set of $2\ell + 2$ vertices where we have cleaners. They will stay there for most of the j -th round.

We then clean the whole j -th column plus the edge gadget of e_j using the remaining four cleaners. For every $i \in [\ell]$, we define $\mathcal{G}(i, j) = U(i, j) \cup P_e(i, j) \cup Z(i, j)$, the intersection of the j -th column with the i -th row. For every i going from 1 to ℓ , we place a $2\ell + 3$ -rd cleaner at $c(i, j)$. All the connected components of $\mathcal{G}(i, j)$ in $G - (X_j \cup \{c(i, j)\})$ are paths or subdivisions of a claw (i.e., a star with three leaves). These graphs, and more generally, disjoint unions of subdivided stars can be decontaminated with three cleaners. Place a first cleaner at the center of the first subdivided star. This disconnects it into a disjoint union of paths, that can be cleaned with two additional cleaners. When all the paths are cleaned, move the first cleaner to the center of the second subdivided star and proceed similarly with the remaining connected components. When all the paths and subdivided claws of $\mathcal{G}(i, j)$ have been decontaminated, we remove the $2\ell + 3$ -rd cleaner (to place it at $c(i + 1, j)$ at the next iteration). We proceed similarly, up to $\mathcal{G}(\ell, j)$. Finally we observe that the edge gadget is also a subdivided claw, so it can be handled with three additional cleaners.

At this point, the remaining toxic gas is confined to the right of the ℓ cleaners placed at $b(1, j), \dots, b(\ell, j)$, that is, on some edge incident to some vertex defined by $j' > j$. If $j = m$, we are done since there can no longer be contaminated edges. Otherwise, we move the cleaners from $a(1, j), \dots, a(\ell, j)$ to $a(1, j + 1), \dots, a(\ell, j + 1)$, then the cleaners from $b(1, j), \dots, b(\ell, j)$ to $b(1, j + 1), \dots, b(\ell, j + 1)$, and start the next round. \square

We now show the correctness of the reduction.

Lemma 10. *If H has a multicolored independent set of size ℓ , then there is a set $\Sigma \subseteq V(G)$ such that $\mathbb{V}_G(\Sigma) = \mathbb{U}$.*

Proof. Assume there is a multicolored independent set X of size ℓ in H . We define the set Σ of sites as follows.

- For each $i \in [\ell]$, we place a site on the m vertices $v(i, j, h_i)$ for all $j \in [m]$, where h_i is the index of the vertex of X in the partite set V_i .
- For every $i \in [\ell]$ and $j \in [m]$, we place two sites at $e(i, j)$ and $z(i, j)$.
- For each edge e_j of H with no endpoint in X , we place a site at $s_j = f(j)$.

- For each edge e_j of H with exactly one endpoint in X , we place a site on the vertex s_j of $R(j)$ adjacent to $v(i, j, h_i)$.

Note that, since X is an independent set, there cannot be an edge e_j with two endpoints in X . Therefore we have covered all cases. This finishes the placement of the sites.

For every $i \in [\ell]$ and $j \in [m]$, we have $\text{cell}(z(i, j), \Sigma) = Z(i, j)$ because $e(i, j)$, the only neighbor of $z(i, j)$, is also a site. It also holds that $\text{cell}(e(i, j), \Sigma) = P_e(i, j)$ because $c(i, j)$ is at distance $t + 1$ from $e(i, j)$ and at distance t from the site $v(i, j, h_i)$. For every $i \in [\ell]$ and $j \in [m - 1]$, the vertices $v(i, j, h_i)$ and $v(i, j + 1, h_i)$ are compatible since $d_G(v(i, j, h_i), b(i, j)) = t + h_i = d_G(v(i, j + 1, h_i), a(i, j))$ and $d_G(v(i, j, h_i), a(i, j + 1)) = t + h_i + 1 = d_G(v(i, j + 1, h_i), b(i, j))$. Here it is relevant to choose the lengths of the paths $P_a(i, j, h)$, $P_b(i, j, h)$ and $P_c(i, j, h)$ to ensure that the shortest path from vertex $v(i, j, h)$ to $a(i, j)$ is $P_a(i, j, h)$, instead of passing through $b(i, j)$ or $c(i, j)$. (Similar statements hold for the shortest paths from $v(i, j, h)$ to $b(i, j)$ and to $c(i, j)$.)

We now only need to check that the site s_j in the edge gadget of e_j —the edge, say, between the h -th vertex of V_i and the h' -th vertex of $V_{i'}$ —is compatible with the sites chosen in Σ for $U(i, j)$ and $U(i', j)$. The nice property that makes everything work is that, for every $i \in [\ell]$, $j \in [m]$, $h \neq h' \in [t]$, $d_G(v(i, j, h), v(i, j, h'))$ is always equal to $2t$. Indeed the shortest path between $v(i, j, h)$ and $v(i, j, h')$ goes through $c(i, j)$, which is at distance t of both vertices.

