

The Inverse Voronoi Problem in Graphs II: Trees

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

May 1, 2020

Abstract

We consider the inverse Voronoi diagram problem in trees: given a tree T with positive edge-lengths and a collection \mathbb{U} of subsets of vertices of $V(T)$, decide whether \mathbb{U} is a Voronoi diagram in T with respect to the shortest-path metric. We show that the problem can be solved in $O(N + n \log^2 n)$ time, where n is the number of vertices in T and $N = n + \sum_{U \in \mathbb{U}} |U|$ is the size of the description of the input. We also provide a lower bound of $\Omega(n \log n)$ time for trees with n vertices.

Keywords: Voronoi diagram in graphs, inverse Voronoi problem, trees, applications of binary search trees, dynamic programming in trees, lower bounds.

1 Introduction

Let T be a tree with n vertices and abstract, positive edge-lengths $\lambda: E(T) \rightarrow \mathbb{R}_{>0}$. The length of a path in T is the sum of the edge-lengths along the path. The (shortest-path) **distance** between two vertices x and y of T , denoted by $d_T(x, y)$, is the length of the unique path in T from x to y .

Let Σ be a subset of $V(T)$. We refer to each element of Σ as a **site**, to distinguish it from an arbitrary vertex of T . The **Voronoi cell** of each site $s \in \Sigma$ is then defined by

$$\text{cell}_T(s, \Sigma) = \{x \in V(T) \mid \forall s' \in \Sigma : d_T(s, x) \leq d_T(s', x)\}.$$

The **Voronoi diagram** of Σ in T is

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

When the tree is clear from the context, we remove the subindex and thus just talk about $d(\cdot, \cdot)$, $\text{cell}(s, \Sigma)$ and $\mathbb{V}(\Sigma)$. It is easy to see that, for each set Σ of sites, each vertex of T belongs to some Voronoi cell. Therefore, the sets in $\mathbb{V}_T(\Sigma)$ cover all vertices of T . On the other hand, the Voronoi cells do not need to be pairwise disjoint. In particular, when some vertex of T is closest to two sites, then it is in both Voronoi cells.

*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, J1-1693. Email address: sergio.cabello@mf.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.

27 In this paper we consider computational aspects of the *inverse* Voronoi problem in trees. This
 28 means that we are given a collection of candidate Voronoi cells in a tree, and we would like to
 29 decide whether they form indeed a Voronoi diagram. Let us describe the problem more formally.

30 GRAPHIC INVERSE VORONOI IN TREES

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

33 Question: Are there sites $s_1, \dots, s_k \in V(T)$ such that $\text{cell}_T(s_i, \{s_1, \dots, s_k\}) = U_i$ for each
 34 $i \in \{1, \dots, k\}$? When the answer is positive, provide a solution: sites $s_1, \dots, s_k \in V(T)$
 35 that certify the positive answer.

36 The *inverse Voronoi* problem can be considered also in arbitrary graphs and metric spaces. In
 37 the accompanying paper [1], we provide NP-hardness and $W[1]$ -hardness for several different
 38 scenarios. The problem is related to questions in pattern recognition; we refer to the discussion
 39 therein. Most notably for our work, we use the framework of parameterized complexity to show
 40 that, assuming the Exponential Time Hypothesis (ETH), the inverse Voronoi problem cannot be
 41 solved for graphs G of pathwidth $p(G)$ in time $f(p(G))|V(G)|^{o(p(G))}$, for any computable function
 42 f . This result justifies considering trees as a special case.

43 **Our results.** One has to be careful with the size of the description of the input because the size of
 44 the Voronoi diagram may be quadratic in the size of the tree. For example, in a star with $2n$ leaves
 45 and sites in n of the leaves, each Voronoi cell has size $\Theta(n)$, and thus an explicit description of the
 46 Voronoi diagram has size $\Theta(n^2)$. Motivated by this, we define the **description size** of an instance
 47 $I = (T, (U_1, \dots, U_k))$ for the GRAPHIC INVERSE VORONOI IN TREES to be $N = N(I) = |V(T)| + \sum_i |U_i|$.
 48 We use n for the number of vertices in the tree T , which is potentially smaller than N .

49 We show that the problem GRAPHIC INVERSE VORONOI IN TREES can be solved in $O(N + n \log^2 n)$
 50 time for arbitrary trees. We also show a lower bound of $\Omega(n \log n)$ in the algebraic computation
 51 tree model.

52 One may be tempted to think that the problem is easy for trees. Our near-linear time algorithm
 53 for arbitrary trees is far from trivial. Of course we cannot exclude the existence of a simpler
 54 algorithm running in near-linear time, but we do think that the problem is more complex than it
 55 may seem at first glance. Figure 1 may help understanding that the interaction between different
 56 Voronoi cells may be more complex than it seems.

57 To obtain our algorithm, we consider the following more general problem, where the input
 58 also specifies, for each Voronoi cell, a subset of vertices where the site has to be placed.

59 GENERALIZED GRAPHIC INVERSE VORONOI IN TREES

60 Input: (T, \mathbb{U}) , where T is a tree with positive edge-lengths and $\mathbb{U} = ((U_1, S_1), \dots, (U_k, S_k))$
 61 is a sequence of pairs of subsets of vertices of G .

62 Question: are there sites $s_1, \dots, s_k \in V(T)$ such that $s_i \in S_i$ and $U_i = \text{cell}_T(s_i, \{s_1, \dots, s_k\})$
 63 for each $i \in \{1, \dots, k\}$? When the answer is positive, provide a solution: sites
 64 $s_1, \dots, s_k \in V(T)$ that certify the positive answer.

65 Following the analogy with GRAPHIC INVERSE VORONOI IN TREES, we define the **description size**
 66 of an instance $I = (T, ((U_1, S_1), \dots, (U_k, S_k)))$ to be $N(I) = |V(T)| + \sum_i |U_i| + \sum_i |S_i|$.

67 Clearly, the problem GRAPHIC INVERSE VORONOI IN TREES can be reduced to the problem
 68 GENERALIZED GRAPHIC INVERSE VORONOI IN TREES by taking $S_i = U_i$ for all $i \in \{1, \dots, k\}$. This
 69 transformation can be done in linear time (in the size of the instance). Thus, for the rest of the
 70 paper our algorithms will be for the problem GENERALIZED GRAPHIC INVERSE VORONOI IN TREES.
 71 (The lower bound holds for the original problem.)

72 In our solution we first make a reduction to the same problem in which Voronoi cells are
 73 disjoint, and then we make another transformation to an instance having maximum degree 3.

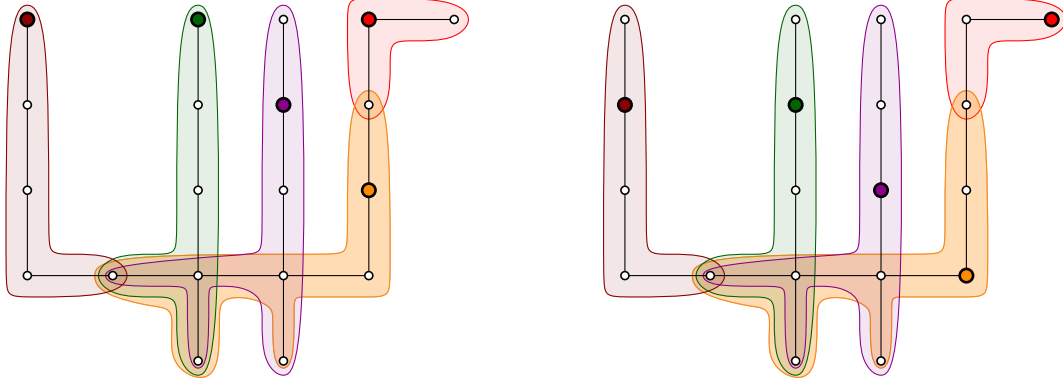


Figure 1: An instance with two solutions. The edges have unit length and the larger, filled dots represent the sites.

74 Finally, we employ a bottom-up dynamic programming procedure that, to achieve near-linear
 75 time, merges the information from the subproblems in time almost proportional to the smallest of
 76 the subproblems. For this, we employ dynamic binary search trees to manipulate sets of intervals.

77 **Roadmap.** In Section 2 we provide some basic tools. In Section 3 we show how to reduce the
 78 problem to a special instance where the candidate Voronoi cells are disjoint and the tree has
 79 maximum degree 3. In Section 4 we describe how to solve the problem, after the transformation,
 80 using dynamic programming. In Section 5 we provide a lower bound.

81 2 Basics

82 For a positive integer k we use the notation $[k] = \{1, \dots, k\}$.

83 In the following results we use T as the ground tree that defines the metric. Note that in the
 84 following claims it is important that T has positive edge-lengths. An alternative way to define
 85 cells is using strict inequalities. More precisely, for a set of sites Σ , the **open Voronoi cell** of each
 86 site $s \in \Sigma$ is then defined by

$$87 \quad \text{cell}^<(s, \Sigma) = \{x \in X \mid \forall s' \in \Sigma \setminus \{s\} : d(s, x) < d(s', x)\}.$$

88 In this case, the cells are disjoint but they do not necessarily form a partition of $V(T)$. The
 89 following two lemmas are straightforward folklore and we omit their proofs.

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

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

90 **Lemma 2.** *For each set Σ of sites, each site $s \in \Sigma$, and each vertex $v \in \text{cell}(s, \Sigma)$, the path in T from
 91 s to v is contained in $T[\text{cell}(s, \Sigma)]$, the subgraph of T induced by $\text{cell}(s, \Sigma)$. The same statement is
 92 true for $\text{cell}^<(s, \Sigma)$.*

93 A consequence of this Lemma is that the shortest path from s to $v \in \text{cell}(s, \Sigma) \setminus \text{cell}^<(s, \Sigma)$ has
 94 a part with vertices inside $\text{cell}^<(s, \Sigma)$ followed by a part with vertices of $\text{cell}(s, \Sigma) \setminus \text{cell}^<(s, \Sigma)$.

95 **Lemma 3.** *Given an instance for the problem GRAPHIC INVERSE VORONOI IN TREES or the GENER-
 96 ALIZED GRAPHIC INVERSE VORONOI IN TREES, and a candidate solution s_1, \dots, s_k , we can check in
 97 $O(N)$ time whether s_1, \dots, s_k is indeed a solution.*




Figure 2: Construction of T_a (left) and the directed acyclic graph D_a (right).

98 *Proof.* Let T be the underlying tree defining the instance. We add a new vertex a (called the *apex*)
 99 to T and connect it to each candidate site s_1, \dots, s_k with edges of the same positive length. See
 100 the left drawing in Figure 2. The resulting graph, denoted by T_a , has treewidth 2, and thus we
 101 can compute shortest paths from a to all vertices in linear time [3]. Let $d_a[v]$ be the distance in
 102 T_a from a to v .

103 Next we build a digraph D_a describing the shortest paths from a to all other vertices. The
 104 vertex set of D_a is $V(T) \cup \{a\} = V(T_a)$. For each arc $u \rightarrow v$, where $uv \in E(T_a)$, we add $u \rightarrow v$
 105 to D_a if and only if $d_a[v] = d_a[u] + \lambda(uv)$. With this we obtain a directed acyclic graph D_a that
 106 contains *all* shortest paths from a to every $v \in V(T)$ and, moreover, each directed path in D_a is
 107 indeed a shortest path in T_a . See Figure 2 right.

108 Now we label each vertex v with the indices i of those sites s_i , whose Voronoi cells contain v ,
 109 as follows. We start setting $L(s_i) = \{i\}$ for each site s_i . Then we consider the vertices $v \in V(T)$ in
 110 topological order with respect to D_a . For each vertex v , we set $L(v)$ to be the union of $L(u)$, where
 111 u iterates over the vertices of $V(T)$ with arcs in D pointing to v . It is easy to see by induction that
 112 $L(v) = \{i \in [k] \mid v \in \text{cell}_T(s_i, \{s_1, \dots, s_k\})\}$. During the process we keep a counter for $\sum_v |L(v)|$,
 113 and if at some moment we detect that the counter exceeds N , we stop and report that s_1, \dots, s_k is
 114 not a solution. Otherwise, we finish the process when we computed the sets $L(v)$.

115 Now we compute the sets $V_i = \{v \in V(T) \mid s_i \in L(v)\}$ for $i = 1, \dots, k$. This is done iterating
 116 over the vertices $v \in V(T)$ and adding v to each site of $L(v)$. This takes $O(N + \sum_v |L(v)|) = O(N)$
 117 time. Note that $V_i = \text{cell}_T(s_i, \{s_1, \dots, s_k\})$. It remains to check that $U_i = V_i$ for all $i \in [k]$. For this
 118 we add flags to $V(T)$ that are initially set to false. Then, for each $i \in [k]$, we do the following:
 119 check that $|U_i| = |V_i|$, iterate over the vertices of U_i setting the flags to true, iterate over the
 120 vertices of V_i checking that the flags are true, iterate over the vertices of U_i setting the flags back
 121 to false. The procedure takes $O(N + \sum_v |L(v)|) = O(N)$ time and, if all the checks were correct,
 122 we have $U_i = V_i = \text{cell}_T(s_i, \{s_1, \dots, s_k\})$ for all $i \in [k]$. \square

123 3 Arbitrary trees – Transforming to nicer instances

124 In this section we provide a transformation to reduce the problem GENERALIZED GRAPHIC INVERSE
 125 VORONOI IN TREES to instances where the tree has maximum degree 3 and the candidate Voronoi
 126 regions are disjoint. First we show how to transform it into disjoint Voronoi regions, and then
 127 we handle the degree. In our description, we first discuss the transformation without paying
 128 attention to its efficiency. At the end of the section we discuss how the transformation can be
 129 done in linear time.

3.1 Transforming to disjoint cells

In this section we explain how to decrease the overlap between different Voronoi regions. The procedure is iterative: we consider one edge of the tree at a time and transform the instance. When there are no edges to process, we can conclude that the original instance has no solution or we can find a solution to the original instance.

Consider an instance $I = (T, ((U_1, S_1), \dots, (U_k, S_k)))$ for the problem GENERALIZED GRAPHIC INVERSE VORONOI IN TREES. See Figure 1 for an example of such an instance.

For each index $i \in [k]$ we define

$$W_i = U_i \setminus \bigcup_{j \neq i} U_j,$$

$$E_i = \{uv \in E(T) \mid u \in W_i, v \in U_i \setminus W_i\}.$$

The intuition is that each W_i should be the open Voronoi cell defined by the (unknown) site s_i , that is, the vertices of T with s_i as unique closest site; see Lemma 1. Each E_i is then the set of edges with one vertex in W_i and another vertex in $U_i \cap U_j$ for some $j \neq i$. The following result is easy to prove using Lemma 2.

