

Compromise Solutions for Robust Combinatorial Optimization with Variable-Sized Uncertainty

André Chassein^{*1} and Marc Goerigk^{†2}

¹Fachbereich Mathematik, Technische Universität Kaiserslautern, Germany

²Department of Management Science, Lancaster University, United Kingdom

Abstract

In classic robust optimization, it is assumed that a set of possible parameter realizations, the uncertainty set, is modeled in a previous step and part of the input. As recent work has shown, finding the most suitable uncertainty set is in itself already a difficult task. We consider robust problems where the uncertainty set is not completely defined. Only the shape is known, but not its size. Such a setting is known as variable-sized uncertainty.

In this work we present an approach how to find a single robust solution, that performs well on average over all possible uncertainty set sizes. We demonstrate that this approach can be solved efficiently for min-max robust optimization, but is more involved in the case of min-max regret, where positive and negative complexity results for the selection problem, the minimum spanning tree problem, and the shortest path problem are provided. We introduce an iterative solution procedure, and evaluate its performance in an experimental comparison.

Keywords: robustness and sensitivity analysis; robust combinatorial optimization; min-max regret; variable-sized uncertainty

1 Introduction

Classic optimization settings assume that the problem data are known exactly. Robust optimization, like stochastic optimization, instead assumes some degree of uncertainty in the problem formulation. Based on the seminal papers [BTN98, BTN00, BTN02], most approaches in robust optimization formalize this uncertainty by assuming that all uncertain parameters ξ are described by a set of possible outcomes \mathcal{U} , the uncertainty set.

For general overviews on robust optimization, we refer to [BTGN09, BBC11, GMT14]. Other surveys focus on robust combinatorial optimization [ABV09, KZ16], algorithmic developments [GS16] or present tutorials to the field [GYdH15, CG16].

*Email: chassein@mathematik.uni-kl.de

†Corresponding author. Email: m.goerigk@lancaster.ac.uk

While the discussion of properties of the robust problem for different types of uncertainty sets \mathcal{U} has always played a major role in the research community, only recently the data-driven design of useful sets \mathcal{U} has become a focus of research. In [BGK13], the authors discuss the design of \mathcal{U} taking problem tractability and probabilistic guarantees of feasibility into account. The paper [BB09] discusses the relationship between risk measures and uncertainty sets, and [YdH12] constructs uncertainty sets by data-driven approximations of ambiguous chance constraints.

In distributionally robust optimization, one assumes that a probability distribution on the data is roughly known; however, this distribution itself is subject to an uncertainty set \mathcal{U} of possible outcomes (see [GS10, BTdHdW⁺13, WKS14]).

Another related approach is the globalized robust counterpart, see [BTGN09]. The idea of this approach is that a relaxed feasibility should be maintained, even if a scenario occurs that is not specified in the uncertainty set. The larger the distance of ξ to \mathcal{U} , the further relaxed becomes the feasibility requirement of the robust solution.

In this work we present an alternative to constructing a specific uncertainty set \mathcal{U} . Instead, we only assume knowledge of a nominal (undisturbed) scenario, and consider a set of possible uncertainty sets of varying size based on this scenario. That is, a decision maker does not need to determine the size of uncertainty, but only its shape. Our goal is to construct a solution for which the worst-case objective with respect to any possible uncertainty set performs well on average over all uncertainty sets parameterized this way.

The general idea of variable-sized uncertainty that this work is based upon was recently introduced in [CG18]. There, the aim is to construct a set of robust candidate solutions that requires the decision maker to choose one that suits him best. This is inspired by parametric optimization, where one traces the change in the optimal solution while problem parameters vary. In our setting, we consider all uncertainty sizes simultaneously, and generate a single solution as a compromise approach to the unknown uncertainty. It can hence be regarded as a goal programming approach for variable-sized uncertainty. We call this setting the *compromise approach to variable-sized uncertainty*.

We focus on combinatorial optimization problems with uncertainty in the objective function, and consider both min-max and min-max regret robustness (see [KZ16]).

This work is structured as follows. In Section 2, we briefly formalize the setting of variable-sized uncertainty. We then introduce our new compromise approach for min-max regret robustness in Section 3. We present complexity results for the selection problem, the minimum spanning tree problem, and the shortest path problem in Section 4, before discussing the case of min-max robustness in Section 3. In Section 6, we evaluate our approach in a computation experiment, and conclude this paper in Section 7.

2 Variable-Sized Uncertainty

In the following, we use the notation $[n] := \{1, \dots, n\}$ and write vectors and matrices in bold, e.g., $\mathbf{x} = (x_i)_{i \in [n]}$. We briefly summarize the setting of [CG18], where the term "variable-sized uncertainty" was coined. Consider an uncertain

combinatorial problem of the form

$$\min \{ \mathbf{c}\mathbf{x} : \mathbf{x} \in \mathcal{X} \} \quad (\text{P}(\mathbf{c}))$$

with $\mathcal{X} \subseteq \{0, 1\}^n$, and an uncertainty set $\mathcal{U}(\lambda) \subseteq \mathbb{R}_+^n$ that is parameterized by some value $\lambda \in \Lambda$. For example,

- interval-based uncertainty $\mathcal{U}(\lambda) = \prod_{i \in [n]} [(1 - \lambda)\hat{c}_i, (1 + \lambda)\hat{c}_i]$ with $\Lambda \subseteq [0, 1]$,
- general interval-based uncertainty $\mathcal{U}(\lambda) = \prod_{i \in [n]} [\hat{c}_i - \lambda d_i, \hat{c}_i + \lambda d_i]$ with $\mathbf{d} \in \mathbb{R}_+^n$, or
- ellipsoidal uncertainty $\mathcal{U}(\lambda) = \{ \mathbf{c} : \mathbf{c} = \hat{\mathbf{c}} + \mathbf{C}\boldsymbol{\xi}, \|\boldsymbol{\xi}\|_2 \leq \lambda \}$ with $\Lambda \subseteq \mathbb{R}_+$, $\mathbf{C} \in \mathbb{R}^{n \times m}$, $\boldsymbol{\xi} \in \mathbb{R}^m$.

We call $\hat{\mathbf{c}}$ the *nominal scenario*, and any $\hat{\mathbf{x}} \in \mathcal{X}$ that is a minimizer of $\text{P}(\hat{\mathbf{c}})$ a *nominal solution*.

