

Approximate and Exact Geometric Generalized Minimum Spanning Trees

Majid Mirzanezhad*

Arash Rafiey †

Abstract

We revisit the geometric variant of the generalized minimum spanning tree (GGMST) problem: we are given a set of n points in \mathbb{R}^d inside an integer grid with N non-empty clusters where each cluster is a unit hypercube, the aim is to find a minimum spanning tree that contains exactly one point from each cluster while respecting the adjacencies of the clusters. For any $\epsilon > 0$, we give an $(1 + \epsilon)$ approximation algorithm whose running time is $O^*(N \cdot (\frac{1}{\epsilon})^{O(\frac{1}{\epsilon})} + N^{d+1} + n \log n)$, where O^* hides factors polynomial in $1/\epsilon$. We also give an exact FPT algorithm that runs in $O(N^2 \rho^2 (2^{O(\kappa^4)}))$, where κ is an implicit parameter that relates to the input integer grid containing the non-empty clusters. Here, ρ is the maximum number of points across all clusters. When a GGMST instance is given in the plane, we present $(1 + \epsilon)$ -approximation with running time $O(N \rho^2 280^{\frac{1}{\epsilon}})$. In the case that for any non-empty cluster, either the row or column it belongs to, consists of at most $\gamma < N$ many non-empty clusters, we give a faster FPT algorithm that runs in $O(N \rho^2 560^\gamma)$ time.

1 Introduction

In *Generalized Minimum Spanning Tree* (GMST) problem, we are given an undirected abstract graph $G = (V, E)$ with vertex set V and edge set $E \subseteq V \times V$ together with a weight function $w : E \rightarrow \mathbb{R}^+$ and a partition of vertices V into non-intersecting clusters. The goal is to find a minimum weight tree containing exactly one vertex from each cluster. This problem is a generalization of the well-known *Minimum Spanning Tree* (MST) Problem where each vertex of the graph is a cluster itself. In the geometric variant of the problem, we are given N non-empty d -dimensional axis-aligned unit hypercubes, which we refer to as *clusters*, and a set of n points in \mathbb{R}^d spread across the N clusters. Each cluster contains at least two points. The aim is to pick exactly one point from each cluster and connect the chosen N points by straight-line segments (with Euclidean weights) to form an MST. We only allow edges between points chosen from *adjacent* clusters, where adjacency

means the clusters share at least a corner. This problem is also known as *Geometric Generalized Minimum Spanning Tree* (GGMST).

The GGMST appears in telecommunications network planning, where a network of node clusters needs to be connected via a tree architecture using exactly one node per cluster [7] and has applications in energy distribution and agricultural irrigation [9]. The GMST problem was first introduced by Myung et al. (1995) [13] and known to be NP-hard [6]. While there is no constant factor approximation unless $P = NP$, some heuristic algorithms have been developed [7, 9]. Pop et al. [15] lay out a 2ρ -approximation algorithm for GGMST where the size of each cluster is bounded by ρ . They take advantage of Integer programming and LP relaxation to tackle the problem. Feremans, Grigoriev, and Sitters consider GGMST [5] and show a strong NP-hardness, even if we restrict to instances in which all non-empty clusters are connected and each cluster contains at most two points (Theorem 1 in [5]). They design a dynamic programming algorithm that solves the problem with connected non-empty clusters in time $O(l \rho^{6k} 2^{34k^2} k^2)$ where the grid has size $l \times k$ and it is polynomial when k is fixed (Theorem 2 in [5]).

Bhattacharya et al. designed a polynomial-time approximation algorithm with a guaranteed ratio of $1 + 4\sqrt{2} + \epsilon$ [4]. Later, Kachooei et al. [10] showed that the problem is still NP-hard when there are at most two points in each cluster even with equal y -coordinates, and it is unlikely to have a fully polynomial time approximation scheme (FPTAS) unless $P = NP$. It is also known that the GGMST is NP-hard even if the cluster graph forms a tree [14].

Our results. We obtained two primary results in higher dimensions and then we show how they can be improved when the instance is given in the plane. We propose two approximation and FPT algorithms for each case as follows:

(1) We give a $(1 + \epsilon)$ -approximation algorithm that runs in $O(n \log n + n/\epsilon^d + (N/\epsilon)^{d+1} + N(\frac{1}{\epsilon})^{O(1/\epsilon)})$, in any constant dimension d (Thm. 5). Our approach utilizes Arora’s PTAS technique for TSP in [2]. Arora’s technique requires portal *pairings* (entering and exiting a quadtree cell in matching pairs) to form a tour. Our GGMST variant tracks *connected* subsets of portals (partitions) to encode how GGMST edges unite bound-

*School of Electrical Engineering and Computer Science, Ohio University, USA, miirza@ohio.edu

†Department of Computer Science, Indiana State University, IN, USA, arash.rafiey@indstate.edu

ary portals to form connected components.

(2) We also give an exact FPT algorithm that runs in $O(\rho^2 N^2 (2^{O(\kappa^4)}))$, where κ is the treewidth of the hypergraph induced by the non-empty clusters of the integer grid in \mathbb{R}^d , for any constant dimension d (Thm. 7). In telecom networks, especially large-scale ones, each local subnetwork (e.g., LAN, building, or regional cluster) designates exactly one node (its gateway or hub) to connect to external links. This is precisely the setting where each cluster has to select one hub.

(3) We show that the optimal GGMST in \mathbb{R}^2 can be computed in time $O(N\rho^2 280^k)$ where the instance $(\mathcal{G}, \mathcal{N}, P)$ lies in a grid of size $k \times l$ and ρ is the maximum number of points inside each cluster (Thm. 8). As a byproduct of this result, we can obtain the following PTAS. Using a similar approximation argument that appeared in [5], one can obtain a $(1 + \epsilon)$ -approximation for instance $(\mathcal{G}, \mathcal{N}, P)$ in time $O(N\rho^2 280^{\frac{1}{\epsilon}})$ (Corollary 9) which significantly improves the $(1 + \epsilon)$ -approximation algorithm with running time $O(\frac{l}{\epsilon^2} \rho^{\frac{6}{\epsilon}} 2^{\frac{34}{\epsilon^2}})$ in [5].

(4) In the case that for any non-empty cluster, either the row or column it belongs to, consists of at most $\gamma < N$ many non-empty clusters, we give another FPT algorithm in \mathbb{R}^2 that runs in $O(N\rho^2 560^\gamma)$ which is much faster than our FPT algorithm in \mathbb{R}^d (Thm. 10).

2 Preliminaries

Given two points $x, y \in \mathbb{R}^d$, the Euclidean distance between them is denoted by $\|x - y\|$. Computing minimum spanning trees requires setting a fixed point as the root of a tree. We denote a subtree rooted at a point x by T_x , consisting of all nodes and edges descending from x , encapsulating the hierarchical structure induced by the root. The tree's weight $W(T)$ represents the total edge lengths of T . When dealing with substructures, we use $W(T_x)$ to specify the weight of the subtree rooted at x .

