



Optimization problems in graphs with locational uncertainty

Marin Bougeret, Jérémy Omer, Michael Poss

► To cite this version:

Marin Bougeret, Jérémy Omer, Michael Poss. Optimization problems in graphs with locational uncertainty. 2021. hal-03331166v2

HAL Id: hal-03331166

<https://hal.science/hal-03331166v2>

Preprint submitted on 3 Sep 2021 (v2), last revised 20 Dec 2023 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimization problems in graphs with locational uncertainty

Marin Bougeret

LIRMM, University of Montpellier, CNRS, France, marin.bougeret@lirmm.fr

Jérémy Omer

IRMAR, INSA de Rennes, Rennes, France, jeremy.omer@insa-rennes.fr

Michael Poss

LIRMM, University of Montpellier, CNRS, France, michael.poss@lirmm.fr

Many discrete optimization problems amount to select a feasible subgraph of least weight. We consider in this paper the context of spatial graphs where the positions of the vertices are uncertain and belong to known uncertainty sets. The objective is to minimize the sum of the distances in the chosen subgraph for the worst positions of the vertices in their uncertainty sets. We first prove that these problems are \mathcal{NP} -hard even when the feasible subgraphs consist either of all spanning trees or of all $s - t$ paths. In view of this, we propose an exact solution algorithm combining integer programming formulations with a cutting plane algorithm, identifying the cases where the separation problem can be solved efficiently. We also propose two types of polynomial-time approximation algorithms. The first one relies on solving a nominal counterpart of the problem considering pairwise worst-case distances. We study in details the resulting approximation ratio, which depends on the structure of the metric space and of the feasible subgraphs. The second algorithm considers the special case of $s - t$ paths and leads to a fully-polynomial time approximation scheme. Our algorithms are numerically illustrated on a subway network design problem and a facility location problem.

Key words: combinatorial optimization, robust optimization, \mathcal{NP} -hardness, approximation algorithms, cutting plane algorithms

1. Introduction

Research in combinatorial optimization has provided efficient algorithms to solve a large variety of complex discrete decision problems, providing exact or near-optimal solutions in reasonable amounts of time. The applications are countless, ranging from logistics (network design, facility location, ...) to scheduling, including even important data science applications such as clustering. Many of these applications amount to select a subset of edges of a graph $G = (V, E)$ among a family of feasible subsets \mathcal{F} and that minimizes its total weight. Among those, we focus on spatial graphs on a given metric space (\mathcal{M}, d) , where each vertex i is assigned a position $u_i \in \mathcal{M}$ and the cost of set $F \in \mathcal{F}$ is given by $\sum_{\{i,j\} \in F} d(u_i, u_j)$, leading

to the combinatorial optimization problem

$$\min_{F \in \mathcal{F}} \sum_{\{i,j\} \in F} d(u_i, u_j). \quad (1)$$

Problem (1) encompasses many applications, such as network design, facility location, and clustering. These are typically subject to data uncertainty, be it because of the duration of the decision process, measurement errors, or simply lack of information. One successful framework that has emerged to address uncertainty is robust optimization (Ben-Tal and Nemirovski 1998), modeling the uncertain parameters with convex sets, such as polytopes, or finite sets of points, among which combinatorial robust optimization focuses on discrete robust optimization problems (Bertsimas and Sim 2003, Buchheim and Kurtz 2018, Kasperski and Zieliński 2016, Kouvelis and Yu 2013). We enter this framework by considering the model where the positions of the vertices are subject to uncertainty, therefore impacting the distances among the vertices. The resulting problem thus seeks to find the feasible subgraph that minimizes its worst-case sum of distances. Formally, we introduce for each vertex $i \in V$ the set of possible locations as the uncertainty set $\mathcal{U}_i \subseteq \mathcal{M}$. Using the notations $u = (u_1, \dots, u_{|V|})$ and $\mathcal{U} = \times_{i \in V} \mathcal{U}_i$, the general problem considered in this paper can be cast as

$$\min_{F \in \mathcal{F}} \max_{u \in \mathcal{U}} \sum_{\{i,j\} \in F} d(u_i, u_j). \quad (2)$$

As an illustration, the following three applications fall into the context of problem (2).

Example 1 (Subway network design). *Designing and expanding a subway network forms an important optimization problem faced by large cities. The new lines should efficiently cover dense areas of the city, while interacting well with the existing transportation lines. A key aspect of this problem amounts to locate the new subway stations. In addition to the technical considerations inherent to any construction, these also involve political considerations as local officials are never happy to let their citizens face the inconvenience of heavy civil engineering. This political lever is particularly complex in cities like Brussels having multiple local governments that must all agree before the stations can finally be constructed. As the overall process takes years, facing local government changes, the exact locations of the metro stations typically evolve between the first draft and the final implementation. Now, the exact locations of the stations impact the lengths of the resulting lines, the construction cost of which is typically proportional to their lengths (Gutiérrez-Jarpa et al. 2013). The cost of*

digging the new lines can therefore be modeled as a network design problem with locational uncertainty on the position of the vertices, usually including additional technical and environmental constraints. When a single line is of interest, \mathcal{F} will consist of paths, cycles, or trees.

Example 2 (Strategic facility location). *A production company wishes to expand its activities in a new region, locating additional facilities. We consider the context of a strategic level where the company may only choose approximate locations, as the exact locations will be known later, after all technical and legal considerations have been studied. Furthermore, the assignment of facilities to clients also needs to be considered in advance as each client may request different products, which require different installations at the facilities. As always in such facility location problems, the distances between the future clients and facilities lead to significant transportation costs that need to be kept as low as possible. In this particular case, these distances depend on locations that are uncertain at the time planning decisions are made. This leads to the problem of covering the vertices representing the clients with a certain number of disjoint stars, the root of each star representing a facility. Importantly, the distances are provided by the underlying road network (Melkote and Daskin 2001), which yields a graph-induced metric (\mathcal{M}, d) .*

Example 3 (Data clustering). *Given a set of data points, a fundamental problem in classification seeks to partition the points into a given number of subsets so as to minimize the dissimilarities among the points grouped into the same subset. The data points may consist of real values, for instance when they model physical measurements, as well as ordinal values, for instance, representing the results of medical or psychological surveys. In both cases, uncertainty on the values is common, be it because of a measurement error, or because of a lack of information – in which case the corresponding coordinate is replaced by the full interval (Masson et al. 2020). One must thus group the data in a way that is robust against these uncertainties. Various objective functions and metric spaces may be considered. With real values, it is common to minimize the sum of squared Euclidean distances. With ordinal values instead, a linear function may be more appropriate, coupled with ad hoc graph-induced distances. In all cases, the resulting problem seeks to partition the set of vertices with cliques.*

Traditionally, robust optimization problems with an objective function that is concave in the uncertain parameters are reformulated as monolithic models using conic duality (Ben-Tal and Nemirovski 1998). These techniques do not readily extend to function $d(u_i, u_j)$ as the

latter is non-concave in general. Actually, for Euclidean metric spaces based on the vector space $\mathbb{R}^\ell, \ell \in \mathbb{Z}^+$, $d(u_i, u_j) = \|u_i - u_j\|_2$ is convex in u_i and u_j . Function $\|u_i - u_j\|_2$ is closely related to the second-order cone (SOC) constraints considered by Zhen et al. (2021) for robust problems with polyhedral uncertainty sets. Zhen et al. (2021) linearize such robust SOC constraints by introducing adjustable variables, turning the problem into an adjustable robust optimization problem that can be tackled exactly (Ayoub and Poss 2016, Zhen et al. 2018, Zeng and Zhao 2013) or approximately using affine decision rules (Ben-Tal et al. 2004) or finite adaptability approaches (Bertsimas and Dunning 2016, Hanasusanto et al. 2015, Postek and Hertog 2016, Subramanyam et al. 2019), among others. Interestingly, the approach of Zhen et al. (2021), extended in Roos et al. (2018) to more general convex functions, makes no particular assumption on the feasibility set of the decision variables, herein represented by the set \mathcal{F} . A second work closely related to (2) is that of Citovsky et al. (2017), who rely on computational geometry techniques to provide constant-factor approximation algorithms in the special case where \mathcal{F} contains all Hamiltonian cycles of G . They propose in particular to solve a deterministic counterpart of (2) where the uncertain distances are replaced by the maximum pairwise distances $d_{ij}^{max} = \max_{u_i \in \mathcal{U}_i, u_j \in \mathcal{U}_j} d(u_i, u_j)$, for each $(i, j) \in V^2, i \neq j$.

To summarize, we see that while Zhen et al. (2021) provide valuable tools for addressing problems defined in Euclidean metric spaces considering uncertainty polytopes, their approaches cannot be used for graph-induced metric spaces, such as those mentioned in Examples 2 and 3. On the other hand, Citovsky et al. (2017) focused on the case where \mathcal{F} contains all Hamiltonian cycles of G . The main purpose of the present paper is thus to provide unified solution algorithms that are valid for any set \mathcal{F} and metric space (\mathcal{M}, d) . We only assume that \mathcal{U} is finite, which encompasses the two aforementioned cases. Specifically, when (\mathcal{M}, d) is a Euclidean space, the distance is a convex function implying that maximizing over polytope $\mathcal{U}^{pol} = \times_{i \in V} \mathcal{U}_i^{pol}$ is equivalent to maximizing over the extreme points of that polytope. Then, in the case of graph-induced metrics, the set \mathcal{M} is the set of nodes of a finite graph, meaning that each $\mathcal{U}_i \subseteq \mathcal{M}$ must be finite as well.

In this context, our contributions can be summarized as follows:

- We prove that problem (2) is \mathcal{NP} -hard even when \mathcal{F} consists of all $s - t$ paths and (\mathcal{M}, d) is the one-dimensional Euclidean metric space or when \mathcal{F} consists of all spanning trees of G . These results illustrate how the nature of problem (2) fundamentally differs

from the classical min-max robust problem with cost uncertainty, which is known to be polynomially solvable whenever the costs lie in independent uncertainty sets (Aissi et al. 2009).

- We provide a general cutting-plane algorithm for problem (2) that relies on integer programming formulations for \mathcal{F} . We further show that the separation problem $c(F) = \max_{u \in \mathcal{U}} \sum_{\{i,j\} \in F} d(u_i, u_j)$ is \mathcal{NP} -hard and provide two algorithms for computing $c(F)$. One is based on integer programming formulations while the other one relies on a dynamic programming algorithm that involves the treewidth of F .
- We extend the approximation algorithm based on d^{max} and suggested by Citovsky et al. (2017) to general sets \mathcal{F} and metric spaces different from the Euclidean one. We study in depth the resulting approximation ratios, which depend on the structure of \mathcal{F} and (\mathcal{M}, d) . Regarding \mathcal{F} , we obtain strong results for cycles, stars, trees, and graphs that can be composed as unions of simpler graphs. As a special case of our results, we find the ratio of 3 for Hamiltonian cycles, as previously proposed by Citovsky et al. (2017). Concerning the structure of (\mathcal{M}, d) , we show that metric spaces that satisfy Ptolemy’s inequality (Apostol 1967) benefit from stronger results.
- We provide a dynamic programming algorithm for the special case where \mathcal{F} consists of all $s - t$ paths, which is turned into a fully-polynomial time approximation scheme by rounding data appropriately.
- We compare numerically the exact cutting plane algorithm with the approximation algorithm that relies on d^{max} . The benchmark is composed of two families of instances. The first family includes Steiner tree instances that illustrate subway network design. The second one is composed of strategic facility location instances. The former application relies on two-dimensional Euclidean metric spaces so we can further include the affine decision rule reformulation from Zhen et al. (2021) to the comparison.

The rest of the paper is structured as follows. In Section 2, we introduce the definitions and notations that will be used throughout the paper. We then study the hardness of problem (2) in Section 3. In Section 4, we develop the exact solution algorithm, including the details of the computations of $c(F)$. We provide and study the approximation algorithm based on d^{max} in Section 5 while the fully-polynomial time approximation scheme

for $s - t$ paths is depicted in Section 6. In Section 7, we present our numerical experiments. Several appendices are annexed to the paper, containing most of the proofs, details on fixed-parameter tractable algorithms (\mathcal{FPT}), and the affine decision rules reformulation from Zhen et al. (2021).

2. Notations and definitions

In this section, we provide notations and definitions that will be used throughout the manuscript.

2.1 Graphs

We consider simple undirected graphs G , where $V(G)$ denotes the set of vertices of G and $E(G)$ its set of edges, and we assume that G has no isolated vertices. When clear from context we use notations $n = |V(G)|$ and $m = |E(G)|$. For any $i \in V(G)$, we denote by $N(i) = \{j \in V(G) \mid \{i, j\} \in E(G)\}$. We say that a graph G is a clique if for any two disjoint vertices i, j , $\{i, j\} \in E$, and we say that G is a star if G is a tree of maximum diameter 2. For any positive integer k , we denote $[k] = \{1, \dots, k\}$.

2.2 Cost functions

Given any position vector $u \in \mathcal{M}^n$ and subset of edges $F \subseteq E(G)$, we define $c(u, F) = \sum_{i,j \in F} d(u_i, u_j)$, and $c(F) = \max_{u \in \mathcal{U}} c(u, F)$. We extend these notations to any subgraph G' by $c(G') = c(E(G'))$ and $c(u, G') = c(u, E(G'))$. Moreover, we define $\text{diam}(\mathcal{U}_i) = \max_{u, u' \in \mathcal{U}_i} d(u, u')$.

2.3 Robust problems

For the statements and proofs of complexity and approximation results, we give more formal definitions of the inputs of the optimization problems we consider.

We are interested in the class \mathcal{S} of deterministic combinatorial optimization problems that can be formulated as $\min_{F \in \mathcal{F}} c(u, F)$. Any $\Pi \in \mathcal{S}$ represents a specific problem, such as the shortest path or the minimum spanning tree. An instance of problem Π is defined by its input $I = (G, \alpha, u, D)$, where G is the aforementioned graph, α contains problem-specific additional input, $u \in \mathcal{M}^n$ denotes the position of the nodes, and D is a (symmetric) matrix recording the distance between each pair $u, u' \in \bigcup_{i \in V(G)} \{u_i\}$. Thus, \mathcal{F} contains the subsets

of edges of G that satisfy the constraints specific to Π for input I . As an example, for the Δ -Shortest Path problem (Δ -SP), the additional input $\alpha = \{s, t\}$ consists of the origin and destination nodes, so the objective is to find an $s - t$ path P minimizing $c(u, P)$. In the Δ -Minimum Spanning Tree problem (Δ -MST), $\alpha = \emptyset$, and the objective is to find a spanning tree T minimizing $c(u, T)$.

For any problem $\Pi \in \mathcal{S}$, we define its robust counterpart ROBUST- Π . The input of an instance of ROBUST- Π is $I = (G, \alpha, \mathcal{U}, \mathcal{D})$ where G and α are as before, $\mathcal{U}_i \subseteq \mathcal{M}$ contains the possible positions for node i , and \mathcal{D} is a (symmetric) matrix recording the distance between each pair $u, u' \in \bigcup_{i \in V(G)} \mathcal{U}_i$. The objective of ROBUST- Π is to find $F \in \mathcal{F}$ that minimizes $\max_{u \in \mathcal{U}} c(u, F) = c(F)$. Observe that ROBUST- Π is a generalization of Π as it corresponds to the case where $|\mathcal{U}_i| = 1$ for any i .

We are also interested in the problem of playing the role of the adversary, corresponding to computing $c(G)$ for a given graph G . For problem EVAL-C, an input $I = (G, \mathcal{U}, \mathcal{D})$ where G, \mathcal{U} and \mathcal{D} are as before.

For all these optimization problems, we define $\text{OPT}(I)$ as the optimal solution cost of instance I . When clear from the context, we may also use the shorter notation OPT .

3. Hardness of the robust problem

Let Π be the assignment problem. Any instance of Π involves a partition of V in two sets of equal size, denoted V^1 and V^2 , such that any $F \in \mathcal{F}$ contains exactly $n/2$ edges covering V^1 and V^1 . Because the edges belonging to any $F \in \mathcal{F}$ are disjoint, we have that

$$c(F) = \sum_{\{i,j\} \in F} \max_{u_i \in \mathcal{U}_i, u_j \in \mathcal{U}_j} d(u_i, u_j) = \sum_{\{i,j\} \in F} d_{ij}^{max}.$$

Hence, ROBUST- Π can be solved in polynomial time by solving the deterministic assignment problem defined by the partition $V_1 \cup V_2$ and the vector of weights d^{max} . In spite of this easy example, we show in this section that ROBUST- Π can in general not be reduced to Π as ROBUST- Π is typically harder than Π .

We illustrate the hardness of ROBUST- Π by focusing on the two polynomially solvable problems Π mentioned in Section 2.3, namely the shortest path problem Δ -SP and the minimum spanning tree problem Δ -MST. These problems have been largely studied in the robust combinatorial optimization literature under cost uncertainty (e.g. Kasperski and Zieliński

(2009), Yaman et al. (2001)), namely

$$\min_{F \in \mathcal{F}} \max_{c \in \mathcal{C}} \sum_{\{i,j\} \in F} c_{ij}, \quad (3)$$

where \mathcal{C} is a given uncertainty set. It is folklore that when \mathcal{C} is the Cartesian product of intervals, $\mathcal{C} = \times_{e \in E(G)} [\underline{c}_e, \bar{c}_e]$ for $\underline{c}_e \leq \bar{c}_e$, problem (3) can be reformulated as $\min_{F \in \mathcal{F}} \sum_{\{i,j\} \in F} \bar{c}_{ij}$ making the robust problems as easy as their nominal counterparts. The result below, proved in Appendix A.1, shows that such is not the case for ROBUST-II, as Δ -SP turns \mathcal{NP} -hard even in the simple case where each \mathcal{U}_i is a subset of \mathbb{R} . Notice that in the 1-dimensional Euclidean space, the convexity of $d(u_i, u_j) = |u_i - u_j|$ implies that \mathcal{U}_i is equivalent to a segment $[\underline{u}_i, \bar{u}_i]$, for some $\underline{u}_i \leq \bar{u}_i$.