In their setting of variable-sized uncertainty, the aim is to find a minimal set of solutions $\mathcal{S} \subseteq \mathcal{X}$ that contains an optimal solution to each robust problem over all λ . Here, the robust problem is either given by the min-max counterpart

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x}$$

or the min-max regret counterpart

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{c} \in \mathcal{U}(\lambda)} \left(\mathbf{c}\mathbf{x} - \min_{\mathbf{y} \in \mathcal{X}} \mathbf{c}\mathbf{y} \right).$$

In the case of min-max robustness, such a set can be found through methods from multi-objective optimization in $\mathcal{O}(|\mathcal{S}| \cdot T)$, where T denotes the complexity of the nominal problem, for many reasonable uncertainty sets. However, \mathcal{S} may be exponentially large.

This setting is related to two other approaches from the optimization literature. The first is fuzzy optimization, which draws on possibility distributions to describe the problem uncertainty. A fuzzy set \tilde{A} consists of a reference set Ω and a membership function $\mu_{\tilde{A}} : \Omega \rightarrow [0, 1]$. The value of the membership function can be interpreted as the degree of membership of an element in \tilde{A} . A λ -cut is then defined as all elements with membership at least λ , i.e., the set $\tilde{A}^\lambda = \{v \in \Omega : \mu_{\tilde{A}}(v) \geq \lambda\}$. The λ -cut with $\lambda = 1$ is called the core of a fuzzy set. For details on possibility theory we refer to [DP88]; its relationship to min-max regret is discussed in [KZ10]. In variable-sized uncertainty, one may consider the uncertainty set \mathcal{U} as being fuzzy. The set $\prod_{i \in [n]} [\hat{c}_i, \hat{c}_i]$ is then the core of the uncertainty, and each set $\mathcal{U}(\lambda)$ corresponds to a λ -cut in possibility theory. More general ways to model possibility distributions exist, which may constitute an interesting way to extend variable-sized uncertainty in the future.

The second related approach is parametric optimization. In this setting, one considers a family of optimization problems that are parameterized through some value λ . The general goal is to compute regions where the optimal solution does not change, meaning that all possible problems with respect to λ are solved simultaneously. Variable-sized uncertainty can be seen as a parametric problem,

where the parameter defines the uncertainty set. We refer to [WOBD13] for a discussion of robustness and parametric (multi-objective) optimization.

The drawback of variable-sized uncertainty is that the solution set \mathcal{S} may be of exponential size, which would require some processing of solutions before they can be presented to the decision maker. The idea of our compromise approach is to present only one solution with a good overall performance instead. Furthermore, this solution is not necessarily in \mathcal{S} , which means that the previous approach might not be able to find it. We introduce our new approach in the following section.

3 Compromise Solutions in the Min-Max Regret Model

In this paper we are interested in finding one single solution that performs well on average over all possible uncertainty sizes $\lambda \in \Lambda$. Recall that in classic min-max regret, one considers the problem

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - \text{opt}(\mathbf{c})$$

with $\text{opt}(\mathbf{c}) = \min_{\mathbf{y} \in \mathcal{X}} \mathbf{c}\mathbf{y}$. We define the compromise approach to variable-sized uncertainty as the following problem:

$$\min \text{val}(\mathbf{x}) \quad \text{with} \quad \text{val}(\mathbf{x}) = \int_{\Lambda} \left(\max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - \text{opt}(\mathbf{c}) \right) d\lambda \quad (\text{CMMR})$$

In the following, we focus our analysis on the classic interval uncertainty sets $\mathcal{U}(\lambda) = \prod_{i \in [n]} [(1 - \lambda)\hat{c}_i, (1 + \lambda)\hat{c}_i]$. To simplify the presentation, we further assume $\Lambda = [0, 1]$. The previous work on variable-sized uncertainty aims at presenting the decision maker with a set of solutions, and assumes that the decision maker will then choose the solution that suits his requirements best. Not presenting a single solution but a set of solution respectively a probability distribution of solutions was also proposed in [MJC15]. There the authors introduce the concept of randomized min-max regret in which the goal is to find a probability distribution over solutions such that the expected maximum regret is minimal. Note that in this concept the uncertainty set size is assumed to be fixed.

In contrast, the concept of compromise min-max regret produces only a single solution that represents a good overall trade-off for all uncertainty set sizes. More precise, the compromise min-max regret solution minimizes the average of the maximum regret over all considered uncertainty set sizes. It can therefore be seen in the tradition of goal programming for multi-objective optimization. As a motivation for our approach, consider the example shown in Figure 1.

We calculated the solution to problem (CMMR) and a minimizer of the min-max regret problem with uncertainty set $\mathcal{U}(0.3)$. Plotted is the difference in the regret objective function between both solutions for different values of λ , normalized such that the zero-line stands for the (CMMR) solution. By construction, the optimal regret solution for $\mathcal{U}(0.3)$ has a good performance for $\lambda = 0.3$, but it has a higher regret on average. Furthermore, it may be that the




Figure 1: Example regret difference for a shortest path instance.

optimal solution to (CMMR) is not an optimal solution to any min-max regret problem with fixed parameter λ ; hence, it would not be "visible" to the decision maker when using the previous approach for variable-sized uncertainty. On the other hand, if a single solution turns out to be optimal for all uncertainty set sizes, it is also optimal for (CMMR).

3.1 Structure of the Objective Function

We first discuss the objective function $val(\mathbf{x})$ for some fixed $\mathbf{x} \in \mathcal{X}$. A well-known result (see, e.g., [ABV09]) states that for fixed \mathbf{x} , the scenario that maximizes the regret is given as

$$c_i(\mathbf{x}, \lambda) = \begin{cases} (1 + \lambda)\hat{c}_i & \text{if } x_i = 1 \\ (1 - \lambda)\hat{c}_i & \text{if } x_i = 0 \end{cases}.$$

Using this result we find that

$$\begin{aligned} reg(\mathbf{x}, \lambda) &:= \max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - opt(\mathbf{c}) \\ &= \mathbf{c}(\mathbf{x}, \lambda)\mathbf{x} - opt(\mathbf{c}(\mathbf{x}, \lambda)) \\ &= \max_{\mathbf{y} \in \mathcal{X}} (1 + \lambda)\mathbf{c}\mathbf{x} - \sum_{i \in [n]} \hat{c}_i (1 - \lambda + 2\lambda x_i) y_i. \end{aligned}$$

Hence, $reg(\mathbf{x}, \lambda)$ is a piecewise linear function in λ , where every possible regret solution \mathbf{y} defines an affine linear regret function $\mathbf{c}(\mathbf{x}, \lambda)(\mathbf{x} - \mathbf{y})$. Note that $reg(\mathbf{x}, \lambda)$ is convex in λ .