A “*Hypergraph*” of an integer grid is a graph whose vertices are the non-empty clusters in the grid and two vertices share an edge iff the clusters are adjacent. Let $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ be an instance of the GGMST where \mathcal{G} is an HG, \mathcal{N} is the set of clusters with $|\mathcal{N}| = N$, and P is the set of points spread across the clusters in \mathcal{N} , with $|P| = n$. An HG \mathcal{G} in \mathcal{I} can attain a spanning tree whose nodes span all the clusters in \mathcal{I} . The said spanning tree is defined as follows: A “*Backbone Spanning Tree (BBST)*” w.r.t. some integer grid clusters is a spanning tree if (1) each vertex corresponds to a cluster, (2) all vertices span the entire non-empty clusters, and (3) two vertices share an edge if their corresponding clusters are adjacent to each other.

Let $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ be an instance of the GGMST problem. For a given BBST T , let $\text{OptMST}(T)$ denote the optimal tree corresponding to the BBST T . Note that every backbone spanning tree can induce a

GGMST whose adjacency of nodes follows the adjacency of nodes in the BBST. Let \mathcal{B} be the set of all possible BBSTs. Then: $T_{\text{opt}} := \min_{T \in \mathcal{B}} \text{OptMST}(T)$. In other words, a branching tree of the BBST T describes the adjacency of the vertices of a GGMST T^* where T^* selects the points of every cluster marked by T and connects them with respect to the adjacency of the clusters in T . While most of the notions are utilized in the main presentation, we included the omitted proofs in Appendix for the sake of brevity.

3 Approximation algorithm for GGMST

As mentioned earlier, the GGMST problem is NP-hard in general, but some underlying geometric structural properties can be utilized to achieve approximation schemes. Our goal is to provide a PTAS, albeit potentially with a high dependence on $1/\epsilon$.

3.1 Algorithmic sketch

We begin with discretizing each cluster into a finite set of candidate points independent of n . By choosing a sufficiently fine grid inside each cluster, we can guarantee that picking the nearest candidate point to the optimal solution's chosen point will not increase the weight (W) of optimal GGMST by more than $\epsilon \cdot W(T_{\text{opt}})$. Next, we construct a *WSPD-based* $(1 + \epsilon)$ -spanner to reduce the potential $O(n^2)$ edge set down to $O(n)$ edges without significantly increasing GGMST cost.

We then build a *quadtrees* on the clusters, placing $O(1/\epsilon)$ portals on each cell boundary to constrain how GGMST edges may cross those boundaries. Applying a bottom-up *dynamic programming* (DP) merges partial GGMST solutions from child cells by matching their portal-connectivity patterns, similar to Arora's technique for TSP that appeared in [2]. By carefully bounding the “overhead”, i.e., inflated weight added to the solution due to the approximation via the portal placement introduced at each scale (from forcing edges to go through portals) and leveraging the spanner's limited set of edges, the algorithm achieves a total GGMST cost within $(1 + \epsilon)$ of the true optimum, all in polynomial time for a fixed ϵ . Let $\delta > 0$, we partition each cluster uniformly into grid cells each of side length δ . This yields $O(\frac{1}{\delta^d})$ candidate points per cluster. Let P be the resulting set of all cluster grid points. The naïve algorithm below is a dynamic programming algorithm that can compute the exact GGMST given a BBST T .

Lemma 1 *For a given BBST T , one can compute the $\text{OptMST}(T)$ in polynomial time.*

Lemma 2 (Discretization) *Let $\epsilon > 0$. We can choose $\delta = \epsilon/(\sqrt{d}(N - 1))$ sufficiently small so that*

there exists a $(1 + \epsilon)$ -approximate GGMST solution selecting from these candidate points. Moreover, the number of grid points overall is $O(N^d/\epsilon^d)$

3.2 The dynamic program

We present the algorithm and analysis in four main steps: (1) quadtree construction and cluster isolation, (2) building a $(1 + \epsilon)$ -Spanner using WSPD, (3) portal placement and the DP scheme, and finally, (4) analysis of approximation guarantee and running time.

The first two steps are relatively straightforward to handle as follows: our first goal is to spatially organize the clusters and points so that we can apply a dynamic program to propagate the minimum weight of a spanning forest from each “zone” to the adjacent ones. We use a *quadtree* to recursively partition the plane until each leaf cell of the quadtree contains at most one cluster (thus fully containing that cluster). By ensuring each cluster resides in precisely one leaf cell, picking a representative grid point from that cluster is a local decision at the leaf. We utilize further to ensure a constant number of points per leaf limits state explosion in the DP state. Without further techniques, considering an MST among n points might require looking at $O(n^2)$ potential edges. To make the problem tractable, we reduce complexity by constructing a $(1 + \epsilon)$ -spanner. Given n points in the plane and $\epsilon > 0$, it is known that one can construct a $(1 + \epsilon)$ -spanner G with $O(n/\epsilon^d)$ edges in $O(n \log n + n/\epsilon^d)$ time [8]. Furthermore: $\text{MST}(G) \leq (1 + \epsilon)\text{MST}(P)$, where P is the set of points and $\text{MST}(P)$ is the optimal MST cost on these n points.

Portal placement and the DP scheme. Even with a spanner, we face the challenge of ensuring we pick exactly one point per cluster to form a GGMST. We approach this as follows: In the Arora-style PTAS algorithm, portals are used to discretize how solutions cross cell boundaries. Place $p = O(1/\epsilon)$ portals uniformly along each side of every quadtree cell. Any MST edge crossing the boundary must “snap” to the nearest portal. By introducing portals: i) we limit the complexity of boundary interactions in the DP states and ii) we ensure that any additional cost from snapping edges to portals is at most $\epsilon \cdot W(T_{\text{opt}})$ after appropriate scaling.

Lemma 3 *Let T_{opt} be an GGMST of cost $W(T_{\text{opt}})$ in a bounding box of side-length M . Suppose at quadtree level i , each cell has side length $\frac{M}{2^i}$. Let X_i be the number of GGMST edges that cross cell boundaries at level i . If every crossing edge is at least $\alpha \frac{M}{2^i}$ in length, then: $X_i \leq \frac{\beta W(T_{\text{opt}})}{\alpha (M/2^i)}$, for some constant $\alpha > 0$ and $\beta > 0$.*

Lemma 4 (Portal Approximation) *There is a way to choose $O(1/\epsilon)$ portals per cell boundary so that rerouting GGMST edges through these portals increases the GGMST cost by at most $\epsilon \cdot W(T_{\text{opt}})$.*

DP state definition and encoding connectivity. Consider a cell C in the quadtree. The DP state $D(C, S, B)$ might encode: (i) C : The cell for which a GGMST is being computed. (ii) S : Which clusters fully contained in C have chosen their representative point. This controls the GGMST to be a forest and then connects to other forests at a higher-level cell. Since each leaf cell has at most one cluster, and these sets merge as we go up, S is manageable. (iii) B : A description of how GGMST connections that leave C through its boundary portals are arranged. This is a pattern representing how portals are connected internally.