Proposition 1. *ROBUST- Δ -SP is \mathcal{NP} -hard, even when (\mathcal{M}, d) is the 1-dimensional Euclidean space.*

For Δ -MST we can prove the hardness of the problem only for a more general metric space. See Appendix A.2 for the proof.

Proposition 2. *ROBUST- Δ -MST is \mathcal{NP} -hard.*

We also detail in Appendix A.2 how, for any positive integer ℓ , the metric space (\mathcal{M}, d) used in the proof of Proposition 2 cannot be embedded isometrically into \mathbb{R}^ℓ . As a consequence, the hardness of ROBUST- Δ -MST in Euclidean spaces remains an open problem.

Having proved that ROBUST-II is \mathcal{NP} -hard, we will develop two types of solution algorithms for the problem. In the next section, we provide a cutting plane algorithm based on integer programming formulations for \mathcal{F} , which has an exponential running-time in general. Then, in Sections 5 and 6, we design polynomial time algorithms that yield approximate solutions to ROBUST-II.

4. Exact solution of the general robust problem

A popular type of algorithms solving exactly difficult robust optimization problems replaces the large uncertainty set by an approximation of small cardinality, leading to a relaxation of the original problem. Then, these algorithms iterate between solving integer programming formulations for the robust problem with small uncertainty set, and checking the optimality of the solution for the relaxation by solving an adversarial separation problem. This process

leads to cutting plane algorithms (e.g., Bertsimas et al. (2016), Fischetti and Monaci (2012), Naoum-Sawaya and Buchheim (2016)). Such algorithms involve frequent calls to computing the objective function EVAL-C, so we start this section by studying how to efficiently solve this problem. Then, we detail in Section 4.2 the overall cutting plane algorithm for ROBUST-II.

4.1 Hardness and exact solution of the cost evaluation problem

4.1.1 Negative results

Our first result (Proposition 3 below) is that EVAL-C is hard, even when the metric space is reduced to two points, or the input graph is a clique. For this, we consider particularly simple metric spaces, and rely on a reduction from problem MAX-CUT. We recall that MAX-CUT is a famous problem in combinatorial optimization that, given any input graph G , seeks a partition $\{V_1, V_2\}$ of $V(G)$ such that $|\{e \in E(G) : |e \cap V_1| = 1\}|$ is maximized. The detailed proof is provided in Appendix B.1.

Proposition 3. *Even when $|\mathcal{M}| = 2$, there is no \mathcal{PTAS} for EVAL-C unless $\mathcal{P} = \mathcal{NP}$.*

Our second result assumes instead that G is a clique, and relies on a reduction from problem MAX-DIVERSITY. Given a ground set X , a natural number k and a diversity function $div : 2^X \rightarrow \mathbb{R}$, MAX-DIVERSITY seeks to output a subset $S \subseteq X$ of size k that maximizes $div(S)$. The special case of interest here, proved \mathcal{NP} -hard in Cevallos et al. (2018), considers that the ground set X is the three-dimensional Euclidean space, and that $div(S) = \sum_{i,j \in S} \|i - j\|_2$. We obtain immediately the following result.

Proposition 4. *Even when restricted to input graphs that are cliques, to metric space defined by squared Euclidean distance in the three-dimensional Euclidean space, and to instances where $\mathcal{U}_i = \mathcal{M}$ for all i , EVAL-C is \mathcal{NP} -hard.*

Let us now turn to parameterized complexity, and let tw be the treewidth of G (see Appendix F.2 for the definition of treewidth). Informally, tw measures the thickness of a tree structure describing G . In particular, $tw(G) = 1$ for any tree G . As we show in the next section that computing EVAL-C is polynomial on trees, a natural question is to determine if we can extend this result by proving that EVAL-C/ tw admits an \mathcal{FPT} algorithm, meaning an algorithm running in $f(tw) \cdot |I|^c$ for some computable function f and constant c (see Appendix F.1 for details). The following proposition implies that it is very unlikely, and thus places EVAL-C with the few problems that are not \mathcal{FPT} by treewidth. See Appendix B.2 for the proof of the proposition.

Proposition 5. EVAL-C/ tw is $\mathcal{W}[1]$ -hard.

4.1.2 Positive results

Given the aforementioned hardness results, we propose two approaches to computing $c(G)$. The first approach relies on an integer programming formulation. For each $i \in V$ and $k \in \{1, \dots, |\mathcal{U}_i|\}$, binary variable y_i^k takes value 1 if and only if vertex i is located at position u_i^k . Therefore, $c(G)$ is equal to

$$\begin{aligned} \max \quad & \sum_{\{i,j\} \in E} \sum_{k=1}^{|\mathcal{U}_i|} \sum_{\ell=1}^{|\mathcal{U}_j|} d(u_i^k, u_j^\ell) y_i^k y_j^\ell \\ \text{s.t.} \quad & \sum_{k=1}^{|\mathcal{U}_i|} y_i^k = 1, \quad \forall i \in V \\ & y_i \in \{0, 1\}^{|\mathcal{U}_i|}, \quad \forall i \in V \end{aligned}$$

which can be linearized using classical techniques.

It is also possible to compute $c(G)$ efficiently whenever G has small treewidth $tw(G)$ using a dynamic programming algorithm. Let us detail the algorithm whenever G is a tree rooted at vertex r , which we assume oriented from r to its leaves L . We denote by $D(i)$ the set that contains the direct descendants of i , which is empty if i is a leaf. Let $\text{OPT}(i, u_i)$ be the maximum value obtained for the subtree starting at i given that node i is located at u_i . We obtain the following recursion:

$$\text{OPT}(i, u_i) = \begin{cases} \sum_{j \in D(i)} \max_{u_j \in \mathcal{U}_j} d(u_i, u_j) + \text{OPT}(j, u_j), & i \in V(G) \setminus L \\ 0, & i \in L \end{cases} \quad (4)$$

and the optimal solution cost is given by $\max_{u_r \in \mathcal{U}_r} \text{OPT}(r, u_r)$. Dynamic programming recursion (4) will be used in our numerical experiments, which involve trees and stars.

Recall that $tw = tw(G)$ and let us further denote $\sigma = \max_{i \in V(G)} |\mathcal{U}_i|$. Using dynamic programming on a nice tree decomposition of G (see Appendix F.2 for the definition), one can readily extend the above idea to any graph of bounded treewidth, leading to Theorem 1, whose proof is deferred to Appendix G. We point out that according to Proposition 5 we cannot (unless $\mathcal{W}[1] = \mathcal{FPT}$) remove the dependency in σ to get for example a $\mathcal{O}(\text{poly}(n) \times f(tw))$, and this holds for any computable function f .

Theorem 1. EVAL-C/ $tw + \sigma$ is \mathcal{FPT} . More precisely, we can compute an optimal solution of EVAL-C in time $\mathcal{O}(n \times tw \times \sigma^{\mathcal{O}(tw)})$.

4.2 Cutting plane algorithm for the robust problem

Now that we have depicted numerical methods for computing EVAL-C, we wish to make the extra step towards the exact solution of the complete problem, ROBUST-II. For this, we design an exact solution algorithm that generates scenarios of \mathcal{U} on the fly in the course of a branch-and-cut algorithm. We thus choose to associate to each $F \in \mathcal{F}$ the binary vector $x^F \in \{0, 1\}^m$ such that $x_{ij}^F = 1$ if and only if $\{i, j\} \in F$. As a result, the counterpart of \mathcal{F} is the set of binary vectors $\mathcal{X} = \{x^F, F \in \mathcal{F}\}$. With these notations, ROBUST-II can be reformulated as

$$\min_{x \in \mathcal{X}} \max_{u \in \mathcal{U}} \sum_{\{i,j\} \in E} x_{ij} d(u_i, u_j). \quad (5)$$

Let $\tilde{\mathcal{U}}$ be a finite subset of \mathcal{U} . An exact algorithm for (5), described in Algorithm 1, relies on the following relaxed formulation

$$\min \left\{ \omega \mid \omega \geq \sum_{\{i,j\} \in E} x_{ij} d(u_i, u_j), \forall u \in \tilde{\mathcal{U}}, x \in \mathcal{X} \right\}. \quad (6)$$

Algorithm 1 describes an iterative cutting-plane implementation, alternating between the solution of the relaxed master problem (6) and the adversarial separation problem $c(F)$. Practical implementation of these algorithms typically rely instead on branch-and-cut algorithms, where the adversarial separation problem is solved at each integer node of the branch-and-bound-tree.

Algorithm 1: Cutting-plane algorithm for (2)

repeat

 Let $(\tilde{\omega}, \tilde{x})$ be an optimal solution of (6)

 Let F be the support of \tilde{x}

 Compute $c(F) = \max_{u \in \mathcal{U}} \sum_{\{i,j\} \in F} d(\tilde{u}_i, \tilde{u}_j)$ and let \tilde{u} be a maximizer

if $c(F) > \tilde{\omega}$ **then** $\tilde{\mathcal{U}} \leftarrow \tilde{\mathcal{U}} \cup \{\tilde{u}\}$

until $c(F) \leq \tilde{\omega}$

return F

For large and complex problems, one can hardly expect Algorithm 1 to run quickly, so it may be wiser to first try a quick approximation algorithm. We provide in the next section such an algorithm and study in depth the worst-case bound between the solution cost returned by the algorithm and the optimal solution cost.

5. Approximation algorithm for the general robust problem

5.1 Reduction to a deterministic problem by using worst case distances

Algorithm 2: Solving a deterministic counterpart based on some representative location u^*

Given an instance of ROBUST-II: $\min_{F \in \mathcal{F}} \max_{u \in \mathcal{U}} \sum_{\{i,j\} \in F} d(u_i, u_j)$
 Select $u^* \in \mathcal{U}$
 Compute F using an approximation algorithm for $\min_{F \in \mathcal{F}} \sum_{\{i,j\} \in F} d(u_i^*, u_j^*)$
return F

A simple approach to ROBUST-II would be to choose a relevant vector u^* to obtain the corresponding problem II, given by $\min_{F \in \mathcal{F}} \sum_{\{i,j\} \in F} d(u_i^*, u_j^*)$, and use any known approximation algorithm for the deterministic problem II. This approach is formalized by Algorithm 2.

Unfortunately, choosing such a representative u^* is not easy. For instance, a natural choice might be to consider the barycenters of each set, i.e., $u^* = u^{bc}$ where $u_i^{bc} \in \arg \min_{u_1 \in \mathcal{U}_i} \sum_{u_2 \in \mathcal{U}_i} d(u_1, u_2)$. Although the choice of barycenters may appear natural at first glance, the cost of the solution F^{bc} returned by Algorithm 2 for $u^* = u^{bc}$ may actually be arbitrarily larger than the optimal solution cost.

Observation 1. Let $\Pi \in \mathcal{S}$ such that $\mathcal{F} = E(G)$. Let e^{bc} be the solution returned by Algorithm 2 for $u^* = u^{bc}$. The ratio $c(e^{bc})/\text{OPT}$ is unbounded, even when (\mathcal{M}, d) is included in the one-dimensional Euclidean metric space.

Proof. For any $\epsilon > 0$ small enough, consider $V = \{1, 2, 3\}$ and $E = \{\{1, 2\}, \{2, 3\}\}$ with $\mathcal{U}_1 = \{\epsilon\}$, $\mathcal{U}_2 = \{0\}$, and $\mathcal{U}_3 = \{-1, 1\}$. We have that $u_1^{bc} = \epsilon$ and $u_2^{bc} = u_3^{bc} = 0$ so Algorithm 2 picks edge $\{2, 3\}$ having a cost of $c(\{2, 3\}) = 1$. In contrast edge $\{1, 2\}$ has a cost of ϵ , yielding a ratio of $1/\epsilon$. \square

We proceed by using a different approach for constructing the deterministic problem II. For each $\{i, j\} \in E$, we define the worst-case distances as $d_{ij}^{max} = \max_{u_i \in \mathcal{U}_i, u_j \in \mathcal{U}_j} d(u_i, u_j)$, and define similarly the value $c^{max}(F) = \sum_{\{i,j\} \in F} d_{ij}^{max}$ for any $F \in \mathcal{F}$.

Notice that, in general, there is no $u^* \in \mathcal{U}$ such that $d(u_i^*, u_j^*) = d_{ij}^{max}$ for any i and j : the worst-case position $u_i \in \mathcal{U}_i$ of a node i typically varies with the edge $\{i, j\} \in E$

depending on the uncertainty set \mathcal{U}_j . Then, rather than choose a particular vector $u^* \in \mathcal{U}$, we define a new metric space (\mathcal{M}', d') where $\mathcal{M}' = \{\mathcal{U}_i, i \in V(G)\}$, and for any vertices $i \neq j$, $d'(\mathcal{U}_i, \mathcal{U}_j) = d_{ij}^{max}$. Observe that d' is indeed a distance. In particular, $d'(\mathcal{U}_i, \mathcal{U}_j) = 0$ means that $\mathcal{U}_i = \mathcal{U}_j$ and that both sets are reduced to a singleton. Consequently, $\min_{F \in \mathcal{F}} c^{max}(F) = \sum_{\{i,j\} \in F} d'(\mathcal{U}_i, \mathcal{U}_j)$ is an instance of Π . The resulting approach is summarized in Algorithm 3, which transfers approximability of Π to ROBUST- Π .

Theorem 2. *Let $\Pi \in \mathcal{S}$ where the output F belongs to a graph class \mathcal{F} . Suppose that there is a polynomial ρ_1 -approximation for Π , and that $c^{max}(F) \leq \rho_2 c(F)$ for any $F \in \mathcal{F}$. Then, using this approximation in Algorithm 3 yields a polynomial $\rho_1 \rho_2$ -approximation for ROBUST- Π .*

Proof. Let F be the output of Algorithm 3, F^* be an optimal solution of the instance of ROBUST- Π , and $F^{max} \in \arg \min_{F \in \mathcal{F}} c^{max}(F)$. Then $c(F) \leq c^{max}(F) \leq \rho_1 c^{max}(F^{max}) \leq \rho_1 c^{max}(F^*) \leq \rho_1 \rho_2 c(F^*)$. \square

Algorithm 3: Solving a deterministic counterpart based on d^{max} distances

Given an instance of ROBUST- Π : $\min_{F \in \mathcal{F}} \max_{u \in \mathcal{U}} \sum_{\{i,j\} \in F} d(u_i, u_j)$
 Compute F using an approximation algorithm for $\min_{F \in \mathcal{F}} c^{max}(F)$
return F

In Section 5.2 and Section 5.3 we prove the existence of constant upper bounds on the ratio $c^{max}(F)/c(F)$ for different families \mathcal{F} and metric spaces. A summary of our results is given in Table 1. In particular, for any graph F , Theorem 4 states that $c^{max}(F) \leq 4c(F)$ for any Ptolemaic metric space, and Theorem 5 states that $c^{max}(F) \leq 9c(F)$ for any metric space. These imply that, up to a constant factor, ROBUST- Π is not harder to approximate than Π , as formalized below.

Theorem 3. *Let Π be a problem of \mathcal{S} . Suppose that there is a ρ_1 -approximation for Π . Then, using this approximation in Algorithm 3, it is a $4\rho_1$ -approximation for ROBUST- Π restricted to Ptolemaic metric space, and a $9\rho_1$ -approximation for ROBUST- Π .*

5.2 Bounding the approximation ratio on general graphs

We divide our study of $c^{max}(F)/c(F)$ for arbitrary sets of edges F into special metric spaces that satisfy an important property, called Ptolemy's inequality, and metric spaces that do not satisfy this property.

Problem (graph family)	deterministic version	$c(F)/c^{max}(F)$	robust counterpart
$\Pi \in \mathcal{S}$	ρ -approx.	9 (Thm 5) 4 (Ptolemaic, Thm 4)	9ρ -approx. 4ρ -approx. (Ptolemaic)
MIN-EWCP (clique)	2-approx. (Eremin et al. 2014)	2 (Prop 10)	4-approx.
Δ -SP (path)	polynomial	2 (Cor 3)	2 approx \mathcal{NP} -hard (Prop 1) \mathcal{FPTAS} (Thm 7)
Δ -MST (tree)	polynomial	6 (Prop 14) 4 (Ptolemaic, Thm 4) $2\sqrt{2}$ (planar Eucl., Prop 15)	6-approx. 4-approx. (Ptolemaic) $2\sqrt{2}$ -approx. (planar Eucl.) \mathcal{NP} -hard (Prop 2)
Δ -TSP (cycle)	$\frac{3}{2}$ -approx. (Christofides 1976)	2 (Citovsky et al. (2017) and Cor 3)	3-approx. (Citovsky et al. 2017) \mathcal{PTAS} (planar Eucl.) (Citovsky et al. 2017))

Table 1: Overview of our results. When no reference is given, the ratios in the "robust counterpart" column are obtained by Theorem 2. All our results are for finite \mathcal{U}_i , except the $2\sqrt{2}$ ratio for trees, which holds for balls in the Euclidean plane, and the \mathcal{PTAS} for Δ -TSP holding for disjoint unit balls in the Euclidean plane.