In the following, we use a theoretical example to illustrate how a general formula for $val(\mathbf{x})$ can be derived. Figure 2 illustrates the objective function with two brake points at λ_1 and λ_2 . In red is the maximum over all regret functions, which defines $val(\mathbf{x})$. On the interval $[0, \lambda_1]$, the regret of some solution \mathbf{y}^2 , while solution \mathbf{y}^3 defines the regret on $[\lambda_1, \lambda_2]$, and \mathbf{y}^4 defines the regret on $[\lambda_2, 1]$. In this case, we can hence compute

$$val(\mathbf{x}) = \int_0^1 reg(\mathbf{x}, \lambda) d\lambda = \int_0^1 \left(\max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - \min_{\mathbf{y} \in \mathcal{X}} \mathbf{c}\mathbf{y} \right) d\lambda$$




Figure 2: Illustration of the structure of $reg(\mathbf{x}, \lambda)$.

$$\begin{aligned}
&= \int_0^{\lambda_1} \max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - \mathbf{c}\mathbf{y}^2 d\lambda + \int_{\lambda_1}^{\lambda_2} \max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - \mathbf{c}\mathbf{y}^3 d\lambda + \int_{\lambda_2}^1 \max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - \mathbf{c}\mathbf{y}^4 d\lambda \\
&= \int_0^{\lambda_1} \mathbf{c}(\mathbf{x}, \lambda)(\mathbf{x} - \mathbf{y}^2) d\lambda + \int_{\lambda_1}^{\lambda_2} \mathbf{c}(\mathbf{x}, \lambda)(\mathbf{x} - \mathbf{y}^3) d\lambda + \int_{\lambda_2}^1 \mathbf{c}(\mathbf{x}, \lambda)(\mathbf{x} - \mathbf{y}^4) d\lambda
\end{aligned}$$

Using that the functions $\mathbf{c}(\mathbf{x}, \lambda)(\mathbf{x} - \mathbf{y}^i)$ are linear in λ , we can now rewrite the integrals using the midpoint in each interval to find

$$\begin{aligned}
val(\mathbf{x}) &= \lambda_1 \mathbf{c}\left(\mathbf{x}, \frac{\lambda_1}{2}\right)(\mathbf{x} - \mathbf{y}^2) + (\lambda_2 - \lambda_1) \mathbf{c}\left(\mathbf{x}, \frac{\lambda_1 + \lambda_2}{2}\right)(\mathbf{x} - \mathbf{y}^3) \\
&\quad + (1 - \lambda_2) \mathbf{c}\left(\mathbf{x}, \frac{\lambda_2 + 1}{2}\right)(\mathbf{x} - \mathbf{y}^4)
\end{aligned}$$

In general, to compute $val(\mathbf{x})$, we need to determine all relevant regret solutions \mathbf{y} , and the intersections of the resulting regret functions. Because \mathcal{X} is a finite set, every calculation of $val(\mathbf{x})$ can be written in the above form.

3.2 Problem Formulation

Let $\bar{\Lambda}(\mathbf{x}) \subseteq \Lambda$ be the set of change points of the piecewise linear function $reg(\mathbf{x}, \cdot)$. To formulate problem (CMMR) as a linear integer program, we use a set $\bar{\Lambda} \supseteq \cup_{\mathbf{x} \in \mathcal{X}} \bar{\Lambda}(\mathbf{x})$. As \mathcal{X} is a finite set, there always exists a set $\bar{\Lambda}$ that is finite. In general, it may contain exponentially many elements.

For the ease of notation, we assume $\bar{\Lambda} = \{\lambda_1, \dots, \lambda_K\}$ with $\lambda_i \leq \lambda_{i+1}$, $\lambda_1 = 0$ and $\lambda_{K+1} := 1$. As a first approach, we model $val(\mathbf{x})$ by using all possible regret solutions $\mathbf{y} \in \mathcal{X}$.

$$\begin{aligned}
\min \quad & \sum_{j \in [K]} (\lambda_{j+1} - \lambda_j) z_j \tag{1} \\
\text{s.t.} \quad & z_j \geq \sum_{i \in [n]} (1 + \bar{\lambda}_j) \hat{c}_i x_i - \sum_{i \in [n]} (1 - \bar{\lambda}_j + 2\bar{\lambda}_j x_i) \hat{c}_i y_i^\ell \quad \forall j \in [K], \mathbf{y}^\ell \in \mathcal{X} \\
& \mathbf{x} \in \mathcal{X}
\end{aligned}$$

where $\bar{\lambda}_j = \frac{1}{2}(\lambda_j + \lambda_{j+1})$ for all $j \in [K]$. Problem (1) calculates $val(\mathbf{x})$ by using the rectangle rule for integration on each of its linear parts, i.e.,

$$val(\mathbf{x}) = \int_{\Lambda} reg(x, \lambda) d\lambda = \sum_{j \in [K]} \int_{\lambda_j}^{\lambda_{j+1}} reg(x, \lambda) d\lambda = \sum_{j \in [K]} (\lambda_{j+1} - \lambda_j) reg(x, \bar{\lambda}_j)$$

as $\bar{\Lambda} \supseteq \cup_{\mathbf{x} \in \mathcal{X}} \bar{\Lambda}(\mathbf{x})$.

Problem (1) can also be used with any arbitrary set $\bar{\Lambda}$. In that case, the optimal objective value is still a lower bound on $val(\mathbf{x})$ (as the rectangle rule underestimates convex functions). This property will be used in our solution algorithm.

If the problems $P(\mathbf{c})$ have a duality gap of zero, i.e., if solving the dual of the linear relaxation also gives an optimal objective value to $P(\mathbf{c})$, this formulation can be simplified. Examples where this is the case include the shortest path or the minimum spanning tree problem. Let us assume that the linear relaxation of \mathcal{X} is given by

$$\bar{\mathcal{X}} = \{\mathbf{x} \in \mathbb{R}_+^n : \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$$

For the regret problem $\min_{\mathbf{x} \in \mathcal{X}} \{\mathbf{c}(\mathbf{x}, \lambda)\mathbf{x} - opt(\mathbf{c}(\mathbf{x}, \lambda))\}$ we may then write the following equivalent problem using the dual (see [ABV09]):

$$\min\{\mathbf{c}(\mathbf{x}, \lambda)\mathbf{x} - \mathbf{b}^t \mathbf{u} : \mathbf{x} \in \bar{\mathcal{X}}, \mathbf{u} \in \bar{\mathcal{Y}}\} \text{ with } \bar{\mathcal{Y}} = \{\mathbf{u} \geq 0 : \mathbf{A}^t \mathbf{u} \leq \mathbf{c}(\mathbf{x}, \lambda)\}$$