Lemma 4. *Supposing that there is a solution to GENERALIZED GRAPHIC INVERSE VORONOI IN TREES with input I , the following hold.*

(a) *Each set U_i ($i \in [k]$) and each set W_i ($i \in [k]$) induces a connected subgraph of T .*

(a) *If two sets U_i and U_j ($i \neq j$) intersect, then $E_i \neq \emptyset$ and $E_j \neq \emptyset$.*

Proof. Consider a solution s_1, \dots, s_k to GENERALIZED GRAPHIC INVERSE VORONOI IN TREES with input I , and define $\Sigma = \{s_1, \dots, s_k\}$. This means that, for each $i \in [k]$, we have $s_i \in S_i$ and $U_i = \text{cell}_T(s_i, \Sigma)$. Note that because of Lemma 1, we have

$$\forall i \in [k]: W_i = U_i \setminus \bigcup_{j \neq i} U_j = \text{cell}_T(s_i, \Sigma) \setminus \bigcup_{j \neq i} \text{cell}_T(s_j, \Sigma) = \text{cell}_T^<(s_i, \Sigma).$$

If there are distinct indices $i, j \in [k]$ such that U_i and U_j intersect, then $W_i \subsetneq U_i$. Because of Lemma 1, we have $s_i \in W_i$, and therefore W_i is nonempty. Because of Lemma 2, the sets U_i and W_i induce subtrees of T . Since $W_i \subsetneq U_i$, it follows that T has some edge from W_i to $U_i \setminus W_i$, and therefore E_i is nonempty. \square

As a preprocessing step, we replace S_i by $S_i \cap W_i$ for each $i \in [k]$. Since a site cannot belong to two Voronoi regions, this replacement does not reduce the set of feasible solutions for I . To simplify notation, we keep using I for the new instance. We check that, for each $i \in [k]$, the set S_i is nonempty and the sets U_i and W_i induce a connected subgraph of T . If any of those checks fail, we correctly report that there is no solution to I .

If the sets U_1, \dots, U_k are pairwise disjoint, we do not need to do anything. If at least two of them overlap but the sets E_1, \dots, E_k are empty, then Lemma 4 implies that there is no solution. In the remaining case some E_i is nonempty, and we transform the instance as follows.

In the transformations we will need “short” edges. To quantify this, we introduce the **resolution** $\text{res}(I)$ of an instance I , defined by

$$\text{res}(I) = \min(\mathbb{R}_{>0} \cap \{d_T(s_i, u) - d_T(s_j, u) \mid u \in U_i \cap U_j, s_i \in S_i, s_j \in S_j, i, j \in [k]\}).$$

Here we take the convention that $\min(\emptyset) = +\infty$. From the definition we have the following property:

$$\forall i, j \in [k], u \in U_i \cap U_j, s_i \in S_i, s_j \in S_j: |d_T(s_i, u) - d_T(s_j, u)| < \text{res}(I) \implies d_T(s_i, u) = d_T(s_j, u). \quad (1)$$




Figure 3: The transformation from the instance I in Figure 1 to I' for two different choices of the set U_1 and $xy \in E_1$. The new vertex y' appearing because of the subdivision is marked with a square. The “shorter” edges in the drawing have length ε ; all other edges have unit length.

172 Consider any value $\varepsilon > 0$. Fix any index $i \in [k]$ such that $E_i \neq \emptyset$ and consider an edge $xy \in E_i$
173 with $x \in W_i$ and $y \in U_i \setminus W_i$. By renaming the sets, if needed, we assume henceforth that $i = 1$,
174 that is, $E_1 \neq \emptyset$, $x \in W_1$ and $y \in U_1 \setminus W_1$. We build a tree T' with edge-lengths λ' and a new set U'_1
175 as follows. We obtain T' from T by subdividing xy with a new vertex y' . We define U'_1 to be the
176 subset of vertices of U_1 that belong to the component of $T - y$ that contains x , and then we also
177 add y' into U'_1 . Note that $u \in U_1$ belongs to U'_1 if and only if $d_T(u, x) < d_T(u, y)$. In particular,
178 $y \notin U'_1$. Finally, we set the edge-lengths $\lambda'(xy') = \lambda(xy)$ and $\lambda'(y'y) = \varepsilon$, and the remaining
179 edges have the same length as in T . This completes the description of the transformation. Note
180 that T' is just a subdivision of T and, effectively, the edge xy became a 2-edge path $xy'y$ that is
181 longer by ε . All distances in T' are larger or equal than in T , and the difference is at most ε .

182 Let I' be the new instance, where we use T' , λ' and U'_1 , instead of T , λ and U_1 , respectively.
183 (We leave U_i unchanged for each $i \in [k] \setminus \{1\}$ and S_i unchanged for each $i \in [k]$.) See Figure 3
184 for two examples of this transformation and Figure 4 for a schematic view. We call I' the instance
185 obtained from I by **expanding the edge xy from E_1 by ε** . Note that y' is not a valid placement
186 for a site in I' , since $y' \notin S_1$.

187 Our definition of $\text{res}(I)$ is carefully chosen so that it does not decrease with the expansion of
188 an edge. That is, $\text{res}(I') \geq \text{res}(I)$. (This property is exploited in the proof of Lemma 8.) This is an
189 important but subtle point needed to achieve efficiency. It will permit that all the short edges that
190 are introduced during the transformations have the same small length ε , and we will be able to
191 treat ε symbolically.

192 The next two lemmas show the relation between solutions to the instances I and I' .

193 **Lemma 5.** *Suppose that $\varepsilon > 0$. If Σ is a solution to GENERALIZED GRAPHIC INVERSE VORONOI IN*
194 *TREES with input I , then Σ is also a solution to GENERALIZED GRAPHIC INVERSE VORONOI IN TREES*
195 *with input I' .*

196 *Proof.* We first introduce some notation. Let V_x be the vertex set of the component of $T' - y'$ that
197 contains x and let V_y be the vertex set of the component of $T' - y'$ that contains y . See Figure 4.
198 Note that $x \in V_x$ and $y \in V_y$, while y' is neither in V_x nor in V_y . From the definition of U'_1 , we
199 have $U'_1 = \{y'\} \cup (V_x \cap U_1)$ and $U_1 \setminus U'_1 = V_y \cap U_1$.

200 We have the following easy relations between distances in T and T' ; we will use them often




Figure 4: Notation in the proof of Lemma 5.

201 without explicit reference.

$$\begin{aligned}
202 \quad & \forall u, v \in V_x : d_{T'}(u, v) = d_T(u, v) \\
203 \quad & \forall u, v \in V_y : d_{T'}(u, v) = d_T(u, v) \\
204 \quad & \forall u \in V_x, v \in V_y : d_{T'}(u, v) = d_T(u, v) + \varepsilon \\
205 \quad & \forall u \in V_x : d_{T'}(u, y') = d_T(u, y) \\
206 \quad & \forall u \in V_y : d_{T'}(u, y') = d_T(u, y) + \varepsilon.
\end{aligned}$$

208 Consider a solution s_1, \dots, s_k to GENERALIZED GRAPHIC INVERSE VORONOI IN TREES with input
209 I , and define $\Sigma = \{s_1, \dots, s_k\}$. This means that, for all $i \in [k]$, we have $s_i \in S_i$ and $U_i = \text{cell}_T(s_i, \Sigma)$.
210 Our objective is to show that $U'_1 = \text{cell}_{T'}(s_1, \Sigma)$ and $U_i = \text{cell}_{T'}(s_i, \Sigma)$ for all $i \in [k] \setminus \{1\}$.

211 Since $U_i = \text{cell}_T(s_i, \Sigma)$ for all $i \in [k]$, Lemma 1 implies that $W_1 = \text{cell}_T^{\leq}(s_1, \Sigma)$ and $s_1 \in W_1$.
212 Since $x \in W_1$, $y \notin W_1$, and W_1 induces a connected subgraph of T because of Lemma 4(a), the
213 set W_1 is contained in V_x . Since $W_1 \subseteq V_x$ and $W_1 \subseteq U_1$, we have $W_1 \subseteq V_x \cap U_1$ and we conclude
214 that $W_1 \subseteq U'_1$. Furthermore, because $s_1 \in \text{cell}_T^{\leq}(s_1, \Sigma) = W_1$ and $W_1 \subseteq V_x$, we obtain that $s_1 \in V_x$.

215 For each $i \in [k] \setminus \{1\}$, we have $x \notin U_i$ because $x \in W_1$, and Lemma 4(a) implies that the set
216 $U_i = \text{cell}_T(s_i, \Sigma)$ is fully contained either in V_x or in V_y .

217 Consider any index $\ell \in [k] \setminus \{1\}$ with the property that $y \in U_1 \cap U_\ell$. Since U_ℓ contains y , it
218 cannot be that $U_\ell \subseteq V_x$, and therefore $U_\ell \subseteq V_y$. In particular, $s_\ell \in V_y$.

219 We first note that the sets U'_1, U_2, \dots, U_k cover $V(T')$. Indeed, since $y \in U_1 \cap U_\ell$, the sites s_1 and
220 s_ℓ are closest sites to y in T , and using that $s_1 \in V_x$ and $s_\ell \in V_y$, we obtain that $U_1 \setminus U'_1$ is contained
221 in U_ℓ . Further, since U_1, \dots, U_k cover $V(T)$, $y' \in U'_1$ by construction, and $V(T') = V(T) \cup \{y'\}$,
222 we conclude that indeed U'_1, U_2, \dots, U_k cover $V(T')$.

223 Next, we make the following two claims.

224 **Claim 5.1.** $y' \in \text{cell}_{T'}(s_1, \Sigma)$ and $y' \notin \text{cell}_{T'}(s_i, \Sigma)$ for any $i \in [k] \setminus \{1\}$.

225 *Proof.* Fix any index $i \in [k] \setminus \{1\}$. Consider first the case when $s_i \in V_x$. In this case the path from
226 s_i to y' passes through x , which is a vertex in $\text{cell}_T^{\leq}(s_1, \Sigma)$. It follows that $d_T(s_1, x) < d_T(s_i, x)$,
227 which implies

$$228 \quad d_{T'}(s_1, y') = d_T(s_1, y) < d_T(s_i, y) = d_{T'}(s_i, y').$$

229 Consider now the case when $s_i \in V_y$. Because $y \in U_1 = \text{cell}(s_1, \Sigma)$, we have $d_T(s_1, y) \leq$
230 $d_T(s_i, y)$ and we conclude that

$$231 \quad d_{T'}(s_i, y') = d_T(s_i, y) + \varepsilon \geq d_T(s_1, y) + \varepsilon = d_{T'}(s_1, y') + \varepsilon > d_{T'}(s_1, y').$$

232 In each case we get $d_{T'}(s_1, y') < d_{T'}(s_i, y')$, and the claim follows. \square

233 **Claim 5.2.** $y \notin \text{cell}_{T'}(s_1, \Sigma)$.

234 *Proof.* Since y belongs to $U_1 \cap U_\ell$, we have $d_T(s_1, y) = d_T(s_\ell, y)$. Using that U_ℓ is contained in
235 V_y , and thus $s_\ell \in V_y$, we have

$$236 \quad d_{T'}(s_\ell, y) = d_T(s_\ell, y) = d_T(s_1, y) = d_{T'}(s_1, y) - \varepsilon < d_{T'}(s_1, y).$$

237 We conclude that y is not an element of $\text{cell}_{T'}(s_1, \Sigma)$. \square

238 Claims 5.1 and 5.2 imply that y' belongs *only* to the Voronoi region $\text{cell}_{T'}(s_1, \Sigma)$ and y does not
 239 belong to $\text{cell}_{T'}(s_1, \Sigma)$. This means that each vertex of V_x belongs only to some regions $\text{cell}_{T'}(s_i, \Sigma)$
 240 with $s_i \in V_x$ and each vertex of V_y belongs to some regions $\text{cell}_{T'}(s_i, \Sigma)$ with $s_i \in V_y$. That is,
 241 it cannot be that some vertex $u \in V_x$ belongs to $\text{cell}_{T'}(s_i, \Sigma)$ with $s_i \in V_y$ and it cannot be that
 242 some vertex $u \in V_y$ belongs to $\text{cell}_{T'}(s_i, \Sigma)$ with $s_i \in V_x$. Effectively, this means that y' splits the
 243 Voronoi diagram $\mathbb{V}_{T'}(\Sigma)$ into the part within $T'[V_x]$ and the part within $T'[V_y]$, with the gluing
 244 property that $y' \in \text{cell}_{T'}(s_1, \Sigma)$. Since $U'_1 \setminus \{y'\} = U_1 \cap V_x$ and the distances within $T'[V_x]$ and
 245 within $T'[V_y]$ are the same as in T , the result follows. \square

246 The converse property is more complicated. We need ε to be small enough and we also have
 247 to assume that I has a solution. It is this tiny technicality that makes the reduction nontrivial.

248 **Lemma 6.** *Suppose that $0 < \varepsilon < \text{res}(I)$ and the answer to GENERALIZED GRAPHIC INVERSE VORONOI
 249 IN TREES with input I is “yes”. If Σ' is a solution to GENERALIZED GRAPHIC INVERSE VORONOI
 250 IN TREES with input I' , then Σ' is also a solution to GENERALIZED GRAPHIC INVERSE VORONOI IN
 251 TREES with input I .*

252 *Proof.* When the instance I has *some* solution, then the properties discussed in Lemmas 4 and 5
 253 hold. We keep using the notation and the properties established earlier. In particular, each set U_i
 254 ($i \in [k] \setminus \{1\}$) is contained either in V_x or in V_y , and we have $W_1 \subseteq U'_1 \subseteq V_x \cup \{y'\}$.

255 Consider a solution s_1, \dots, s_k to GENERALIZED GRAPHIC INVERSE VORONOI IN TREES with input
 256 I' , and set $\Sigma = \{s_1, \dots, s_k\}$. This means that $U'_1 = \text{cell}_{T'}(s_1, \Sigma)$ and, for all $i \in [k] \setminus \{1\}$, we have
 257 $U_i = \text{cell}_{T'}(s_i, \Sigma)$. We have to show that, for all $i \in [k]$, we have $U_i = \text{cell}_T(s_i, \Sigma)$, which implies
 258 that Σ is a solution to input I .

259 Like before, we split the proof into claims that show that Σ is a solution to GENERALIZED
 260 GRAPHIC INVERSE VORONOI IN TREES with input I . We start with an auxiliary property that plays
 261 a key role.

262 **Claim 6.1.** *For each $i \in [k]$, we have $y \in U_i$ if and only if $y \in \text{cell}_T(s_i, \Sigma)$.*

263 *Proof.* Suppose first that $y \in U_i$ and $i \neq 1$. Then $U_i \subseteq V_y$. Since $y \in U_i = \text{cell}_{T'}(s_i, \Sigma)$ and
 264 $y \notin U'_1 = \text{cell}_{T'}(s_1, \Sigma)$, we have

$$265 \quad d_T(s_i, y) = d_{T'}(s_i, y) < d_{T'}(s_1, y) = d_T(s_1, y) + \varepsilon. \quad (2)$$

267 Since $y' \notin U_i = \text{cell}_{T'}(s_i, \Sigma)$ and $y' \in U'_1 = \text{cell}_{T'}(s_1, \Sigma)$, we have

$$268 \quad d_T(s_1, y) = d_{T'}(s_1, y') < d_{T'}(s_i, y') = d_T(s_i, y) + \varepsilon. \quad (3)$$

270 Combining (2) and (3) we get

$$271 \quad |d_T(s_i, y) - d_T(s_1, y)| < \varepsilon < \text{res}(I).$$

272 From property (1) and since $y \in U_1 \cap U_i$, we conclude that $d_T(s_1, y) = d_T(s_i, y)$. For each $s_j \in V_y$
 273 we use that $y \in U_i = \text{cell}_{T'}(s_i, \Sigma)$ to obtain

$$274 \quad d_T(s_j, y) = d_{T'}(s_j, y) \geq d_{T'}(s_i, y) = d_T(s_i, y).$$

275 For each $s_j \in V_x$ we use that the path from s_j to y goes through $x \in U'_1 = \text{cell}_{T'}(s_1, \Sigma)$ to obtain

$$276 \quad d_T(s_j, y) \geq d_T(s_1, y) = d_T(s_i, y).$$

277 We conclude that for each $j \in [k]$ we have $d_T(s_j, y) \geq d_T(s_i, y)$, and therefore $y \in \text{cell}_T(s_i, \Sigma)$.