There are two cases: s_j is on the path $P(e_j)$ or $s_j = f(j)$. If $s_j \in P(e_j)$, it means that one of the endpoints of e_j is in the multicolored independent set X . Without loss of generality, we assume that it is the h -th vertex of V_i (hence, $h = h_i$). In that case, the sites s_j and $v(i, j, h_i)$ are adjacent vertices, and therefore they are compatible. The sites s_j and $v(i', j, h_{i'})$ are also compatible since $d_G(s_j, v(i', j, h_{i'})) = 2t + 1$ and $d_G(v(i', j, h_{i'}), v(i, j, h_i)) = 2t$.

Now, if $s_j = f(j)$, it means that e_j does not touch any vertex of Σ . Hence, $h \neq h_i$ and $h' \neq h_{i'}$. Then we have $d_G(s_j, v(i, j, h_i)) = d_G(s_j, v(i', j, h_{i'})) = 2t + 1$, $d_G(v(i, j, h), v(i, j, h_i)) = 2t$ and $d_G(v(i', j, h'), v(i', j, h_{i'})) = 2t$. It follows that also in this case the site s_j is compatible with $v(i, j, h_i)$ and $v(i', j, h_{i'})$.

Therefore, we showed that each site $v(i, j, h_i) \in \Sigma$ is compatible with every other site of Σ . This implies that for every $i \in [\ell]$ and $j \in [m]$, we have $\text{cell}(v(i, j, h_i), T) = U(i, j)$. In turn, it implies that $\text{cell}(s_j, T) = R(j)$ for each $j \in [m]$, and therefore $\mathbb{V}_G(\Sigma) = \mathbb{U}$. \square

Lemma 11. *If H has no multicolored independent set of size ℓ , then there is no set $\Sigma \subseteq V(G)$ such that $\mathbb{V}_G(\Sigma) = \mathbb{U}$.*

Proof. A solution for the GRAPHIC INVERSE VORONOI has to put sites on every $e(i, j)$ and $z(i, j)$, otherwise the Voronoi cell $Z(i, j)$ would not appear in the set of cells. As $e(i, j)$ is at distance $t + 1$ of $c(i, j)$, the site chosen for the cell $U(i, j)$ has to be at distance exactly t of $c(i, j)$ (otherwise, this site would not be compatible with $e(i, j)$). So, the site chosen for $U(i, j)$ has to be in $I(i, j)$.

Then we prove that if a site is placed on $v(i, j, h)$, a site should be placed consistently on $v(i, j + 1, h)$. This is immediate by construction, since the only vertex of $U(i, j + 1)$ which has a distance to $a(i, j + 1)$ equal to $d_G(v(i, j, h), b(i, j)) = t + h$ is $v(i, j + 1, h)$. Here we are using again that the shortest path from $v(i, j, h)$ to $b(i, j)$ is indeed $P_b(i, j, h)$, and does not detour through $a(i, j)$ or $c(i, j)$. This implies that, for each $i \in [\ell]$, all the choices of sites for the cells $\{U(i, j)\}_{j \in [m]}$ have to be consistent to the same vertex, say of index h_i in V_i . This defines a (consistent) set X of ℓ vertices of H .

As by assumption X cannot be an independent set, there is an edge e_j with both endpoints in X . Say those endpoints are the vertices in partite sets V_i and $V_{i'}$. Then, the site for $R(j)$ cannot be closer to the two vertices of $R(j)$ that are adjacent to $v(i, j, h_i)$ and $v(i', j, h_{i'})$. Hence there is no $\Sigma \subseteq V(G)$ such that $\mathbb{V}_G(\Sigma) = \mathbb{U}$. \square

Proof of Theorem 8. Because of Lemmas 10 and 11, solving GRAPHIC INVERSE VORONOI for (G, \mathbb{U}) also solves MULTICOLORED INDEPENDENT SET for H . The graph G has $O(m\ell t^2) = O(|V(H)|^5)$

vertices and pathwidth $p \leq 2\ell + 5$ because of Lemma 9. An algorithm for the GRAPHIC INVERSE VORONOI with running time $f(p)|V(G)|^{o(p)}$ (for some computable function f) would imply that we can solve MULTICOLORED INDEPENDENT SET in time $f(2\ell + 5)(|V(H)|^5)^{o(2\ell + 5)} = g(\ell)n^{o(\ell)}$ for a computable function g . This would disprove the Exponential Time Hypothesis. \square

We show an almost matching upper bound when the potential Voronoi cells form a partition of the vertex set.