Using this reformulation, we rewrite problem (1). Variables z_j represent the regret of solution \mathbf{x} with respect to the uncertainty set $\mathcal{U}(\lambda_j)$. By substituting these variables with the dual problems to calculate the regret, we find the following program for (CMMR)

$$\begin{aligned} \min \quad & \sum_{j \in [K]} (\lambda_{j+1} - \lambda_j) (\mathbf{c}(\mathbf{x}, \bar{\lambda}_j)\mathbf{x} - \mathbf{b}^t \mathbf{u}^j) & (2) \\ \text{s.t.} \quad & \mathbf{A}^t \mathbf{u}^j \leq \mathbf{c}(\mathbf{x}, \bar{\lambda}_j) & \forall j \in [K] \\ & \mathbf{x} \in \bar{\mathcal{X}} \\ & \mathbf{u}^j \in \bar{\mathcal{Y}} & \forall j \in [K] \end{aligned}$$

For binary variables \mathbf{x} , the product $\mathbf{c}(\mathbf{x}, \bar{\lambda}_j)\mathbf{x}$ can then be linearized. If a set $\bar{\Lambda}$ can be found that is of polynomial size, this is a compact formulation. We analyze set $\bar{\Lambda}$ in more detail for the selection, minimum spanning tree, and shortest path problem in Section 4. In general, constraints and variables can be added in an iterative algorithm that generates new candidate values for λ in Problem (2). This approach is explained in the following section. If the zero duality gap assumption does not hold, we can use Formulation (1), where both values for λ and regret solutions \mathbf{y} need to be generated.

3.3 General Algorithm

In Algorithm 1 we describe how to compute the set $\bar{\Lambda}(\mathbf{x})$ of change points and the corresponding set $\mathcal{Y}(\mathbf{x})$ of regret solutions for $reg(\mathbf{x}, \cdot)$.

We begin with only two regret functions, for the extreme points $\bar{\Lambda}(\mathbf{x}) = \{0, 1\}$. The resulting two regret functions will intersect at one new candidate change point λ . We find the regret solution \mathbf{y} maximizing the regret at this point

Algorithm 1 Algorithm to compute $\bar{\Lambda}(\mathbf{x})$ and $\mathcal{Y}(\mathbf{x})$

Input: An instance of (CMMR), a fixed solution $\mathbf{x} \in \mathcal{X}$.

- 1: $\bar{\Lambda}(\mathbf{x}) \leftarrow \{0, 1\}$, $\bar{\Lambda}^{\text{new}}(\mathbf{x}) \leftarrow \bar{\Lambda}(\mathbf{x})$
 - 2: $\mathcal{Y}(\mathbf{x}) \leftarrow \emptyset$
 - 3: **for all** $\lambda \in \bar{\Lambda}^{\text{new}}(\mathbf{x})$ **do**
 - 4: Solve $P(\mathbf{c}(\mathbf{x}, \lambda))$. Let \mathbf{y} be the resulting solution.
 - 5: $\mathcal{Y}(\mathbf{x}) \leftarrow \mathcal{Y}(\mathbf{x}) \cup \{\mathbf{y}\}$
 - 6: $\bar{\Lambda}^{\text{new}}(\mathbf{x}) \leftarrow \bar{\Lambda}^{\text{new}}(\mathbf{x}) \setminus \{\lambda\}$
 - 7: **end for**
 - 8: change \leftarrow false
 - 9: **for all** $\mathbf{y}^i, \mathbf{y}^j \in \mathcal{Y}(\mathbf{x})$ **do**
 - 10: Calculate λ as the point where the affine linear regret functions defined by \mathbf{y}^i and \mathbf{y}^j intersect.
 - 11: **if** $\lambda \notin \bar{\Lambda}(\mathbf{x})$ **then**
 - 12: $\bar{\Lambda}(\mathbf{x}) \leftarrow \bar{\Lambda}(\mathbf{x}) \cup \{\lambda\}$
 - 13: $\bar{\Lambda}^{\text{new}}(\mathbf{x}) \leftarrow \bar{\Lambda}^{\text{new}}(\mathbf{x}) \cup \{\lambda\}$
 - 14: change \leftarrow true
 - 15: **end if**
 - 16: **end for**
 - 17: **if** change = true **then**
 - 18: Go to 3
 - 19: **end if**
 - 20: Reduce $\bar{\Lambda}(\mathbf{x})$ and $\mathcal{Y}(\mathbf{x})$ such that only change points and regret functions remain that define the maximum over all affine linear functions.
 - 21: **return** $\bar{\Lambda}(\mathbf{x})$, $\mathcal{Y}(\mathbf{x})$
-

by solving the problem $P(\mathbf{c}(\mathbf{x}, \lambda))$. We then repeat this process by iteratively calculating the regret at all current intersection points. Note that if there exists any regret function that is larger than all current regret functions at some point λ , then it is also larger than all current functions at the intersection point between two of them. Hence, Algorithm 1 finds all relevant change points λ . As it may also produce unnecessary candidates λ , we reduce the solution sets at the end to contain only those change points and regret functions that define the maximum.

Algorithm 1 is used to solve Formulation (2) in the case of a zero duality gap for $P(\mathbf{c})$ as described in Algorithm 2.

The algorithm begins with a starting set $\bar{\Lambda} = \{0, 1\}$ as a guess for relevant change points. Using the current set $\bar{\Lambda}$, it then solves Formulation (2). Recall that the objective value that is found is only a lower bound LB^k on the true optimal objective value of the problem. To evaluate the resulting candidate solution \mathbf{x}^k , we compute $val(\mathbf{x}^k)$ in Step 3. As \mathbf{x}^k is a feasible solution to (CMMR), $val(\mathbf{x}^k)$ gives an upper bound on the optimal objective value. Hence, if lower and upper bound coincide, an optimal solution has been found. Otherwise, we extend the set $\bar{\Lambda}$ in Step 7 and repeat the procedure.

If $P(\mathbf{c})$ has a duality gap larger than zero, the same algorithm can be used with the adjustment that Problem (1) is solved in Step 2. To this end, also the regret solutions $\mathcal{Y}(\mathbf{x}^k)$ generated in the computation of $val(\mathbf{x}^k)$ in Step 3 have to be collected and used.

Algorithm 2 Exact algorithm for (CMMR)

Input: An instance of (CMMR).

- 1: $\bar{\Lambda} \leftarrow \{0, 1\}$, $k \leftarrow 0$
 - 2: Solve Formulation (2) using $\bar{\Lambda}$. Let the solution be \mathbf{x}^k , and the objective value LB^k .
 - 3: Use Algorithm 1 to find $\bar{\Lambda}(\mathbf{x}^k)$. Compute $val(\mathbf{x}^k)$.
 - 4: **if** $val(\mathbf{x}^k) = LB^k$ **then**
 - 5: **END:** \mathbf{x}^k is an optimal solution.
 - 6: **else**
 - 7: $\bar{\Lambda} \leftarrow \bar{\Lambda} \cup \bar{\Lambda}(\mathbf{x}^k)$
 - 8: $k \leftarrow k + 1$
 - 9: Go to 2
 - 10: **end if**
-

3.4 Extensions

We briefly mention possible extensions of the (CMMR) approach as introduced at the beginning of this section.