- *Portal Partitions:* To represent how portals connect inside C , we consider a partition of the set of portals on C 's boundary into connected components. Each connected component represents a partial GGMST structure inside C that connects some subset of these boundary portals. The number of partitions of a set of m elements is given by the m th Bell number B_m , which grows super-exponentially in k . For $m = p = O(1/\epsilon)$, we know that Bell number B_m is: $B_m = \Theta(m^m / (e^m \cdot m^{3/2}))$. Hence $B_m \leq m^m = O(1/\epsilon)^{O(1/\epsilon)}$. This encoding is crucial. By storing a “minimal structure” that indicates which portals are connected, we can later merge these configurations at higher-level cells.

- *DP Recurrence:* For a leaf cell C , we know exactly which points it contains (at most a constant number). We try all possible ways of picking one point for the cluster (if there is one) and then determine how this affects portal connectivity. In particular, for each choice, we compute the partial GGMST cost inside C (often trivial since it is just one cluster plus some boundary edges to portals). Then, we encode the resulting portal connectivity in a partition B .

$$D(C, S, B) = \underbrace{\text{PortalsCost}(p, B)}_{\text{how } p \text{ connects to boundary portals}}$$

where $\text{PortalsCost}(p, B)$ is the cost of hooking up p to the boundary portals in a way that yields the partition B . In other words, how does p connect to (some or all) of the portals in B ? This ensures that the internal MST structure is consistent with B . For an internal cell C with children C_1, C_2, C_3, C_4 , we combine their solutions:

$$D(C, S, B) = \min_{\substack{S_1 \cup S_2 \cup S_3 \cup S_4 = S \\ B_1, B_2, B_3, B_4}} \left(\sum_{i=1}^4 D(C_i, S_i, B_i) + \text{MERGECOST}(B_1, B_2, B_3, B_4, B) \right)$$




Figure 1: Connected components in cell S have formed through a detour through the adjacent cell S' reckoned by MERGECOST. The dotted lines display the edges of the underlying spanner. A partition of active (bold) portals on the boundary can be stored in B .

where MERGECOST involves adding edges from the spanner G between portals to ensure the combined configuration at C respects partition B , see Figure 1.

Analysis: approximation and running time. The approximation factors set to $\epsilon/3$ for various steps such as cluster discretization, spanner construction, and portal partitioning combine to yield a $(1 + \epsilon)$ -approximation overall. We now give a detailed view of the running time analysis, including how the complexity of handling portals is controlled. The runtime analysis is as follows:

(1) We discretize each of N cell into $O(N^d/\epsilon^d)$ cells, hence it takes $O(N^{d+1}/\epsilon^d)$ time. (3) We build a $(1 + \epsilon/3)$ -spanner G in time of $O(n \log n + n/\epsilon^d)$. (2) We build a quadtree on N cells in time $O(N \log N)$. The quadtree has $O(N)$ cells since we stop subdividing when each cell has at most a constant number of points and at most one cluster. Each leaf cell is thus simple to handle. Suppose each leaf has $l = O(\frac{N^d}{\epsilon^d})$ candidate points. For each candidate grid point, we check edges to up to $p = O(\frac{1}{\epsilon})$ portals. That's $O(l \cdot k) = O(\frac{N^d}{\epsilon^d} \cdot \frac{1}{\epsilon})$. Summation across all leaves, since we have $O(N)$ leaves, the total is $N \times O(\frac{N^d}{\epsilon^d} \cdot \frac{1}{\epsilon}) = O(\frac{N^{d+1}}{\epsilon^{d+1}})$. (4) At internal cells, we use a dynamic programming (DP) approach. The DP states consider: (4a) Which of the four subcells are represented in the cell (chosen at lower levels). (4b) How the GGMST inside the subcells connects to one another outside via a set of boundary portals. The challenge is that at each cell, we must consider different ways to connect portals. The number of ways to partition a set of m elements into connected subsets is at most the m th Bell number $B_m = (1/\epsilon)^{O(1/\epsilon)}$. Since the number of the cells in the quadtree is $O(N)$, the total would be $O(N(1/\epsilon)^{O(1/\epsilon)})$. Hence, the total runtime is $O(N^{d+1}/\epsilon^d) + O(n \log n + n/\epsilon^d) + O(N \log N) + O((N/\epsilon)^{d+1}) + O(N(1/\epsilon)^{O(1/\epsilon)})$, which is: $O(n \log n + n/\epsilon^d + (N/\epsilon)^{d+1} + N(1/\epsilon)^{O(1/\epsilon)})$.

Theorem 5 Given an instance $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ with $|\mathcal{N}| = N$ and $|P| = n$, where clusters are axis-aligned

unit hypersquares in \mathbb{R}^d , for any $\epsilon > 0$, one can compute a $(1 + \epsilon)$ -GGMST of \mathcal{I} in $O^*((1/\epsilon)^{O(1/\epsilon)}N + N^{d+1} + n \log n)$ time, where O^* hides factors polynomial in $1/\epsilon$.

4 An FPT algorithm for exact GGMST

Our algorithm breaks down into two levels by mixing a coarse dynamic program to enumerate all BBSTs like T and a fine-grained dynamic programming algorithm to obtain $\text{OptMST}(T)$ (using the algorithm in Lemma 1). However, the number of backbone spanning trees can grow exponentially with the size of the graph, making enumeration computationally challenging. Given an instance $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ with $|\mathcal{N}| = N$ and $|P| = n$, we present an algorithm that efficiently enumerates all spanning trees of our $\mathcal{G} = (V, E)$ with bounded treewidth κ , exploiting the structure provided by a nice tree decomposition, where V is the set of clusters and E is the set of edges representing the adjacency between the hypernode in V . Computing the exact treewidth of a graph is NP-complete [1]. However, Korhonen [12] presented an algorithm that computes a tree decomposition of width k' satisfying $\kappa' \leq 2\kappa + 1$, where κ is the optimal treewidth of \mathcal{G} . Given a graph G with treewidth κ , one can compute a tree decomposition of \mathcal{G} with width $k' \leq 2\kappa + 1$ in time $O(2^{O(k')} \cdot N)$ [12].

In our algorithm, we use the tree decomposition obtained from Korhonen's algorithm and κ' will be used as the treewidth parameter in our analysis. Every tree decomposition can be transformed into a "nice" tree decomposition with at most $O(N)$ nodes [11]. Our goal is to enumerate all combinatorially different spanning trees of \mathcal{G} . We achieve this by performing dynamic programming over the nice tree decomposition of \mathcal{G} . The key idea is to represent partial solutions at each node of the decomposition and combine them according to specific rules that ensure correctness and avoid redundancy.

Our algorithm consists of several components that update the state of connectivity of vertices in each bag and the acyclicity of the spanning tree induced by the vertices in the bag.

State Representation: At each node t of the nice tree decomposition, we maintain a set of DP states. Each state represents a set of partial spanning trees consistent with the subgraph induced by the vertices seen so far.

Definition 6 (DP State) A DP state at node t is a tuple (σ, δ) where:

(1) σ is a partition of the vertices in X_t , representing the connected components among them in the partial spanning tree.

(2) $\delta \subseteq E_t$ is the set of edges included in the partial spanning tree, where E_t is the set of edges between vertices in X_t .