278 Since $d_T(s_1, y) = d_T(s_i, y)$ whenever $y \in U_1 \cap U_i$, and $y \in U_\ell$ for some $\ell \in [k] \setminus \{1\}$, we also
 279 obtain $y \in \text{cell}_T(s_1, \Sigma)$. With this we have shown one direction of the implication.

280 To show the other implication, consider some index $i \in [k]$ such that $y \in \text{cell}_T(s_i, \Sigma)$. If $i = 1$,
 281 then $y \in U_1$ by construction, and the implication holds. So we consider the case when $i \neq 1$.
 282 First we show that it cannot be that $s_i \in V_x$. Assume, for the sake of reaching a contradiction,
 283 that $s_i \in V_x$. Because of the implication left-to-right that we showed, we have $y \in \text{cell}_T(s_1, \Sigma)$.
 284 Since we have $y \in \text{cell}_T(s_i, T)$ and $y \in \text{cell}_T(s_1, \Sigma)$, we obtain $d_T(s_i, y) = d_T(s_1, y)$. Because
 285 $s_1, s_i \in V_x$, we obtain $d_T(s_i, x) = d_T(s_1, x)$ and therefore $d_{T'}(s_i, x) = d_{T'}(s_1, x)$. Further, since
 286 $x \in U'_1 = \text{cell}'_T(s_1, \Sigma)$, we get $x \in \text{cell}_{T'}(s_i, \Sigma) = U_i$, which implies $x \notin W_1$. We conclude that it
 287 must be $s_i \notin V_x$, and thus $s_i \in V_y$.

288 Take an index $\ell \in [k] \setminus \{1\}$ such that $y \in U_\ell$. Such an index exists because $y \notin W_1$. We have
 289 $U_\ell \subseteq V_y$ and thus $s_\ell \in V_y$. Because of the implication left-to-right that we showed, we have $y \in$
 290 $\text{cell}_T(s_\ell, \Sigma)$. Since we have $y \in \text{cell}_T(s_i, T)$ and $y \in \text{cell}_T(s_\ell, \Sigma)$, we obtain $d_T(s_i, y) = d_T(s_\ell, y)$.
 291 Because $s_i, s_\ell \in V_y$ we then have

$$292 \quad d_{T'}(s_i, y) = d_T(s_i, y) = d_T(s_\ell, y) = d_{T'}(s_\ell, y).$$

293 Since $d_{T'}(s_i, y) = d_{T'}(s_\ell, y)$ and $y \in U_\ell = \text{cell}_{T'}(s_\ell, \Sigma)$, we conclude that $y \in \text{cell}_{T'}(s_i, \Sigma) = U_i$. \square

294 **Claim 6.2.** $x \in \text{cell}_T(s_1, \Sigma)$ and $x \notin \text{cell}_T(s_i, \Sigma)$ for any $i \in [k] \setminus \{1\}$.

295 *Proof.* Since $x \in U'_1 = \text{cell}'_T(s_1, \Sigma)$ and $x \notin U_i = \text{cell}_{T'}(s_i, \Sigma)$ for any $i \in [k] \setminus \{1\}$, we have

$$296 \quad \forall i \in [k] \setminus \{1\} : d_{T'}(s_1, x) < d_{T'}(s_i, x).$$

297 We then have

$$298 \quad \forall s_i \in V_x, s_i \neq s_1 : d_T(s_1, x) = d_{T'}(s_1, x) < d_{T'}(s_i, x) = d_T(s_i, x). \quad (4)$$

300 For each $s_i \in V_y$, note that the path from s_i to x passes through y , and $y \in \text{cell}_T(s_1, \Sigma)$ because
 301 of Claim 6.1. Using that $s_1 \in V_x$, we have

$$302 \quad \forall s_i \in V_y : d_T(s_1, x) < d_T(s_i, x). \quad (5)$$

304 Combining (4) and (5), the claim follows. \square

305 **Claim 6.3.** For each $u \in V_y$, we have $u \in U_1$ if and only if $u \in \text{cell}_T(s_1, \Sigma)$.

306 *Proof.* Consider some solution s_1^*, \dots, s_k^* to GENERALIZED GRAPHIC INVERSE VORONOI IN TREES
 307 with input I , and set $\Sigma^* = \{s_1^*, \dots, s_k^*\}$. This means that $U_i = \text{cell}_T(s_i^*, \Sigma^*)$ for each $i \in [k]$. We also
 308 fix an index $\ell \in [k] \setminus \{1\}$ such that $y \in U_\ell \cap U_1$. Recall that $U_\ell \subseteq V_y$ because $x \notin U_\ell$, and $W_1 \subseteq V_x$
 309 because $x \in W_1$ and $y \notin W_1$. Using Claim 6.1 and using that Σ^* is a solution to I we have

$$310 \quad d_T(s_1, y) = d_T(s_\ell, y) \quad \text{and} \quad d_T(s_1^*, y) = d_T(s_\ell^*, y). \quad (6)$$

312 Consider some $u \in U_1 \cap V_y$. We will show that $u \in \text{cell}_T(s_1, \Sigma)$. Consider the subtree \tilde{T} defined
 313 by the paths connecting the vertices $s_1, s_1^*, s_\ell, s_\ell^*, u$. See Figure 5. The path from u to s_1^* attaches to
 314 the path from s_ℓ^* to y at the vertex y . Indeed, if it attaches at another vertex $a \neq y$, then we would
 315 have $d_T(s_\ell^*, a) < d_T(s_1^*, a)$ because of (6), which would imply $d_T(s_\ell^*, u) < d_T(s_1^*, u)$, contradicting
 316 the assumption that $u \in \text{cell}_T(s_1^*, \Sigma^*) = U_1$. Because W_ℓ does not contain y and W_ℓ is a connected
 317 subgraph of T (applying Lemma 2), W_ℓ is contained in a connected component of $T - y$. Further
 318 since W_ℓ contains s_ℓ and s_ℓ^* , and we replaced S_ℓ with $S_\ell \cap W_\ell$ in the preprocessing step¹, s_ℓ and s_ℓ^*
 319 are in the same component of $T - y$. Therefore, the (u, s_1) -path attaches to the (s_ℓ, y) -path at the
 320 vertex y .

¹ Without the replacement S_ℓ with $S_\ell \cap W_\ell$, the lemma is actually not true because it can happen that $s_\ell \in U_1 \cap U_\ell$.
 Indeed, we could have $s_\ell \in S_\ell \cap U_\ell \cap U_1$, which is not a valid placement in I but would be a valid placement in I' .




Figure 5: Situation in the proof of Claim 6.3.

321 Since each path from s_1, s_1^*, s_ℓ and s_ℓ^* to u passes through y , from (6) we get

$$322 \quad d_T(s_1, u) = d_T(s_\ell, u) \quad \text{and} \quad d_T(s_1^*, u) = d_T(s_\ell^*, u). \quad (7)$$

324 Together with $u \in U_1 = \text{cell}_T(s_1^*, \Sigma^*)$ we conclude that $u \in \text{cell}_T(s_\ell^*, \Sigma^*) = U_\ell$. Since $u \in U_\ell =$
325 $\text{cell}_{T'}(s_\ell, \Sigma)$ we have

$$326 \quad \forall s_j \in V_y : \quad d_T(s_1, u) = d_T(s_\ell, u) \leq d_T(s_j, u).$$

327 Together with the fact that each $s_j \in V_x$ is no closer to u than s_1 because $x \in \text{cell}_T(s_1, \Sigma)$, we
328 conclude that $u \in \text{cell}_T(s_1, \Sigma)$. This finishes the left-to-right direction of the implication.

329 Consider now a vertex $u \in V_y \cap \text{cell}_T(s_1, \Sigma)$. Since y is on the path from s_1 to u , we obtain
330 from (6) that $d_T(s_\ell, u) \leq d_T(s_1, u)$, and therefore $u \in \text{cell}_T(s_\ell, \Sigma)$. Because $u \in V_y$, $d_{T'}(s_\ell, u) =$
331 $d_T(s_\ell, u)$, and distances in T' can only be larger than in T , we have $u \in \text{cell}_{T'}(s_\ell, \Sigma) = U_\ell =$
332 $\text{cell}_T(s_\ell^*, \Sigma^*)$. This means that

$$333 \quad \forall i \in [k] : \quad d_T(s_\ell^*, u) \leq d_T(s_i^*, u). \quad (8)$$

335 Since $u \in \text{cell}_T(s_1, \Sigma)$, $u \in \text{cell}_T(s_\ell, \Sigma)$ and $d_T(s_1, y) = d_T(s_\ell, y)$, the vertex y is on the path from
336 s_ℓ to u . Note that the vertices s_ℓ and s_ℓ^* must be contained in the same component of $T - y$ because
337 the (s_ℓ, s_ℓ^*) -path must be contained W_ℓ (Lemma 2 and footnote 1), but $y \notin W_\ell$. This implies that y
338 is also on the path from s_ℓ^* to u . Since y is also on the path from s_1^* to u , we get from (6) and (8)
339 that

$$340 \quad \forall i \in [k] : \quad d_T(s_1^*, u) = d_T(s_\ell^*, u) \leq d_T(s_i^*, u).$$

341 It follows that $u \in \text{cell}_T(s_1^*, \Sigma^*) = U_1$. This finishes the proof of the claim. \square

342 We are now ready to prove Lemma 6: for all $i \in [k]$ we have $U_i = \text{cell}_T(s_i, \Sigma)$. Consider first
343 the case $i = 1$. Because of Claim 6.3 we have $V_y \cap U_1 = V_y \cap \text{cell}_T(s_1, \Sigma)$. It remains to show
344 that $U_1' = V_x \cap U_1 = V_x \cap \text{cell}_T(s_1, \Sigma)$. Consider any vertex $u \in U_1'$. Because of Claim 6.2 we have
345 $x \in \text{cell}_T^<(s_1, \Sigma)$, and therefore $u \in V_x$ implies

$$346 \quad \forall s_j \in V_y : \quad d_T(s_1, u) < d_T(s_j, u).$$

347 On the other hand, since $u \in U_1' = \text{cell}_{T'}(s_1, \Sigma)$ we have

$$348 \quad \forall s_j \in V_x : \quad d_T(s_1, u) = d_{T'}(s_1, u) \leq d_{T'}(s_j, u) = d_T(s_j, u).$$

349 We conclude that $d_T(s_1, u) \leq d_T(s_j, u)$ for all $s_j \in \Sigma$, and therefore $u \in \text{cell}_T(s_1, \Sigma)$. This shows
350 that $U_1' \subseteq V_x \cap \text{cell}_T(s_1, \Sigma)$. To show the inclusion in the other direction, consider any $u \in$
351 $V_x \cap \text{cell}_T(s_1, \Sigma)$. We then have

$$352 \quad \forall s_j \in \Sigma : \quad d_{T'}(s_1, u) = d_T(s_1, u) \leq d_T(s_j, u) \leq d_{T'}(s_j, u),$$

353 which implies $u \in \text{cell}_{T'}(s_1, \Sigma) = U_1'$. This finishes the proof of $U_1 = \text{cell}_T(s_1, \Sigma)$, that is, the case
354 $i = 1$.



Figure 6: A similar transformation for arbitrary graphs does not work. On the right side we have the transformed instance with a feasible solution that does not correspond to a solution in the original setting.

355 Consider now the indices $i \in [k] \setminus \{1\}$ with $s_i \in V_y$. Recall that we have $U_i = \text{cell}_{T'}(s_i, \Sigma)$
 356 and $U_i \subseteq V_y$. Fix an index $\ell \in [k] \setminus \{1\}$ such that $y \in U_\ell \cap U_1$. Such an index exists because
 357 $y \notin W_1$. We must have $U_\ell \subseteq V_y$ because $x \notin U_\ell$, and thus $s_\ell \in V_y$. Because of Claim 6.1 we have
 358 $d_T(s_1, y) = d_T(s_\ell, y)$, and using that $x \in \text{cell}_T^<(s_1, \Sigma)$, implied by Claim 6.2, we get

$$359 \quad \forall u \in V_y, s_j \in V_x : d_T(s_\ell, u) \leq d_T(s_1, u) \leq d_T(s_j, u).$$

360 This implies that in T each vertex of V_y has at least one closest site (from Σ) that belongs to V_y .
 361 Therefore, for each $s_i \in V_y$, we have

$$362 \quad \text{cell}_T(s_i, \Sigma) = \text{cell}_{T[V_y]}(s_i, \Sigma \cap V_y).$$

363 A similar argument can be used for T' : no site in V_x is the closest site to any vertex of V_y and the
 364 closest site to y' is s_1 . Therefore, for each $s_i \in V_y$, we have

$$365 \quad \text{cell}_{T'}(s_i, \Sigma) = \text{cell}_{T'[V_y]}(s_i, \Sigma \cap V_y).$$

366 Noting that $T[V_y] = T'[V_y]$ we obtain, for each $s_i \in V_y$,

$$367 \quad U_i = \text{cell}_{T'}(s_i, \Sigma) = \text{cell}_{T'[V_y]}(s_i, \Sigma \cap V_y) = \text{cell}_{T[V_y]}(s_i, \Sigma \cap V_y) = \text{cell}_T(s_i, \Sigma).$$

368 It remains to consider the indices $i \in [k] \setminus \{1\}$ with $s_i \in V_x$. The approach is similar, and
 369 actually simpler because $x \in \text{cell}_T^<(s_1, T)$ implies that there is no influences from the sites V_y . (No
 370 care is needed for y' because it belongs to $\text{cell}_{T'}^<(s_1, \Sigma)$). Therefore, for each $s_i \in V_x \setminus \{s_1\}$,

$$371 \quad U_i = \text{cell}_{T'}(s_i, \Sigma) = \text{cell}_{T'[V_x]}(s_i, \Sigma \cap V_x) = \text{cell}_{T[V_x]}(s_i, \Sigma \cap V_x) = \text{cell}_T(s_i, \Sigma).$$

372 We have covered all the cases: $s_i = s_1$, $s_i \in V_y$, and $s_i \in V_x \setminus \{s_1\}$. This finishes the proof of
 373 the Lemma. \square

374 It is important to note that the transformation described above only works for trees. A similar
 375 transformation for arbitrary graphs may have feasible solutions that do not correspond to solutions
 376 in the original problem. See Figure 6 for a simple example.

377 Another important point is that we need the assumption that I had a solution. This means
 378 that, any solution Σ' we obtain after making a sequence of expansions, has to be tested in the
 379 original instance. However, if Σ' is not a valid solution in I , then I has no solution.

380 We are going to make a sequence of edge expansions. The replacement of S_i with $S_i \cap W_i$ (for
 381 $i \in [k]$) needs to be made only at the preprocessing step and it is important for correctness (see
 382 footnote 1). It is not needed later on because with each edge expansion the sets W_i (for $i \in [k]$)
 383 can only increase.