One possible direction to extend the concept is to consider different uncertainty sets. The more general interval uncertainty sets $\mathcal{U}(\lambda) = \prod_{i \in [n]} [\hat{c}_i - \lambda d_i, \hat{c}_i + \lambda d_i]$ are a direct generalization of the classic interval uncertainty sets. For these sets the scenario maximizing the regret for a fixed solution \mathbf{x} is given by

$$c_i(\mathbf{x}, \lambda) = \begin{cases} \hat{c}_i + \lambda d_i & \text{if } x_i = 1 \\ \hat{c}_i - \lambda d_i & \text{if } x_i = 0 \end{cases}.$$

We assume again that $\lambda \in [0, 1]$ and that $0 \leq d_i \leq \hat{c}_i$. This ensures that $c_i(\mathbf{x}, \lambda) \geq 0$. In this case, the previous arguments of this section are still valid and the same algorithms can be used to solve the resulting problems. Note that the problem formulations need to be adapted accordingly.

Using ellipsoidal uncertainty sets which have been introduced in min-max regret only recently (see [CG17]) turns out to pose more difficult problems. In this setting, $c_i(\mathbf{x}, \lambda)$ does not obey the simple structure as in the case of interval sets. Hence, the presented methods are not applicable for this case. Note that using ellipsoidal uncertainty sets for min-max regret problems is problematic since already evaluating the regret of a given solution turns out to be an NP-complete problem even for basic problems like the minimum spanning tree or shortest path problem (see [CG17]).

A further extension is to use a weight function $w : \Lambda \rightarrow [0, 1]$ that reflects the decision maker's preferences, e.g., if smaller or larger degrees of uncertainty should be focussed on. The resulting variant of (CMMR) is then given as

$$\min val(\mathbf{x}) \quad \text{with} \quad val(\mathbf{x}) = \int_{\Lambda} w(\lambda) \left(\max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - opt(\mathbf{c}) \right) d\lambda$$

Note that the solution algorithms from this section can be extended to this model if w is given as a piecewise constant function. To this end, one needs to ensure that $\bar{\Lambda} = \{\lambda_1, \dots, \lambda_K\}$ contains all change points of w .

Furthermore, the previous discussion focussed on continuous sets Λ . For a discrete set Λ we may hence be interested in the problem

$$\min_{\lambda \in \Lambda} \sum w(\lambda) \left(\max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} - \text{opt}(\mathbf{c}) \right)$$

instead. Our concept is then a direct extension of the classic min-max regret approach with Λ being a singleton. The algorithms presented here can still be applied in a simplified way, as all relevant value for λ are already given.

3.5 Approximation

We now consider the approximability of problem (CMMR). For classic min-max regret problems with interval uncertainty, the following result is known:

Theorem 1. [KZ06] *Let $\mathcal{U} = \prod_{i \in [n]} [a_i, b_i]$ be an interval-based uncertainty set, and let $\hat{\mathbf{x}}$ be an optimal solution to the scenario $\hat{\mathbf{c}} := \frac{1}{2}(\mathbf{a} + \mathbf{b})$. Then*

$$\max_{\mathbf{c} \in \mathcal{U}} (\mathbf{c}\hat{\mathbf{x}} - \text{opt}(\mathbf{c})) \leq 2 \min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{c} \in \mathcal{U}} (\mathbf{c}\mathbf{x} - \text{opt}(\mathbf{c})),$$

i.e., the midpoint solution is a 2-approximation for the min-max regret problem.

We can hence also approximate the compromise problem.

Corollary 1. *Any optimal solution $\hat{\mathbf{x}}$ with respect to $P(\hat{\mathbf{c}})$ is also a 2-approximation for problem (CMMR).*

Proof. Using Theorem 1, we find:

$$\text{val}(\hat{\mathbf{x}}) = \int_{\Lambda} \text{reg}(\hat{\mathbf{x}}, \lambda) d\lambda \leq 2 \int_{\Lambda} \min_{\mathbf{x} \in \mathcal{X}} \text{reg}(\mathbf{x}, \lambda) d\lambda \leq 2 \min_{\mathbf{x} \in \mathcal{X}} \int_{\Lambda} \text{reg}(\mathbf{x}, \lambda) d\lambda$$

□

Note that the example instance used in [KZ06] uses uncertainty intervals of the form $[1, 1]$ and $[0, 2]$, which is more general than our setting $\mathcal{U}(\lambda) = \prod_{i \in [n]} [(1 - \lambda)\hat{c}_i, (1 + \lambda)\hat{c}_i]$ allows. Therefore one might hope that a better approximation ratio can be achieved in our setting. However, Figure 3 shows a shortest path instance with costs $\hat{\mathbf{c}}$, where the midpoint solution also performs as a 2-approximation for any $\lambda \in (0, 1]$. One can assume that the midpoint solution chooses the path along the top two edges with regret 4λ . The path using the middle edge has a regret of only 2λ , which shows that the approximation guarantee is still tight.

4 Min-Max Regret Compromise Solutions for Specific Problems

4.1 Minimum Selection

The minimum selection problem is given by

$$\min \left\{ \mathbf{c}\mathbf{x} : \sum_{i \in [n]} x_i = p, \mathbf{x} \in \{0, 1\}^n \right\}$$




Figure 3: Instance where approximation guarantee of midpoint solution is tight.

and has been frequently studied in the literature on min-max regret optimization (see [KZ16]). While the classic selection problem can be solved in $\mathcal{O}(n)$, the min-max regret version can be solved in $\mathcal{O}(n \cdot \min\{n, n-p\})$ time, see [Con04]. We show that also problem (CMMR) can be solved in polynomial time.

Theorem 2. *Let $\mathcal{U}(\lambda) = \prod_{i \in [n]} [(1-\lambda)\hat{c}_i, (1+\lambda)\hat{c}_i]$ for a fixed $\lambda \in [0, 1]$. Then the nominal solution $\tilde{\mathbf{x}}$ is an optimal solution to the min-max regret selection problem.*

Proof. We assume that items are sorted with respect to $\hat{\mathbf{c}}$. Let $\tilde{\mathbf{x}}$ be an optimal solution with $\tilde{x}_i = 0$ for an item $i \leq p$. We assume i is the smallest such item. Then there exists some $j > p$ with $\tilde{x}_j = 1$. Consider the solution \mathbf{x}' with $x'_k = \tilde{x}_k$ for $k \neq i, j$ and $x'_i = 1, x'_j = 0$.