Note that each state must satisfy the following invariants: (1) *Acyclicity*: The edges in δ do not form cycles, and (2) *Consistency*: The partition σ accurately reflects the connectivity induced by δ . Let DP_t denote the set of DP states at node t . We define recursive formulas for computing DP_t based on the type of node. We have the following breakdowns for our DP recurrences:

- *Leaf Nodes*: $\text{DP}_t = \{(\emptyset, \emptyset)\}$.

- *Introduce Nodes*: Let t be some node which is an introduce node introducing vertex v , with child t' . The DP states are computed as: $\text{DP}_t = \bigcup_{(\sigma', \delta') \in \text{DP}_{t'}} \bigcup_{\delta_v \subseteq E_v} \{(\sigma_v, \delta_v \cup \delta') \mid \text{Valid}(\sigma_v, \delta_v \cup \delta')\}$, where: E_v is the set of edges between v and vertices in $X_{t'}$, δ_v is a subset of E_v (possible ways v can connect to $X_{t'}$), σ_v is the updated partition obtained by adding v to σ' and merging parts if v is connected to vertices in δ_v , and $\text{Valid}(\sigma_v, \delta_v \cup \delta')$ ensures acyclicity and consistency.

- *Forget Nodes*: Let t be a forget node forgetting vertex v , with child t' . The DP states are computed as: $\text{DP}_t = \{(\sigma|_{X_t}, \delta') \mid (\sigma', \delta') \in \text{DP}_{t'}\}$, where $\sigma|_{X_t}$ is the restriction of the partition σ' to X_t .

- *Join Nodes*: Let t be a join node with children t_1 and t_2 . The DP states are computed as:

$$\text{DP}_t = \left\{ (\sigma, \delta_1 \cup \delta_2) \mid (\sigma, \delta_1) \in \text{DP}_{t_1}, (\sigma, \delta_2) \in \text{DP}_{t_2}, \right. \\ \left. \text{Valid}(\sigma, \delta_1 \cup \delta_2) \right\}.$$

The function $\text{Valid}(\sigma, \delta)$ returns true if the edges in δ do not form cycles. The partition σ correctly represents the connected components induced by δ . See Appendix for further explanation on the recursive formulas above.

To analyze the runtime, let $N = |V|$ be the number of vertices in \mathcal{G} and κ' be the width of the tree decomposition obtained from Korhonen's algorithm. At each node t , the size of the bag is $|X_t| \leq \kappa' + 1$. The number of possible partitions σ of X_t is given by the Bell number $B_{\kappa'+1}$ and the number of possible edge subsets δ among vertices in X_t is $2^{\binom{\kappa'+1}{2}}$. Since for any $m > 0$, the m th Bell number is: $B_m = \left((m/e)^m m^{3/2} \right)$, the total number of states at each node, and hence the required space is: $S(N, \kappa) = B_{\kappa'+1} \cdot 2^{\binom{\kappa'+1}{2}} = O\left(2^{O(\kappa^2)} \left(O(\kappa/e)^{2\kappa+2} O(\kappa)^{3/2} \right)\right) = O\left(2^{O(\kappa^2)}\right)$. Since $\kappa' \leq 2\kappa + 1$. Regarding time complexity, we need to explore each:

(I) *Introduce Nodes*: For each state in $\text{DP}_{t'}$, we consider $2^{|E_v|}$ subsets δ_v . Note that $|E_v| \leq \kappa'$, thus we have $2^{\kappa'}$ possibilities. The runtime per introduce node: $O(S \cdot 2^{\kappa'}) = O(S \cdot 2^{O(\kappa)})$. (ii) *Forget Nodes*: For each state in $\text{DP}_{t'}$, we perform operations in $O(1)$ time. Hence The runtime per forget node: $O(S)$. (iii) *Join Nodes*: For each pair of states from the two children, we check for compatibility and merge. Number of state pairs is $O(S^2)$. To obtain the total run-

time: The number of nodes in the nice tree decomposition is $O(N)$. Hence, the total time complexity is: $T(N, \kappa) = O(N(S^2 + S \cdot 2^{O(\kappa)})) = O\left(N(2^{O(\kappa^4)})\right)$.

Theorem 7 *Given $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ with $|\mathcal{N}| = N$, $|P| = n$, and κ as the treewidth of \mathcal{G} , one can compute the exact GGMST of \mathcal{I} in time $O\left(\rho^2 N^2 (2^{O(\kappa^4)})\right)$, where ρ is the maximum number of points in across all clusters in \mathcal{N} .*

5 Approximation and FPT algorithms in \mathbb{R}^2

We view the grid column-by-column. For any two consecutive columns only the cells that appear in those columns can be joined, so a *backbone forest* needs to be chosen inside every $k \times 2$ window. A direct enumeration shows that at most 280^k distinct forests can occur in a $k \times 3$ strip; hence we can exhaustively list every candidate backbone for the first three columns, and for each of them run a dynamic program that selects the best point from every non-empty cell and wires the choices together. It takes $O(\rho^2 N_3)$ time on a single forest, where ρ bounds the points per cluster and N_3 is the number of non-empty cells in the first three columns. We store the optimal value together with the induced forest on the last two columns and then slide the $k \times 2$ window one column to the right. At step i we match every stored forest on columns $(i-2, i-1)$ against each of the at most 280^k possibilities on $(i-1, i)$; the same dynamic program, now charged to the non-empty cells in $(i-2, i-1, i)$, updates the table. After the last column is processed the cheapest entry whose backbone is a tree is the optimum. Overall running time: $O(N \rho^2 280^k)$.

Theorem 8 *Given $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ in \mathbb{R}^2 with $|\mathcal{N}| = N$ and $|P| = n$ and with k fixed rows. The optimal GGMST of \mathcal{I} can be computed in $O(N \rho^2 280^k)$ time.*

Replacing the exact dynamic programme inside the sliding window by the PTAS of [5] yields a $(1 + \varepsilon)$ -approximation in $O(N \rho^2 280^{1/\varepsilon})$ time.

Corollary 9 *There is $(1 + \varepsilon)$ -approximation algorithm for GGMST of \mathcal{I} in \mathbb{R}^2 that runs in $O(N \rho^2 280^{1/\varepsilon})$ time.*

We say $(\mathcal{G}, \mathcal{N}, P)$ is γ -bounded if every non-empty cluster (i, j) has at most γ non-empty clusters in row i or column j . For a γ -bounded instance every column holds at most γ non-empty cells, so within any 3 consecutive columns we must decide how the vertices of a graph on $\leq 3\gamma$ labelled cells are joined. Lemma 14 (see Appendix) shows that the number of legal patterns is bounded by 280^γ (the constant 280 appears exactly as in Lemma 12, but raised only to γ instead of k). During the sweep we must remember, for every cell that sits on

the border between the two columns of the sliding window, whether its incident backbone edge continues left or right; this doubles the local state space, whence the factor $(2 \cdot 280)^\gamma = 560^\gamma$. Apart from this bookkeeping tweak, the algorithm is identical to the one sketched above: enumerate, run the same $O(\rho^2)$ dynamic program per pattern, and slide the window. The result is an FPT–algorithm as state in the following theorem:

Theorem 10 *Given $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ in \mathbb{R}^2 that is γ -bounded with $|\mathcal{N}| = N$ and $|P| = n$. The GGMST of \mathcal{I} can be computed in $O(N\rho^2 560^\gamma)$ time.*

6 Concluding remarks

In this paper we proposed some exact and approximation algorithms improving the runtime on the latest results on the GGMST problem. Our algorithms work in any dimension d and make use of some known techniques, the one by Arora appeared in [2]. Our main contribution is indeed the $(1 + \epsilon)$ -approximation for GGMST in \mathbb{R}^d . Arora’s original PTAS for TSP relies on portal pairings—ensuring entry and exit in matching pairs—to construct a tour, whereas our GMST variant instead tracks connected subsets (partitions) of portals to capture how GMST edges merge boundary portals into forests, achieving an approximation error of $\epsilon \cdot W(T_{opt})$ via a geometric series and Bell number analysis. Unlike TSP, where every point must appear in the tour, our formulation requires exactly one point per cluster; these activations are encoded at the leaf cells of our dynamic program to prevent redundancy at higher levels. Furthermore, since edges can only connect adjacent unit squares, the complexity associated with long-range edges is reduced. While Arora’s method typically places $O(\frac{\log n}{\epsilon})$ portals per cell side—resulting in runtime factors such as $n^{\mathcal{O}(1/\epsilon)}$ —our approach utilizes only $O(\frac{1}{\epsilon})$ portals, distributing overhead hierarchically across the quadtree levels, thereby avoiding the extra $\log n$ factor in portal spacing. As with Arora’s technique, the overall runtime depends super-polynomially on $\frac{1}{\epsilon}$, confirming that our method is a PTAS: it is polynomial in n for any fixed ϵ , but not fully polynomial in $\frac{1}{\epsilon}$. Our other algorithms are also improving on some known results by [5] in \mathbb{R}^2 and provide fixed parameter tractability for the case where the GMST instances behave naturally such as having bounded treewidth and/or γ -bounded as we introduced in this paper.

References

[1] S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM Journal on Algebraic and Discrete Methods*, 8(2):277–284, 1987.

[2] S. Arora. Polynomial-time approximation schemes for euclidean tsp and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.

[3] B. Bhattacharya, A. Čustić, A. Rafiey, A. Rafiey, and V. Sokol. Approximation algorithms for generalized mst and tsp in grid clusters. In Z. Lu, D. Kim, W. Wu, W. Li, and D. Z. Du, editors, *Combinatorial Optimization and Applications*, volume 9486 of *Lecture Notes in Computer Science*, pages 117–129. Springer, Cham, 2015.

[4] B. K. Bhattacharya, A. Custic, A. Rafiey, A. Rafiey, and V. Sokol. Approximation algorithms for generalized MST and TSP in grid clusters. In *Combinatorial Optimization and Applications - 9th International Conference, COCOA 2015, Houston, TX, USA*, volume 9486 of *Lecture Notes in Computer Science*, pages 110–125. Springer, 2015.

[5] C. Feremans, A. Grigoriev, and R. Sitters. The geometric generalized minimum spanning tree problem with grid clustering. *JOR*, 4:319–329, 2006.

[6] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[7] B. L. Golden, S. Raghavan, and D. Stanojevic. Heuristic search for the generalized minimum spanning tree problem. *INFORMS Journal on Computing*, 17(3):290–304, 2005.

[8] S. Har-peled. *Geometric Approximation Algorithms*. American Mathematical Society, USA, 2011.

[9] H. Jiang and Y. Chen. An efficient algorithm for generalized minimum spanning tree problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pages 217–224, 2010.

[10] H. A. Kachoei, M. Davoodi, and D. Tayebi. On the generalized minimum spanning tree in the euclidean plane. In *Proceedings of the 1st Iranian Conference on Computational Geometry (ICCG)*, pages 19–23, 2018.

[11] T. Kloks. *Treewidth: computations and approximations*, volume 842 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1994.

[12] T. Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2021.

[13] Y.-S. Myung, C. ho Lee, and D. wan Tcha. On the generalized minimum spanning tree problem. *Networks*, 26(4):231–241, 1995.

[14] P. C. Pop. The generalized minimum spanning tree problem. 2002.

[15] P. C. Pop, G. Still, and W. Kern. An approximation algorithm for the generalized minimum spanning tree problem with bounded cluster size. In *Proceedings of the Algorithms and Complexity in Discrete Mathematics (ACiD)*, pages 115–121, 2005.

Appendix

This appendix discloses the omitted proofs and algorithmic details skipped in each dedicated section.

Proof of lemmas and theorems in Section 3

Lemma 1 *For a given BBST T , one can compute the $\text{OptMST}(T)$ in polynomial time.*

Proof. Select an internal vertex of T corresponding to a cluster C as the root. For any leaf cluster $l \in T$, and for any point $p \in l$, set $W(T_p) = 0$. For any internal node $I \in T$, and any point $q \in I$, let $\text{Chd}(I)$ denote the set of children of I in BBST T . Then we have: $W(T_q) = \sum_{c \in \text{Chd}(I)} \min_{p \in c} \{W(T_p) + \|p - q\|\}$. The algorithm runs in $O(N\rho^2)$ time using a forward and then backward breadth-first search on T since the number of children per node is $O(2^{d+1}) = O(1)$, where $d = O(1)$ and $\rho < n$ is the maximum number of points in a cluster. \square

Lemma 2 (Discretization) *Let $\epsilon > 0$. We can choose $\delta = \epsilon/(\sqrt{d}(N-1))$ sufficiently small so that there exists a $(1 + \epsilon)$ -approximate GGMST solution selecting from these candidate points. Moreover, the number of grid points overall is $O(N^d/\epsilon^d)$.*

Proof. Consider an optimal GGMST solution T_{opt} that selects a point q_i in each cluster C_i . By construction of the grid inside C_i , there is a grid point p_i no more than $\delta\sqrt{d}$ away from q_i (since the cluster is unit size and we have a δ -spaced lattice).

Replace each q_i by p_i to obtain a new (approximate) solution T' . The increase in the length of any edge (q_i, q_j) in T_{opt} is at most proportional to δ because shifting points can increase distances by at most $\sqrt{d}\delta$. Thus, the total additional cost incurred by this rounding is at most $\sqrt{d}\delta N$, since the MST has at most $N-1$ edges and each edge length changes by at most $\sqrt{d}\delta$.