5.2.1 Ptolemaic metric spaces

Let us recall Ptolemy's inequality (see e.g. Apostol (1967)). A metric space (\mathcal{M}, d) is *Ptolemaic* if for any four points A, B, C, D in \mathcal{M} ,

$$d(A, C) \cdot d(B, D) \leq d(A, B) \cdot d(C, D) + d(B, C) \cdot d(A, D).$$

We consider throughout the section that (\mathcal{M}, d) is Ptolemaic, which includes for instance Euclidean spaces. A direct consequence of the definition is given by the following lemma.

Lemma 1. *Let ABC be a triangle in a Ptolemaic metric space (\mathcal{M}, d) such that $d(B, C) \geq \max\{d(A, B), d(A, C)\}$. Then, for any $O \in \mathcal{M}$,*

$$d(O, A) \leq d(O, B) + d(O, C). \tag{7}$$

Proof. Using $d(B, C) \geq \max\{d(A, B), d(A, C)\}$ and Ptolemy's inequality, we get

$$d(B, C) \cdot d(O, A) \leq d(B, C) \cdot d(O, B) + d(B, C) \cdot d(O, C),$$

and the result follows. \square

Using the above inequality, we can get a constant bound on the approximation ratio by focusing on the extremities of a diameter of each uncertainty set $\mathcal{U}_i \in V$. The detailed proof is in Appendix C.1. It results in the following ratio.

Theorem 4. *Let (\mathcal{M}, d) be a Ptolemaic metric space. Then, $c^{\max}(F) \leq 4 c(F)$.*

5.2.2 Arbitrary metric spaces

Lemma 1 does not apply to non-Ptolemaic metric spaces, as illustrated in the following example.




Figure 1: Counter example of Lemma 1 for non-Ptolemaic metric spaces.

Example 4. *Let \mathcal{M} be a metric space on four points A, B, C, O , and consider a distance d such that (see also Figure 1)*

$$d(X, Y) = \begin{cases} 1 & \text{if } \{X, Y\} \subseteq \{A, B, C\} \\ 0.5 & \text{if } \{X, Y\} \in \{\{O, B\}, \{O, C\}\} \\ 1.5 & \text{if } \{X, Y\} = \{O, A\} \end{cases} ,$$

One readily verifies that $d(O, A) = \frac{3}{2}(d(O, B) + d(O, C))$.

In fact, multiplying the right-hand-side of (7) by $3/2$, as in Example 4, is enough for any metric space.

Lemma 2. *Let ABC be a triangle of \mathcal{M} such that $d(B, C) \geq \max\{d(A, B), d(A, C)\}$. Then, for any $O \in \mathcal{M}$,*

$$d(O, A) \leq \frac{3}{2}(d(O, B) + d(O, C)).$$

Proof. Using $d(B, C) \geq \max\{d(A, B), d(A, C)\}$ and applying the triangular inequality at each step, we get

$$d(O, A) \leq d(O, B) + d(A, B) \leq d(O, B) + d(B, C) \leq 2d(O, B) + d(O, C)$$

$$d(O, A) \leq d(O, C) + d(A, C) \leq d(O, C) + d(B, C) \leq d(O, B) + 2d(O, C)$$

Adding the above two inequalities provides the result. □

Using the above result, we can obtain a weaker counterpart of Theorem 4 for non-Ptolemaic metric spaces. See Appendix C.2 for the detailed proof.

Theorem 5. *For any metric space (\mathcal{M}, d) , $c^{max}(F) \leq 9 c(F)$.*

5.3 Bounding the approximation ratio on specific structures

In what follows, we assume that the structure of the subgraph induced by F can be leveraged to obtain stronger bounds than in the previous section. We first describe how graph decomposition can be used to obtain such bounds. We then address the special graphs that have been singled out in our introductory applications, namely: paths, cycles, trees (subway network design), cliques (clustering), and stars (facility location). Unless stated otherwise, we assume throughout the section that (\mathcal{M}, d) is an arbitrary metric space, non-necessarily Ptolemaic.

5.3.1 Building blocks

We study below how the bounds obtained for distinct subsets of edges can be combined to obtain a bound on their union.

Proposition 6. *Let G be a graph, $F_t \subseteq E(G)$ and $\rho_t \geq 1$ such that $c^{max}(F_t) \leq \rho_t c(F_t)$ for each $t = 1, \dots, T$. Then:*

- $c^{max}(\cup_{t=1}^T F_t) \leq T \times \max_{t=1, \dots, T} \rho_t c(\cup_{t=1}^T F_t)$, and
- $c^{max}(\cup_{t=1}^T F_t) \leq \max_{t=1, \dots, T} \rho_t c(\cup_{t=1}^T F_t)$ if, in addition, $F_t \cap F_{t'} = \emptyset$ for each $t \neq t' \in [T]$.

Proof. In the first case we have

$$T \cdot c(F) \geq \sum_{t=1}^T c(F_t) \geq \sum_{t=1}^T \frac{1}{\rho_t} c^{max}(F_t) \geq \frac{c^{max}(F)}{\max_{t=1, \dots, T} \rho_t}.$$

In the second case we have

$$c(F) = \max_{u \in \mathcal{U}} \sum_{t=1}^T \sum_{\{i, j\} \in F_t} d(u_i, u_j) = \sum_{t=1}^T c(F_t) \geq \sum_{t=1}^T \frac{1}{\rho_t} c^{max}(F_t) \geq \frac{c^{max}(F)}{\max_{t=1, \dots, T} \rho_t}. \quad \square$$

The above results are particularly useful when combining sets F_t having low values of ρ_t . The simplest example of such a set is a single edge.

Observation 2. *For any $e \in E(G)$, $c^{max}(e) = c(e)$.*

From the above observation and Proposition 6, we obtain immediately that matchings also satisfy the equality.

Corollary 1. *For any matching F , $c^{max}(F) = c(F)$.*

Matchings can be further combined to obtain general bounds that depend on the characteristics of F . In the remainder, $\chi(F)$ denotes the edge chromatic number of F and $\Delta(F)$ denotes its maximum degree.

Corollary 2. $c^{max}(F) \leq \chi(F) c(F)$.

Proof. We can cover F by $\chi(F)$ disjoint matchings. □

5.3.2 Graphs with small maximum degree

Recall that Vizing's theorem states that $\chi(F) \leq \Delta(F) + 1$. Combining this with Corollary 2 implies that $c^{max}(F) \leq (\Delta(F) + 1)c(F)$. Actually, the bound can be decreased to $\Delta(F)$, as stated below. See Appendix C.3 for the proof.

Proposition 7. *For any graph F , $c^{max}(F) \leq \Delta(F) c(F)$.*

Proposition 7 immediately implies the following Corollary, which had an ad-hoc proof for cycles in Citovsky et al. (2017).

Corollary 3. *Let F be a path or a cycle, then $c^{max}(F) \leq 2 c(F)$.*

We verify in the two propositions below that the above bound is tight. See Appendices D.1 and D.2 for the proofs.

Proposition 8. *For any path F of length at least three, there exists an uncertainty set \mathcal{U} such that $c^{max}(F) = 2 c(F)$.*

Proposition 9. *For any cycle F of length at least four, there exists an uncertainty set \mathcal{U} such that $c^{max}(F) = 2 c(F)$.*

The above proposition considers cycles that contain at least four vertices, so one can wonder what happens in the case of smaller cycles. We show next that for cycles that contain only 3 vertices, the bound can be reduced to $3/2$.

Remark 1. Consider the 3-cycle F . Applying the triangular inequality three times yields:

$$\begin{aligned} c(F) &= \max_{u \in \mathcal{U}} (d(u_1, u_2) + d(u_2, u_3) + d(u_1, u_3)) \geq \max_{u \in \mathcal{U}} (2 \max\{d(u_1, u_2), d(u_2, u_3), d(u_1, u_3)\}) \\ &\geq 2 \max\{d_{12}^{max}, d_{23}^{max}, d_{13}^{max}\} \\ &\geq \frac{2}{3} c^{max}(F), \end{aligned}$$

so the maximum worst-case factor is bounded by $3/2$. This bound is tight. To see this, one can look at the case exhibited in the proof of Proposition 8 ($\mathcal{U}_1 = \{0\}$, $\mathcal{U}_2 = \{0, 1\}$ and $\mathcal{U}_3 = \{1\}$). For the 3-cycle, $c^{max}(F) = 3$ and $c(F) = 2$.

5.3.3 Cliques

We now turn to the special case where \mathcal{F} contains only cliques. One specificity of a clique F is that for any matching M of size $\lfloor \frac{n}{2} \rfloor$, every edge of F belongs to a triangle including one edge of M . Applying the triangle inequality repeatedly for a well chosen matching provides the following ratio. The detailed proof is given in Appendix C.4.

Proposition 10. Let F be a clique. Then, $c^{max}(F) \leq 2 c(F)$.

We show below that the above bound is asymptotically tight, even for very simple metric spaces. See Appendix D.3 for a proof.

Proposition 11. If F is a k -clique, there exists an uncertainty set \mathcal{U} such that $c^{max}(F) = \frac{2(k-1)}{k} c(F)$ if k is odd and $c^{max}(F) = \frac{2k}{k+1} c(F)$ if k is even.

5.3.4 Stars

In what follows, we consider stars whose center is vertex 1 (meaning that for any $i \neq 1$, $|N(i)| = 1$).

Proposition 12. Let F be a star. Then, assuming that (\mathcal{M}, d) is a Ptolemaic space, $c^{max}(F) \leq 2 c(F)$.

Proof. Let $\{u_1^1, u_1^2\}$ such that $d(u_1^1, u_1^2) = \text{diam}(\mathcal{U}_1)$. Let $i \in \{2, \dots, n\}$ and $\bar{u}_1 \in \mathcal{U}_1, \bar{u}_i \in \mathcal{U}_i$ such that $d(\bar{u}_1, \bar{u}_i) = d_{1i}^{max}$. We follow the same approach as in the proof of Theorem 4 (see Appendix C.1). We thus set $\bar{\mathcal{U}}_1 = \{u_1^1, u_1^2\}$ and $\bar{\mathcal{U}}_i = \{\bar{u}_i\}$ for $i > 1$, and for $i \in V$, we consider the random variables \tilde{u}_i uniformly distributed on \mathcal{U}_i . Since $\{u_1^1, u_1^2\}$ is a diameter of \mathcal{U}_1 , we can apply Lemma 1 to get

$$d(u_1^1, \bar{u}_i) + d(u_1^2, \bar{u}_i) \geq d(\bar{u}_1, \bar{u}_i) = d_{1i}^{max}.$$

This implies $\mathbb{E}[d(\tilde{u}_1, \tilde{u}_i)] \geq \frac{1}{2}d_{1i}^{max}$ for any $i > 1$, and thus the claimed ratio. \square

Corollary 4. *Let F be a star. Then, $c^{max}(F) \leq 3 c(F)$.*

Proof. The result is obtained with the exact same proof as Proposition 12 where we apply Lemma 2 instead of Lemma 1. \square

We show below that the bound from Corollary 4 is asymptotically tight. See Appendix D.4 for a proof.

Proposition 13. *Let F be a star on n vertices. There is an uncertainty set \mathcal{U} such that $c^{max}(F) = \frac{3(n-1)}{n+1} c(F)$.*

5.3.5 Trees

We conclude our study of specific structures with trees. Our first result combines the bounds obtained for stars in the previous section with the composition results presented in Section 5.3.1 to improve the ratio of 9 obtained in Theorem 5 for general graphs and metric spaces. See Appendix C.5 for a detailed proof.

Proposition 14. *Let F be a tree. Then $c^{max}(F) \leq 6 c(F)$.*

Next, we show how the bound can be further tightened when considering Euclidean metric spaces and spherical uncertainty sets, see Appendix C.6 for a proof.

Proposition 15. *Assume that \mathcal{M} is a Euclidean space with d the associated Euclidean distance, and assume that for all $i \in V$, \mathcal{U}_i is a closed ball with radius r_i . Then, for any tree F , $c^{max}(F) \leq 2\sqrt{2} c(F)$.*

6. FPTAS for robust shortest path

Up to now, we have mostly focused on general problems ROBUST- Π and provided either exact algorithms or constant-factor approximation algorithms. The purpose of this section is to focus on a specific problem, ROBUST- Δ -SP. Recall that we proved in Section 3 that ROBUST- Δ -SP is \mathcal{NP} -hard. The main goal of the section is to show that ROBUST- Δ -SP admits an *FPTAS*. We do this by providing a dynamic programming algorithm for ROBUST-SP (a generalization of ROBUST- Δ -SP where d is not required to verify the triangle inequality) in Section 6.1, which can also be used on its own when the appropriated parameters have small values. The algorithm is then used in Section 6.2 to derive the *FPTAS*.

6.1 Dynamic programming algorithm

In the ROBUST-SP problem, the input is $I = (G, s, t, \mathcal{U}, \mathcal{D})$ is the same as in the ROBUST- Δ -SP problem, except that \mathcal{D} is now a non-negative matrix that is only assumed to be symmetric, e.g., $\mathcal{D}(u, v) = \mathcal{D}(v, u)$ for any $u, v \in \bigcup_{i \in V(G)} \mathcal{U}_i$. In particular, the function associated with \mathcal{D} may not respect the triangle inequality. This level of generality will be useful in the next section when deriving the \mathcal{FPTAS} for ROBUST- Δ -SP.

Next, we introduce further notations that are used to derive the dynamic programming algorithm. Let $\mathcal{F}^{i \rightarrow t}$ denote the set of all $i - t$ simple paths in G , and $\mathcal{F}_\kappa^{i \rightarrow t}$ those having at most κ edges. Given $i \in V(G)$ and a path $P \in \mathcal{F}^{i \rightarrow t}$, we define the worst-case cost given that $u_i = u_i^\ell$ for $\ell \in [|\mathcal{U}_i|]$ as

$$c^\ell(P) = \max\{c(u, P) \mid u \in \mathcal{U}, u_i = u_i^\ell\},$$

and $Pr(P) = (c^\ell(P))_{\ell \in [|\mathcal{U}_i|]}$ as the profile of P . We also introduce for any $\kappa \in [n]$ the set of profiles of all $i - t$ paths with at most κ edges as $\mathcal{P}^{(i, \kappa)} = \{Pr(P) \mid P \in \mathcal{F}_\kappa^{i \rightarrow t}\}$. We denote

$$\text{Val}(I) = \{c^\ell(P) \mid i \in V(G), P \in \mathcal{F}^{i \rightarrow t}, \ell \in [|\mathcal{U}_i|]\},$$

$n_{val} = |\text{Val}(I)|$, and $n_{\mathcal{P}} = |\bigcup_{i \in V(G), \kappa \in [n]} \mathcal{P}^{(i, \kappa)}|$ the total number of different profiles.

Our objective is to define an algorithm $A(i, \kappa)$ that, given any $i \in V(G)$ and $\kappa \in [n]$ computes $(\mathcal{P}^{(i, \kappa)}, Q^{(i, \kappa)})$ such that

- $Q^{(i, \kappa)} \subseteq \mathcal{F}_\kappa^{i \rightarrow t}$,
- $|Q^{(i, \kappa)}| = |\mathcal{P}^{(i, \kappa)}|$,
- for any $p \in \mathcal{P}^{(i, \kappa)}$, there exists $P \in Q^{(i, \kappa)}$ such that $Pr(P) = p$.

Informally, $A(i, \kappa)$ computes all profiles associated to (i, κ) as well as a representative path for each one of these profiles. Let us first verify that computing this is enough to solve ROBUST-SP optimally.

Lemma 3. *Given $(\mathcal{P}^{(s, n)}, Q^{(s, n)})$, we can find an optimal solution in time polynomial in n and linear in $n_{\mathcal{P}}$.*

Proof. For any $p \in \mathcal{P}^{(s, n)}$, $p = (p_\ell)_{\ell \in [|\mathcal{U}_s|]}$, let $x_p = \max_{\ell \in [|\mathcal{U}_s|]} p_\ell$. We define $p^{min} = \arg \min_{p \in \mathcal{P}^{(s, n)}} x_p$ and output $P^{min} \in Q^{(s, n)}$ such that $Pr(P^{min}) = p^{min}$. Let P^* be an optimal solution and $p^* = Pr(P^*)$. As $p^* \in \mathcal{P}^{(s, n)}$, we have $c(P^{min}) = x_{p^{min}} \leq x_{p^*} = c(P^*)$. \square

We provide next the dynamic programming recursion for $\mathcal{P}^{(i,\kappa)}$, leaving aside the computation of $Q^{(i,\kappa)}$ to simplify the presentation. Given $i \in V(G)$, $\kappa \in [n]$, $j \in N(i)$, $P' \in \mathcal{F}_\kappa^{j \rightarrow t}$, and $p' = Pr(P')$, we consider the $i - t$ path $P = iP'$ obtained by concatenating i with P' . One readily verifies that $Pr(P) = p(i, \kappa, j, p')$, where $p(i, \kappa, j, p') = (y_\ell)_{\ell \in [\mathcal{U}_i]}$, with $y_\ell = \max_{\ell' \in [\mathcal{U}_j]} d(u_i^\ell, u_j^{\ell'}) + p'_{\ell'}$. We obtain that for any $i \neq s$ and $\kappa > 0$

$$\mathcal{P}^{(i,\kappa)} = \{p(i, \kappa, j, p') \mid j \in N(i), p' \in \mathcal{P}^{(j,\kappa-1)}\}, \quad (8)$$

and $\mathcal{P}^{(s,0)} = (0)_{\ell \in [\mathcal{U}_s]}$. Recall that $\sigma = \max_{i \in V(G)} |\mathcal{U}_i|$. We provide in the next lemma the complexity of the resulting dynamic programming algorithm. See Appendix E.1 for the proof.