Let $\tilde{\mathbf{y}}$ be a regret solution for $\tilde{\mathbf{x}}$. We can assume that $\tilde{y}_i = 1$, as $(1-\lambda)\hat{c}_i$ must be one of the p cheapest items. We can also assume $\tilde{y}_j = 0$, as $(1+\lambda)\hat{c}_j$ is not among the p cheapest items. Let \mathbf{y}' be the regret solution for \mathbf{x}' .

The solutions $\tilde{\mathbf{x}}$ and \mathbf{x}' differ only on the two items i and j . Hence, the following cases are possible:

- Case $y'_i = 1$ and $y'_j = 0$, i.e., $\tilde{\mathbf{y}} = \mathbf{y}'$. We have

$$\begin{aligned} \text{reg}(\tilde{\mathbf{x}}, \lambda) - \text{reg}(\mathbf{x}', \lambda) &= (1+\lambda)\hat{c}_j - (1+\lambda)\hat{c}_i - (1-\lambda)\hat{c}_i + (1+\lambda)\hat{c}_i \\ &= (1+\lambda)\hat{c}_j - (1-\lambda)\hat{c}_i \geq 0 \end{aligned}$$

- Case $y'_i = 1$ and $y'_j = 1, y'_k = 0$ for some $k > i$ with $\tilde{y}_k = 1$. Note that this means $(1+\lambda)\hat{c}_j \geq (1-\lambda + 2\tilde{x}_k\lambda)\hat{c}_k$, as otherwise, the regret solution $\tilde{\mathbf{y}}$ could be improved. Hence,

$$\begin{aligned} \text{reg}(\tilde{\mathbf{x}}, \lambda) - \text{reg}(\mathbf{x}', \lambda) &= (1+\lambda)\hat{c}_j - (1+\lambda)\hat{c}_i - (1-\lambda)\hat{c}_i + (1+\lambda)\hat{c}_i \\ &\quad - (1-\lambda + 2\lambda\tilde{x}_k)\hat{c}_k + (1-\lambda)\hat{c}_j \\ &= (1-\lambda)(\hat{c}_j - \hat{c}_i) + (1+\lambda)\hat{c}_j - (1-\lambda + 2\lambda\tilde{x}_k)\hat{c}_k \geq 0 \end{aligned}$$

- Case $y_i = 0$ and $y'_\ell = 1$ for some $\ell > i$ with $\tilde{y}_\ell = 0$. As the costs of item i have increased by using solution \mathbf{x}' instead of $\tilde{\mathbf{x}}$, the resulting two cases are analogue to the two cases above.

Overall, solution \mathbf{x}' has regret less than or equal to the regret of $\tilde{\mathbf{x}}$. \square

Note that this result does not hold for general interval uncertainty sets, where the problem is NP-hard. It also does not necessarily hold for other combinatorial optimization problems; e.g., a counter-example for the assignment problem can be found in [CG18].

Finally, it remains to show that $val(\mathbf{x})$ can also be computed in polynomial time.

Theorem 3. *For the compromise min-max regret problem of minimum selection it holds that $|\bar{\Lambda}(\mathbf{x})| \in \mathcal{O}(\min\{p, n - p\})$ for any fixed $\mathbf{x} \in \mathcal{X}$.*

Proof. If \mathbf{x} is fixed, then there are p items i with costs $(1 + \lambda)\hat{c}_i$, and $(n - p)$ items i with costs $(1 - \lambda)\hat{c}_i$. The regret solution is determined by the p smallest items. Accordingly, when λ increases, the regret solution only changes if an item i with $x_i = 1$, that used to be among the p smallest items, moves to the $(n - p)$ largest items, and another item j with $x_j = 0$ becomes part of the p smallest items. There are at most $\min\{p, n - p\}$ values for λ where this is the case. □

Note that $\bar{\Lambda} = \cup_{\mathbf{x} \in \mathcal{X}} \bar{\Lambda}(\mathbf{x})$ has $\mathcal{O}(n^2)$ many elements. As the size of $\bar{\Lambda}(\mathbf{x})$ is polynomially bounded, $val(\mathbf{x})$ can be computed in polynomial time, and we get the following conclusion.

Corollary 2. *The compromise min-max regret problem of minimum selection can be solved in $\mathcal{O}(n)$, and evaluated in $\mathcal{O}(n \min\{p, n - p\})$.*

Proof. As the nominal solution can be computed in $\mathcal{O}(n)$, the first claim follows. For the second claim, we need to evaluate the regret of $\hat{\mathbf{x}}$ at $\mathcal{O}(\min\{p, n - p\})$ many points. Each evaluation can be done in $\mathcal{O}(n)$. □

4.2 Minimum Spanning Tree

The min-max regret spanning tree problem in a graph $G = (V, E)$ has previously been considered, e.g., in [YKP01, KZ11]. The regret of a fixed solution can be computed in polynomial time, but it is NP-hard to find an optimal solution. We now consider the compromise min-max regret counterpart (CMMR).

Let any spanning tree \mathbf{x} be fixed. To compute $val(\mathbf{x})$, we begin with $\lambda = 0$ and calculate a regret spanning tree by solving a nominal problem with costs $\hat{\mathbf{c}}$. Recall that this can be done using Kruskal's algorithm that considers edges successively according to an increasing sorting

$$\hat{c}_{e_1} \leq \dots \leq \hat{c}_{e_{|E|}}$$

with respect to costs. If λ increases, edges that are included in \mathbf{x} have costs $(1 + \lambda)\hat{c}_e$ (i.e., their costs increase) and edges not in \mathbf{x} have costs $(1 - \lambda)\hat{c}_e$ (i.e., their costs decrease). Kruskal's algorithm will only find a different solution if the sorting of edges change. As there are $|V| - 1$ edges with increasing costs, and $|E| - |V| + 1$ edges with decreasing costs, the sorting can change at most $(|V| - 1)(|E| - |V| + 1) = \mathcal{O}(|E|^2)$ times (note that two edges with increasing costs or two edges with decreasing costs cannot change relative positions). We have therefore shown:

Theorem 4. *A solution to the compromise min-max regret problem of minimum spanning tree can be evaluated in polynomial time.*

If the solution \mathbf{x} is not known, we can still construct a set $\bar{\Lambda}$ with size $\mathcal{O}(|E|^2)$ that contains all possible change points along the same principle. We can conclude:

Theorem 5. *There exists a compact mixed-integer programming formulation for the compromise min-max regret problem of minimum spanning tree.*

Note that Theorem 5 does also hold in the more general setting of general interval uncertainty sets. However, we show in the following that solving the compromise problem is NP-hard. To this end, we use the following result:

Theorem 6. *[AL04] The min-max regret spanning tree problem is NP-hard, even if all intervals of uncertainty are equal to $[0, 1]$.*

Note that if all intervals are of the form $[a, b]$ (i.e., $\mathcal{U} := \prod_{i \in [n]} [a, b]$), then

$$\begin{aligned} \max_{\mathbf{c} \in \mathcal{U}} (\mathbf{c}\mathbf{x} - \text{opt}(\mathbf{c})) &= \sum_{e \in E} bx_e - \min_{\mathbf{y} \in \mathcal{X}} \left(\sum_{\substack{e \in E \\ x_e=1}} by_e + \sum_{\substack{e \in E \\ x_e=0}} ay_e \right) \\ &= (|V| - 1)b - \min_{\mathbf{y} \in \mathcal{X}} \left((|V| - 1)b - \sum_{\substack{e \in E \\ x_e=0}} (b - a)y_e \right) \\ &= (b - a) \max_{\substack{\mathbf{y} \in \mathcal{X} \\ x_e=0}} \sum_{e \in E} y_e \end{aligned}$$

Therefore, the min-max regret problem with costs $[0, 1]$ is equivalent to the min-max regret problem with any other costs $[a, b]$, in the sense that objective values only differ by a constant factor and both problems have the same set of optimal solutions. In particular, a solution \mathbf{y} that maximizes the regret of \mathbf{x} with respect to cost intervals $[a, b]$ is also a maximizer of the regret for any other cost intervals $[a', b']$. We can conclude:

Theorem 7. *The compromise problem of min-max regret minimum spanning tree is NP-hard.*

Proof. Let an instance of the min-max regret spanning tree problem with cost intervals $[0, 1]$ be given. Consider an instance of the compromise problem with $\hat{c}_e = 1$ for all $e \in E$. Then

$$\text{val}(\mathbf{x}) = \int_0^1 \text{reg}(\mathbf{x}, \lambda) d\lambda = \int_0^1 \left(2\lambda \max_{\substack{\mathbf{y} \in \mathcal{X} \\ x_e=0}} \sum_{e \in E} y_e \right) d\lambda = \max_{\substack{\mathbf{y} \in \mathcal{X} \\ x_e=0}} \sum_{e \in E} y_e$$

Hence, any minimizer of $\text{val}(\mathbf{x})$ is also an optimal solution to the min-max regret spanning tree problem. \square

4.3 Shortest Path

For the shortest path problem, we consider \mathcal{X} as the set of all simple $s-t$ paths in a graph $G = (V, E)$ (for the min-max regret problem, see, e.g., [AL04]). As for the minimum spanning tree problem, the regret of a fixed solution can be computed in polynomial time, but it is NP-hard to find an optimal solution. For the compromise problem (CMMR), we have:

$$\begin{aligned} \text{reg}(\mathbf{x}, \lambda) &= \sum_{\substack{e \in E \\ x_e=1}} (1 + \lambda) \hat{c}_e - \left(\min_{\mathbf{y} \in \mathcal{X}} \sum_{\substack{e \in E \\ x_e=1}} (1 + \lambda) \hat{c}_e y_e + \sum_{\substack{e \in E \\ x_e=0}} (1 - \lambda) \hat{c}_e y_e \right) \\ &= \sum_{\substack{e \in E \\ x_e=1}} (1 + \lambda) \hat{c}_e + \left(\min_{\mathbf{y} \in \mathcal{X}} \lambda \sum_{\substack{e \in E \\ x_e=1}} 2\hat{c}_e y_e + (1 - \lambda) \sum_{e \in E} \hat{c}_e y_e \right) \end{aligned}$$

We can interpret the minimization problem as a weighted sum of the bicriteria problem

$$\min \left\{ \left(\sum_{\substack{e \in E \\ x_e=1}} 2\hat{c}_e y_e \right) : \mathbf{y} \in \mathcal{X} \right\}$$

The number of solutions we need to generate to compute $\text{val}(\mathbf{x})$ can therefore be bounded by the number of solutions we can find through such weighted sum computations (the set of efficient extreme solutions \mathcal{E} , see, e.g., [Ehr06]).

Lemma 1. *For the compromise shortest path problem, it holds that $|\bar{\Lambda}(\mathbf{x})| \leq |\mathcal{E}|$.*

Depending on the graph G , the following bounds on the number of efficient extreme solutions \mathcal{E} can be taken from the literature [Car83, CG18]:

- for series-parallel graphs, $\mathcal{E} \in \mathcal{O}(|E|)$
- for layered graphs with width w and length ℓ , $\mathcal{E} \in \mathcal{O}(2^{\log w \log(\ell+1)})$
- for acyclic graphs, $\mathcal{E} \in \mathcal{O}(|V|^{\log |V|})$
- for general graphs $\mathcal{E} \in 2^{\Omega(\log^2 |V|)}$

We conclude:

Corollary 3. *A solution to the compromise min-max regret problem of shortest path can be evaluated in polynomial time on series-parallel graphs and layered graphs with fixed width or length.*

Note that Lemma 1 and Corollary 3 hold also in the case of general interval uncertainty sets. Note that the number of extreme efficient solutions is only an upper bound on $|\bar{\Lambda}(\mathbf{x})|$. Unfortunately, we cannot hope to find a better performance than this bound on general graphs, as the following result demonstrates.

Theorem 8. *For any bicriteria shortest path instance with costs (\mathbf{a}, \mathbf{b}) , $a_e > 0$ for all $e \in E$, there is an instance of (CMMR) and a solution \mathbf{x} where $|\bar{\Lambda}(\mathbf{x})| = |\mathcal{E}|$.*

Proof. Let an instance of the bicriteria shortest path problem be given, i.e., a directed graph $G = (V, E)$ with arc costs $c_e = (a_e, b_e)$ for all $e \in E$. As $a_e > 0$ for all $e \in E$, we rescale costs such that $2a_e \geq b_e$ for all $e \in E$ without changing the set of efficient extreme solutions. We create the following instance of (CMMR).

Every arc $e = (i, j) \in E$ is substituted by three arcs $e' = (i, i'(e))$, $e'' = (i'(e), j'(e))$ and $e''' = (j'(e), j)$. We set $\hat{c}_{e'} = a_e - \frac{b_e}{2}$, $\hat{c}_{e''} = \frac{b_e}{2}$ and $\hat{c}_{e'''} = 0$ (see Figure 4 for an example of such a transformation). Let E' , E'' and E''' contain all edges of the respective type. Additionally, we choose an arbitrary order of edges (e_1, \dots, e_m) , and create arcs $E_M = \{(s, i'(e_1)), (j'(e_1), i'(e_2)), \dots, (j'(e_m), t)\}$. We set costs of these arcs to be a sufficiently large constant M . Finally, let \mathbf{x} be the path that follows all edges in E_M and E'' . Note that edges in E''' can be contracted, but are shown for better readability.