In the case that $N \leq 2^d$ and assuming $d = O(1)$, one can enumerate all $N^{N-2} = 2^{(d2^d)} = O(1)$ many BBSTs and compute GGMST with a minimum total weight in polynomial time in an exact manner utilizing the algorithm described in Lemma 1. For all $N > 2^d$, following the packing argument demonstrated in Lemma 1 in [3], we derive that $W(T_{\text{opt}}) > 1$. Hence: $W(T') \leq W(T_{\text{opt}}) + (N-1) \cdot \sqrt{d}\delta = W(T_{\text{opt}}) + (N-1) \cdot \frac{\epsilon}{N-1} = W(T_{\text{opt}}) + \epsilon \leq W(T_{\text{opt}}) + \epsilon \cdot W(T_{\text{opt}}) = (1 + \epsilon) \cdot W(T_{\text{opt}})$. \square

Lemma 3 *Let T_{opt} be an GGMST of cost $W(T_{\text{opt}})$ in a bounding box of side-length M . Suppose at quadtree level i , each cell has side length $\frac{M}{2^i}$. Let X_i be the number of GGMST edges that cross cell boundaries at level i . If every crossing edge is at least $\alpha \frac{M}{2^i}$ in length, then: $X_i \leq \frac{\beta W(T_{\text{opt}})}{\alpha (M/2^i)}$, for some constant $\alpha > 0$ and $\beta > 0$.*

Proof. Any GGMST edge e crossing a cell boundary at level i has length at least $\alpha \frac{M}{2^i}$, because it spans distinct cells of side $\frac{M}{2^i}$, for some $\alpha > 0$. Suppose there are X_i such crossing edges. Then these X_i edges alone have total length

$$\sum_{e \text{ crosses at level } i} \|e\| \geq X_i \left(\alpha \frac{M}{2^i} \right).$$

Applying a standard geometric pigeonhole argument, the GGMST cannot afford a set of edges whose combined length exceeds $\beta \cdot W(T_{\text{opt}})$, for some constant $\beta > 0$. Concretely, if the total length of all crossing edges at level i were allowed to surpass $\beta \cdot W(T_{\text{opt}})$, one could reconstruct a substructure heavier than $W(T_{\text{opt}})$, contradicting T_{opt} 's minimality.

Consider the MST edges crossing cell boundaries at level i , each having length at least $\alpha \frac{M}{2^i}$. The *pigeons* here are precisely those boundary-crossing edges themselves, while the *holes* represent ‘‘slots’’ in the MST's total cost $W(T_{\text{opt}})$. In other words, if each crossing edge (pigeon) demands at least an $\alpha \frac{M}{2^i}$ share of the MST length budget (a hole), and the MST cannot allocate more than $W(T_{\text{opt}})$ total, then having ‘‘too many’’ such edges (pigeons) would exceed the MST's entire weight (holes). Hence:

$$\sum_{e \text{ crosses at level } i} \|e\| \leq \beta W(T_{\text{opt}}).$$

Combining both inequalities yields:

$$X_i \left(\alpha \frac{M}{2^i} \right) \leq \beta W(T_{\text{opt}}),$$

which rearranges to

$$X_i \leq \frac{\beta W(T_{\text{opt}})}{\alpha \left(\frac{M}{2^i} \right)} = \frac{2^i \beta W(T_{\text{opt}})}{\alpha M}.$$

If X_i were any larger, the total cost of those crossing edges alone would exceed $W(T_{\text{opt}})$, which is a contradiction. \square

Lemma 4 (Portal Approximation) *There is a way to choose $O(1/\epsilon)$ portals per cell boundary so that rerouting GGMST edges through these portals increases the GGMST cost by at most $\epsilon \cdot W(T_{\text{opt}})$.*

Proof. Following the geometric scaling argument from Arora's PTAS, at each level of the quadtree, the cost added by snapping edges to portals is a small fraction of the GGMST's cost *at that scale*. Summing these fractions over all levels (cells at different scales) yields at most $\epsilon \cdot W(T_{\text{opt}})$ in total overhead. Let the bounding box of the entire cluster set have side length M . The quadtree is built by recursively splitting into four equal squares. At level 0 (the root), cells have side length M . At level i , cells have side length $\frac{M}{2^i}$. We place $p = O\left(\frac{1}{\epsilon}\right)$ portals on each cell boundary, so the portal spacing at level i is on the order of $\frac{M}{2^i \cdot p} = O\left(\epsilon \frac{M}{2^i}\right)$. The cost analysis unfolds as follows:

- *Overhead per Crossing at Level i :* Consider an GGMST edge e that crosses a boundary at level i . Snapping e from its original crossing point to the nearest portal adds at most $O\left(\epsilon \frac{M}{2^i}\right)$ extra length, since the portal spacing is $O\left(\epsilon \frac{M}{2^i}\right)$.

- *Bounding the Number of Crossings:* We now focus on all GGMST edges crossing cell boundaries at level i . Each such crossing indicates an edge of length at least $\Theta\left(\frac{M}{2^i}\right)$, since it spans across different cells at that scale. Following Lemma 3, if there were too many such ‘‘long’’ edges, their combined length would exceed $W(T_{\text{opt}})$, contradicting the GGMST's minimality. Hence, we conclude that the *total length* of edges crossing boundaries at level i is bounded by some constant factor of $W(T_{\text{opt}})$. Since each crossing can add up to $O\left(\epsilon \frac{M}{2^i}\right)$ extra length, the total *overhead* at this

level is at most $(\epsilon \cdot c) \cdot W(T_{\text{opt}})$ for some universal constant $c > 0$.

- *Summation Over All Levels:* The quadtree may have $O(\log N)$ levels (or $O(\log M)$). If we allocated a full $\epsilon W(T_{\text{opt}})$ overhead *per level*, we could end up with $O(\log N) \cdot \epsilon W(T_{\text{opt}})$. To avoid this, we use a *geometric distribution* of overhead: $\text{Overhead}_i \leq \frac{\epsilon}{2^i} W(T_{\text{opt}})$, at each level i . By placing portals more densely at coarser levels, the detour at scale i is forced to be only $\frac{\epsilon}{2^i} W(T_{\text{opt}})$. Summing these over $i = 0$ to $\log N$ gives $\text{Total Overhead} = \sum_{i=0}^{\log N} \text{Overhead}_i \leq \epsilon W(T_{\text{opt}}) \sum_{i=0}^{\log N} \frac{1}{2^i} < 2\epsilon W(T_{\text{opt}})$. Adjusting constants or using $\epsilon' = \frac{\epsilon}{2}$ in the construction ensures the total overhead does not exceed $\epsilon \cdot W(T_{\text{opt}})$. This completes the proof. \square

Proof of lemmas and theorems in Section 4

The key idea in Section 4 is to represent partial solutions at each node of the decomposition and combine them according to specific rules that ensure correctness and avoid redundancy.

DP formulation

Our algorithm consists of several components that update the state of connectivity of vertices in each bag and the acyclicity of the spanning tree induced by the vertices in the bag.

- *State Representation:* At each node t of the nice tree decomposition, we maintain a set of DP states. Each state represents a set of partial spanning trees consistent with the subgraph induced by the vertices seen so far.

Definition 11 (DP State) *A DP state at node t is a tuple (σ, δ) where:*

1. σ is a partition of the vertices in X_t , representing the connected components among them in the partial spanning tree.
2. $\delta \subseteq E_t$ is the set of edges included in the partial spanning tree, where E_t is the set of edges between vertices in X_t .