Lemma 4. *Let $i \in V(G)$, $\kappa \in [n]$. Given $\mathcal{P}^{(j,\kappa-1)}$ and $Q^{(j,\kappa-1)}$ for any $j \in N(i)$, we can compute $(\mathcal{P}^{(i,\kappa)}, Q^{(i,\kappa)})$ in time $\mathcal{O}(n \times \sigma^2 \times n_{\mathcal{P}})$ and space $\mathcal{O}(n_{\mathcal{P}} n (\log(n)))$.*

We are now ready to state the main result of this section.

Theorem 6. *ROBUST-SP can be solved in time $\mathcal{O}(n^3 \sigma^2 n_{\mathcal{P}})$ and space $\mathcal{O}(n_{\mathcal{P}} n^3 \log(n))$.*

Proof. We compute $(\mathcal{P}^{(s,n)}, Q^{(s,n)})$ using a DP algorithm based on (8), and obtain an optimal solution following Lemma 3. As the total number of entries of the associated memoization table has size $\mathcal{O}(n^2)$, we get the claimed complexity. \square

Let us further elaborate on the value of $n_{\mathcal{P}}$ that arises in Theorem 6. First of all, we see that $n_{\mathcal{P}} \leq (n_{val})^\sigma$, leading to the observation below, used in the next section to derive the FPTAS.

Observation 3. *ROBUST-SP can be solved in time $\mathcal{O}(n^3 \sigma^2 (n_{val})^\sigma)$*

From a more theoretical viewpoint, recall that the reduction of Proposition 1 proving the hardness of ROBUST-SP involves a “large” metric space, so a natural question is whether ROBUST- Δ -SP becomes polynomial for “small” metric space, either in terms of diameter or number of elements. It so happens that the two questions can be answered positively. Namely, observe that for any value $c^\ell(P)$, we can find some $n_{v,v'} \in [n]$ for each $(v, v') \in \mathcal{M}^2$ such that $c^\ell(P) = \sum_{(v,v') \in \mathcal{M}^2} n_{v,v'} d(v, v')$. Therefore, n_{val} can be bounded by $n^{\frac{|\mathcal{M}|(|\mathcal{M}|-1)}{2}}$, so that ROBUST-SP is polynomially solvable if $|\mathcal{M}|$ is constant. Alternatively, if all distances are integer, meaning d has integer values, then $n_{\mathcal{P}} \leq n \times \text{diam}(\mathcal{M})$, so ROBUST-SP can be solved in $\mathcal{O}(n^3 \sigma^2 (n \times \text{diam}(\mathcal{M}))^\sigma)$ in that case.

6.2 FPTAS

We now consider solving ROBUST- Δ -SP approximately. More precisely, given any $\epsilon > 0$, we want to provide an $(1 + \epsilon)$ approximated solution. Let $I = (G, s, t, \mathcal{U}, \mathcal{D})$ be an input to ROBUST- Δ -SP and A be an upper bound to $\text{OPT}(I)$ and $\epsilon' > 0$. We define a matrix \mathcal{D}' by rounding each element of \mathcal{D} to the closest value of the form $\ell\epsilon'A$ for some $\ell \in \mathbb{N}$. We obtain an instance $I' = (G, s, t, \mathcal{U}, \mathcal{D}')$ to ROBUST-SP. Having Observation 3 in mind, a straightforward application of the DP from the previous section to I' would yield too many values n_{val} . Hence, we show in Appendix E.2 how to adapt the DP, and choose A and ϵ' (depending of ϵ) appropriately to obtain the result below.

Theorem 7. *ROBUST- Δ -SP admits an FPTAS for fixed σ : for any ϵ , we can compute an $(1 + \epsilon)$ approximated solution in time $\mathcal{O}(n^3\sigma^2(\frac{n^2}{\epsilon})^\sigma)$.*

7. Computational experiments

In this section, we compare numerically the exact algorithm from Section 4, the heuristic algorithms from Section 5 and the affine decision rules (ADR) approach inspired by Zhen et al. (2021) and detailed in Appendix H. We compare these algorithms on two of the applications mentioned in the introduction: a subway network design problem (ND) and a simple plant location problem (PL). The ADR reformulation is assessed only on ND, because PL relies on a graph-induced metric, which is not compatible with the reformulation proposed by Zhen et al. (2021).

The purpose of our experiments is two-fold. First, we wish to assess the numerical efficiency of the different algorithms in terms of solution times. Second, we measure the approximation ratios obtained by the heuristic algorithms in practice, by comparing the cost of their solutions to the optimal solution costs.

The algorithms have been coded in Julia (Bezanson et al. 2012), using JuMP (Dunning et al. 2017) to interface the mixed integer linear programming (MILP) solver CPLEX. They have been carried out on a processor Intel(R) Core(TM) i7-10510U CPU 1.80GHz. The source code of every algorithm is publicly available at https://github.com/mjposs/locational_uncertainty.

In the remainder of the section, we refer to the cutting plane algorithm as **exact**, the heuristic based on the barycenters as **center**, that based on the maximum pairwise distances

as **dmax**, and that based on affine decision rules as **adr**. Since the applications involve stars and trees, the separation problems of **exact** can be solved using the dynamic programming recurrence presented in (4).

7.1 Subway network design

We consider the problem of expanding the subway network of a city, modeled as a Euclidean Steiner tree problem. The compulsory points model the future stops of the subway, while the other points model the possible knickpoints of the lines. Designing such an expansion plan is a complex urban planning problem that involve multiple political and economical layers. In particular, even when the topological design of the network has been decided, the exact physical location of the stations and knickpoints may still change, be it because of political constraints (inability to buy the rights of a given location) or physical ones (impossibility to dig as planned). The lengths of the lines are directly impacted by these uncertainties, and so is their cost, which can reasonably be assumed proportional to their Euclidean lengths.

Let us model the problem on an undirected graph $G = (V, E)$ where $T \subseteq V$ is the set of compulsory vertices; we consider an arbitrary root $t_0 \in T$ and define $T_0 = T \setminus \{t_0\}$. Set \mathcal{F} thus contains all trees of G that cover the vertices of T . Sets $\mathcal{U}_i \subseteq \mathbb{R}^2$ model the possible locations for the vertices, which we assume to be polyhedral sets, and we assume that the distance $d(u_i, u_j) = \|u_i - u_j\|_2$ is the Euclidean distance. We consider the classical disaggregated MILP formulation for the problem involving two sets of variables (Magnanti and Wong 1984). For each undirected edge $\{i, j\} \in E$, binary variable x_{ij} takes value 1 if and only if the edge is used. Then, for each $t \in T_0$ and $e = \{i, j\} \in E$, we consider a fractional variable f_{ij}^t stating how much flow related to t is sent on the directed arc (i, j) . Defining the incoming and outgoing stars of i as $\delta^-(i) = \{(j, i) : \{i, j\} \in E\}$ and $\delta^+(i) = \{(i, j) : \{i, j\} \in E\}$, respectively, and the balance of vertex i as $b_i^t = 0$ for $i \in T_0 \setminus \{t\}$, $b_{t_0}^t = -1$ and $b_t^t = 1$, we obtain

$$\min \left\{ \max_{u \in \mathcal{U}} \sum_{\{i, j\} \in E} x_{ij} \|u_i - u_j\|_2 \mid \begin{array}{l} \sum_{(j, i) \in \delta^-(i)} f_{ji}^t - \sum_{(i, j) \in \delta^+(i)} f_{ij}^t = b_i^t, \quad \forall i \in V, t \in T_0 \\ f_{ij}^t + f_{ji}^t \leq x_{ij}, \quad \forall \{i, j\} \in E, t \in T_0 \\ f \geq 0, x \text{ binary} \end{array} \right\}.$$

7.1.1 Instances

We assess the different solution algorithms on two groups of instances based on Euclidean distances that are publicly available at <http://steinlib.zib.de/testset.php>. Since **adr**




Figure 2: Small instances inspired by the *format* instance from SteinLib.

does not scale well, we first consider the *format* instance which includes 7 vertices and 9 edges (the instance is available at <http://steinlib.zib.de/format.php>). To get larger instances from the *format* instance, we remove the central terminal and add layered copies of the instance. Figure 2 depicts the original structure of the *format* instance and that obtained by adding one copy. We denote as $\text{format}(\kappa)$ the instance with κ copies of the original graph. Then, we consider the smallest instances with Euclidean costs, namely P6E, among which we keep only those having 100 vertices and 5 terminals. As a result, we get three sparse instances P6E (p619, p620, and 621), each having 180 edges.

The position of the vertices, denoted \bar{u}_i hereafter, are not available in the data files P6E, so we estimate them using a variant of the MDS-MAP algorithm from Shang et al. (2003). Specifically, we apply classical multidimensional scaling (MDS) from the Julia package MultivariateStats (see <https://github.com/JuliaStats/MultivariateStats.jl>) to compute the positions \bar{u} from the distances, completing the distance matrix with the shortest path values. The uncertainty sets $\mathcal{U}_i, i \in V$ are then computed randomly based on two parameters: μ that scales the diameter of each set \mathcal{U}_i , and σ the common number of elements of all $\mathcal{U}_i, i \in V$. To be more precise, we first compute the average distance among pairs of points in V , $\bar{d} = \sum_{i < j} \frac{\|\bar{u}_i - \bar{u}_j\|}{n(n-1)/2}$. For each $i \in V$, we then uniformly draw one random value in $\rho_i \in [0, \mu\bar{d}]$ and define

$$\mathcal{U}_i = \left\{ \left(\bar{u}_{i1} + \rho_i \cos\left(\frac{2k\pi}{\sigma}\right), \bar{u}_{i2} + \rho_i \sin\left(\frac{2k\pi}{\sigma}\right) \right), k = 1, \dots, \sigma \right\}.$$

Following the above procedure, we create 2000 random instances $\text{format}(\kappa)$ for each $\kappa \in [3]$, and 20 instances for each P6E instance. Notice that **adr** is only assessed on the small instances as it was not able to solve any of the larger ones within hours of computing time.

7.1.2 Results

Figures 3 and 4 report the average solution times for the small and large instances, respectively, illustrating the impact of each parameter. More specifically, Figure 3 highlights the rapidly increasing solution times of **adr** along with the values of σ and κ . The three other approaches, including **exact**, can solve all these instances within a fraction of second. Figure 4 further depicts how **exact** is sensitive to μ for the larger instances, while **center** and **dmax** seem less affected by the value of that parameter.

Figure 5 reports the cumulative distributions of the cost increase of each of the three heuristic, relatively to the cost of the exact solution. Formally, let $z(H)$ denote the cost of the solution returned by $H \in \{\mathbf{adr}, \mathbf{dmax}, \mathbf{center}\}$ and z^* denote the optimal solution cost. For each $H \in \{\mathbf{adr}, \mathbf{dmax}, \mathbf{center}\}$, the corresponding curve reports

$$f(x) = 100 \frac{\#\{\text{instances for which } z(H) \leq (1+x) \cdot z^*\}}{\#\{\text{all instances}\}}. \quad (9)$$

These results confirm that **dmax** and **adr** provide solutions with values very close to the optimal one. In contrast, the quality of **center** becomes rather poor as μ increases, ranging up to an extra cost 60% for some of these instances, and with nearly half of the instances with $\mu = 1$ having an extra cost of at least 20%.




Figure 3: ND: Average solution times in seconds on instances $\text{format}(\kappa)$ for each algorithm when varying one of the parameters.

7.2 Simple plant location

We consider a strategic facility location problem where the exact location of the facility may be perturbed due to local political and technical considerations, while the exact position of the clients themselves is subject to uncertainty (Correia and da Gama 2015). The distances between the facilities and the clients are computed from the shortest path distance on a




Figure 4: ND: Average solution times in seconds on large instances (P6E) for each value of Δ .




Figure 5: ND: For each heuristic algorithm, the curve plots (9), the cumulative distribution of the % of instances for which the returned solution has an additional (relative) cost less than the value of the abscissa.

weighted graph that represents the underlying road network. The problem can then be modeled with the weighted graph $G = (V, E, l)$, the vertices of which represent the possible locations for the facilities and clients, while each edge and its weight represent the existence of a road between two vertices together with its length. The metric is induced by graph G , so $\mathcal{M} = V$ and $d(u, v)$ is equal to the shortest path between u and v for every $u, v \in V$.

Let $I \subseteq V$ and $J \subseteq V$ represent the set of clients and possible locations for the facilities. We consider the problem of choosing p facilities among J and assigning every client to its closest facility so as to minimize the total assignment cost. For each $j \in J$, let y_j be a binary variable indicating whether a facility is located at j , and for each $i \in I, j \in J$, let $x_{ij} \in \{0, 1\}$ indicate whether client i is assigned to facility j . The robust problem can then be formulated

as

$$\min \left\{ \max_{u \in \mathcal{U}} \sum_{i \in I, j \in J} x_{ij} d(u_i, u_j) \mid \begin{array}{l} \sum_{j \in J} x_{ij} = 1, \quad \forall i \in I \\ x_{ij} \leq y_j, \quad \forall i \in I, j \in J \\ x, y \text{ binary} \end{array} \right\}.$$

7.2.1 Instances

We construct the graph $G = (V, E, l)$ as follows. For each vertex i , we generate its position u_i uniformly in the square $[0, 1]^2$ and we select edges so that the resulting graph is planar and connected and shorter edges are more likely to appear. This procedure allows to mimic real transportation networks (Daskin 1993). More precisely, we first compute a minimum cost spanning tree based on the weights $\{w_{ij} = \|u_i - u_j\|_2^{-2}\}$ to ensure the graph is connected. Then, we iteratively select $m - n + 1$ additional edges following the probability distribution $p_{ij} = \frac{w_{ij}}{\sum_{\{i', j'\}} w_{i' j'}}$ for each $i \neq j \in V$ while ensuring the resulting graph is planar. The length l_{ij} of each edge $\{i, j\} \in E$ is then given by $\|u_i - u_j\|_2$ and the distance between every pair of vertices is given by the shortest path between them in G . For each $i \in V$, we finally define \mathcal{U}_i as the σ vertices that are closest to i .

7.2.2 Results

Figure 6 reports the average solution times, showing that **exact** is able to solve every instance to optimality within the time limit. The figure further underlines the impact of parameters m and σ : larger values for these parameters impact significantly the solution time of **exact**, whereas **center** and **dmax** solve all the instances in a fraction of a second. Then, following again formula (9), Figure 7 reports the cumulative distributions of the cost increase of each of the two heuristics, relatively to the cost of the exact solution. Notice that **adr** is not included in this comparison, because this algorithm can handle only the Euclidean distance whereas these results rely on graph-induced distances. The results illustrate that when $\sigma = 3$, $m = 120$ and $m = 160$, **dmax** returns almost always slightly better quality solutions than **center**. The distinction between the two is less marked for the three remaining cases.

Overall, these results illustrate that **exact** is a useful tool for solving this problem exactly, while larger instances may be solved to near-optimality by either **center** or **dmax**, the latter being slightly preferable over the former.




Figure 6: PL: Average solution times in seconds for each algorithm when varying one of the parameters.

8. Concluding remarks

This paper has been devoted to the study of general combinatorial optimization problems defined in spatial graphs with locational uncertainty, thus encompassing applications arising in transportation, location, and clustering, among others. We have developed a suite of solution techniques to handle these problems: an exact solution algorithm based on scenario generation and approximation schemes with theoretical guarantees. We have compared these algorithms numerically to the affine decision rules conservative approximation proposed by Zhen et al. (2021) on Steiner tree instances inspired by the scientific literature. Our results illustrate that, while the approximation provided by Zhen et al. (2021) is particularly tight, it does not scale well and is limited to solving small problems. In contrast, our exact algorithm is capable of solving medium-size instances. Moreover, the approximation algorithm based on maximum pairwise distances scales up very well, and it provides solutions of good quality. In fact, the costs of these solutions are nearly as good as those provided by Zhen et al. (2021) and much closer to the optimal value than suggested by theoretical ratios.

Our results also illustrate that while the naive heuristic considering the barycenters can produce near-optimal solutions on some applications (as for the simple plant location), it may also be far worse than the heuristic based on maximum pairwise distances (as for the Steiner tree problem). As a consequence, in the absence of extra knowledge, our tests suggest that using the worst-case distances is more appropriate if an exact solution cannot be computed in acceptable time.