384 Consider an instance $I = (T, ((U_1, S_1), \dots, (U_k, S_k)))$. Set $I_0 = I$ and define, for $t \geq 1$,
 385 the instance I_t by transforming I_{t-1} using an expansion of some edge. For all expansions
 386 we use the same parameter ε . We finish the sequence when we obtain the first instance




Figure 7: The behavior of the reduction to obtain maximum degree 3. Left: part of an instance with a tree of arbitrary degrees. Right: result after the reduction for the left instance. The edges between different candidate Voronoi cells are shortened by δ .

387 $\tilde{I} = (\tilde{T}, ((\tilde{U}_1, \tilde{S}_1), \dots, (\tilde{U}_k, \tilde{S}_k)))$ such that the sets $\tilde{U}_1, \dots, \tilde{U}_k$ are pairwise disjoint. Note that
 388 this procedure stops because the number of pairs (i, j) with $U_i \cap U_j \neq \emptyset$ decreases with each
 389 expansion. This implies that the number of steps is at most $\binom{k}{2}$. In fact, the number of steps is
 390 even smaller.

391 **Lemma 7.** \tilde{I} is reached after at most $k - 1$ edge expansions.

392 *Proof.* We prove this by induction on k . There is nothing to show if $k = 1$. Otherwise, note that
 393 the sets U_i in V_x and those in V_y (respectively) give rise to two independent subproblems with k_x
 394 and k_y sites (respectively), where $k_x + k_y = k$. By induction, the number of edge expansions is at
 395 most $1 + (k_x - 1) + (k_y - 1) = k - 1$. \square

396 The next lemma shows that using the same parameter ε for all edge expansions is a correct
 397 choice. This is due to our careful definition of resolution $\text{res}(\cdot)$.

398 **Lemma 8.** Assume that $0 < \varepsilon < \text{res}(I)$ and the answer to GENERALIZED GRAPHIC INVERSE VORONOI
 399 IN TREES with input I is “yes”. Then Σ is a solution to GENERALIZED GRAPHIC INVERSE VORONOI IN
 400 TREES with input I if and only if Σ is also a solution to GENERALIZED GRAPHIC INVERSE VORONOI
 401 IN TREES with input \tilde{I} .

402 *Proof.* Note that, by construction, $\text{res}(I_{t-1}) \leq \text{res}(I_t)$ for all $t \geq 1$. Indeed, when we expand the
 403 edge xy inserting y' , then there is no set U_i that is on both sides of $T' - y'$. This means that for all
 404 the parameters s_i, s_j, u_i, u_j considered in the definition of $\text{res}(I_t)$ we have $d_{T'}(s_i, u) - d_{T'}(s_j, u) =$
 405 $d_T(s_i, u) - d_T(s_j, u)$. Therefore, $\varepsilon < \text{res}(I_t)$ for all t . The claim now follows easily from Lemmas 5
 406 and 6 by induction on t . \square

407 3.2 Transforming to maximum degree 3

408 Consider an instance $I = (T, ((U_1, S_1), \dots, (U_k, S_k)))$ for the problem GENERALIZED GRAPHIC
 409 INVERSE VORONOI IN TREES, where T is a tree and the sets U_1, \dots, U_k are pairwise disjoint. We
 410 assume that each U_i induces a connected subgraph in T . See Figure 7 for an example of such an
 411 instance viewed around a vertex of degree > 3 . We want to transform it into another instance
 412 $I' = (T', ((U'_1, S'_1), \dots, (U'_k, S'_k)))$ where the maximum degree of T' is 3, the sets U'_1, \dots, U'_k
 413 are pairwise disjoint, and a solution to I' corresponds to a solution of I .

414 In the transformations we will need “short” edges again and we will shorten some edges. We
 415 need *another* version of the resolution:

$$416 \text{res}'(I) = \min(\mathbb{R}_{>0} \cap \{d_T(v, u) - d_T(v', u) \mid v, v', u \in V(T)\}).$$

417 In particular, $\text{res}'(I) \leq \lambda(uv)$ for all edges uv of T . From the definition we have the following
 418 property:

$$419 \quad \forall v, v', u \in V(T) : \quad d_T(v, u) < d_T(v', u) \implies d_T(v, u) + \frac{\text{res}'(I)}{2} < d_T(v', u). \quad (9)$$

421 We explain how to transform the instance into one where all vertices have maximum degree 3.
 422 We will use T' and λ' for the new graph and its edge-lengths. The construction uses two values δ
 423 and δ' , where

$$424 \quad 0 < \delta < \frac{\text{res}'(I)}{6n} \quad \text{and} \quad \delta' = \frac{\delta}{4n}.$$

425 The intuition is that edges connecting different candidate Voronoi cells are shortened by δ , and
 426 then we split the vertices of degree larger than three using short edges of length δ' , where
 427 $0 < \delta' \ll \delta \ll \text{res}'(I)$.

428 For each edge uv of T we place two vertices $a_{u,v}$ and $a_{v,u}$ in T' , and connect them with an
 429 edge. If u and v belong to the same set U_i , then the length λ' of such an edge $a_{u,v}a_{v,u}$ is set
 430 to $\lambda(uv)$. If $u \in U_i$ and $v \in U_j$ with $i \neq j$, then the length λ' of such an edge $a_{u,v}a_{v,u}$ is set to
 431 $\lambda(uv) - \delta > 0$. For each vertex u of T , we connect the vertices $\{a_{u,v} \mid uv \in E(T)\}$ with a path. The
 432 length λ' of the edges on these $|V(T)|$ paths is set to δ' . Finally, for each $i \in [k]$ we define the sets

$$433 \quad U'_i = \{a_{u,v} \mid u \in U_i, uv \in E(T)\},$$

$$434 \quad S'_i = \{a_{u,v} \mid u \in S_i, uv \in E(T)\}.$$

436 Note that the sets U'_1, \dots, U'_k are pairwise disjoint. For an example of the whole process see
 437 Figure 7.

438 To recover the solutions, we define the projection map $\pi(a_{u,v}) = u$. Thus, π sends each vertex
 439 of T' to the corresponding vertex of T that was used to create it. Note that for each $i \in [k]$ we
 440 have $\pi(S'_i) = S_i$ and $\pi(U'_i) = U_i$.

441 The distances in T' and T are closely related. Using that the tree T' has fewer than $2n$ new
 442 short edges of length δ' and the path connecting any two vertices of U'_i is contained in $T'[U'_i]$ we
 443 get

$$444 \quad \forall i \in [k], u, v \in U'_i : \quad d_T(\pi(u), \pi(v)) \leq d_{T'}(u, v) < d_T(\pi(u), \pi(v)) + 2n\delta'$$

$$445 \quad < d_T(\pi(u), \pi(v)) + \delta. \quad (10)$$

446 Using that the path between two vertices in different sets U_i and U_j , $i \neq j$, uses at least one edge
 447 and at most $n - 1$ edges that have been shortened by δ , we get

$$448 \quad \forall i \neq j \in [k], u \in U'_i, v \in U'_j : \quad d_T(\pi(u), \pi(v)) - n\delta < d_{T'}(u, v) < d_T(\pi(u), \pi(v)) - \delta + 2n\delta'$$

$$449 \quad < d_T(\pi(u), \pi(v)). \quad (11)$$

450 **Lemma 9.** *Suppose that $0 < \delta < \text{res}'(I)/6n$ and the sets U_1, \dots, U_k are pairwise disjoint subsets of*
 451 *$V(T)$ that induce connected subtrees of T . The answer to $(T, ((U_1, S_1), \dots, (U_k, S_k)))$ is “yes” if and*
 452 *only if the answer to $(T', ((U'_1, S'_1), \dots, (U'_k, S'_k)))$ is “yes”.*

453 *Proof.* We first show the “if” part. Suppose that the answer to I' is “yes”. Then, there exist
 454 s'_1, \dots, s'_k with $s'_i \in S'_i$ and $U'_i = \text{cell}_{T'}(s'_i, \{s'_1, \dots, s'_k\})$ for each $i \in [k]$. Set $s_i = \pi(s'_i)$ for all $i \in [k]$,
 455 $\Sigma' = \{s'_1, \dots, s'_k\}$ and $\Sigma = \{s_1, \dots, s_k\}$.

456 Consider any fixed $i \in [k]$ and any vertex $u \in U_i$. There exists some vertex $u' \in U'_i$ such that
 457 $u = \pi(u')$. Since $u' \in U'_i = \text{cell}_{T'}(s'_i, \Sigma')$ and $u' \notin U'_j = \text{cell}_{T'}(s'_j, \Sigma')$ for all $j \neq i$, we have

$$458 \quad \forall j \in [k] \setminus \{i\} : \quad d_{T'}(s'_i, u') < d_{T'}(s'_j, u').$$

459 For $j \neq i$, since $u, s_i \in U_i$ and $s_j \notin U_i$ we use the relations (10) and (11) to get

$$460 \quad \forall j \in [k] \setminus \{i\} : d_T(s_i, u) \leq d_{T'}(s'_i, u') < d_{T'}(s'_j, u') < d_T(s_j, u).$$

461 We conclude that $u \in \text{cell}_T(s_i, \Sigma)$ and $u \notin \text{cell}_T(s_j, \Sigma)$ for all $j \in [k] \setminus \{i\}$. It follows that $U_i =$
 462 $\text{cell}_T(s_i, \Sigma)$ for all $i \in [k]$, and the answer to the instance I “yes”.

463 Now we turn to the “only if” part. Then, there exist s_1, \dots, s_k with $s_i \in S_i$ and $U_i =$
 464 $\text{cell}_T(s_i, \{s_1, \dots, s_k\})$ for each $i \in [k]$. Take a vertex $s'_i \in \pi^{-1}(s_i)$ for each $i \in [k]$, $\Sigma = \{s_1, \dots, s_k\}$
 465 and $\Sigma' = \{s'_1, \dots, s'_k\}$.

466 Consider any fixed index $i \in [k]$ and any vertex $u' \in U'_i$. Set $u = \pi(u') \in U_i$. Since $u \in U_i =$
 467 $\text{cell}_T(s_i, \Sigma)$ and $u \notin U_j = \text{cell}_T(s_j, \Sigma)$ for all $j \neq i$, we have

$$468 \quad \forall j \in [k] \setminus \{i\} : d_T(s_i, u) < d_T(s_j, u).$$

469 Because of property (9) we have

$$470 \quad \forall j \in [k] \setminus \{i\} : d_T(s_i, u) + \frac{\text{res}'(I)}{2} < d_T(s_j, u),$$

471 and thus

$$472 \quad \forall j \in [k] \setminus \{i\} : d_T(s_i, u) + 3n \cdot \delta < d_T(s_j, u).$$

473 Using the relations (10) and (11) we get

$$474 \quad \forall j \in [k] \setminus \{i\} : d_{T'}(s'_i, u') < d_T(s_i, u) + 2n\delta < d_T(s_j, u) - n\delta < d_{T'}(s'_j, u').$$

475 This implies that $u' \in \text{cell}_{T'}(s'_i, \Sigma')$ and $u' \notin \text{cell}_{T'}(s'_j, \Sigma')$ for all $j \in [k] \setminus \{i\}$. It follows that
 476 $U'_i = \text{cell}_{T'}(s'_i, \Sigma')$. Since this holds for all $i \in [k]$, it follows that the answer to the instance I'
 477 “yes”. \square

478 3.3 Algorithm to transform

479 We are now ready to explain algorithmic details of the whole transformation and explain its
 480 efficient implementation.

481 Suppose that we have an instance $I = (T, (U_1, \dots, U_k))$ for the problem GRAPHIC INVERSE
 482 VORONOI IN TREES. Let us use n for the number of vertices in T and $N = N(I) = |V(T)| + \sum_i |U_i|$
 483 for the description size of I . As mentioned earlier, we can convert in $O(N)$ time this to an
 484 equivalent instance $(T, ((U_1, S_1), \dots, (U_k, S_k)))$ for the problem GENERALIZED GRAPHIC INVERSE
 485 VORONOI IN TREES. Let I' be this new instance and note that its description size is $O(N)$.

486 First, we root the tree T at an arbitrary vertex r and store for each vertex v of T its parent
 487 node $\text{pa}(v)$. (The parent of r is set to NULL.) We add to each vertex a flag to indicate whether it
 488 belongs to a subset of vertices under consideration. Initially all flags are set to false. This takes
 489 $O(|V(T)|) = O(N)$ time.

490 With this representation of T we can check whether any given subset U of vertices of T induces
 491 a connected subgraph in $O(|U|)$ time. The key observation is that the subgraph $T[U]$ induced by
 492 U is connected if and only if there is exactly one vertex in U whose parent does not belong to U .
 493 (Here we use the convention that for the root $\text{pa}(r) = \text{NULL} \notin U$.) To check this condition, we
 494 set the flag of the vertices of U to true, count how many vertices $v \in U$ have the property that
 495 $\text{pa}(v) \notin U$, decide the connectivity of $T[U]$ depending on the counter, and at the end set the flags
 496 of vertices of U back to false.

497 For each vertex $v \in V(T)$ we make a list $L(v)$ that contains the indices $i \in [k]$ with $v \in U_i$.
 498 The lists $L(v)$, for all $v \in V(T)$, can be computed in $O(N)$ time by scanning the sets U_1, \dots, U_k :
 499 for each $v \in U_i$ we add i to $L(v)$. Note that a vertex $v \in V(T)$ belongs to W_i if and only if i is

500 the only index in the list $L(v)$. Thus, for any given $v \in U_i$, we can decide in $O(1)$ time whether
501 $v \in W_i$. With this we can compute the sets W_1, \dots, W_k in $O(\sum_i |U_i|) = O(N)$ time. Scanning the
502 sets S_1, \dots, S_k , we can replace each set S_i with the set $S_i \cap W_i$. Together we have spent $O(N)$ time
503 and we have made the preprocessing step described after Lemma 4.

504 During the algorithm, as we make the edge expansions, we maintain the lists $L(v)$ for each
505 vertex v and the rooted representation of the tree.

506 Now we explain how to make the expansions of the edges in *batches*: we iterate over the
507 indices $i \in [k]$ and, for each fixed i , we identify E_i and make all the edge expansions for E_i in
508 $O(|U_i|)$ time. Assume for the time being that ε is already known. We will discuss its choice below.

509 Consider any fixed index $i \in [k]$. We compute W_i in $O(|U_i|)$ time using the lists $L(v)$ for
510 $v \in U_i$. (The set W_i may have changed because of expansions for E_j , $j \neq i$, and thus has to be
511 computed again.) We also check in $O(|U_i|)$ time that U_i and W_i induce connected subgraphs of T
512 using the representation of T . (If any of them fails the test, then we correctly report that there is
513 no solution.) We construct the induced tree $T[U_i]$ explicitly and store it using adjacency lists:
514 for each vertex $v \in U_i$ we can find its neighbors in $T[U_i]$ in time proportional to the number of
515 neighbors. From this point, we will use the representation of $T[U_i]$.

516 Next, we compute E_i , for the fixed index $i \in [k]$, in the obvious way. For each edge xy of
517 $T[U_i]$, we check whether $x \in W_i$ and $y \notin W_i$ or whether $y \in W_i$ and $x \notin W_i$ to decide whether
518 $xy \in E_i$. This procedure to compute E_i takes $O(|U_i|)$ time.