Note that each state must satisfy the following invariants:

(1) *Acyclicity:* The edges in δ do not form cycles, and (2) *Consistency:* The partition σ accurately reflects the connectivity induced by δ . Let DP_t denote the set of DP states at node t . We define recursive formulas for computing DP_t based on the type of node. We have the following breakdowns for our DP recurrences:

- *Leaf Nodes:* $\text{DP}_t = \{(\emptyset, \emptyset)\}$.
- *Introduce Nodes:* Let t be an introduce node introducing vertex v , with child t' . The DP states are computed as:

$$\text{DP}_t = \bigcup_{(\sigma', \delta') \in \text{DP}_{t'}} \bigcup_{\delta_v \subseteq E_v} \{(\sigma_v, \delta_v \cup \delta') \mid \text{Valid}(\sigma_v, \delta_v \cup \delta')\}$$

where: E_v is the set of edges between v and vertices in $X_{t'}$, δ_v is a subset of E_v (possible ways v can connect to $X_{t'}$), σ_v is the updated partition obtained by adding v to σ' and merging parts if v is connected to

vertices in δ_v , and $\text{Valid}(\sigma_v, \delta_v \cup \delta')$ ensures acyclicity and consistency.

- *Forget Nodes:*

Let t be a forget node forgetting vertex v , with child t' . The DP states are computed as:

$$\text{DP}_t = \{(\sigma|_{X_t}, \delta') \mid (\sigma', \delta') \in \text{DP}_{t'}\}$$

where $\sigma|_{X_t}$ is the restriction of the partition σ' to X_t .

- *Join Nodes:* Let t be a join node with children t_1 and t_2 . The DP states are computed as:

$$\text{DP}_t = \{(\sigma, \delta_1 \cup \delta_2) \mid (\sigma, \delta_1) \in \text{DP}_{t_1}, (\sigma, \delta_2) \in \text{DP}_{t_2}, \text{Valid}(\sigma, \delta_1 \cup \delta_2)\}$$

The function $\text{Valid}(\sigma, \delta)$ returns true if: The edges in δ do not form cycles. The partition σ correctly represents the connected components induced by δ .

Runtime and space analysis

Let $N = |V|$ be the number of vertices in \mathcal{G} and κ' be the width of the tree decomposition obtained from Korhonen's algorithm. At each node t , the size of the bag is $|X_t| \leq \kappa' + 1$. The number of possible partitions σ of X_t is given by the Bell number $B_{\kappa'+1}$ and the number of possible edge subsets δ among vertices in X_t is $2^{\binom{\kappa'+1}{2}}$. Since for any $m > 0$, the m th Bell number is: $B_m = \left(\frac{m}{e}\right)^m m^{3/2}$, the total number of states at each node, and hence the required space is:

$$\begin{aligned} S(N, \kappa) &= B_{\kappa'+1} \cdot 2^{\binom{\kappa'+1}{2}} \\ &= O\left(2^{O(\kappa^2)} \cdot \left(O\left(\frac{\kappa}{e}\right)^{2\kappa+2} \cdot O(\kappa)^{3/2}\right)\right) \\ &= O\left(2^{O(\kappa^2)}\right) \end{aligned}$$

Since $\kappa' \leq 2\kappa + 1$. Regarding time complexity, we need to explore each:

- *Introduce Nodes:* For each state in $\text{DP}_{t'}$, we consider $2^{|E_v|}$ subsets δ_v . Note that $|E_v| \leq \kappa'$, thus we have $2^{\kappa'}$ possibilities. The runtime per introduce node: $O(S \cdot 2^{\kappa'}) = O(S \cdot 2^{O(\kappa)})$.
- *Forget Nodes:* For each state in $\text{DP}_{t'}$, we perform operations in $O(1)$ time. Hence The runtime per forget node: $O(S)$.
- *Join Nodes:* For each pair of states from the two children, we check for compatibility and merge. Number of state pairs is $O(S^2)$.

To obtain the total runtime: The number of nodes in the nice tree decomposition is $O(N)$. Hence, the total time complexity is:

$$T(N, \kappa) = O(N(S^2 + S \cdot 2^{O(\kappa)})) = O\left(N\left(2^{O(\kappa^4)}\right)\right)$$

Theorem 7 *Given $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ with $|\mathcal{N}| = N$, $|P| = n$, and κ as the treewidth of \mathcal{G} , one can compute the exact GGMST of \mathcal{I} in time $O\left(\rho^2 N^2 (2^{O(\kappa^4)})\right)$, where ρ is the maximum number of points in across all clusters in \mathcal{N} .*

Proof. As shown above, given an instance \mathcal{I} , with κ as the treewidth of \mathcal{G} , one can enumerate all possible BBSTs in time $T(N, \kappa) = O\left(N\left(2^{O(\kappa^4)}\right)\right)$. Once all BBSTs like T are obtained, we can explicitly pick points from the selected clusters that result in the optimal GGMST, i.e., $\text{OptMST}(T)$ using Lemma 1. During enumeration of all BBSTs, we fix each $T \in \text{BBST}$ and utilize the algorithm described in Lemma 1 that runs in $O(N\rho^2)$. Therefore, the total runtime would be: $T(N, n, \kappa) = O\left(\rho^2 N^2\left(2^{O(\kappa^4)}\right)\right)$. \square

Further skipped details in Section 5

Let $(\mathcal{G}, \mathcal{N}, P)$ be an example of the GGMST problem with N non-empty cells where the non-empty cells are connected. Let ρ be the maximum number of points inside each cluster. We break the algorithms down into a separate lemma below:

Lemma 12 *The number of backbone forests in a grid $k \times 3$ (with some non-empty cells that) is at most 280^k .*

Proof. Let $F(k)$ denote the number of such forests. Let C_1, C_2, C_3 denote the clusters in row $k - 1$ and C_4, C_5, C_6 be the clusters in row k . Let H be a graph with vertices C_1, \dots, C_6 and edges $(C_1, C_4), (C_1, C_5), (C_2, C_4), (C_2, C_5), (C_2, C_6), (C_3, C_5), (C_3, C_6), (C_4, C_5), (C_5, C_6)$. Here, the set of edges of H is the adjacency of each of the clusters C_1, \dots, C_6 without considering the adjacency between C_1, C_2 and C_2, C_3 cells. $F(k) \leq h \times F(k-1)$ where h is the number of forests that are sub-graphs of G . It is not difficult to design a computer simulation and verify that $h = 280$. Therefore, $F(k) \leq 280^k$. \square

Theorem 8 *Given $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ in \mathbb{R}^2 with $|\mathcal{N}| = N$ and $|P| = n$ and with k fixed rows. The optimal GGMST of \mathcal{I} can be computed in $O(N\rho^2 280^k)$ time.*

Proof. Let $\mathcal{F}(i, j)$, where $i \leq j$, denote the set of possible forests restricted to the columns from i to j , covering all k rows of \mathcal{G} . By Lemma 12, there exist at most 280^k trees from $\mathcal{F}(1, 3)$ that can serve as the backbone of \mathcal{G} . For each such tree T , we apply the dynamic programming approach described in Lemma 1 to obtain $\text{OptMST}(T)$. This requires $O(N_3\rho^2 280^k)$ time to determine the optimal forest, where N_3 represents the number of non-empty clusters up to column 3. We store the backbone forest restricted to columns c_2 and c_3 along with the weight of the minimum spanning forest, computed using the dynamic programming approach from Lemma 1. Notice that no BBST connect two clusters which are two columns (rows) apart from each other. Thus, we only need to look at two consecutive columns and enumerate possible forests as backbones.