From a theoretical point of view, we have determined constant factor approximation ratios for a variety of problems, singling out metric spaces that are Ptolemaic. We have also understood the \mathcal{NP} -harness of basic problems, including the shortest path problem for




Figure 7: PL: For each heuristic algorithm, the curve plots (9), the cumulative distribution of the % of instances for which the returned solution has an additional (relative) cost less than the value of the abscissa.

which we have provided an \mathcal{FPTAS} . These theoretical results pave the way for several interesting open questions. First, it is unknown whether the robust minimum spanning tree (ROBUST- Δ -MST) is \mathcal{NP} -hard in the 1-dimensional Euclidean space. Furthermore, does the problem become polynomial whenever $|\mathcal{M}|$ is constant, or admits a \mathcal{PTAS} when σ is constant? Another independent open question concerns closing the gaps for the ratios $c^{max}(F)/c(F)$ for arbitrary finite sets \mathcal{U}_i . Specifically, for general graphs and arbitrary metric spaces, we have devised an upper bound of 9 for this ratio while the highest lower bound we could identify, given by the tightness result obtained for stars, is equal to 3. In the case of Ptolemaic metric spaces, our upper bound is 4 while our strongest lower bound is equal to 2 for paths.

References

- Aissi H, Bazgan C, Vanderpooten D (2009) Min-max and min-max regret versions of combinatorial optimization problems: A survey. *Eur. J. Oper. Res.* 197(2):427–438, URL <http://dx.doi.org/10.1016/j.ejor.2008.09.012>.
- Apostol TM (1967) Ptolemy’s Inequality and the Chordal Metric. *Mathematics Magazine* 40(5):233–235, URL <http://dx.doi.org/10.1080/0025570X.1967.11975804>.
- Ayoub J, Poss M (2016) Decomposition for adjustable robust linear optimization subject to uncertainty polytope. *Comput. Manag. Science* 13(2):219–239, URL <http://dx.doi.org/10.1007/s10287-016-0249-2>.

- Ben-Tal A, Goryashko AP, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math. Program.* 99(2):351–376, URL <http://dx.doi.org/10.1007/s10107-003-0454-y>.
- Ben-Tal A, Nemirovski A (1998) Robust convex optimization. *Mathematics of Operations Research* 23(4):769–805.
- Bertsimas D, Dunning I (2016) Multistage robust mixed-integer optimization with adaptive partitions. *Oper. Res.* 64(4):980–998, URL <http://dx.doi.org/10.1287/opre.2016.1515>.
- Bertsimas D, Dunning I, Lubin M (2016) Reformulation versus cutting-planes for robust optimization. *Comput. Manag. Science* 13(2):195–217.
- Bertsimas D, Sim M (2003) Robust discrete optimization and network flows. *Mathematical Programming* 98(1-3):49–71.
- Bezanson J, Karpinski S, Shah VB, Edelman A (2012) Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*.
- Bodlaender HL, Drange PG, Dregi MS, Fomin FV, Lokshtanov D, Pilipczuk M (2013) A $O(c^k n)$ 5-approximation algorithm for treewidth. *CoRR* abs/1304.6321, URL <http://arxiv.org/abs/1304.6321>.
- Buchheim C, Kurtz J (2018) Robust combinatorial optimization under convex and discrete cost uncertainty. *EURO J. Computational Optimization* 6(3):211–238, URL <http://dx.doi.org/10.1007/s13675-018-0103-0>.
- Cevallos A, Eisenbrand F, Morell S (2018) Diversity maximization in doubling metrics. *CoRR* abs/1809.09521, URL <http://arxiv.org/abs/1809.09521>.
- Christofides N (1976) Worst-case analysis of a new heuristic for the traveling salesman problem. *Symposium on new directions and recent results in algorithms and complexity*, 441.
- Citovsky G, Mayer T, Mitchell JSB (2017) TSP With Locational Uncertainty: The Adversarial Model. Aronov B, Katz MJ, eds., *33rd International Symposium on Computational Geometry (SoCG 2017)*, volume 77 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 32:1–32:16 (Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik), ISBN 978-3-95977-038-5, ISSN 1868-8969, URL <http://dx.doi.org/10.4230/LIPIcs.SocG.2017.32>.
- Correia I, da Gama FS (2015) Facility location under uncertainty. *Location science*, 177–203 (Springer).
- Crescenzi P (1997) A short guide to approximation preserving reductions. *Proceedings of Computational Complexity. Twelfth Annual IEEE Conference*, 262–273, URL <http://dx.doi.org/10.1109/CCC.1997.612321>.
- Cygan M, Fomin FV, Kowalik L, Lokshtanov D, Marx D, Pilipczuk M, Pilipczuk M, Saurabh S (2015) *Parameterized Algorithms* (Springer), URL <http://dx.doi.org/10.1007/978-3-319-21275-3>.
- Daskin M (1993) Genrand2: A random network generator. *Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL, USA*.
- Downey RG, Fellows MR (2013) *Fundamentals of Parameterized Complexity*. Texts in Computer Science (Springer), URL <http://dx.doi.org/10.1007/978-1-4471-5559-1>.
- Dunning I, Huchette J, Lubin M (2017) Jump: A modeling language for mathematical optimization. *SIAM Review* 59(2):295–320, URL <http://dx.doi.org/10.1137/15M1020575>.
- Eremin I, Gimadi EK, Kel'manov A, Pyatkin A, Khachai MY (2014) 2-approximation algorithm for finding a clique with minimum weight of vertices and edges. *Proceedings of the Steklov Institute of Mathematics* 284(1):87–95.

- Fischetti M, Monaci M (2012) Cutting plane versus compact formulations for uncertain (integer) linear programs. *Math. Program. Comput.* 4(3):239–273.
- Gutiérrez-Jarpa G, Obreque C, Laporte G, Marianov V (2013) Rapid transit network design for optimal cost and origin–destination demand capture. *Computers & Operations Research* 40(12):3000–3009, ISSN 0305-0548, URL <http://dx.doi.org/https://doi.org/10.1016/j.cor.2013.06.013>.
- Hanasusanto GA, Kuhn D, Wiesemann W (2015) K-adaptability in two-stage robust binary programming. *Operations Research* 63(4):877–891.
- Kasperski A, Zieliński P (2009) On the approximability of minmax (regret) network optimization problems. *Information Processing Letters* 109(5):262–266.
- Kasperski A, Zieliński P (2016) Robust discrete optimization under discrete and interval uncertainty: A survey. *Robustness analysis in decision aiding, optimization, and analytics*, 113–143 (Springer).
- Kloks T (1994) *Treewidth, Computations and Approximations*, volume 842 of *Lecture Notes in Computer Science* (Springer-Verlag), URL <http://dx.doi.org/10.1007/BFb0045375>.
- Kouvelis P, Yu G (2013) *Robust discrete optimization and its applications*, volume 14 (Springer Science & Business Media).
- Magnanti TL, Wong RT (1984) Network design and transportation planning: Models and algorithms. *Transportation science* 18(1):1–55.
- Masson MH, Quost B, Destercke S (2020) Cautious relational clustering: A thresholding approach. *Expert Systems with Applications* 139:112837.
- Melkote S, Daskin MS (2001) An integrated model of facility location and transportation network design. *Transportation Research Part A: Policy and Practice* 35(6):515–538, ISSN 0965-8564, URL [http://dx.doi.org/https://doi.org/10.1016/S0965-8564\(00\)00005-7](http://dx.doi.org/https://doi.org/10.1016/S0965-8564(00)00005-7).
- Naoum-Sawaya J, Buchheim C (2016) Robust critical node selection by benders decomposition. *INFORMS J. Comput.* 28(1):162–174, URL <http://dx.doi.org/10.1287/ijoc.2015.0671>.
- Postek K, Hertog Dd (2016) Multistage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. *INFORMS Journal on Computing* 28(3):553–574.
- Roos E, den Hertog D, Ben-Tal A, de Ruiter FJ, Zhen J (2018) Approximation of hard uncertain convex inequalities. *Optimization Online* URL http://www.optimization-online.org/DB_HTML/2018/06/6679.html.
- Shang Y, Ruml W, Zhang Y, Fromherz MP (2003) Localization from mere connectivity. *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, 201–212.
- Subramanyam A, Gounaris CE, Wiesemann W (2019) K-adaptability in two-stage mixed-integer robust optimization. *Mathematical Programming Computation* 1–32.
- Yaman H, Karasan OE, Pinar MÇ (2001) The robust spanning tree problem with interval data. *Oper. Res. Lett.* 29(1):31–40, URL [http://dx.doi.org/10.1016/S0167-6377\(01\)00078-5](http://dx.doi.org/10.1016/S0167-6377(01)00078-5).
- Zeng B, Zhao L (2013) Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters* 41(5):457–461.
- Zhen J, de Ruiter FJ, Roos E, den Hertog D (2021) Robust optimization for models with uncertain second-order cone and semidefinite programming constraints. *INFORMS Journal on Computing*.
- Zhen J, den Hertog D, Sim M (2018) Adjustable robust optimization via fourier-motzkin elimination. *Operations Research* 66(4):1086–1100, URL <http://dx.doi.org/10.1287/opre.2017.1714>.

A. Hardness proofs for the robust problem

A.1 Proof of Proposition 1

Given a set of integers $\{a_1, \dots, a_n\}$, with $A = \sum_{i=1}^n a_i$, the \mathcal{NP} -complete decision problem partition asks for a subset $S \subset \{1, \dots, n\}$ such that $\sum_{i \in S} a_i = A/2$. Let $K > 0$ be a large enough integer. The reduction considers the graph G with $2n+2$ vertices and $4n$ edges as illustrated Figure 8; the regions \mathcal{U}_i are translated away from vertex o for visibility. Specifically, our reduction locates vertices s and t at 0 while $\mathcal{U}_i = [-u_i^-, u_i^+]$ for each vertex i different from s and t . The definition of u^+ and u^- alternates along the vertices $v_i, v_{i+1}, v_{i+2}, \dots$ and similarly for vertices w_i : for each $i = 2k+1$, we define $u_{v_i}^+ = K + a_i$, $u_{v_i}^- = K + \frac{A}{n} - a_i$, $u_{w_i}^+ = K$ and $u_{w_i}^- = K + \frac{A}{n}$, while for each $i = 2k$, we define $u_{v_i}^+ = K + \frac{A}{n} - a_i$, $u_{v_i}^- = K + a_i$, $u_{w_i}^+ = K + \frac{A}{n}$ and $u_{w_i}^- = K$.






Figure 8: Reduction from partition when \mathcal{U} is a Cartesian product of segments.

We first show that for K large enough, the worst-case $u \in \mathcal{U}$ for any path F from s to t alternates from the top of an interval to the bottom of the subsequent interval along the path. To prove this, notice that for any vertex $v \in G \setminus \{s, t\}$, $u_v^- \in [K - A, K + A]$ and $u_v^+ \in [K - A, K + A]$ and the same holds for any vertex w . Hence, if u alternates for the entire path, the resulting cost is not smaller than $c = 2n(K - A)$. On the contrary, if u misses one alternation, its cost cannot be greater than $c' = 2(n - 1)(K + A) + 2A$. Hence, taking $K > 2nA$ ensures $c > c'$.

The reduction works as follow. Let $S \subseteq \{1, \dots, n\}$ be a subset of integers and \bar{S} its complement. We associate to S the path F_S from s to t that contains v_i for each $i \in S$ and w_i for each $i \in \bar{S}$. From the above, only two scenarios in \mathcal{U} must be considered in the worst-case and each vertex $i \in \{w_1, v_1, \dots, w_n, v_n\}$ contributes to the total length with either



(a) G



(b) $G_{\mathcal{M}}$

Figure 9: Graphs used in the reduction for the minimum spanning tree problem.

$2u_i^+$ or $2u_i^-$, depending on the scenario considered. We have

$$\begin{aligned}
 c(F_S) &= \max_{u \in \mathcal{U}} \sum_{\{i,j\} \in F_S} \|u_i - u_j\|_2 \\
 &= 2 \max \left(nK + \sum_{i \in S} a_i, n(K + \frac{A}{n}) - \sum_{i \in S} a_i \right) \\
 &= 2 \max \left(nK + \sum_{i \in S} a_i, A + nK - \sum_{i \in S} a_i \right) \\
 &= 2nK + 2 \max \left(\sum_{i \in S} a_i, \sum_{i \in \bar{S}} a_i \right).
 \end{aligned}$$

Hence, there exists a path $F_S \in \mathcal{F}$ with minimum cost of $2nK + A$ if and only if there exists a set S such that $\sum_{i \in S} a_i = \sum_{i \in \bar{S}} a_i = A/2$.

A.2 Proof of Proposition 2

We consider the same partition problem as in the proof of Proposition 1. Now G contains the $2n + 2$ vertices and $3n + 1$ edges depicted on Figure 9a. We consider the metric space (\mathcal{M}, d) induced by the weighted graph $G_{\mathcal{M}} = (V_{\mathcal{M}}, E_{\mathcal{M}}, \omega)$ depicted on Figure 9b. Let $K > 0$ be a number large enough. The dashed edges and thin edges have their weights equal to K and $2K$, respectively, while $\omega_{v_{i-1}^1 v_i^1} = 3K + a_i$, $\omega_{v_{i-1}^2 v_i^2} = 3K + \frac{A}{n} - a_i$, $\omega_{w_{i-1}^1 w_i^1} = 3K$, and $\omega_{w_{i-1}^2 w_i^2} = 3K + \frac{A}{n}$ for each $i = 1, \dots, n$. Weights vector ω satisfies the triangle inequalities,

so the metric d induced on $V_{\mathcal{M}}$ by the shortest paths in $G_{\mathcal{M}}$ satisfies $d_{ij} = \omega_{ij}$ for each $\{i, j\} \in E_{\mathcal{M}}$. Finally, we define $\mathcal{U}_{v_i} = \{v_i^1, v_i^2\}$ and $\mathcal{U}_{w_i} = \{w_i^1, w_i^2\}$ for $i = 0, \dots, n$.

We first observe that the cost of a vertical edge $\{v_i, w_i\}$ is equal to K for all positions of $(v_i, w_i) \in \mathcal{U}_{v_i} \times \mathcal{U}_{w_i}$. Let us consider any tree F in G that contains n^v vertical edges and n^h horizontal edges, where $n \leq n^h \leq 2n$. For K large enough, we claim that the worst-case $u \in \mathcal{U}$ locates all vertices either in the bottom layer of $G_{\mathcal{M}}$ that consists of vertices v_i^1 and w_i^1 for $i = 0, \dots, n$, or in the top layer that consists of the remaining vertices. To prove the claim, notice that the weight of any horizontal edge in $G_{\mathcal{M}}$ is comprised between $3K - A$ and $3K + A$, while the weight of any diagonal edge is $2K$. Hence, if u locates all its vertices either in the bottom or in the top layer, the resulting cost is not smaller than $c = n^h(3K - A) + n^v K$. On the contrary, if u alternates at least once between the layers, its cost cannot be greater than $c' = (n^h - 1)(3K + A) + 2K + n^v K$. Hence, taking $K > (4n - 1)A \geq (2n^h - 1)A$ ensures $c > c'$, proving the claim.

We prove next that for K large enough, any optimal tree F in G must contain $n+1$ vertical edges and n horizontal ones. Following the above claim, the cost of a horizontal edge $\{v_i, v_{i+1}\}$ or $\{w_i, w_{i+1}\}$ for a worst-case $u \in \mathcal{U}$ is comprised between $3K - A$ and $3K + A$. Hence, any tree F with $n^h \in \{n+1, \dots, 2n\}$ horizontal edges costs at least $c = n^h(3K - A) + n^v K$ while any tree having $n^h - 1$ horizontal edges costs at most $c = (n^h - 1)(3K + A) + (n^v + 1)K$. Hence, taking $K > 2nA \geq n^h A$ ensures $c > c'$, proving $n^h = n$ in any optimal solution.

As in the proof of Proposition 1, we let $S \subseteq \{1, \dots, n\}$ be a subset of integers and \bar{S} its complement. We associate to S the tree F_S that contains $\{v_{i-1}, v_i\}$ for each $i \in S$ and $\{w_{i-1}, w_i\}$ for each $i \in \bar{S}$. Following the claim above, only two scenarios in \mathcal{U} must be considered, and following again the reasoning used in the proof of Proposition 1, we have

$$c(F_S) = \max \left(3nK + \sum_{i \in S} a_i, 3nK + A - \sum_{i \in S} a_i \right) = 3nK + \max \left(\sum_{i \in S} a_i, \sum_{i \in \bar{S}} a_i \right).$$

Hence, there exists a spanning tree $F_S \in \mathcal{F}$ with minimum cost of $3nK + A/2$ if and only if there exists a set S such that $\sum_{i \in S} a_i = \sum_{i \in \bar{S}} a_i = A/2$.

Remark 2. The graph $G_{\mathcal{M}}$ described in the above proof cannot be embedded isometrically into an Euclidean space, as can be seen by considering the triangle $w_0^1 w_1^1 w_0^2$ and the fourth point v_0^2 . The sides of the triangle have length $d(w_0^1, w_1^1) = 3K$, $d(w_1^1, w_0^2) = 2K$, and $d(w_0^2, w_0^1) = 2K$. Hence, since $d(v_0^2, w_0^1) = d(v_0^2, w_0^2) = K$, any isometric embedding maps v_0^2 to the midpoint of segment $w_0^1 w_0^2$, so its Euclidean distance to w_1^1 must be $\sqrt{\frac{11}{2}}K$. This is

in contradiction with $d(v_0^2, w_1^1) = \min(\omega_{v_0^2 w_1^1} + \omega_{w_0^1 w_1^1}, \omega_{v_0^2 v_1^2} + \omega_{v_1^2 w_1^1}) = \min(4K, 4K + \frac{A}{n} - a_1)$. The above illustrates that the complexity of problem (2) when \mathcal{F} contains all spanning trees of G is still open when one considers only Euclidean metric spaces.