519 We keep considering the fixed index $i \in [k]$. Now we make the expansion for each edge of E_i .
520 Here it is important that the expansion of different edges of E_i are independent: each expansion
521 affects to U_i in a different connected component of $T - W_i$. We make the expansion of an edge
522 $xy \in E_i$ with $x \in W_i$ and $y \in U_i \setminus W_i$ as follows: edit T by inserting y' , set the new edge-lengths
523 for the edges yy' and xy' , remove from U_i the subset R_{xy} of elements of U_i that are closer to y
524 than to x , and insert y' in U_i . The set R_{xy} of elements to be removed from U_i is obtained using
525 the representation of $T[U_i]$ in $O(|R_{xy}|)$ time. We correct the lists $L(v)$ by removing i from $L(v)$
526 for each $v \in R_{xy}$. (We do not need to update $T[U_i]$ because the sets R_{xy} are pairwise
527 disjoint for all $xy \in E_i$.) We conclude that expanding an edge $xy \in E_i$ takes $O(|R_{xy}|)$. Since
528 each element of U_i can be deleted at most once from U_i , and the elements y' we insert cannot be
529 deleted because they belong only to (the new) U_i , the expansions for the edges in E_i takes $O(|U_i|)$
530 time all together. This finishes the description of the work carried out for a fixed $i \in [k]$.

531 We iterate over all $i \in [k]$ making the expansions for (the current) edges in E_i . Since for each
532 $i \in [k]$ we spend $O(|U_i|)$ time, all the expansions required for Lemma 8 are carried out in $O(N)$
533 time. All this was assuming that the value ε is available, which remains to be discussed. Let \tilde{I} be
534 the resulting instance with the disjoint sets.

535 Now we can make the transformation from \tilde{I} to an instance with maximum degree 3. Assume
536 for the time being that we have the parameter δ available. Then the transformation described
537 in Section 3.2 can be easily carried out in linear time. Thus, in $O(N)$ time we obtain the final
538 instance with pairwise disjoint sets U_1, \dots, U_k and the tree T of maximum degree 3.

539 It remains to discuss how to choose the values of ε and δ for the transformations. It is
540 unclear whether ε or δ can be computed in $O(N)$ time when the edges have arbitrary lengths.
541 (If, for example, all edges have integral lengths, then we could take $\varepsilon = 1/4$, $\delta = 1/10n$ and
542 $\delta' = 1/40n^2$.) We will handle this using composite lengths. The length of each edge e is going to
543 be described by a triple (a, b, c) that represents the number $a + b\varepsilon + c\delta'$ for infinitesimals $\delta' \ll \varepsilon$.
544 (Recall that $4n\delta' = \delta$.) Thus the length encoded by (a, b, c) is smaller than the length encoded
545 by (a', b', c') if and only if (a, b, c) is lexicographically smaller than (a', b', c') . In the original
546 graph we replace the length of each edge e by $(\lambda(e), 0, 0)$. In the expansion, the new edges yy'
547 get length $(0, 1, 0)$, and in converting the tree to maximum degree 3 we introduce new edges of
548 length $(0, 0, 1) \equiv \delta'$ and we replace some edges of length $(a, b, 0)$ by $(a, b, -4n)$. The length of a
549 path becomes a triple (a, b, c) that is obtained as the vector sum of the triples over its edges. Each

550 comparison and addition of edge-lengths costs $O(1)$ time. We summarize.

551 **Theorem 10.** *Suppose that we are given an instance I for the problem GRAPHIC INVERSE VORONOI*
 552 *IN TREES or for the problem GENERALIZED GRAPHIC INVERSE VORONOI IN TREES of description size*
 553 *$N = N(I)$ over a tree T with n vertices. In $O(N)$ time we can either detect that I has no solutions, or*
 554 *construct another instance I' for the problem GENERALIZED GRAPHIC INVERSE VORONOI IN TREES*
 555 *over a tree T' with the following properties:*

- 556 • *the tree T' in the instance I' has maximum degree 3,*
- 557 • *the sets in the instance I' are pairwise disjoint,*
- 558 • *the description size of I' and the number of vertices in T' is $O(n)$,*
- 559 • *if the answer to I is “yes”, then any solution to I' is also a solution to I .*

560 *Proof.* It remains only to bound the size of T' and the description size of I' . If $k > n$, then the
 561 instance I has no solution and we report it. Otherwise, Lemma 7 implies that we are making
 562 $k - 1$ expansions, which means that the resulting tree T' has $n + k - 1 = O(n)$ vertices. The size
 563 of the instance I' is $O(n)$ because the sets in the instance are pairwise disjoint and there are $O(n)$
 564 vertices in total. □

565 4 Algorithm for subcubic trees with disjoint Voronoi cells

566 In this section we consider the problem GENERALIZED GRAPHIC INVERSE VORONOI IN TREES for an
 567 input (T, \mathbb{U}) , with the following properties:

- 568 • T is a tree of maximum degree 3
- 569 • \mathbb{U} is a sequence of pairs $(U_1, S_1), \dots, (U_k, S_k)$ where the sets U_1, \dots, U_k are pairwise disjoint.

570 Our task is to find sites s_1, \dots, s_k such that, for each $i \in [k]$, we have $U_i = \text{cell}_T(s_i, \{s_1, \dots, s_k\})$
 571 and $s_i \in S_i$. We may assume that $V(T) = \bigcup_{i \in [k]} U_i$, that $T[U_i]$ is connected for each $i \in [k]$, and
 572 that $S_i \subseteq U_i$ for each $i \in [k]$, as otherwise it is clear that there is no solution. These conditions
 573 can easily be checked in linear time.

574 First, we describe an approach to decide whether there is a solution without paying much atten-
 575 tion to the running time. Then, we describe its efficient implementation taking time $O(N \log^2 N)$,
 576 where N is the description size of the instance.

577 4.1 Characterization

578 For each vertex v , let $i(v)$ be the unique index such that $v \in U_{i(v)}$. We choose a leaf r of T as a
 579 root and henceforth consider the tree T rooted at r . We do this so that each vertex of T has at
 580 most two children. For each vertex v of T , let $T(v)$ be the subtree of T rooted at v , and define
 581 also

$$582 \quad J(v) = \{j \in [k] \mid U_j \cap T(v) \neq \emptyset\}.$$

584 Note that $i(v) \in J(v)$. Since each U_j defines a connected subset of $T(v)$, for each $j \in J(v)$, $j \neq i(v)$,
 585 we have $U_j \subseteq T(v)$ and therefore it must be that $s_j \in T(v)$.

586 Consider a fixed vertex v of T and the corresponding subtree $T(v)$. We want to parameterize
 587 possible distances from v to the site $s_{i(v)}$, that is, the site whose cell contains the vertex v , that
 588 provide the desired Voronoi diagram restricted to $T(v)$. A more careful description is below. We
 589 distinguish possible placements of $s_{i(v)}$ within $T(v)$, which we refer as “below” (or on) v and for




Figure 8: The tree $T_\alpha^+(v)$ used to define $A(v)$.

590 which we use the notation $B(v)$, and possible placements outside $T(v)$, which we refer as “above”
 591 and for which we use the notation $A(v)$.

592 First we deal with the placements where $s_{i(v)}$ is “below” v . In this case we start defining $X(v)$
 593 as the set of tuples $(s_j)_{j \in J(v)}$ that satisfy the following two conditions:

$$594 \quad \forall j \in J(v) : s_j \in S_j,$$

$$595 \quad \forall j \in J(v) : \text{cell}_{T(v)}(s_j, \{s_t \mid t \in J(v)\}) \cap T(v) = U_j \cap T(v).$$

597 Note that $X(v) \subseteq \prod_{j \in J(v)} S_j$. Finally, we define

$$598 \quad B(v) = \{d_T(s_{i(v)}, v) \mid (s_j)_{j \in J(v)} \in X(v)\}.$$

600 The set $B(v)$ represents the valid distances at which we can place $s_{i(v)}$ inside $T(v)$ such that $s_{i(v)}$
 601 is the closest site to v , and still complete the rest of the placements of the sites to get the correct
 602 portion of \mathbb{U} inside $T(v)$.

603 Now we deal with the placements “above” v . For $\alpha > 0$, let $T_\alpha^+(v)$ be the tree obtained from
 604 $T(v)$ by adding an edge vv_{new} , where v_{new} is a new vertex, and setting the length of vv_{new} to α .
 605 The role of v_{new} is the placement of the site closest to v , when it is outside $T(v)$. See Figure 8 for
 606 an illustration. In the following discussion we also use Voronoi diagrams with respect to $T_\alpha^+(v)$.
 607 Let $Y_\alpha(v)$ be the set of tuples $(s_j)_{j \in J(v)}$ that satisfy all of the following conditions:

$$608 \quad s_{i(v)} = v_{\text{new}},$$

$$609 \quad \forall j \in J(v) \setminus \{i(v)\} : s_j \in S_j,$$

$$610 \quad \forall j \in J(v) : \text{cell}_{T_\alpha^+(v)}(s_j, \{s_t \mid t \in J(v)\}) \cap T(v) = U_j \cap T(v).$$

612 Finally we define

$$613 \quad A(v) = \{\alpha \in \mathbb{R}_{>0} \mid Y_\alpha(v) \neq \emptyset\}.$$

615 We are interested in deciding whether $B(r)$ is nonempty. Indeed, for the root r we have
 616 $J(r) = [k]$ and $T(r) = T$ by construction. The definition of $X(v)$ implies that $B(r)$ is nonempty if
 617 and only if there is some tuple $(s_1, \dots, s_k) \in S_1 \times \dots \times S_k$ such that

$$618 \quad \forall i \in J(r) = [k] : \text{cell}_T(s_i, \{s_1, \dots, s_k\}) = \text{cell}_{T(r)}(s_i, \{s_1, \dots, s_k\}) = U_i \cap T(r) = U_i.$$

619 This is precisely the condition we have to check to solve GENERALIZED GRAPHIC INVERSE VORONOI
 620 IN TREES.




Figure 9: Different cases in the computation of $A(v)$ and $B(v)$ when v has children v_1 and v_2 . (The case $i(v) = i(v_2) \neq i(v_1)$ is symmetric to the case $i(v) = i(v_1) \neq i(v_2)$.)

621 We are going to compute $A(v)$ and $B(v)$ in a bottom-up fashion along the tree T . If v is leaf of
 622 T , then $J(v) = \{i(v)\}$ and clearly we have

$$623 \quad A(v) = \mathbb{R}_{>0} \quad \text{and} \quad B(v) = \begin{cases} \{0\} & \text{if } v \in S_{i(v)}, \\ \emptyset & \text{if } v \notin S_{i(v)}. \end{cases}$$

624 Consider now a vertex v of T that has two children v_1 and v_2 . Assume that we already have
 625 $A(v_j)$ and $B(v_j)$ for $j = 1, 2$. For $j = 1, 2$ define the sets

$$\begin{aligned}
 626 \quad A'(v_j) &= \{x - \lambda(vv_j) \mid x \in A(v_j)\}, \\
 627 \quad B'(v_j) &= \{x + \lambda(vv_j) \mid x \in B(v_j)\}, \\
 628 \quad C'(v_j) &= \{\alpha \mid \exists x \in B(v_j) \text{ such that } x - \lambda(vv_j) < \alpha < x + \lambda(vv_j)\}.
 \end{aligned}$$

630 This is the offset we obtain when we take into account the length of the edge vv_j . The set $C'(v_j)$
 631 will be relevant for the case when $i(v) \neq i(v_j)$. The following lemmas show how to compute $A(v)$
 632 and $B(v)$ from its children. Figure 9 is useful to understand the different cases.

633 **Lemma 11.** *If the vertex v has two children v_1 and v_2 , then*

$$634 \quad A(v) = \mathbb{R}_{>0} \cap \begin{cases} A'(v_1) \cap A'(v_2) & \text{if } i(v) = i(v_1) = i(v_2), \\ A'(v_1) \cap C'(v_2) & \text{if } i(v) = i(v_1) \neq i(v_2), \\ A'(v_2) \cap C'(v_1) & \text{if } i(v) = i(v_2) \neq i(v_1), \\ C'(v_1) \cap C'(v_2) & \text{if } i(v) \neq i(v_1) \text{ and } i(v) \neq i(v_2). \end{cases}$$

635 *Proof.* This is a standard proof in dynamic programming. We only point out the main insight
 636 showing the role of $A'(v_j)$ and $C'(v_j)$ for $j \in \{1, 2\}$.

637 When $i(v) = i(v_j)$, placing $s_{i(v)}$ at v_{new} of the tree $T_{\alpha}^+(v)$ is the same as placing it at v_{new} of
 638 $T_{\alpha + \lambda(vv_j)}^+(v_j)$. The valid values α for $T_{\alpha + \lambda(vv_j)}^+(v_j)$ are described by $A'(v_j)$, a shifted version of
 639 $A(v_j)$.

640 When $i(v) \neq i(v_j)$, there has to be a compatible placement of $s_{i(v_j)}$ inside $T(v_j)$ such that v is
 641 closer to $s_{i(v)} = v_{\text{new}}$ than to $s_{i(v_j)}$, while v_j is closer to $s_{i(v_j)}$ than to $s_{i(v)}$. That is, we must have

$$642 \quad d_T(v_{\text{new}}, v) < d_T(s_{i(v_j)}, v) \quad \text{and} \quad d_T(s_{i(v_j)}, v_j) < d_T(v_{\text{new}}, v_j),$$

643 or equivalently, α must satisfy

$$644 \quad \alpha < d_T(s_{i(v_j)}, v_j) + \lambda(vv_j) \quad \text{and} \quad d_T(s_{i(v_j)}, v_j) < \alpha + \lambda(vv_j).$$

645 Thus, each possible value x of $d_T(s_{i(v_j)}, v_j)$, that is, each $x \in B(v_j)$, gives the interval $(x -$
 646 $\lambda(vv_j), x + \lambda(vv_j))$ of possible values for α . The union of these intervals over $x \in B(v_j)$ is precisely
 647 $C'(v_j)$. \square

648 To construct $B(v)$ it is useful to have a function that tells whether v is a valid placement for
 649 $s_{i(v)}$. For this matter we define the following function:

$$650 \quad \chi(v) = \begin{cases} \{0\} & \text{if } i(v) = i(v_1) = i(v_2), v \in S_{i(v)}, 0 \in A'(v_1) \text{ and } 0 \in A'(v_2), \\ \{0\} & \text{if } i(v) = i(v_1) \neq i(v_2), v \in S_{i(v)}, 0 \in A'(v_1) \text{ and } 0 \in C'(v_2), \\ \{0\} & \text{if } i(v) = i(v_2) \neq i(v_1), v \in S_{i(v)}, 0 \in A'(v_2) \text{ and } 0 \in C'(v_1), \\ \{0\} & \text{if } i(v) \neq i(v_1), i(v) \neq i(v_2), v \in S_{i(v)}, 0 \in C'(v_1) \text{ and } 0 \in C'(v_2), \\ \emptyset & \text{otherwise.} \end{cases}$$

651 **Lemma 12.** *If the vertex v has two children v_1 and v_2 , then*

$$652 \quad B(v) = \chi(v) \cup \begin{cases} (B'(v_1) \cap A'(v_2)) \cup (B'(v_2) \cap A'(v_1)) & \text{if } i(v) = i(v_1) = i(v_2), \\ B'(v_1) \cap C'(v_2) & \text{if } i(v) = i(v_1) \neq i(v_2), \\ B'(v_2) \cap C'(v_1) & \text{if } i(v) = i(v_2) \neq i(v_1), \\ \emptyset & \text{if } i(v) \neq i(v_1) \text{ and } i(v) \neq i(v_2). \end{cases}$$

653 *Proof.* First we note that $\chi(v) = \{0\}$ if and only if v is a valid placement for $s_{i(v)}$. Indeed, the
 654 formula is the same that was used for $A(v)$, but for the value $\alpha = 0$, and it takes into account
 655 whether $v \in S_{i(v)}$.