Theorem 12. *Instances (G, \mathbb{U}) of GRAPHIC INVERSE VORONOI can be solved in time $|V(G)|^{O(w \log k)}$, when the k cells of \mathbb{U} are pairwise disjoint and w is the treewidth of G .*

Proof. We solve a more general problem where each potential Voronoi cell of \mathbb{U} comes with a prescribed subset, specifying where one can actually place its site. Let H be the graph on k vertices obtained by contracting each cell of \mathbb{U} into a single vertex. Contracting edges can only decrease the treewidth (see for instance [11]). Thus the treewidth of H is at most w . We exhaustively guess in time k^{w+1} a balanced vertex-separator S of size $w + 1$ in the graph H . Each connected component of $H - S$ has thus less than $2k/3$ vertices. We further guess in time at most $|V(G)|^{w+1}$ the $w + 1$ corresponding sites, say, s'_1, \dots, s'_{w+1} (for the cells of S) in a fixed solution. For each guess, we remove the $w + 1$ corresponding cells – say, their union is U – from G , and update the prescribed subsets of the remaining cells to those placements compatible with the sites that are already fixed. More precisely, if s is a placement of the site of U' (not included in U) incompatible with a site s'_i (for some $i \in [w + 1]$), then we remove s from the prescribed set for cell U' . We then solve recursively each connected component of $G - U$. Thus, we get $k|V(G)|^{w+1} \leq |V(G)|^{2(w+1)}$ independent subproblems, each of them with at most $2k/3$ candidate Voronoi regions (and restricted subset of possible placements). Since the depth of the branching tree is $O(\log k)$, the total running time is bounded by $|V(G)|^{O(w \log k)}$. \square

6 Conclusions

We have introduced the inverse Voronoi problem for graphs and we have shown several different hardness results, also within the framework of parameterized complexity. In an accompanying paper [6], we consider the problem GRAPHIC INVERSE VORONOI when the underlying graph is a tree.

Here we list some possible directions for further research.

- Is there an algorithm to solve the problem in $n^{O(w)}$ time for graphs with n vertices and treewidth w when the candidate Voronoi cells intersect? Perhaps one can also use some treewidth associated to the candidate Voronoi regions. In particular, for planar graphs a running time of $n^{O(\sqrt{k})}$ seems plausible but challenging when the Voronoi cells overlap.
- Considering cells defined by additively weighted sites.
- Following the analogy to problems considered in the Euclidean case [1, 4], find the smallest set Σ such that each U_i is the union of some Voronoi cells in $\mathbb{V}(\Sigma)$. Taking $\Sigma = V(G)$ gives a feasible solution, and our hardness implies that the problem is NP-hard. Can one get approximation algorithms?

Acknowledgments

Part of this work was done at the 21st Korean Workshop on Computational Geometry, held in Rogla, Slovenia, in June 2018. We thank all workshop participants for their helpful comments.

References

- [1] G. Aloupis, H. Pérez-Rosés, G. Pineda-Villavicencio, P. Taslakian, and D. Trinchet-Almaguer. Fitting Voronoi Diagrams to Planar Tessellations. *Combinatorial Algorithms - 24th International Workshop, IWOCA*, pp. 349–361. Springer, Lecture Notes in Computer Science 8288, 2013, https://doi.org/10.1007/978-3-642-45278-9_30.
- [2] P. F. Ash and E. D. Bolker. Recognizing Dirichlet tessellations. *Geometriae Dedicata* 19(2):175–206, Nov 1985, <http://dx.doi.org/10.1007/BF00181470>.
- [3] S. Bandyapadhyay, A. Banik, S. Das, and H. Sarkar. Voronoi game on graphs. *Theor. Comput. Sci.* 562:270–282, 2015, <https://doi.org/10.1016/j.tcs.2014.10.003>.
- [4] S. Banerjee, B. B. Bhattacharya, S. Das, A. Karmakar, A. Maheshwari, and S. Roy. On the Construction of Generalized Voronoi Inverse of a Rectangular Tessellation. *Trans. Computational Science*, pp. 22–38. Springer, Lecture Notes in Computer Science 8110, 2013, https://doi.org/10.1007/978-3-642-41905-8_3.
- [5] A. Biniiaz, S. Cabello, P. Carmi, J. De Carufel, A. Maheshwari, S. Mehrabi, and M. Smid. On the minimum consistent subset problem. *CoRR* abs/1810.09232, 2018, <http://arxiv.org/abs/1810.09232>.
- [6] É. Bonnet, S. Cabello, B. Mohar, and H. Pérez-Rosés. The inverse Voronoi problem in graphs II: trees, 2018. Manuscript. Its content is included in <http://arxiv.org/abs/1811.12547>.
- [7] S. Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. *ACM Trans. Algorithms*, to appear. Preliminary version presented at SODA 2017. Full version available at <http://arxiv.org/abs/1702.07815>.
- [8] É. Colin de Verdière. Shortest cut graph of a surface with prescribed vertex set. *Algorithms - ESA 2010, 18th Annual European Symposium, Part II*, pp. 100–111. Springer, Lecture Notes in Computer Science 6347, 2010, http://dx.doi.org/10.1007/978-3-642-15781-3_9.
- [9] M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015, <http://dx.doi.org/10.1007/978-3-319-21275-3>.
- [10] T. K. Dey, F. Fan, and Y. Wang. Graph induced complex on point data. *Comput. Geom.* 48(8):575–588, 2015, <https://doi.org/10.1016/j.comgeo.2015.04.003>.
- [11] R. Diestel. *Graph Theory, 4th Edition*. Graduate texts in mathematics 173. Springer, 2012.
- [12] M. Erwig. The graph Voronoi diagram with applications. *Networks* 36(3):156–163, 2000, [http://dx.doi.org/10.1002/1097-0037\(200010\)36:3<156::AID-NET2>3.0.CO;2-L](http://dx.doi.org/10.1002/1097-0037(200010)36:3<156::AID-NET2>3.0.CO;2-L).
- [13] P. Gawrychowski, H. Kaplan, S. Mozes, M. Sharir, and O. Weimann. Voronoi diagrams on planar graphs, and computing the diameter in deterministic $\tilde{O}(n^{5/3})$ time. *Proc. 29th ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pp. 495–514. SIAM, 2018, <http://dx.doi.org/10.1137/1.9781611975031.33>. Full version available at <http://arxiv.org/abs/1704.02793>.
- [14] P. Gawrychowski, S. Mozes, O. Weimann, and C. Wulff-Nilsen. Better tradeoffs for exact distance oracles in planar graphs. *Proc. 29th ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pp. 515–529, 2018, <http://dx.doi.org/10.1137/1.9781611975031.34>.