B. Hardness proofs for the cost evaluation problem

B.1 Proof of Proposition 3

Let us denote the objective function of MAXCUT as $c^{MAX-CUT}(V_1, V_2) = |\{e \in E(G) \mid |e \cap V_1| = 1\}|$. Further, we denote by $\text{OPT}^{MAX-CUT}(G)$ the value of an optimal solution for graph G .

Lemma 5. *For any input graph G of MAX-CUT, we can compute in polynomial time an instance I of EVAL-C such that*

- *for any solution $\{V_1, V_2\}$ of MAX-CUT on G , there exists a solution u of EVAL-C on I such that $c(u, G) = c^{MAX-CUT}(V_1, V_2)$,*
- *for any solution u of EVAL-C on I such that $c(u, G) = c^{MAX-CUT}(V_1, V_2)$, there exists a solution $\{V_1, V_2\}$ of MAX-CUT on G such that $c^{MAX-CUT}(V_1, V_2) = c(u, G)$ (and this solution can be computed in polynomial time in $|V(G)|$),*
- *the metric space of I is $\mathcal{M} = \{0, 1\}$, the distance d is defined by $d(x, y) = |x - y|$, and for any $i \in V(G)$, $\mathcal{U}_i = \mathcal{M}$.*

Proof. Given G , we simply define I as the graph G itself, and the metric space, distance and \mathcal{U}_i as claimed. Given a solution $\{V_1, V_2\}$ (which is a partition) of MAX-CUT, we define $u_i = 0$ if $i \in V_1$, and 1 if $i \in V_2$. This implies $c(u, I) = c^{MAX-CUT}(V_1, V_2)$. For the reverse direction, given a solution u of EVAL-C, we define $V_1 = \{i \mid u_i = 0\}$ and $V_2 = V(G) \setminus V_1$, and we also have $c(u, I) = c^{MAX-CUT}(V_1, V_2)$. \square

Let us recall that S -reduction (see Crescenzi (1997)) is one of the most restrictive form of reduction, that preserves in particular approximability. Lemma 5 immediatly implies that there is an S -reduction from MAX-CUT to EVAL-C, implying the following Observation and the proof of Proposition 3.

Observation 4. *Let G be a graph. Using $\mathcal{M} = \{0, 1\}$, d defined by $d(x, y) = |x - y|$, and for any $i \in V(G)$, $\mathcal{U}_i = \mathcal{M}$, we get $c(G) = \text{OPT}^{MAX-CUT}(G)$.*

B.2 Proof of Proposition 5

Given a graph G , and a set of integers (called colors) $L(i)$ for any $i \in V(G)$, problem LIST-COL aims at deciding whether we can find a color $f(i) \in L(i)$ for any $i \in V(G)$ such that for any edge $\{ij\} \in E(G)$, $f(i) \neq f(j)$.

Given a graph G a list of colors $L(i)$ for any $i \in V(G)$, we define $\mathcal{M} = \bigcup_{i \in V(G)} L(i)$, and $d(c_1, c_2) = 0$ if $c_1 = c_2$, and 1 otherwise. We define the uncertainty set of G as follows: for any i , we let $\mathcal{U}_i = L(i)$. It is now straightforward to verify that we have a YES-instance of LIST-COL iff $c(G) = m$. As the reduction can be computed in polynomial time, and the graph (and thus its treewidth) is unchanged, we get the desired result.

C. Proofs related to the approximation ratio

C.1 Proof of Theorem 4

For all $i \in \{1, \dots, n\}$, let $[u_i^1, u_i^2]$ be a diameter of \mathcal{U}_i , i.e., $u_i^1 \in \mathcal{U}_i, u_i^2 \in \mathcal{U}_i$ and $d(u_i^1, u_i^2) = \text{diam}(\mathcal{U}_i)$, and let $\tilde{\mathcal{U}}_i = \{u_i^1, u_i^2\}$ and $\tilde{\mathcal{U}}$ be the cross product of the $\tilde{\mathcal{U}}_i$. Let $\tilde{c}(F) = \max_{u \in \tilde{\mathcal{U}}} \sum_{\{i,j\} \in F} d(u_i, u_j)$. As $\tilde{\mathcal{U}} \subseteq \mathcal{U}$, we get $c(F) \geq \tilde{c}(F)$. Let us now prove that $\tilde{c}(F) \geq \frac{c^{max}(F)}{4}$.

Now, define the random variable \tilde{u}_i taking any value of $\tilde{\mathcal{U}}_i$ with equal probability $1/3$. The worst-case length of the graph is not smaller than its expected edge length, i.e.,

$$c(F) = \max_{u \in \tilde{\mathcal{U}}} \sum_{\{i,j\} \in F} d(u_i, u_j) \geq \mathbb{E} \left[\sum_{\{i,j\} \in F} d(\tilde{u}_i, \tilde{u}_j) \right],$$

where, by linearity of expectation,

$$\mathbb{E} \left[\sum_{\{i,j\} \in F} d(\tilde{u}_i, \tilde{u}_j) \right] = \sum_{\{i,j\} \in F} \mathbb{E} [d(\tilde{u}_i, \tilde{u}_j)].$$

We then consider some arbitrary edge $\{i, j\} \in F$:

$$\mathbb{E} [d(\tilde{u}_i, \tilde{u}_j)] = \frac{1}{4} (d(u_i^1, u_j^1) + d(u_i^1, u_j^2) + d(u_i^2, u_j^1) + d(u_i^2, u_j^2)).$$

Let $\bar{u}_i \in \mathcal{U}_i$ and $\bar{u}_j \in \mathcal{U}_j$ such that $d(\bar{u}_i, \bar{u}_j) = d_{ij}^{max}$. As $u_j^1 u_j^2$ is a diameter of \mathcal{U}_j , we have $d(u_j^1, u_j^2) \geq \max(d(u_j^1, \bar{u}_j), d(u_j^2, \bar{u}_j))$, and we can apply Lemma 1 twice in triangle $u_j^1 u_j^2 \bar{u}_j$ to get

$$\begin{cases} d(u_i^1, u_j^1) + d(u_i^1, u_j^2) \geq d(u_i^1, \bar{u}_j) \\ d(u_i^2, u_j^1) + d(u_i^2, u_j^2) \geq d(u_i^2, \bar{u}_j). \end{cases}$$

One last application of the lemma in $u_i^1 u_i^2 \bar{u}_i$ then yields

$$d(u_i^1, \bar{u}_j) + d(u_i^2, \bar{u}_j) \geq d(\bar{u}_i, \bar{u}_j) = d_{ij}^{max}.$$

Summarizing the above, we get to $\tilde{c}(F) \geq \sum_{\{i,j\} \in F} \mathbb{E} [d(\tilde{u}_i, \tilde{u}_j)] \geq \frac{1}{4} c^{max}(F) \geq \frac{c^{max}(F)}{4}$.

C.2 Proof of Theorem 5

The proof follows exactly the approach followed in the proof of Theorem 4, but we use Lemma 2 instead of Lemma 1 when needed. We thus use the same notations as in the proof of Theorem 4. As $u_j^1 u_j^2$ is a diameter of \mathcal{U}_j , we have $d(u_j^1, u_j^2) \geq \max(d(u_j^1, \bar{u}_j), d(u_j^2, \bar{u}_j))$, and we can apply Lemma 2 twice in triangle $u_j^1 u_j^2 \bar{u}_j$ to get

$$\begin{cases} d(u_i^1, u_j^1) + d(u_i^1, u_j^2) \geq \frac{2}{3} d(u_i^1, \bar{u}_j) \\ d(u_i^2, u_j^1) + d(u_i^2, u_j^2) \geq \frac{2}{3} d(u_i^2, \bar{u}_j). \end{cases}$$

One last application of Lemma 2 in $u_i^1 u_i^2 \bar{u}_i$ then yields

$$d(u_i^1, \bar{u}_j) + d(u_i^2, \bar{u}_j) \geq \frac{2}{3} d(\bar{u}_i, \bar{u}_j) = \frac{2}{3} d_{ij}^{max}.$$

Summarizing the above, we get to $\tilde{c}(F) \geq \sum_{\{i,j\} \in F} \mathbb{E} [d(\tilde{u}_i, \tilde{u}_j)] \geq \frac{1}{4} \frac{4}{9} c^{max}(F) \geq \frac{c^{max}(F)}{9}$.

C.3 Proof of Proposition 7

We follow the proof of Theorem 4 (and re-use same notations), but using different sets $\tilde{\mathcal{U}}_i$. For any edge $\{i, j\}$ of F , define $u_i^j \in \mathcal{U}_i$ and $u_j^i \in \mathcal{U}_j$ such that $d(u_i^j, u_j^i) = d_{ij}^{max}$. For any vertex i , let $\tilde{\mathcal{U}}_i = \{u_i^\ell, \ell \in N(i)\}$, and $\Delta_i = |N(i)|$ be the degree of i . Let $\tilde{c}(F) = \max_{u \in \tilde{\mathcal{U}}} \sum_{\{i,j\} \in F} d(u_i, u_j)$. As $\tilde{\mathcal{U}} \subseteq \mathcal{U}$, we get $c(F) \geq \tilde{c}(F)$. Let us now prove that $\tilde{c}(F) \geq \frac{c^{max}(F)}{\Delta}$.

Now, consider random variable \tilde{u}_i taking any value of $\tilde{\mathcal{U}}_i$ with equal probability. As in the proof of Theorem 4, it is enough to lower bound the following quantity for an arbitrary edge $\{i, j\} \in F$:

$$\mathbb{E} [d(\tilde{u}_i, \tilde{u}_j)] = \frac{1}{\Delta_i \Delta_j} \left(\sum_{\ell_1 \in N(i), \ell_2 \in N(j)} d(u_i^{\ell_1}, u_j^{\ell_2}) \right).$$

Suppose without loss of generality that $\Delta_j \leq \Delta_i$. Let $\Delta^{min} = \min(\Delta_i, \Delta_j) = \Delta_j$, $X_i = N(i) \setminus \{j\}$, and $X_j = N(j) \setminus \{i\}$ (X_i or X_j may be empty when $\Delta^{min} = 1$). Let us define an

arbitrary injective mapping $\phi : X_j \rightarrow X_i$. Observe that

$$\begin{aligned}
\sum_{\ell_1 \in N(i) \ell_2 \in N(j)} d(u_i^{\ell_1}, u_j^{\ell_2}) &\geq d_{ij}^{max} + \sum_{\ell \in X_j} d(u_i^j, u_j^\ell) + \sum_{\ell_1 \in X_i, \ell_2 \in X_j} d(u_i^{\ell_1}, u_j^{\ell_2}) + \sum_{\ell \in X_i} d(u_i^\ell, u_j^i) \\
&\geq d_{ij}^{max} + \sum_{\ell \in X_j} d(u_i^j, u_j^\ell) + \sum_{\ell \in X_j} d(u_j^\ell, u_i^{\phi(\ell)}) + \sum_{\ell \in X_i} d(u_i^\ell, u_j^i) \\
&\geq d_{ij}^{max} + \sum_{\ell \in X_j} d(u_i^j, u_j^\ell) + \sum_{\ell \in X_j} d(u_j^\ell, u_i^{\phi(\ell)}) + \sum_{\ell \in X_j} d(u_i^{\phi(\ell)}, u_j^i) \\
&= d_{ij}^{max} + \sum_{\ell \in X_j} \left(d(u_i^j, u_j^\ell) + d(u_j^\ell, u_i^{\phi(\ell)}) + d(u_i^{\phi(\ell)}, u_j^i) \right) \\
&\geq d_{ij}^{max} + \sum_{\ell \in X_j} d_{ij}^{max} = \Delta^{min} d_{ij}^{max}.
\end{aligned}$$

We obtain

$$\mathbb{E} [d(\tilde{u}_i, \tilde{u}_j)] \geq \frac{1}{\Delta_i \Delta_j} \Delta^{min} d_{ij}^{max} = \frac{1}{\Delta_i} d_{ij}^{max} \geq \frac{1}{\Delta(F)} d_{ij}^{max}.$$

C.4 Proof of Proposition 10

Recall that n denotes $|V(F)|$. It is folklore that $\chi(F) = n - 1$ when n is even, and $\chi(F) = n$ when n is odd, leading to $\chi(F) \leq n$ for any n . This implies that $E(F)$ can be partitioned into n matchings M_i , each of size $\lfloor n/2 \rfloor$, and thus that $c^{max}(F) = \sum_{i=1}^n c^{max}(M_i)$. Therefore, there is a matching of F , denoted M^* , such that

$$c^{max}(M^*) \geq \frac{1}{n} c^{max}(F). \quad (10)$$

Next, we define u^* as any element from $\arg \max_{u \in \mathcal{U}} \sum_{\{i,j\} \in M^*} d(u_i, u_j)$, and we use the shorter notation $d_{ij}^* = d(u_i^*, u_j^*)$, and $d_e^* = d_{ij}^*$ for any edge $e = \{i, j\}$. Observe that because M^* is a matching

$$d_e^* = d_e^{max} \quad (11)$$

for each $e \in M^*$. For any $E' \subseteq E(F)$, let $d^*(E') = \sum_{e \in E'} d_e^*$. Observe that $c(F) \geq d^*(F)$. Our objective is to prove that $d^*(F) \geq \frac{1}{2} c^{max}(F)$.

Assume without loss of generality that $M^* = \{\{2i - 1, 2i\} \mid i \in [\lfloor \frac{n}{2} \rfloor]\}$, so n is the only vertex not belonging to any edge of M^* when n is odd. For any $i \in [\lfloor \frac{n}{2} \rfloor]$, let $X(2i - 1, 2i) = \{\{2i - 1, l\} \cup \{2i, l\} \mid l \in V(F) \setminus \{2i - 1, 2i\}\}$. Observe that the triangle inequality yields

$$d^*(X(2i - 1, 2i)) = \sum_{l \in V(F) \setminus \{2i - 1, 2i\}} (d_{(2i-1)l}^* + d_{2il}^*) \geq (n - 2) d_{(2i-1)(2i)}^*. \quad (12)$$

Summing up (12) for all $i \in [\lfloor \frac{n}{2} \rfloor]$, we obtain

$$\sum_{i \in [\lfloor \frac{n}{2} \rfloor]} d^*(X(2i-1, 2i)) \geq (n-2)d^*(M^*) = c^{max}(M^*),$$

where the last equality follows from (11). What is more, any edge $e \in E(F) \setminus M^*$ belongs to at most two sets $X(2i-1, 2i)$, so that

$$2d^*(E(F) \setminus M^*) \geq \sum_{i \in [\lfloor \frac{n}{2} \rfloor]} d^*(X(2i-1, 2i)).$$

We obtain

$$2d^*(F) = 2 \left[d^*((E(F) \setminus M^*)) + d^*(M^*) \right] \geq (n-2)c^{max}(M^*) + 2c^{max}(M^*) = nc^{max}(M^*) \geq c^{max}(F),$$

where the last inequality follows from (10).

C.5 Proof of Proposition 14

Observe first that we can partition $E(F)$ into E_1 and E_2 such that each E_i induces a star forest (a graph where any connected component is a star). Indeed, to obtain such a partition we root the tree at vertex 1, and define S^i as the star whose central vertex is i and whose leaves are the children vertices of i in F . Then, we define E_1 (resp. E_2) as the union of $E(S^i)$ for vertices i which are an odd (resp. even) distance from vertex 1. By Proposition 6 and Corollary 4, we get the claimed ratio.

C.6 Proof of Proposition 15

We assume without loss of generality that F is rooted at vertex 1. The proof is made by induction on the height of the tree. For this, we consider the following induction statement:

P(h) : *If F has height h , there are two solutions $u^1, u^2 \in \mathcal{U}$ such that*

- $c^{max}(F) \leq 2\sqrt{2} \sum_{\{i,j\} \in F} d(u_i^1, u_j^1) = \sum_{\{i,j\} \in F} d(u_i^2, u_j^2)$,
- $u_i^1 = u_i^2$ for each vertex i with level $l < h$,
- $[u_i^1, u_i^2]$ is a diameter of \mathcal{U}_i for each vertex i with level h .

If $h = 0$, F has only one vertex which is the root of the tree, so the induction statement is trivially satisfied with any diameter $[u^1, u^2]$ of \mathcal{U}_1 .

Assume that $P(h)$ is true for some $L \geq 1$, and let F be a rooted tree with height $h + 1$. Without loss of generality, we assume that the vertices of V^F are sorted by increasing level and let n_{h-1} , n_h and n_{h+1} be such that the vertices with level h are $V_h^F = \{n_{h-1} + 1, \dots, n_h\}$ and those with level $h + 1$ are $V_{h+1}^F = \{n_h + 1, \dots, n_{h+1}\}$. Let $F_{\leq h} = (V_{\leq h}^F, F_{\leq h})$ be the subtree of F induced by vertices $\{1, \dots, n_h\}$. Tree $F_{\leq h}$ has height h so we can apply the induction statement to get two solutions $u^1, u^2 \in \mathcal{U}$ as described in **P(h)**.