656 The proof for the correctness of $B(v)$ is again based in standard dynamic programming. The
 657 case for $s_{i(v)}$ being placed at v is covered by $\chi(v)$. The main insight for the case when $s_{i(v)}$ is
 658 placed in $T(v_1)$ is that, from the perspective of the other child, v_2 , the vertex is placed ‘‘above’’ v_2 .
 659 That is, only the distance from $s_{i(v)}$ to v_2 is relevant. Thus, we have to combine $B(v_1)$ and $A(v_2)$,
 660 with the appropriate shifts. More precisely, for v_2 we have to use $A'(v_2)$ or $C'(v_2)$ depending on
 661 whether $i(v_2) = i(v)$ or $i(v_2) \neq i(v)$. \square

662 When v has a unique child v' , then the formulas are simpler and the argumentation is similar.
 663 We state them for the sake of completeness without discussing their proof.

$$664 \quad A(v) = \mathbb{R}_{>0} \cap \begin{cases} A'(v') & \text{if } i(v) = i(v'), \\ C'(v') & \text{if } i(v) \neq i(v'). \end{cases}$$

$$665 \quad B(v) = \begin{cases} B'(v') \cup \{0\} & \text{if } i(v) = i(v'), v \in S_{i(v)}, \text{ and } \lambda(vv') \in A(v'), \\ B'(v') & \text{if } i(v) = i(v') \text{ and } (v \notin S_{i(v)} \text{ or } \lambda(vv') \notin A(v')), \\ \{0\} & \text{if } i(v) \neq i(v'), v \in S_{i(v)} \text{ and } 0 \in C'(v'), \\ \emptyset & \text{if } i(v) \neq i(v') \text{ and } (v \notin S_{i(v)} \text{ or } 0 \notin C'(v')). \end{cases}$$

668 4.2 Efficient manipulation of monotonic intervals

669 The efficient algorithm that we will present is based on an efficient representation of the sets $A(v)$
 670 and $B(v)$ using binary search trees. Here we discuss the representation that we will be using.

671 We first consider how to store a set X of real values under the following operations.

- 672 • Copy makes a copy of the data structure storing X ;
- 673 • Report returns the elements of X sorted;
- 674 • Insert(y) adds a new element y in X ;
- 675 • Delete(y) removes the element $y \in X$ from X ;

- 676 • $\text{Succ}(y)$ returns the successor of y in X , defined as the smallest number in X that is at least
677 as large as y ;
- 678 • $\text{Pred}(y)$ returns the predecessor of y in X , defined as the largest number in X that is smaller
679 or equal than y ;
- 680 • $\text{Split}(y)$ returns the representation for $X_{\leq} = \{x \in X \mid x \leq y\}$ and the representation for
681 $X_{>} = \{x \in X \mid x > y\}$; the representation of X is destroyed in the process;
- 682 • $\text{Join}(X_1, X_2)$ returns the representation of $X = X_1 \cup X_2$ if $\max(X_1) < \min(X_2)$, and otherwise
683 it returns an error; the representations of X_1 and X_2 are destroyed in the process;
- 684 • $\text{Shift}(\alpha)$ adds the given value α to all the elements of X .

685 These operations can be done efficiently using dynamic balanced binary search tree with
686 so-called augmentation, that is, with some extra information attached to the nodes. Strictly
687 speaking the following result is not needed, but understanding it will be useful to understand the
688 more involved data structure we eventually employ.

689 **Theorem 13.** *There is an augmented dynamic binary search tree to store sets of m real values with
690 the following time guarantees:*

- 691 • the operations *Copy* and *Report* take $O(m)$ time;
- 692 • the operations *Insert*, *Delete*, *Succ*, *Pred*, *Split*, *Join* and *Shift* take $O(\log m)$ time. (For *Join*
693 the value m is the size of the resulting set.)

694 *Proof.* Let X be the set of values to store. We use a dynamic balanced binary search tree \mathcal{T} where
695 each node represents one element of X . For each node μ of \mathcal{T} , let $x(\mu)$ be the value represented
696 by μ . The tree \mathcal{T} is a binary search tree with respect to the values $x(\mu)$. However, we *do not*
697 store $x(\mu)$ explicitly at μ , but we store it in so-called difference form. At each non-root node μ
698 with parent μ' , we store $\text{diff-val}(\mu) := x(\mu) - x(\mu')$. At the root r we store $\text{diff-val}(r) = x(r)$.
699 (This choice is consistent with using $x(\text{NULL}) = 0$.) This is a standard technique already used by
700 Tarjan [6]. Whenever we want to obtain $x(\mu)$ for a node μ , we have to add $\text{diff-val}(\mu')$ for the
701 nodes μ' along the root-to- μ path. Since operations on a tree are performed always locally, that is,
702 accessing a node from a neighbour, we spend $O(\log m)$ time to compute the first value $x(\mu)$, and
703 from there on each value $x(\cdot)$ is computed in $O(1)$ additional time from the value of its neighbor.
704 Of course, the values $\text{diff-val}(\mu)$ have to be updated through the changes in the tree, including
705 rotations or other balancing operations.

706 With this representation it is trivial to perform the operation $\text{Shift}(\alpha)$ in constant time: at the
707 root r of \mathcal{T} , we just add α to $\text{diff-val}(r)$.

708 For the rest of operations, the time needed to execute them is the same as for the dynamic
709 balanced search trees we employ. Brass [2, Chapter 3] explains dynamic trees with the requested
710 properties; see Section 3.11 of the book for the more complex operations of split and join. (The
711 same time bounds with amortized time bounds, which are sufficient for our application, can be
712 obtained using the classical splay trees [5].) □

713 Consider now a family \mathbb{I} of closed intervals on the real line. The family \mathbb{I} is **monotonic** if no
714 interval contains another interval. In a monotonic family of intervals, the left endpoints have to
715 be distinct and the right endpoints also have to be distinct. Also, for such a family, sorting the
716 intervals by their left endpoints or their right endpoints gives the same result. Because of this, we
717 can talk about the ordering of the intervals, and we can also talk about the rightmost or leftmost
718 interval in \mathbb{I} with a certain property.

719 We want to maintain a set \mathbb{I} of monotonic intervals under the following operations.

- 720 • `IntCopy` makes a copy of the data structure storing \mathbb{I} .
- 721 • `IntReport` returns the elements of \mathbb{I} sorted by their left endpoint.
- 722 • `IntInsert(J)` adds a new interval J in \mathbb{I} ; it assumes that the resulting family keeps being
723 monotonic.
- 724 • `IntDelete(J)` deletes the interval $J \in \mathbb{I}$.
- 725 • `IntHitBy(J)`, for an interval J , returns whether J intersects some interval of \mathbb{I} .
- 726 • `IntContaining(J)`, for an interval J , returns the representation for $\mathbb{I}' = \{I \in \mathbb{I} \mid J \subseteq I\}$ and
727 the representation for $\mathbb{I}'' = \mathbb{I} \setminus \mathbb{I}'$. The representation of \mathbb{I} is destroyed in the process.
- 728 • `IntClip(J)`, for an interval $J = [x, y]$, returns the representation for the intervals $\mathbb{I}' = \{I \cap J \mid$
729 $I \in \mathbb{I}\}$ and for the intervals $\mathbb{I}'' = \{I \cap (-\infty, x] \mid I \in \mathbb{I}\} \cup \{I \cap [y, +\infty) \mid I \in \mathbb{I}\}$. In both cases
730 we remove empty intervals, and remove intervals contained in another one, so that we keep
731 having monotonic families. The representation of \mathbb{I} is destroyed in the process.
- 732 • `IntJoin($\mathbb{I}_1, \mathbb{I}_2$)` returns the representation of $\mathbb{I} = \mathbb{I}_1 \cup \mathbb{I}_2$ if \mathbb{I} is a monotonic family and all the
733 intervals of \mathbb{I}_1 are to the left of all the intervals of \mathbb{I}_2 . Otherwise it returns an error. The
734 representations of \mathbb{I}_1 and \mathbb{I}_2 are destroyed in the process.
- 735 • `IntShift(α)`, for a given real value α , shifts all the intervals by α ; this is, each interval $[a, b]$
736 in \mathbb{I} is replaced by $[a + \alpha, b + \alpha]$.
- 737 • `IntExtend(λ)`, for a given real value $\lambda > 0$, extends all the intervals by λ in both directions;
738 this is, each interval $[a, b]$ in \mathbb{I} is replaced by $[a - \lambda, b + \lambda]$.

739 **Theorem 14.** *There is a data structure to store monotonic families of m intervals with the following*
740 *time guarantees:*

- 741 • *the operations `IntCopy` and `IntReport` take $O(m)$ time;*
- 742 • *the operations `IntInsert`, `IntDelete`, `IntHitBy`, `IntContaining`, `IntClip`, `IntJoin`, `IntShift`, `IntExtend`*
743 *take $O(\log m)$ time. (For `IntJoin` the value m is the size of the resulting set \mathbb{I} .)*

744 *Proof.* Let \mathbb{I} be the family of monotonic intervals to store. We use a dynamic balanced binary
745 search tree \mathcal{T} where each node represents one element of \mathbb{I} . For the node μ of \mathcal{T} that represents
746 the interval I , let $a(\mu)$, $b(\mu)$ and $\ell(\mu) = b(\mu) - a(\mu)$ be the left endpoint, the right endpoint, and
747 the length of I , respectively. Thus, if μ represents $[a_i, b_i]$, we have $a_i = a(\mu)$ and $b_i = a(\mu) + \ell(\mu)$.

748 The tree \mathcal{T} is a binary search tree with respect to the values $a(\mu)$. Because the family of
749 intervals is monotonic, \mathcal{T} is *also* a binary search tree with respect to the values $b(\mu)$. However,
750 the values $a(\mu)$, $b(\mu)$ or $\ell(\mu)$ are *not stored explicitly*. Instead, the values are stored in difference
751 form and implicitly. More precisely, at each node μ of \mathcal{T} we store two values, $\text{diff-val}(\mu)$ and
752 $\text{diff-len}(\mu)$, defined as follows. If μ is the root of the tree and represents the interval $[a, b]$, then
753 $\text{diff-val}(\mu) = a$ and $\text{diff-len}(\mu) = b - a$. If μ is a non-root node of the tree representing $[a, b]$, and
754 μ' is its parent, then $\text{diff-val}(\mu) = a - \text{diff-val}(\mu')$ and $\text{diff-len}(\mu) = (b - a) - \text{diff-len}(\mu')$.

755 This is an extension of the technique employed in the proof of Theorem 13. In fact, \mathcal{T} is just
756 the tree in the proof of Theorem 13 for the left endpoints of the intervals, where additionally
757 each node stores information about the length of the interval, albeit this additional information is
758 stored also in difference form.

759 Whenever we want to obtain $a(\mu)$ or $\ell(\mu)$ for a node μ , we have to add $\text{diff-val}(\mu')$ or
760 $\text{diff-len}(\mu')$ for the nodes μ' along the root-to- μ path, respectively. The right endpoint $b(\mu)$ is
761 obtained from $b(\mu) = a(\mu) + \ell(\mu)$. Since operations in a tree always go from a node to a neighbor,

762 we can assume that the values $a(\mu)$, $b(\mu)$ and $\ell(\mu)$ are available at a cost of $O(1)$ time per node,
 763 after an initial cost of $O(\log m)$ time to compute the values at the first node. Of course, the values
 764 $\text{diff-val}(\mu)$ and $\text{diff-len}(\mu)$ have to be updated through the changes in the tree, including rotations
 765 or other balancing operations.

766 Since \mathcal{T} is a binary search tree with respect to the values $a(\cdot)$ and also with respect to the
 767 values $b(\cdot)$, we can make the usual operations that can be performed in a binary search tree,
 768 such as predecessor or successor, with respect to any of those two keys. For example, we can
 769 get in $O(\log m)$ time the rightmost interval that contains a given value y , which amounts to a
 770 predecessor query with y for the values $a(\cdot)$, or we can get the leftmost interval that contains a
 771 given value y , which amounts to a successor query with y for the values $b(\cdot)$.

772 With this representation, it is trivial to perform the operations $\text{IntShift}(\alpha)$ or $\text{IntExtend}(\lambda)$ in
 773 $O(1)$ time. We just update diff-val or diff-len at the root.

774 The operations IntCopy , IntReport , IntInsert and IntDelete can be carried out as normal
 775 operations in a dynamic binary search tree. The operation IntJoin is also just the join operation
 776 for trees.

777 For the operation $\text{IntHitBy}(J)$ with $J = [x, y]$ we make a predecessor and a successor query
 778 with x for the values $a(\cdot)$. This gives the two intervals $I_1, I_2 \in \mathbb{I}$ such that x is between the left
 779 endpoint of I_1 and I_2 . We then check whether $I_1 \cup I_2$ intersect J , which requires constant time.

780 For the operation $\text{IntContaining}(J)$ we proceed as follows. We find the rightmost interval
 781 $[a_\ell, b_\ell] \in \mathbb{I}$ with the left endpoint outside J . We find the rightmost interval $[a_r, b_r] \in \mathbb{I}$ with the
 782 right endpoint inside J . Because \mathbb{I} is a monotonic family of intervals, the intervals contained in J
 783 are precisely those with the right endpoint in the half-open interval $(a_\ell, a_r]$. We use the operations
 784 $\text{Split}(a_\ell)$ and $\text{Split}(a_r)$ with respect to the values $a(\cdot)$ to obtain the representations of

$$\begin{aligned} 785 \quad \mathbb{I}_1 &= \{[a, b] \in \mathbb{I} \mid a \leq a_\ell\}, \\ 786 \quad \mathbb{I}_2 &= \{[a, b] \in \mathbb{I} \mid a_\ell < a \leq a_r\} = \{I \in \mathbb{I} \mid J \subseteq I\}, \\ 787 \quad \mathbb{I}_3 &= \{[a, b] \in \mathbb{I} \mid a_r < a\}. \end{aligned}$$

789 We then use the Join operation to merge the representations of \mathbb{I}_1 and \mathbb{I}_3 .

790 For the operation $\text{IntClip}(J)$ with the interval $J = [x, y]$ we proceed as follows. We use
 791 $\text{IntContaining}([x, x])$, $\text{IntContaining}([y, y])$, and IntJoin to separate \mathbb{I} into the group \mathbb{I}' of intervals
 792 pierced by x or y , and the rest, \mathbb{I}'' . Then we use $\text{Split}(x)$ (with respect to $a(\cdot)$) and $\text{Split}(y)$ (with
 793 respect to $a(\cdot)$) to split \mathbb{I}'' into three groups: \mathbb{I}_1 containing intervals of \mathbb{I} contained in $(-\infty, x]$, \mathbb{I}_2
 794 containing intervals of \mathbb{I} contained in $[x, y]$, and \mathbb{I}_3 containing intervals of \mathbb{I} contained in $[y, +\infty)$.
 795 In \mathbb{I}' we find the leftmost interval that contains x , clip it with $(-\infty, x]$, and add it to \mathbb{I}_1 . Again in
 796 the same group, \mathbb{I}' , we find the rightmost interval that contains x , clip it with $[x, y]$, and add it to
 797 \mathbb{I}_2 . We do a similar procedure for y : add to \mathbb{I}_2 the leftmost interval of \mathbb{I}' that contains y , clipped
 798 with J , and add to \mathbb{I}_3 the rightmost interval of \mathbb{I}' that contains y , clipped with $[y, +\infty)$. If the
 799 two intervals we added to \mathbb{I}_2 are the same, which means that they both are $[x, y]$, we only add
 800 one of them. The procedure takes $O(\log m)$ time. \square

801 Consider a set $A \subseteq \mathbb{R}$. A **representation** of A is a family \mathbb{I} of monotonic intervals such that
 802 $A = \bigcup_{I \in \mathbb{I}} I$. The intervals in \mathbb{I} may intersect and the representation is not uniquely defined. See
 803 Figure 10 for an example. The *size* of the representation \mathbb{I} is the number of (possibly non-disjoint)
 804 intervals in \mathbb{I} . This is potentially larger than the minimum number of intervals that is needed
 805 because the intervals in \mathbb{I} can intersect.