- [15] D. Gerbner, V. Mészáros, D. Pálvölgyi, A. Pokrovskiy, and G. Rote. Advantage in the discrete Voronoi game. *J. Graph Algorithms Appl.* 18(3):439–457, 2014, <https://doi.org/10.7155/jgaa.00331>.
- [16] L. Gottlieb, A. Kontorovich, and P. Nisnevitch. Near-optimal sample compression for nearest neighbors. *IEEE Trans. Information Theory* 64(6):4120–4128, 2018, doi:10.1109/TIT.2018.2822267, <https://doi.org/10.1109/TIT.2018.2822267>.
- [17] P. E. Hart. The condensed nearest neighbor rule (corresp.). *IEEE Trans. Information Theory* 14(3):515–516, 1968, <https://doi.org/10.1109/TIT.1968.1054155>.
- [18] R. Impagliazzo and R. Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.* 62(2):367–375, 2001, <https://doi.org/10.1006/jcss.2000.1727>.
- [19] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences* 63(4):512–530, 2001, <https://doi.org/10.1006/jcss.2001.1774>.
- [20] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.* 63(4):512–530, 2001, <https://doi.org/10.1006/jcss.2001.1774>.
- [21] S. Kannan, C. Mathieu, and H. Zhou. Graph reconstruction and verification. *ACM Trans. Algorithms* 14(4):40:1–40:30, 2018, <http://doi.acm.org/10.1145/3199606>. Preliminary version in ICALP 2015.
- [22] L. M. Kirousis and C. H. Papadimitriou. Interval graphs and searching. *Discrete Mathematics* 55(2):181–184, 1985, [https://doi.org/10.1016/0012-365X\(85\)90046-9](https://doi.org/10.1016/0012-365X(85)90046-9).
- [23] D. Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.* 11(2):329–343, 1982, <https://doi.org/10.1137/0211025>.
- [24] D. Marx. Can you beat treewidth? *Theory of Computing* 6(1):85–112, 2010, <https://doi.org/10.4086/toc.2010.v006a005>.
- [25] D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using Voronoi diagrams. *CoRR* abs/1504.05476, 2015, <http://arxiv.org/abs/1504.05476>. Preliminary version at ESA 2015.
- [26] W. Mulzer and G. Rote. Minimum-weight triangulation is NP-hard. *J. ACM* 55(2):11:1–11:29, 2008, <http://doi.acm.org/10.1145/1346330.1346336>.
- [27] A. Okabe and K. Sugihara. *Spatial Analysis along Networks. Statistical and Computational Methods*. John Wiley and Sons, 2012.
- [28] G. L. Ritter, H. B. Woodruff, S. R. Lowry, and T. L. Isenhour. An algorithm for a selective nearest neighbor decision rule (corresp.). *IEEE Trans. Information Theory* 21(6):665–669, 1975, <https://doi.org/10.1109/TIT.1975.1055464>.