Now, for all $i \in V_h^F$, let $S_i \subset F$ be the set of edges of the star graph whose internal vertex is i and whose leaves are the children vertices of i (which all belong to V_{h+1}^F). Similarly to what was done in the proof of Proposition 12, for all $\{i, j\} \in S_i$, we can set $u_j^1 \in \arg \max_{u \in \mathcal{U}_j} d(u, u_i^1)$ and $u_j^2 \in \arg \max_{u \in \mathcal{U}_j} d(u, u_i^2)$, which yields $d(u_i^1, u_j^1) + d(u_i^2, u_j^2) \geq d_{ij}^{max}$. We can then assume without loss of generality that $2 \sum_{\{i,j\} \in S_i} d(u_i^1, u_j^1) \geq \sum_{\{i,j\} \in S_i} d(u_i^2, u_j^2)$, i.e.,

$$c^{max}(S_i) \leq 2 \sum_{\{i,j\} \in S_i} d(u_i^1, u_j^1). \quad (13)$$

Given that we are considering the Euclidean distance with spherical uncertainty sets, the above implies that segment $[u_i^1, u_j^1]$ goes through the center o_j of \mathcal{U}_j (direct application of the triangular inequality). Then, let $[\bar{u}_j^1, \bar{u}_j^2]$ be the diameter of \mathcal{U}_j that is orthogonal to $[u_i^1, u_j^1]$; it exists because \mathcal{U}_j is spherical. We then compute,

$$\begin{cases} d(u_i^1, u_j^1)^2 = d(u_i^1, o_j)^2 + r_j^2 + 2r_j d(u_i^1, o_j) \leq 2d(\bar{u}_j^1, o_j)^2 + r_j^2 \\ d(u_i^1, \bar{u}_j^1)^2 = d(u_i^1, \bar{u}_j^2)^2 = d(u_i^1, o_j)^2 + r_j^2 \end{cases}$$

As a consequence, $d(u_i^1, \bar{u}_j^1) = d(u_i^1, \bar{u}_j^2) \geq \sqrt{2}d(u_i^1, u_j^1)$. Using (13), we then get

$$c^{max}(S_i) \leq 2\sqrt{2} \sum_{\{i,j\} \in S_i} d(u_i^1, \bar{u}_j^1) = 2\sqrt{2} \sum_{\{i,j\} \in S_i} d(u_i^1, \bar{u}_j^2)$$

. The same applies to all i , so we build two solutions $\tilde{u}^1, \tilde{u}^2 \in \mathcal{U}$ such that

- $\tilde{u}_i^1 = \tilde{u}_i^2 = u_i^1$ for all $i \in V_{\leq h}^F$,
- $[\tilde{u}_i^1, \tilde{u}_i^2] = [\bar{u}_i^1, \bar{u}_i^2]$ is a diameter of \mathcal{U}_i for all $i \in V_{h+1}^F$.

To conclude, we observe that $F = F_{\leq h} \cup \left(\bigcup_{i \in V_h^F} S_i \right)$, so

$$\begin{aligned}
\sum_{\{i,j\} \in F} d(\bar{u}_i^1, \bar{u}_j^1) &= \sum_{\{i,j\} \in F_{\leq h}} d(\bar{u}_i^1, \bar{u}_j^1) + \sum_{i \in V_h^F} \sum_{\{i,j\} \in S_i} d(\bar{u}_i^1, \bar{u}_j^1) \\
&= \sum_{\{i,j\} \in F_{\leq h}} d(u_i^1, u_j^1) + \sum_{i \in V_h^F} \sum_{\{i,j\} \in S_i} d(u_i^1, \bar{u}_j^1) \\
&\geq \frac{1}{2\sqrt{2}} c^{\max}(F_{\leq h}) + \frac{1}{2\sqrt{2}} \sum_{i \in V_h^F} c^{\max}(S_i) \\
&= \frac{1}{2\sqrt{2}} c^{\max}(F).
\end{aligned}$$

D. Tightness proofs

D.1 Proof of Proposition 8

We consider a path F composed of $n \geq 3$ vertices, where $E(F) = \{\{i, i+1\} : i \in \{1, \dots, n-1\}\}$. The vertices are located on a one-dimensional line where $\mathcal{U}_1 = \{0\}$, $\mathcal{U}_2 = \{0, 1\}$ and $\mathcal{U}_3 = \dots = \mathcal{U}_n = \{1\}$. We have $d_{12}^{\max} = d_{23}^{\max} = 1$ and $d_{i,i+1}^{\max} = 0, \forall i = 3, \dots, n-1$, so $c^{\max}(F) = 2$. There are only two feasible solutions depending on whether $u_2 = 1$ or $u_2 = 0$, and they have the same cost $c(F) = 1$.

D.2 Proof of Proposition 9

We consider a cycle F composed of $n \geq 4$ vertices, where $E(F) = \{\{i, i+1\} : i \in \{1, \dots, n-1\}\} \cup \{\{n, 1\}\}$. The vertices are located on a one-dimensional line where $\mathcal{U}_1 = \{0\}$, $\mathcal{U}_2 = \{0, 1\}$, $\mathcal{U}_3 = \{1\}$, $\mathcal{U}_4 = \{0, 1\}$ and $\mathcal{U}_5 = \dots = \mathcal{U}_n = \{0\}$ if $n \geq 5$. We verify that $c^{\max}(F) = 4$, and there are four feasible solutions depending on whether $u_2 = 1$ or $u_2 = 0$ and $u_4 = 1$ or $u_4 = 0$. These four solutions all have the same cost $c(F) = 2$.

D.3 Proof of Proposition 11

According to Observation 4, we have $c(F) = \text{OPT}^{\text{MAX-CUT}}(F)$, and notice that $c^{\max}(F) = m$, where $m = |E(F)| = \frac{n(n-1)}{2}$. Thus, it only remains to compute $\text{OPT}^{\text{MAX-CUT}}(F)$ in both cases. When n is even, $\text{OPT}^{\text{MAX-CUT}}(F) = \frac{n^2}{4}$, and when n is odd, $\text{OPT}^{\text{MAX-CUT}}(F) = \frac{(n-1)}{2} \frac{(n+1)}{2}$, leading in both cases to the claimed ratio.

D.4 Proof of Proposition 13

We consider the uncertainty set $\mathcal{U} = \times_{i=1}^n \mathcal{U}_i$ where $\mathcal{U}_1 = \{u_1^2, \dots, u_1^n\}$ and $\mathcal{U}_i = \{u_i\}, i = 2, \dots, n$, such that:

- for all $(i, j) \in \{2, \dots, n\}^2, i \neq j, d(u_1^i, u_1^j) = 2/3$,
- for all $(i, j) \in \{2, \dots, n\}^2, i \neq j, d(u_i, u_j) = 2/3$,
- for all $i = 2, \dots, n, d(u_1^i, u_i) = 1$ and $\forall j \neq i, d(u_1^j, u_i) = 1/3$.

The triangular inequality is verified, so d is a distance.

By symmetry of the star graph and of the uncertainty set, every solution $u \in \mathcal{U}$ is optimal and has the same value $1 + (n - 2) \cdot 1/3$. Moreover, the maximum distance between \mathcal{U}_1 and \mathcal{U}_i is equal to 1 for all $i \in \{2, \dots, n\}$. As a result, $c^{max}(F) = 3(n - 1)/(n + 1)c(F)$.

E. Proofs related to the FPTAS

E.1 Proof of Lemma 4

For complexity issues we assume that $\mathcal{P}^{(j, \kappa-1)}$ and $Q^{(j, \kappa-1)}$ are represented as arrays A_P and A_Q indexed by profiles, where given a profile p , $A_P[p]$ is true iff $p \in \mathcal{P}^{(j, \kappa-1)}$, and $A_Q[p]$ contains a path P such that $Pr(P) = p$ if $p \in \mathcal{P}^{(j, \kappa-1)}$, and \emptyset otherwise. This explains the $\mathcal{O}(n_P n (\log(n)))$ required space for storing Q . As profiles are vectors of length at most σ , we consider that it takes $\mathcal{O}(\sigma)$ to obtain the value stored at index p of array A_P or A_Q .

We compute $\mathcal{P}^{(i, \kappa)}$ following the recursion relation (8), and compute $Q^{(i, \kappa)}$ along the way. More precisely, we start by initializing two arrays A'_P and A'_Q of size n_P . Then, for all $j \in N(i)$ and $p' \in \mathcal{P}^{(j, \kappa-1)}$, we compute $p(i, \kappa, j, p')$ in time $\mathcal{O}(\sigma^2)$. Now, we perform the following operations in $\mathcal{O}(1)$. If $p(i, \kappa, j, p')$ is not already in \mathcal{P} , we add it to \mathcal{P} , we find a path P' in $Q^{(j, \kappa-1)}$ such that $Pr(P') = p'$, and we add the path iP' to Q .

E.2 Proof of Theorem 7

Let $I = (G, s, t, \mathcal{U}, \mathcal{D})$ be an instance of ROBUST- Δ -SP. Our objective is to provide a solution of cost at most $(1 + \epsilon)\text{OPT}(I)$. Let $\epsilon' \in \mathbb{R}^+$. Using the 2-approximation obtained combining Theorem 2 and Proposition 7, we start by computing an $s - t$ path P_A of cost $c(P_A) = A$, where $\text{OPT}(I) \leq A \leq 2\text{OPT}(I)$. For any $x, y \in \bigcup_{i \in V(G)} \mathcal{U}_i$, we define $\mathcal{D}'(x, y)$ by rounding up

$\mathcal{D}(x, y)$ to the closest value of the form $\ell\epsilon'A$ for some $\ell \in \mathbb{N}$. For any path P and $u \in \mathcal{U}$ we denote by $c'(u, P) = \sum_{i,j \in P} \mathcal{D}'(u_i, u_j)$, $c'(P) = \max_{u \in \mathcal{U}} c'(u, P)$. Let $I' = (G, s, t, \mathcal{U}, \mathcal{D}')$ be the instance of ROBUST-SP obtained when using \mathcal{D}' instead of \mathcal{D} .

Observe that

- the function $d'(u, v) = \mathcal{D}'(u, v)$ may not be a distance,
- for any $x, y \in \bigcup_{i \in V(G)} \mathcal{U}_i$, we have $\mathcal{D}(u, v) \leq \mathcal{D}'(u, v) \leq \mathcal{D}(u, v) + \epsilon'A$
- for any path P , $c(P) \leq c'(P) \leq c(P) + n\epsilon'A$
- $\text{OPT}' \leq \text{OPT} + n\epsilon'A$.

Let $i \in V(G)$ and P be an $i - t$ path. We say that P is *useless* if it verifies $c'(P) > A(1 + n\epsilon')$; otherwise, P is said to be *good*. According to previous observations, we see that $c'(P_A) \leq A(1 + n\epsilon')$. Thus, for any $i \in V(G)$ and useless $i - t$ path P , we have $c'(P) > c'(P_A)$. This implies that P cannot be the suffix of an optimal solution to input I' (meaning that there is no optimal solution of I' that first goes from s to i , and then uses P). As a consequence, in the DP algorithm provided in Section 6.1, we can restrict our attention to the profiles of good paths, without loosing optimality in I' . More formally, for any $i \in V(G)$ and $\kappa \in [n]$, we adapt the previous definition $\mathcal{P}^{(i, \kappa)}$ to

$$\mathcal{P}_g^{(i, \kappa)} = \{Pr(P) \mid P \in \mathcal{F}_\kappa^{i \rightarrow t}, P \text{ is good}\},$$

and we now consider that the DP algorithm $A_g(i, \kappa)$ computes $(\mathcal{P}_g^{(i, \kappa)}, Q^{(i, \kappa)})$ instead of $(\mathcal{P}^{(i, \kappa)}, Q^{(i, \kappa)})$. We now compute P^* an optimal solution on instance I' (for cost function c') using Theorem 6 (with DP algorithm A_g), and output P^* .

We have $c(P^*) \leq c'(P^*) = \text{OPT}(I') \leq \text{OPT}(I) + n\epsilon'A \leq \text{OPT}(I)(1 + 2n\epsilon')$. Let us now consider the complexity of computing P^* . As $c'(P) = \max_{\ell \in [|\mathcal{U}_i|]} c'^\ell(P)$, observe that for any good path P we have $c'^\ell(P) \leq A(1 + n\epsilon')$. Moreover, as for any $x, y \in \bigcup_{i \in V(G)} \mathcal{U}_i$, $\mathcal{D}'(u, v)$ is a multiple of $\epsilon'A$, we get that for any good path P , $c'^\ell(P) = \ell\epsilon'A$ for $0 \leq \ell \leq n + \lceil \frac{1}{\epsilon'} \rceil$. This implies that $n_{val} \leq (n + 2 + \frac{1}{\epsilon'}) = \mathcal{O}(\frac{n}{\epsilon'})$, so Observation 3 leads to the desired complexity of $\mathcal{O}(n^3 \sigma^2 (\frac{n}{\epsilon'})^\sigma)$ to get a ratio $1 + 2n\epsilon'$. Finally, given a target ratio $1 + \epsilon$, we set $\epsilon' = \frac{\epsilon}{2n}$, obtaining the claimed complexity.

F. Definitions related to parameterized complexity and treewidth

F.1 Parameterized complexity

We refer the reader to Downey and Fellows (2013), Cygan et al. (2015) for basic background on parameterized complexity, and we recall here only some basic definitions. A *parameterized problem* is a language $L \subseteq \Sigma^* \times \mathbb{N}$, where Σ is some fixed alphabet. For an instance $I = (x, k) \in \Sigma^* \times \mathbb{N}$, k is called the *parameter*. Given a classical (non-parameterized) decision problem $L_c \subseteq \Sigma^*$ and a function $\kappa : \Sigma^* \rightarrow \mathbb{N}$, we denote by $L_c/\kappa = \{(x, \kappa(x)) \mid x \in L_c\}$ the associated parameterized problem.

A parameterized problem L is *fixed-parameter tractable* (\mathcal{FPT}) if there exists an algorithm \mathcal{A} , a computable function f , and a constant c such that given an instance $I = (x, k)$, \mathcal{A} (called an \mathcal{FPT} algorithm) correctly decides whether $I \in L$ in time bounded by $f(k) \cdot |I|^c$. For instance, the VERTEX COVER problem parameterized by the size of the solution is \mathcal{FPT} .

Within parameterized problems, the \mathcal{W} -hierarchy may be seen as the parameterized equivalent to the class \mathcal{NP} of classical decision problems. Without entering into details (see Downey and Fellows (2013), Cygan et al. (2015) for the formal definitions), a parameterized problem being $\mathcal{W}[1]$ -hard can be seen as a strong evidence that this problem is *not* \mathcal{FPT} . The canonical example of $\mathcal{W}[1]$ -hard problem is INDEPENDENT SET parameterized by the size of the solution.

The most common way to transfer $\mathcal{W}[1]$ -hardness is via parameterized reductions. A *parameterized reduction* from a parameterized problem L_1 to a parameterized problem L_2 is an algorithm that, given an instance (x, k) of L_1 , outputs an instance (x', k') of L_2 such that

- (x, k) is a yes-instance of L_1 if and only if (x', k') is a yes-instance of L_2 ,
- $k' \leq g(k)$ for some computable function g , and
- the running time is bounded by $f(k) \cdot |x|^{\mathcal{O}(1)}$ for some computable function f

If L_1 is $\mathcal{W}[1]$ -hard and there is a parameterized reduction from L_1 to L_2 , then L_2 is $\mathcal{W}[1]$ -hard as well.

F.2 Tree decompositions and treewidth.

A *tree decomposition* of a graph G is a pair $\mathcal{D} = (T, \mathcal{B})$, where T is a tree and $\mathcal{B} = \{X^w \mid w \in V(T)\}$ is a collection of subsets of $V(G)$, called *bags*, such that:

- $\bigcup_{w \in V(T)} X^w = V(G)$,
- for every edge $\{i, j\} \in E$, there is a $w \in V(T)$ such that $\{i, j\} \subseteq X^w$, and
- for every $\{x, y, z\} \subseteq V(T)$ such that z lies on the unique path between x and y in T , $X^x \cap X^y \subseteq X^z$.

We call the vertices of T *vertices* of \mathcal{D} and the sets in \mathcal{B} *bags* of \mathcal{D} . The *width* of a tree decomposition $\mathcal{D} = (T, \mathcal{B})$ is $\max_{w \in V(T)} |X^w| - 1$. The *treewidth* of a graph G , denoted by $tw(G)$, is the smallest integer t such that there exists a tree decomposition of G of width at most t . We need to introduce nice tree decompositions, which will make the presentation of the algorithm used to proof Theorem 1 much simpler.

Let $\mathcal{D} = (T, \mathcal{B})$ be a rooted tree decomposition of G (meaning that T has a special vertex r called the *root*). As T is rooted, we naturally define an ancestor relation among bags, and say that $X^{w'}$ is a *descendant* of X^w if the vertex set of the unique simple path in T from r to w' contains w . In particular, every vertex w is a descendant of itself. For every $w \in V(T)$, we define $G^w = G[\bigcup\{X^{w'} \mid X^{w'} \text{ is a descendant of } X^w \text{ in } T\}]$.

Such a rooted decomposition is called a *nice tree decomposition* of G if the following conditions hold:

- $X^r = \emptyset$.
- Every vertex of T has at most two children in T .
- For every leaf $\ell \in V(T)$, $X^\ell = \emptyset$. Each such vertex ℓ is called a *leaf vertex*.
- If $w \in V(T)$ has exactly one child w' , then either
 - $X^w = X^{w'} \cup \{i\}$ for some $i \notin X^{w'}$. Each such vertex is called an *introduce vertex*.
 - $X^w = X^{w'} \setminus \{i\}$ for some $i \in X^{w'}$. Each such vertex is called a *forget vertex*.
- If $w \in V(T)$ has exactly two children w_L and w_R , then $X^w = X^{w_L} = X^{w_R}$. Each such vertex w is called a *join vertex*.