806 Consider some set A and its representation \mathbb{I} . If we use the data structure of Theorem 14 to
 807 store \mathbb{I} , the operations reflect operations we do with A . For example, $\text{IntHitBy}(J)$ tells whether J
 808 intersects A , while $\text{IntClip}([x, y])$ returns a representation of $A \cap [x, y]$ and a representation of
 809 $A \cap (-\infty, x] \cup A \cap [y, +\infty)$. The operation $\text{IntContaining}(J)$ will be used only when \mathbb{I} is a set of
 810 zero-length intervals, and in that case it returns a representation of $A \cap J$. When \mathbb{I}_1 and \mathbb{I}_2 are

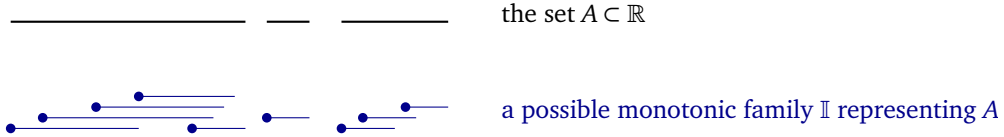


Figure 10: The set A at the top and one possible representation \mathbb{I} of A . The size of this representation is 9.

811 representations of A_1 and A_2 , then $\text{IntJoin}(\mathbb{I}_1, \mathbb{I}_2)$ returns the representation of $A_1 \cup A_2$, assuming
 812 that $\max(A_1) < \min(A_2)$.

813 4.3 Algorithm

814 In this section we present an efficient algorithm based on the characterization of the previous
 815 section. We keep using the same notation. In particular, T keeps being a rooted tree and each
 816 vertex has at most two children. We use n for the number of vertices of T .

817 There are two main ideas used in our approach. The first one is that, for each vertex of the
 818 tree with two children, we want to spend time (roughly) proportional to the size of the smaller
 819 subtree of its children. The second idea is to use representations of $A(v)$ and $B(v)$ and manipulate
 820 them using the data structure of Theorem 14.

821 The following lemma, which is folklore, shows the advantage of the first idea. For each node
 822 v with two children, let v_1 and v_2 be its two children. If v has only one child, we denote it by v_1 .
 823 For each node v , let $n(v)$ be the number of vertices in the subtree $T(v)$. (Thus $n(r) = n$.)

824 **Lemma 15.** *If V_2 denotes the vertices of T with two children, then*

$$825 \sum_{v \in V_2} \min\{n(v_1), n(v_2)\} = O(n \log n).$$

826 *Proof.* For each vertex u of T define

$$827 \sigma(u) = \sum_{v \in V_2 \cap V(T(u))} \min\{n(v_1), n(v_2)\}.$$

828 Thus, we want to bound $\sigma(r)$. We show by induction on $n(u)$ that

$$829 \sigma(u) \leq n(u) \log_2 n(u).$$

830 For the base case note that, when $n(u) = 1$, the vertex u is a leaf and $\sigma(u) = 0$, so the statement
 831 holds.

832 If u has one child u_1 , then we have $V_2 \cap T(u) = V_2 \cap T(u_1)$,

$$833 \sigma(u) = \sigma(u_1) \leq n(u_1) \log_2 n(u_1) \leq n(u) \log_2 n(u),$$

834 and the bound holds. If u has two children u_1 and u_2 , then we can assume without loss of
 835 generality that $n(u_1) \leq n(u_2)$, which implies that $n(u_1) < n(u)/2$. Using the induction hypothesis

836 for $n(u_1)$ and $n(u_2)$, we obtain

$$\begin{aligned}
837 \quad \sigma(u) &= \sum_{v \in V_2 \cap V(T(u))} \min\{n(v_1), n(v_2)\} \\
838 &= \sigma(u_1) + \sigma(u_2) + n(u_1) \\
839 &\leq n(u_1) \log_2 n(u_1) + n(u_2) \log_2 n(u_2) + n(u_1) \\
840 &< n(u_1) \log_2 (n(u)/2) + n(u_2) \log_2 n(u) + n(u_1) \\
841 &= n(u_1) (\log_2 n(u) - 1) + n(u_2) \log_2 n(u) + n(u_1) \\
842 &= (n(u_1) + n(u_2)) \log_2 n(u) \\
843 &< n(u) \log_2 n(u). \quad \square
\end{aligned}$$

845 We manipulate the sets $A(v)$ and $B(v)$ using representations $\mathbb{I}(A(v))$ and $\mathbb{I}(B(v))$, respectively.
846 In the case of $B(v)$, since $B(v)$ is a finite set of values, the family $\mathbb{I}(B(v))$ consists of zero-length
847 monotonic intervals. The reason for this artificial approach to treat $B(v)$, as opposed to using a
848 set of values, is that in our algorithm sometimes we set the lengths of intervals defined by $B(v)$.
849 Thus, there is no real difference between how we treat the representations of $A(\cdot)$ and $B(\cdot)$.

850 The families of intervals $\mathbb{I}(A(v))$ and $\mathbb{I}(B(v))$ are stored and manipulated using the data
851 structure of Theorem 14. Thus, we are using the data structure described in Theorem 14 to
852 represent $A(v)$ and $B(v)$ implicitly, as the union of monotonic intervals. The reason for this choice
853 is technical and reflected in the proof of the next lemma.

854 For each vertex v of T , we use $m_A(v)$ and $m_B(v)$ to denote the sizes of $\mathbb{I}(A(v))$ and $\mathbb{I}(B(v))$,
855 respectively. Although the value $m_A(v)$ actually depends on the family $\mathbb{I}(A(v))$ of intervals that is
856 used, this relaxation of the notation will not lead to confusion.

857 It is clear that $B(v)$ has at most $n(v)$ values because each value corresponds to a vertex of
858 $T(v)$. Thus, $m_B(v) \leq n(v)$. A similar bound will hold for $m_A(v)$ by induction.

859 **Lemma 16.** Consider a vertex v of T with two children v_1 and v_2 , and assume that we have
860 representations $\mathbb{I}(A(v_1))$, $\mathbb{I}(B(v_1))$, $\mathbb{I}(A(v_2))$ and $\mathbb{I}(B(v_2))$ of $A(v_1)$, $B(v_1)$, $A(v_2)$ and $B(v_2)$, respec-
861 tively, each of them stored in the data structure of Theorem 14. Set $m_1 = m_A(v_1) + m_B(v_1)$ and
862 $m_2 = m_A(v_2) + m_B(v_2)$, and assume that $m_1 \leq m_2$. We can compute in $O(m_1 \log m_2)$ time families
863 $\mathbb{I}(A(v))$ and $\mathbb{I}(B(v))$ that represent $A(v)$ and $B(v)$, respectively, each of them stored in the data
864 structure of Theorem 14.² Moreover, the representation $\mathbb{I}(A(v))$ has size at most

$$865 \quad \max\{m_A(v_1) + m_A(v_2), m_A(v_1) + m_B(v_2), m_B(v_1) + m_A(v_2), m_B(v_1) + m_B(v_2)\}.$$

866 *Proof.* First we compute $\chi(v)$. To check whether $0 \in A'(v_j)$, where $j \in \{1, 2\}$, we perform the
867 operation $\text{IntHitBy}([\lambda(vv_j), \lambda(vv_j)])$ in the representation $\mathbb{I}(A(v_j))$. To check whether $0 \in C'(v_j)$,
868 where $j \in \{1, 2\}$, we observe that $0 \in C'(v_j)$ if and only if $[-\lambda(vv_j), +\lambda(vv_j)]$ contains some
869 element of $B(v_j)$. This latter question is answered making the query $\text{IntHitBy}([-\lambda(vv_j), +\lambda(vv_j)])$
870 in the representation $\mathbb{I}(B(v_j))$. We conclude, that $\chi(v)$ can be computed in $O(\log m_2)$ time without
871 changing any of the representations.

872 Next, for each $j \in \{1, 2\}$, we compute the representation $\mathbb{I}(A'(v_j))$ of $A'(v_j)$ applying the
873 operation $\text{IntShift}(-\lambda(vv_j))$ to $\mathbb{I}(A(v_j))$. Similarly, we can compute the representation $\mathbb{I}(B'(v_j))$ of
874 $B'(v_j)$. This takes $O(\log m_1) + O(\log m_2) = O(\log m_2)$ time. More importantly, with an additional
875 cost of $O(\log m_j)$ time we can use indistinctly the representation of $B(v_j)$ or $B'(v_j)$, whatever is
876 more convenient.

877 Note that we cannot afford to make copies of the representations $\mathbb{I}(A'(v_2))$ or $\mathbb{I}(B'(v_2))$ because
878 this would take $\Theta(m_2)$ time, which may be too much. On the other hand, we can manipulate and
879 make explicit copies of $\mathbb{I}(A'(v_1))$ and $\mathbb{I}(B'(v_1))$ because it takes $O(m_1)$ time. Define the *minimal*

²In the process we destroy the data structures for $\mathbb{I}(A(v_2))$ and $\mathbb{I}(B(v_2))$.

880 **representation** of a set $A \subset \mathbb{R}$ to be the maximal intervals (with respect to inclusion) in A . From
881 $\mathbb{I}(A'(v_1))$ we can compute the minimal representation of $A'(v_1)$ in linear time, that is, $O(m_1)$ time.
882 For this we use the operation `IntReport` in $\mathbb{I}(A'(v_1))$, which returns the intervals in $\mathbb{I}(A'(v_1))$ sorted
883 by their left endpoints, and sequentially merge adjacent intervals that intersect. Similarly, we can
884 find a minimal representation of $B'(v_1)$, which is a list of the values in $B'(v_1)$. Thus, after $O(m_1)$
885 time we have the minimal representation of $A'(v_1)$ as a list of (sorted) intervals J_1, \dots, J_s and
886 $B'(v_1)$ as a sorted list of values y_1, \dots, y_t , where $k + t \leq m_1$.

887 Now we distinguish cases depending on the relations between $i(v)$, $i(v_1)$ and $i(v_2)$.

888 Consider the case when $i(v) = i(v_1) = i(v_2)$. We have two parts.

- 889 1. First we compute the representation $\mathbb{I}(B(v))$ of $B(v)$. Because of Lemma 12, we have

$$890 \quad B(v) = \chi(v) \cup (B'(v_1) \cap A'(v_2)) \cup (B'(v_2) \cap A'(v_1)).$$

891 Recall that we have an explicit representation of $B'(v_1)$. For each element y in $B'(v_1)$, we
892 query $\mathbb{I}(A'(v_2))$ using `IntHitBy`($[y, y]$) to decide whether $y \in A'(v_2)$. Thus, we can compute
893 an explicit representation of $B'(v_1) \cap A'(v_2)$ in $O(m_1 \log m_2)$ time.

894 Recall that we also have an explicit minimal representation J_1, \dots, J_s of $A'(v_1)$. For each in-
895 terval J in that representation, we query $\mathbb{I}(B(v_2))$ with `IntClip`(J) to obtain the representation
896 of $J \cap B'(v_2)$. Since the sets J_1, \dots, J_s are pairwise disjoint, we indeed obtain representations
897 of the sets $J_1 \cap B'(v_2), \dots, J_s \cap B'(v_2)$. We then merge them using `IntJoin`. Since the intervals
898 J_1, \dots, J_t are pairwise disjoint, the operation `IntJoin` can be indeed performed. In total we
899 have used $t \leq m_1$ times the operations `IntClip` and `IntJoin`, and thus we spent $O(m_1 \log m_2)$
900 time in total. Inserting in this representation the values (as zero-length intervals) of
901 $B'(v_1) \cap A'(v_2)$, we finally obtain a representation of $(B'(v_1) \cap A'(v_2)) \cup (B'(v_2) \cap A'(v_1))$. If
902 $\chi(v)$ is nonempty, we also insert the interval $[0, 0]$ in the representation. The final result
903 is a representation $\mathbb{I}(B(v))$ of $B(v)$. Note that in this computation we have destroyed the
904 representation of $\mathbb{I}(B'(v_2))$ because of the operations `IntClip`.

- 905 2. Next we compute the representation $\mathbb{I}(A(v))$ of $A(v)$. Because of Lemma 11 we have
906 that $A(v) = \mathbb{R}_{>0} \cap A'(v_1) \cap A'(v_2)$. Recall that we have an explicit minimal representation
907 J_1, \dots, J_s of $A'(v_1)$. For each interval J_i in the minimal representation of $A'(v_1)$, we extract
908 from $\mathbb{I}(A'(v_2))$ a representation of $J_i \cap A'(v_2)$ using `IntClip`(J_i). Then we compute a rep-
909 resentation of $\bigcup_{i \in [s]} J_i \cap A'(v_2) = A'(v_1) \cap A'(v_2)$ using $s - 1$ times the operation `IntJoin`.
910 In both cases it is important that the intervals J_1, \dots, J_s are pairwise disjoint. This takes
911 $O(s \log m_2) = O(m_1 \log m_2)$ time. To obtain $\mathbb{I}(A(v))$ we apply `IntClip`($\mathbb{R}_{>0}$). (Strictly speak-
912 ing, in Theorem 14 we were assuming closed intervals, but this is not an important feature
913 and we can maintain arbitrary intervals.) Note that in this computation we have destroyed
914 the representation $\mathbb{I}(A'(v_2))$ of $A'(v_2)$, because of the `IntClip` operations, and therefore this
915 step has to be made after the computation of $B(v)$, which is also using the representation
916 $\mathbb{I}(A'(v_2))$, but not changing it.