Next, we apply a sliding window technique, moving from left to right with a window of size $k \times 2$. At step $i - 1$, we assume that the minimum spanning forest whose backbone belongs to $\mathcal{F}(i - 2, i - 1)$ has already been computed. To expand to column i , we must consider all possible forests in $\mathcal{F}(i - 1, i)$ that overlap with $\mathcal{F}(i - 2, i - 1)$. This requires evaluating forests in $\mathcal{F}(i - 2, i - 1, i)$ to identify the valid overlaps. The selection of actual points from each cluster is again determined using the dynamic programming approach from Lemma 1, which runs in $O(N_i\rho^2)$ time for each given

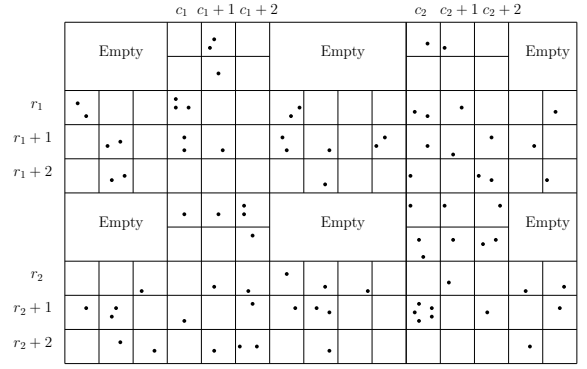


Figure 2: Example of the γ -bounded instance. Note that there could be non-empty cluster in row c_1 as well as in row r_1 , etc. The empty text means the clusters without any point in them.

forest, where N_i represents the number of non-empty cells across the three consecutive columns c_{i-2}, c_{i-1}, c_i . Since there are at most 280^k such forests, the runtime at step i is $O(N_i\rho^2 280^k)$. Thus, upon reaching the last column, the overall running time is $O(N\rho^2 280^k)$. Notice that at the end of the last row, we need to make sure we end up with a tree, eliminating the forests that are not tree after all. \square

Corollary 9 *There is $(1 + \epsilon)$ -approximation algorithm for GGMST of \mathcal{I} in \mathbb{R}^2 that runs in $O(N\rho^2 280^{\frac{1}{\epsilon}})$ time.*

Proof. We apply the approximation algorithm from [5] to obtain an $(1 + \epsilon)$ -approximation. Their approach includes a subroutine that computes OptMST within a grid of size $k \times l$, running in time $O(l\rho^{6k} 2^{34k^2} k^2)$. However, as shown in Theorem 8, this step can be performed in $O(\rho^2 N 280^k)$ time. By carefully analyzing their PTAS, we conclude that a GMST T can be computed such that its weight is at most $(1 + \epsilon) \cdot \text{OptMST}$ while maintaining an overall runtime of $O(N\rho^2 280^{\frac{1}{\epsilon}})$. \square

Definition 13 (γ -bounded) *Let $(\mathcal{G}, \mathcal{N}, P)$ be an instance of the GGMST. We say $(\mathcal{G}, \mathcal{N}, P)$ is γ -bounded if for every non-empty cluster (i, j) , (in row i and column j), the number of non-empty clusters in row i is at most γ or the number of non-empty clusters in column j is at most γ .*

By a simple adaptation of the argument of the Lemma 12 we have the following lemma (proof is straightforward):

Lemma 14 *The number of backbone forests for instance $(\mathcal{G}, \mathcal{N}, P)$ which lies in a $l \times 3$ grids and for each column 1, 2, 3 have at most γ non-empty cluster is at most 280^γ .*

By a simple adaptation of the argument of the Theorem 8 and Lemma 14, we have the following lemma.

Lemma 15 *Let $(\mathcal{G}, \mathcal{N}, P)$ be an instance of the GGMST problem so that the number of non-empty clusters in every column is at most γ . Then optimal GGMST can be computed with running time $O(N\rho^2 280^\gamma)$.*

We are now ready to conclude our main theorem:

Theorem 10 *Given $\mathcal{I} = (\mathcal{G}, \mathcal{N}, P)$ in \mathbb{R}^2 that is γ -bounded with $|\mathcal{N}| = N$ and $|P| = n$. The GGMST of \mathcal{I} can be computed in $O(N\rho^2 560^\gamma)$ time.*

Proof. Consider that \mathcal{G} is embedded within a grid of dimensions $\alpha \times \beta$. We partition this grid as follows : Let $c_1, c_1 + 1, \dots, c_1 + \ell_1$ denote a sequence of consecutive columns where each column contains at most γ non-empty clusters. Similarly, let $c_2, c_2 + 1, \dots, c_2 + \ell_2$ be the next sequence of consecutive columns satisfying where each of them has at most γ non-empty clusters. We continue this partitioning until reaching the last column β (refer to Figure 2 for visualization). Likewise, define $r_1, r_1 + 1, \dots, r_1 + s_1$ as the first set of consecutive rows where each row contains at most γ non-empty clusters. Define $r_2, r_2 + 1, \dots, r_2 + s_2$ as the next set of consecutive rows where each also contains at most γ non-empty clusters. We continue this partitioning until reaching the last row α .

Using Lemma 15, we can compute the minimum spanning forest for sub-grid restricted to columns from c_1 to $c_1 + \ell_1$ and all the α rows. Next, we compute the optimal minimum forest within the rows $r_1, r_1 + 1, \dots, r_1 + s_1$ up to column $c_1 + 1$, denoted as $\text{OptMST}(F_{r_1})$. Observe that the number of possible forests in these rows and the corresponding column c_1 is bounded by 280^γ . Since the dynamic programming approach processes the rows sequentially, we ensure that the constructed forest is compatible with the existing forest on column c_1 . This requires $O(N'\rho^2 280^\gamma \times 2^\gamma)$ time, where N' represents the number of non-empty clusters in the rows from r_1 to $r_1 + s_1$. Observe that there are at most 2^γ possible forests on column c_1 . A similar procedure is applied to the subsequent set of rows, $r_2, r_2 + 1, \dots, r_2 + s_2$, and the next block of columns $c_2, c_2 + 1, \dots, c_2 + s_2$. To track all possible forests restricted to any row and up to column $c_1 + 1$, we spend $O(N_1\rho^2 280^\gamma \times 2^\gamma)$ time, where N_1 represents the number of non-empty clusters in this restricted section. Finally, as we process the columns $c_1 + 1, c_1 + 2, \dots, c_1 + \ell_1$, we maintain a sliding window of size $\alpha \times 2$. Note that for each of these forests there is another super-forest \mathcal{S} that is a potential part of the OptMST . We compute the $\text{OptMST}(\mathcal{S})$ using the dynamic programming approach in Lemma 1. The approach proceeds iteratively:

1. Process the grid portion between columns $c_1 + \ell_1$ and c_2 row by row until reaching c_2 .
2. Continue processing column by column.
3. Repeat 1 and 2 until we reach to the last row and last column in \mathcal{G} .

Thus, the overall runtime complexity of the algorithm is bounded by $O(N\rho^2 560^\gamma)$, as required. \square