We recall that one of the key property of such a nice decomposition is that for any $w \in V(T)$, X^w is a separator of G . This implies in particular that, in a join vertex, there is no edge $\{i, j\}$ in G^w where $i \in V(G^{w_L}) \setminus X^w$ and $j \in V(G^{w_R}) \setminus X^w$.

Given a tree decomposition of a graph G of width t and x vertices, it is possible to transform it in polynomial time into a *nice* one of width t and xt vertices (Kloks 1994). Moreover, it is possible (Bodlaender et al. (2013)) to compute a tree decomposition of width $tw' = \mathcal{O}(tw(G))$ and $\mathcal{O}(n)$ vertices in time $\mathcal{O}(c^{tw(G)}n)$, where $n = |V(G)|$. By using these two results, we can compute in time $\mathcal{O}(c^{tw(G)}n)$ a nice tree decomposition of width $\mathcal{O}(tw(G))$ with $\mathcal{O}(tw(G)n)$ vertices.

G. Computing the objective function on small treewidth graphs

Throughout this section, we denote by $u|_X$ the vector u restricted to components u_i such that $i \in X$, for any $X \subseteq V(G)$.

G.1 Definition of the auxiliary problem

In this section we consider that we are given a fixed input of EVAL-C, and a nice tree decomposition $\mathcal{D} = (T, \mathcal{B})$ of G . Given $w \in V(T)$, we denote $\mathcal{U}^w = \times_{i \in V(G^w)} \mathcal{U}_i$. Let us define the following maximization problem Π . An input of Π is a pair (w, f) where $w \in V(T)$, and f is a function from X^w to \mathcal{M} such that for any $i \in X^w$, $f(i) \in \mathcal{U}_i$. An output is a vector $u \in \mathcal{U}^w$ such that for any $i \in X^w$, $u_i = f(i)$, which we denote by $u \vdash (w, f)$. The objective is to maximize $c(u, G^w)$. We denote by $\text{OPT}(w, f)$ the optimal value for instance (w, f) . As usual in DP algorithms, to simplify the presentation we will define an algorithm A that given an input (w, f) only computes the value $\text{OPT}(w, f)$. This algorithm could be easily modified to get an associated optimal solution.

G.2 Join case

Let w be a join vertex with children w^L and w^R . Given two vectors $u^L \in \mathcal{U}^{w^L}$ and $u^R \in \mathcal{U}^{w^R}$, such that for any $i \in X^w$, $u_i^L = u_i^R$, we define $u = u^L \diamond u^R$ by $u_i = u_i^L$ for any $i \in V(G^{w^L})$, and $u_i = u_i^R$ for any $i \in V(G^{w^R})$. Observe that u is well defined as for $i \in X^w$, $u_i^L = u_i^R$.

Lemma 6. *Let (w, f) be an input of Π such that w is a join vertex with children w^L and w^R . For any $u \in \mathcal{U}^w$, $u \vdash (w, f)$ iff there exists u^L, u^R such that the following conditions hold:*

- $u^L \vdash (w^L, f)$
- $u^R \vdash (w^R, f)$
- $u = u^L \diamond u^R$

Proof. The \Rightarrow direction is immediate by defining $u^L = u|_{V(G^{w_L})}$ (resp. $u^R = u|_{V(G^{w_R})}$). In the \Leftarrow direction, observe that $u^L \diamond u^R$ is well defined as for any $i \in X^w$, $u_i^L = u_i^R = f(i)$, and $u \vdash (w, f)$ is also immediate. \square

Lemma 7. *Let (w, f) be an input of Π such that w is a join vertex with children w^L and w^R . Then, $\text{OPT}(w, f) = \text{OPT}(w^L, f) + \text{OPT}(w^R, f) - d^{(w, f)}$, where $d^{(w, f)} = \sum_{i, j \in X^w, \{i, j\} \in E(G)} d(f(i), f(j))$.*

Proof. Let us start with the \leq inequality. Let u such that $c(u, G^w) = \text{OPT}(w, f)$. Let u^L and u^R as defined by Lemma 6. Observe that $c(u, G^w) = c(u, G^{w_L}) + c(u, G^{w_R}) - d^{(w, f)}$ as edges inside X^w are counted twice in the first two terms. We have $c(u, G^{w_L}) = c(u^L, G^{w_L})$, and $c(u^L, G^{w_L}) \leq \text{OPT}(w^L, f)$ as $u^L \vdash (w^L, f)$, and same properties hold for the right side. This implies $\text{OPT}(w, f) \leq \text{OPT}(w^L, f) + \text{OPT}(w^R, f) - d^{(w, f)}$.

Let us now turn to the other inequality. Let u^L such that $c(u^L, G^{w_L}) = \text{OPT}(w^L, f)$, u^R such that $c(u^R, G^{w_R}) = \text{OPT}(w^R, f)$, and $u = u^L \diamond u^R$. According to Lemma 6, $u \vdash (w, f)$, and again $c(u, G^w) = c(u^L, G^{w_L}) + c(u^R, G^{w_R}) - d^{(w, f)}$, implying the desired inequality. \square

We are now ready to define the DP algorithm A in the join case. Given an input (w, f) of Π such that w is a join vertex with children w^L and w^R , $A(w, f)$ returns $A(w^L, f) + A(w^R, f) - d^{(w, f)}$. It is immediate by induction and using Lemma 7 that $A(w, f) = \text{OPT}(w, f)$.

G.3 Introduce case

Given any input (w, f) of Π and $X \subseteq X^w$, we denote by $f|_X$ function f restricted to X . The following two lemmas are immediate.

Lemma 8. *Let (w, f) be an input of Π such that w is an introduce vertex with children w' . Let i be such that $X^w = X^{w'} \cup \{i\}$. For any $u \in \mathcal{U}^w$, $u \vdash (w, f)$ iff the following conditions hold:*

- $u|_{V(G^{w'})} \vdash (w', f|_{X^{w'}})$
- $u_i = f(i)$

Lemma 9. *Let (w, f) be an input of Π such that w is an introduce vertex with children w' . Let i be such that $X^w = X^{w'} \cup \{i\}$. Then, $\text{OPT}(w, f) = \text{OPT}(w', f|_{X^{w'}}) + d^{(i,w,f)}$, where $d^{(i,w,f)} = \sum_{j \in X^w, \{i,j\} \in E(G)} d(f(i), f(j))$.*

We are now ready to define the DP algorithm A in the introduce case. Given an input (w, f) of Π such that w is an introduce vertex with children w' , where $X^w = X^{w'} \cup \{i\}$, $A(w, f)$ returns $A(w', f|_{X^{w'}}) + d^{(i,w,f)}$. It is immediate by induction and using Lemma 9 that $A(w, f) = \text{OPT}(w, f)$.

G.4 Forget case

Let (w, f) be an input of Π such that w is a forget vertex with children w' . Let i such that $X^{w'} = X^w \cup \{i\}$. For any $x \in \mathcal{M}$, we denote $f^{(i,x)}$ the function from $X^{w'}$ to \mathcal{M} such that $f^{(i,x)}(j) = f(j)$ for any $j \neq i$, and $f^{(i,x)}(i) = x$. The following Lemma is immediate.

Lemma 10. *Let (w, f) be an input of Π such that w is a forget vertex with children w' . Let i be such that $X^{w'} = X^w \cup \{i\}$. For any $u \in \mathcal{U}^w$, $u \vdash (w, f)$ iff $u \vdash (w', f^{(i,u_i)})$.*

Lemma 11. *Let (w, f) be an input of Π such that w is a forget vertex with children w' . Let i be such that $X^{w'} = X^w \cup \{i\}$. Then, $\text{OPT}(w, f) = \max_{x \in \mathcal{U}_i} \text{OPT}(w', f^{(i,x)})$.*

Proof. Observe first that $G^w = G^{w'}$. Let us start with the \leq inequality. Let u such that $c(u, G^w) = \text{OPT}(w, f)$. Notice that $c(u, G^w) = c(u, G^{w'})$. By Lemma 10, $u \vdash (w', f^{(i,u_i)})$, implying $c(u, G^{w'}) \leq \text{OPT}(w', f^{(i,u_i)}) \leq \max_{x \in \mathcal{U}_i} \text{OPT}(w', f^{(i,x)})$.

Let us now turn to the other inequality. Let $x^* \in \mathcal{U}_i$ maximizing the right side. Let u such that $c(u, G^{w'}) = \text{OPT}(w', f^{(i,x^*)})$. Notice that as $u \vdash (w', f^{(i,x^*)})$, $u_i = x^*$, and thus $u \vdash (w', f^{(i,u_i)})$. According to Lemma 10, $u \vdash (w, f)$, implying that $c(u, G^{w'}) = c(u, G^w) \leq \text{OPT}(w, f)$. \square

We are now ready to define the DP algorithm A in the forget case. Given an input (w, f) of Π such that w is a forget vertex with children w' , where $X^{w'} = X^w \cup \{i\}$, $A(w, f)$ returns $\max_{x \in \mathcal{U}_i} A(w', f^{(i,x)})$. It is immediate by induction and using Lemma 11 that $A(w, f) = \text{OPT}(w, f)$.

G.5 Putting pieces together

Theorem 8. $\text{EVAL-C}/tw + \sigma$ is \mathcal{FPT} . More precisely, we can compute an optimal solution of EVAL-C in time $\mathcal{O}(ntw\sigma^{\mathcal{O}(tw)})$, where $n = |V(G)|$, $tw = tw(G)$, and $\sigma = \max_{i \in V(G)} |\mathcal{U}_i|$.

Proof. Given an input $(\mathcal{M}, d, G, \mathcal{U})$ of EVAL-C , we start (see Appendix F.2) by computing in time $\mathcal{O}(c^{tw(G)}n)$ a nice tree decomposition of width $\mathcal{O}(tw(G))$ with $N = \mathcal{O}(ntw(G))$ vertices. Remember that this nice tree decomposition is rooted on a vertex r such that $X^r = \emptyset$. Then, we output $A(r, \emptyset)$. Notice that as $X^r = \emptyset$, the second parameter (the function from X^r to \mathcal{M}) is defined nowhere and denoted \emptyset . As A solves Π optimally, we have $A(r, \emptyset) = \text{OPT}(r, \emptyset)$. Moreover, as $G^r = G$, we have $\text{OPT}(r, \emptyset) = c(G)$.

Let us now consider the running time of A . Given a tree decomposition with N vertices (in the tree of bags) and of width t , the size of the DP table is $\mathcal{O}(N\sigma^t)$, the time to compute one entry is dominated by the forget case where the branching is in $\mathcal{O}(\sigma)$, implying a running time in $\mathcal{O}(N\sigma^{t+1})$. Plugging the corresponding values, we get the claimed running time. \square

H. Affine decision rules for Euclidean metric spaces

In what follows, we assume that (\mathcal{M}, d) is the ℓ -dimensional Euclidean space so applying an epigraphic reformulation to (5) leads to

$$\begin{aligned} \min \quad & \omega \\ \text{s.t.} \quad & \omega \geq \sum_{\{i,j\} \in E} x_{ij} \|u_i - u_j\|_2, \quad \forall u \in \mathcal{U} \\ & x \in \mathcal{X}. \end{aligned}$$

We provide next a conservative reformulation of the above problem, adapting the techniques proposed by Zhen et al. (2021). Let $\mathcal{U}_i = \text{conv}(u_i^1, \dots, u_i^\sigma)$ be defined by its σ extreme points for each $i \in V$, and denote the unit ball of dimension ℓ by \mathcal{W}^ℓ , as well as $\mathcal{W} = \times_{i \in V} \mathcal{W}^\ell$. Let us also direct arbitrarily every edge in E , leading to the set of directed edges \vec{E} . We obtain that the constraint $\forall u \in \mathcal{U} : \sum_{(i,j) \in \vec{E}} x_{ij} \|u_i - u_j\|_2 \leq \omega$ (notice that $x_{ij} = x_{ji}$) is equivalent

to

$$\begin{aligned}
& \forall u \in \mathcal{U} : \sum_{(i,j) \in \vec{E}} x_{ij} \max_{w_{ij} \in \mathcal{W}^\ell} w_{ij}^T (u_i - u_j) \leq \omega \\
& \Leftrightarrow \forall w \in \mathcal{W}, u \in \mathcal{U} : \sum_{(i,j) \in \vec{E}} x_{ij} w_{ij}^T (u_i - u_j) \leq \omega \\
& \Leftrightarrow \forall w \in \mathcal{W} : \max \left\{ \sum_{(i,j) \in \vec{E}} x_{ij} w_{ij}^T \left(\sum_{k=1}^{\sigma} \lambda_i^k u_i^k - \sum_{l=1}^{\sigma} \lambda_j^l u_j^l \right) : \sum_{k=1}^{\sigma} \lambda_i^k = 1, \forall i \in V, \lambda \geq 0 \right\} \leq \omega \\
& \Leftrightarrow \forall w \in \mathcal{W} : \min \left\{ \sum_{i \in V} \mu_i : \mu_i \geq \sum_{(i,j) \in \vec{E}} x_{ij} w_{ij}^T u_i^k - \sum_{(j,i) \in \vec{E}} x_{ji} w_{ji}^T u_i^k, \forall i \in V, k = 1, \dots, \sigma \right\} \leq \omega \\
& \Leftrightarrow \max_{w \in \mathcal{W}} \min \left\{ \sum_{i \in V} \mu_i : \mu_i \geq \sum_{(i,j) \in \vec{E}} x_{ij} w_{ij}^T u_i^k - \sum_{(j,i) \in \vec{E}} x_{ji} w_{ji}^T u_i^k, \forall i \in V, k = 1, \dots, \sigma \right\} \leq \omega.
\end{aligned}$$

Introducing affine decision rules

$$\mu_i(w) = \mu_i^0 + \sum_{e \in \vec{E}} \mu_{ie}^T w_e,$$

we obtain

$$\begin{aligned}
& \min \quad \omega \\
& \text{s.t.} \quad \omega \geq \sum_{i \in V} \left(\mu_i^0 + \sum_{e \in \vec{E}} \mu_{ie}^T w_e \right), \quad \forall w \in \mathcal{W} \\
& \quad \mu_i^0 + \sum_{e \in \vec{E}} \mu_{ie}^T w_e \geq \sum_{(i,j) \in \vec{E}} x_{ij} w_{ij}^T u_i^k - \sum_{(j,i) \in \vec{E}} x_{ji} w_{ji}^T u_i^k, \quad \forall i \in V, k = 1, \dots, \sigma, w \in \mathcal{W} \\
& \quad x \in \mathcal{X},
\end{aligned}$$

and after dualization

$$\begin{aligned}
& \min \quad \omega \\
& \text{s.t.} \quad \omega \geq \sum_{i \in V} \mu_i^0 + \sum_{(i,j) \in \vec{E}} \left\| \mu_{i,ij} + \mu_{j,ij} + \sum_{i' \neq i,j} \mu_{i',ij} \right\|_2 \\
& \quad \mu_i^0 \geq \sum_{(i,j) \in \vec{E}} x_{ij} \|u_i^k - \mu_{i,ij}\|_2 + \sum_{(j,i) \in \vec{E}} x_{ji} \|u_i^k + \mu_{i,ji}\|_2 + \sum_{e \in \vec{E} \setminus \delta(i)} \|\mu_{ie}\|_2, \quad \forall i \in V, k = 1, \dots, \sigma \\
& \quad x \in \mathcal{X}.
\end{aligned}$$

Noticing that we can set $\mu_{ie} = 0$ for all $e \in \vec{E} \setminus \delta(i)$ without affecting optimality, we obtain

$$\begin{aligned}
& \min \quad \omega \\
& \text{s.t.} \quad \omega \geq \sum_{i \in V} \mu_i^0 + \sum_{(i,j) \in \vec{E}} \nu_{ij} \\
& \quad \mu_i^0 \geq \sum_{(i,j) \in \vec{E}} x_{ij} \nu_{i,ij}^k + \sum_{(j,i) \in \vec{E}} x_{ji} \nu_{i,ji}^k, \quad \forall i \in V, k = 1, \dots, \sigma \\
& \quad \nu_{ij} \geq \|\mu_{i,ij} + \mu_{j,ij}\|_2, \quad \forall (i,j) \in \vec{E} \\
& \quad \nu_{i,ij}^k \geq \|u_i^k - \mu_{i,ij}\|_2, \quad \forall i \in V, (i,j) \in \delta^+(i), k = 1, \dots, \sigma \\
& \quad \nu_{i,ji}^k \geq \|u_i^k + \mu_{i,ji}\|_2, \quad \forall i \in V, (j,i) \in \delta^-(i), k = 1, \dots, \sigma \\
& \quad x \in \mathcal{X}.
\end{aligned}$$

We linearize products $x_{ij} \nu_{i,ij}^k$ and $x_{ji} \nu_{i,ji}^k$ by introducing variables $\pi_{i,ij}^k \geq 0$ and $\pi_{i,ji}^k \geq 0$ together with constraints $\pi_{i,ij}^k \geq \nu_{i,ij}^k - M_{ij}(1 - x_{ij})$ and $\pi_{i,ji}^k \geq \nu_{i,ji}^k - M_{ij}(1 - x_{ij})$, with $M_{ij} = d_{ij}^{max}$.