917 Consider now the case when $i(v) = i(v_1) \neq i(v_2)$. We proceed as follows.

- 918 1. First we compute the representation $\mathbb{I}(B(v))$ of $B(v)$. Because of Lemma 12 we have
919 $B(v) = \chi(v) \cup (B'(v_1) \cap C'(v_2))$. Note that, for each $y \in \mathbb{R}$, we have $y \in C'(v_2)$ if and only if
920 the interval $[y - \lambda(vv_2), y + \lambda(vv_2)]$ contains some element of $B(v_2)$. Recall that we have an
921 explicit description y_1, \dots, y_t of $B'(v_1)$. Therefore, for each element $y \in B'(v_1)$, we use the
922 operation `IntHitBy`($[y - \lambda(vv_2), y + \lambda(vv_2)]$) in $\mathbb{I}(B(v_2))$ to detect whether $y \in C'(v_2)$. With
923 this we computed $B'(v_1) \cap C'(v_2)$ explicitly in $O(m_1 \log m_2)$ time and we did not change the
924 representation $\mathbb{I}(B(v_2))$. Finally, we build the data structure for the representation $\mathbb{I}(B(v))$
925 of $B(v)$ by inserting the intervals $[y, y]$ with $y \in B'(v_1) \cap C'(v_2)$ and, if $\chi(v)$ is nonempty,
926 we also insert $[0, 0]$ in the data structure.

927 2. Next we compute the representation $\mathbb{I}(A(v))$ of $A(v)$. Because of Lemma 11 we have
 928 $A(v) = \mathbb{R}_{>0} \cap A'(v_1) \cap C'(v_2)$. Note that we cannot compute $C'(v_2)$ explicitly, since that
 929 would take $\Theta(m_2)$ time. Recall that we have an explicit minimal representation J_1, \dots, J_s of
 930 $A'(v_1)$. For each interval $J_i = [x_i, y_i]$ in the minimal representation of $A'(v_1)$, we use the
 931 operation $\text{IntClip}([x_i - \lambda(vv_2), y_i + \lambda(vv_2)])$ in the representation $\mathbb{I}(B'(v_2))$. Note that the
 932 intervals $[x_i - \lambda(vv_2), y_i + \lambda(vv_2)]$ over J_1, \dots, J_s may be intersecting, and therefore for
 933 index i we are actually obtaining the representation of

$$934 \quad B'(v_2) \cap \left([x_i - \lambda(vv_2), y_i + \lambda(vv_2)] \setminus \bigcup_{j < i} [x_j - \lambda(vv_2), y_j + \lambda(vv_2)] \right).$$

935 Nevertheless, using IntJoin over the representations reported we obtain the representation
 936 of the set (of zero-length intervals)

$$937 \quad X := B'(v_2) \cap \bigcup_{i \in [s]} [x_i - \lambda(vv_2), y_i + \lambda(vv_2)].$$

938 We then have

$$939 \quad A'(v_1) \cap C'(v_2) = \bigcup_{x \in X} [x - \lambda(vv_2), x + \lambda(vv_2)],$$

940 which means that we obtain a representation of $A'(v_1) \cap C'(v_2)$ from the representation of X
 941 using the operation $\text{IntExtend}(\lambda(vv_2))$. To obtain $\mathbb{I}(A(v))$ we apply $\text{IntClip}(\mathbb{R}_{>0})$. Since we
 942 are making $O(m_1)$ operations, we spend $O(m_1 \log m_2)$ time. Note that in this computation
 943 we have destroyed the representation of $\mathbb{I}(B'(v_2))$, and thus this step has to be made after
 944 the computation of $\mathbb{I}(B(v))$, which is also using $\mathbb{I}(B'(v_2))$ (or the equivalent representation
 945 $\mathbb{I}(B(v_2))$).

946 Consider now the case when $i(v) = i(v_2) \neq i(v_1)$. We proceed as follows.

947 1. First we compute the representation $\mathbb{I}(B(v))$ of $B(v)$. Because of Lemma 12 we have
 948 $B(v) = \chi(v) \cup (B'(v_2) \cap C'(v_1))$. We compute explicitly the minimal representation of $C'(v_1)$.
 949 Then, for each interval I in that representation we query for the elements $I \cap B'(v_2)$ using
 950 $\text{IntClip}(I)$ in $\mathbb{I}(B'(v_2))$ and join the answers using IntJoin over all intervals I . This takes
 951 $O(m_1 \log m_2)$ time and changes the data structure of the representation $\mathbb{I}(B'(v_2))$. Finally, if
 952 $\chi(v)$ is nonempty, we also insert $[0, 0]$ in the result. The total time is $O(m_1 \log m_2)$.

953 2. Next we compute the representation $\mathbb{I}(A(v))$ of $A(v)$. Because of Lemma 11 we have
 954 $A(v) = \mathbb{R}_{>0} \cap A'(v_2) \cap C'(v_1)$. Again, we compute explicitly the minimal representation of
 955 $C'(v_1)$. For each interval I in the minimal representation of $C'(v_1)$ we use the operation
 956 $\text{IntClip}(I)$ in $\mathbb{I}(A'(v_2))$ to obtain a representation of $I \cap A'(v_2)$, and then use IntJoin to join
 957 all the answers. With this we obtain a representation of $A'(v_2) \cap C'(v_1)$, to which we apply
 958 $\text{IntClip}(\mathbb{R}_{>0})$. This procedure takes $O(m_1 \log m_2)$ time and changes the representation of
 959 $\mathbb{I}(A'(v_2))$.

960 Consider now the remaining case, when $i(v) \neq i(v_1)$ and $i(v) \neq i(v_2)$. We proceed as follows.

- 961 1. The computation of $B(v)$ is trivial, since $B(v) = \chi(v)$ by Lemma 12.
- 962 2. The computation of the representation of $A(v) = C'(v_1) \cap C'(v_2)$ is similar to the case when
 963 $i(v) = i(v_1) \neq i(v_2)$. We compute explicitly the minimal representation of $C'(v_1)$, and use it
 964 as it was done there (for $\mathbb{I}(A'(v_1))$). This takes $O(m_1 \log m_2)$ time.

965 In each case we spent $O(m_1 \log m_2)$ time, and the time bound follows. For the upper bound on
 966 the representation $\mathbb{I}(A(v))$ of $A(v)$, we note that each left endpoint of each interval in $\mathbb{I}(A(v))$ gives
 967 rise to at most one interval in the representation of $A(v)$. The four terms correspond to the four
 968 possible cases we considered for the indices $i(v)$, $i(v_1)$ and $i(v_2)$. \square

969 **Lemma 17.** *The problem GENERALIZED GRAPHIC INVERSE VORONOI IN TREES for an input (T, \mathbb{U})*
 970 *where T is an n -vertex tree of maximum degree 3 and the candidate Voronoi cells are pairwise disjoint,*
 971 *can be solved in $O(n \log^2 n)$ time.*

972 *Proof.* We root T at a leaf so that each node has at most two descendants. For each vertex v
 973 of T , we compute a representation $\mathbb{I}(A(v))$ and $\mathbb{I}(B(v))$ of the sets $A(v)$ and $B(v)$, respectively.
 974 The computation is bottom-up: we compute $\mathbb{I}(A(v))$ and $\mathbb{I}(B(v))$ when this has been computed
 975 for all the children of v . If v has two children, we use Lemma 16. If v has one child, then the
 976 computation can be done in $O(\log m_A(v) + \log m_B(v))$ time in a straightforward manner. When
 977 we arrive to the root r , we just have to check whether $B(r)$ is nonempty.

978 We can see by induction that, for each vertex v of T , $m_A(v) \leq n(v)$. (We already mentioned
 979 earlier that $B(v)$ has at most $n(v)$ values, one per vertex of $T(v)$.) This is clear for the leaves
 980 because $A(\cdot)$ has only one interval. For the internal nodes v that have one child u it follows because
 981 the representation $\mathbb{I}(A(v))$ of $A(v)$ is obtained from the representation of $\mathbb{I}(A(u))$ by a shift. For
 982 the internal nodes v with two children v_1 and v_2 , the bound on $m_A(v)$ follows by induction from
 983 the bound in Lemma 16. In particular, we have $O(\log m_A(v) + \log m_B(v)) = O(\log n)$ at each node
 984 v of T .

985 For each vertex with one child we spend $O(\log n)$ time. For each vertex v with two children v_1
 986 and v_2 we spend $O(\min\{n(v_1), n(v_2)\} \log n)$ time. Thus, if V_1 and V_2 denote the vertices with one
 987 and two children, respectively, we spend

$$\begin{aligned}
 988 \quad O(n) + \sum_{v \in V_1} O(\log n) + \sum_{v \in V_2} O(\min\{n(v_1), n(v_2)\} \log n) \\
 989 \quad \quad \quad = O(n \log n) + O(\log n) \sum_{v \in V_2} O(\min\{n(v_1), n(v_2)\}) \\
 990
 \end{aligned}$$

991 time. Using Lemma 15, this time is $O(n \log^2 n)$.

992 Standard (but non-trivial) adaptations can be used to recover an actual solution. One option
 993 is to use persistent data structures for the search trees that store families \mathbb{I} of monotonic intervals.
 994 A persistent data structure allows to make queries to any version of the tree in the past. Thus, it
 995 stores implicitly copies of the trees that existed at any time. Sarnak and Tarjan [4] explain how to
 996 make red-black tree persistent (and how the Join and Split operations can also be done). Since
 997 we have access to the past versions of the tree, we can recover how the solution was obtained.
 998 Each operation in the past takes $O(\log m)$ time, where m is the sum of operations that were
 999 performed. In our case this is $O(\log(n \log^2 n)) = O(\log n)$ time per operation/query in the tree,
 1000 and the running time is not modified. Another, conceptually simpler option is to store through
 1001 the algorithm information on how to undo each operation. Then, at the end of the algorithm, we
 1002 can run the whole algorithm backwards and recover the solutions. \square

1003 **Theorem 18.** *The problem GENERALIZED GRAPHIC INVERSE VORONOI IN TREES for instances*
 1004 *$I = (T, ((U_1, S_1), \dots, (U_k, S_k)))$ can be solved in time $O(N + n \log^2 n)$, where T is a tree with n*
 1005 *vertices and $N = |V(T)| + \sum_i (|U_i| + |S_i|)$.*

1006 *Proof.* Because of Theorem 10, we can transform in $O(N)$ time the instance I to another instance
 1007 $I' = (T', ((U'_1, S'_1), \dots, (U'_k, S'_k)))$, where T' has maximum degree 3, the sets U'_1, \dots, U'_k are pairwise
 1008 disjoint, and T' has $O(n)$ vertices. We can compute a solution to instance I' in $O(n \log^2 n)$ time
 1009 using Lemma 17. Then, we have to check whether this solution is actually a solution for I . For
 1010 this we use Lemma 3. \square

1011 **Corollary 19.** *The problem GRAPHIC INVERSE VORONOI IN TREES for instances $I = (T, (U_1, \dots, U_k))$,*
 1012 *can be solved in time $O(N + n \log^2 n)$, where T is a tree with n vertices and $N = |V(T)| + \sum_i |U_i|$.*




Figure 11: Construction to show the lower bound in Theorem 20.

5 Lower bound for trees

We can show the following lower bound on any algorithm based on algebraic operations on the lengths of the edges.

Theorem 20. *In the algebraic computation tree model, solving GRAPHIC INVERSE VORONOI IN TREES with n vertices takes $\Omega(n \log n)$ operations, even when the lengths are integers.*

Proof. Consider an instance X, Y for the decision problem SET INTERSECTION, where $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_n\}$ are sets of integers. The task is to decide whether $X \cap Y$ is nonempty. This problem has a lower bound of $\Omega(n \log n)$ in the algebraic computation tree model [7]. (In particular, this implies the same lower bound for the bounded-degree algebraic decision tree model.) Adding a common value to all the numbers, we may assume that X and Y contain only positive integers.

We construct an instance to the GRAPHIC INVERSE VORONOI IN TREES problem, as follows. See Figure 11. We construct a star S_X with $n + 1$ leaves. The edges of S_X have lengths $x_1, \dots, x_n, 2$. We construct also a star S_Y with $n + 1$ leaves whose edges have lengths $y_1 + 1, \dots, y_n + 1, 1$. Finally, we identify the leaf of S_X incident to the edge of length 2 and the leaf of S_Y incident to the edge of length 1. Let T be the resulting tree. We take the sets U_1 and U_2 to be the vertex sets of S_X and S_Y , respectively. Note that T has $2n + 3$ vertices. The reduction makes $O(n)$ operations.

Since placing the sites on the center of the stars does not produce a solution, it is straightforward to see that the answers to SET INTERSECTION(X, Y) and to GRAPHIC INVERSE VORONOI IN TREES($T, (U_1, U_2)$) are the same. Thus, solving GRAPHIC INVERSE VORONOI IN TREES($T, (U_1, U_2)$) in $o(n \log n)$ time would provide a solution to SET INTERSECTION(X, Y) in $o(n \log n)$ time, and contradict the lower bound. \square

The lower bound also extends to the problem GENERALIZED GRAPHIC INVERSE VORONOI IN TREES with disjoint regions because we can apply the transformation to make the cells disjoint.

6 Conclusions

We have provided an algorithm for the inverse Voronoi problem in trees and a lower bound in a standard computation model. Since the upper bound of our algorithm and the lower bound differ, the main open question is closing this gap. Considering trees with unit edge lengths may also be interesting. Our lower bound does not apply for such instances.

1042 Acknowledgments

1043 We are very grateful to the anonymous reviewers for pointing out an error in the previous version
1044 of Section 3.2 and several other useful corrections.

1045 Part of this work was done at the 21st Korean Workshop on Computational Geometry, held in
1046 Rogla, Slovenia, in June 2018. We thank all workshop participants for their helpful comments.

1047 References

- 1048 [1] É. Bonnet, S. Cabello, B. Mohar, and H. Pérez-Rosés. The inverse Voronoi problem in graphs I:
1049 hardness, 2018. Manuscript. Its content is included in <http://arxiv.org/abs/1811.12547>.
- 1050 [2] P. Brass. *Advanced Data Structures*. Cambridge University Press, 2008, [https://doi.org/10.](https://doi.org/10.1017/CB09780511800191)
1051 [1017/CB09780511800191](https://doi.org/10.1017/CB09780511800191).
- 1052 [3] S. Chaudhuri and C. D. Zaroliagis. Shortest paths in digraphs of small treewidth. part
1053 I: sequential algorithms. *Algorithmica* 27(3):212–226, 2000, [https://doi.org/10.1007/](https://doi.org/10.1007/s004530010016)
1054 [s004530010016](https://doi.org/10.1007/s004530010016).
- 1055 [4] N. Sarnak and R. E. Tarjan. Planar point location using persistent search trees. *Commun.*
1056 *ACM* 29(7):669–679, 1986, <https://doi.org/10.1145/6138.6151>.
- 1057 [5] D. D. Sleator and R. E. Tarjan. Self-adjusting binary search trees. *J. ACM* 32(3):652–686,
1058 1985, <http://doi.acm.org/10.1145/3828.3835>.
- 1059 [6] R. E. Tarjan. Dynamic trees as search trees via euler tours, applied to the network simplex
1060 algorithm. *Math. Program.* 77:169–177, 1997, <https://doi.org/10.1007/BF02614369>.
- 1061 [7] A. C. Yao. Lower bounds for algebraic computation trees with integer inputs. *SIAM J. Comput.*
1062 20(4):655–668, 1991, <https://doi.org/10.1137/0220041>.