Figure 4: Example for the transformation used in the proof of Theorem 8.

Note that M is sufficiently large so that no regret path \mathbf{y} will use any edge in E_M . Hence, if \mathbf{y} uses an edge in E' , it will also have to use the following edges in E'' and E''' , i.e., \mathbf{y} corresponds to a path in the original graph $G = (V, E)$.

The regret of \mathbf{x} is

$$\begin{aligned}
\text{reg}(\mathbf{x}, \lambda) &= \sum_{e \in E_M} (1 + \lambda)M + \sum_{e \in E''} (1 + \lambda)\hat{c}_e - \min_{\mathbf{y} \in \mathcal{X}} \sum_{\substack{e \in E \\ x_e = 1}} (1 + \lambda)\hat{c}_e y_e + \sum_{\substack{e \in E \\ x_e = 0}} (1 - \lambda)\hat{c}_e y_e \\
&= (1 + \lambda) \cdot \text{const.} - \min_{\mathbf{y} \in \mathcal{X}} \sum_{e \in E} \left((1 + \lambda) \frac{b_e}{2} + (1 - \lambda)(a_e - \frac{b_e}{2}) \right) y_e \\
&= (1 + \lambda) \cdot \text{const.} - \min_{\mathbf{y} \in \mathcal{X}} \sum_{e \in E} (a_e + \lambda(b_e - a_e)) y_e
\end{aligned}$$

Therefore, if λ goes from 0 to 1, all extreme efficient paths in the original graph G are used to calculate $\text{reg}(\mathbf{x}, \lambda)$. \square

We now consider the complexity of finding a solution $\mathbf{x} \in \mathcal{X}$ that minimizes $\text{val}(\mathbf{x})$. Note that the reduction in [AL04] uses interval costs of the form $[0, 1]$ and $[1, 1]$, which does not fit into our cost framework $[(1 - \lambda)\hat{c}_e, (1 + \lambda)\hat{c}_e]$. Instead, we make use of the following result:

Theorem 9. [CG18] *The min-max regret shortest path problem is NP-hard for layered graphs with interval costs $[0, 1]$.*

Note that for layered graphs, all paths have the same cardinality. Hence, $\text{reg}(\mathbf{x}) = (b - a) \max_{\mathbf{y} \in \mathcal{X}} \sum_{\substack{e \in E \\ x_e = 0}} y_e$ (see Section 4.2), and the problem with costs $[0, 1]$ is equivalent to the problem with costs $[a, b]$ for any $a < b$. Analog to the last section, we can therefore conclude:

Theorem 10. *Finding an optimal solution to the compromise shortest path problem is NP-hard on layered graphs.*

5 Compromise Solutions in the Min-Max Model

We now consider the compromise approach for min-max problems. Using the extension presented in Section 3.4, we apply a weight function $w : \Lambda \rightarrow [0, 1]$ and consider the compromise problem

$$\min_{\mathbf{x} \in \mathcal{X}} \text{val}(\mathbf{x}) \quad \text{with} \quad \text{val}(\mathbf{x}) = \int_{\Lambda} w(\lambda) \left(\max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} \right) d\lambda \quad (\text{CMM})$$

The solution of this problem minimizes the weighted average of the worst cast performance over all considered uncertainty set sizes. In the following, we assume that $\text{val}(x)$ is well-defined (e.g., Λ has bounded support) and consider (CMM) for different shapes of $\mathcal{U}(\lambda)$.

Theorem 11. *Let $\mathcal{U}(\lambda) = \prod_{i \in [n]} [(1 - \lambda)\hat{c}_i, (1 + \lambda)\hat{c}_i]$ be an interval-based uncertainty set with $\lambda \in \Lambda = [0, 1]$. Then, a nominal solution $\hat{\mathbf{x}}$ is an optimal solution of (CMM).*

Proof. As $\max_{\mathbf{c} \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} = (1 + \lambda)\hat{\mathbf{c}}\mathbf{x}$, we get

$$\text{val}(\mathbf{x}) = \int_0^1 w(\lambda) ((1 + \lambda)\hat{\mathbf{c}}\mathbf{x}) d\lambda = \left(\int_0^1 (1 + \lambda)w(\lambda) d\lambda \right) \hat{\mathbf{c}}\mathbf{x}$$

Therefore, a minimizer of the nominal problem with costs $\hat{\mathbf{c}}$ is also a minimizer of (CMM). \square

Theorem 12. Let $\mathcal{U}(\lambda) = \prod_{i \in [n]} [\hat{c}_i - \lambda d_i, \hat{c}_i + \lambda d_i]$ be an general interval uncertainty set with $\lambda \in \Lambda \subseteq \mathbb{R}_+$. Then, an optimal solution to (CMM) can be found by solving a single robust problem with general interval uncertainty.

Proof. Using that $\max_{c \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} = \hat{\mathbf{c}}\mathbf{x} + \lambda \mathbf{d}\mathbf{x}$, we find that

$$\begin{aligned} \text{val}(\mathbf{x}) &= \int_{\Lambda} w(\lambda) (\hat{\mathbf{c}}\mathbf{x} + \lambda \mathbf{d}\mathbf{x}) d\lambda \\ &= \left(\int_{\Lambda} w(\lambda) d\lambda \right) \hat{\mathbf{c}}\mathbf{x} + \left(\int_{\Lambda} \lambda w(\lambda) d\lambda \right) \mathbf{d}\mathbf{x} \end{aligned}$$

To find a minimizer of $\text{val}(\mathbf{x})$, we can solve the min-max robust counterpart of $P(\mathbf{c})$ using an uncertainty set $\mathcal{U}(\lambda')$ with $\lambda' = (\int_{\Lambda} \lambda w(\lambda) d\lambda) / (\int_{\Lambda} w(\lambda) d\lambda)$. \square

Lemma 2. For an ellipsoidal uncertainty set $\mathcal{U}(\lambda) = \{\hat{\mathbf{c}} + \mathbf{C}\boldsymbol{\xi} : \|\boldsymbol{\xi}\|_2 \leq \lambda\}$ with $\lambda \in \mathbb{R}_+$, it holds that

$$\max_{c \in \mathcal{U}(\lambda)} \mathbf{c}\mathbf{x} = \hat{\mathbf{c}}\mathbf{x} + \lambda \|\mathbf{C}^t \mathbf{x}\|_2$$

Proof. This result has been shown in [BTN99] for $\lambda = 1$. The proof holds analogously. \square

Theorem 13. Let $\mathcal{U}(\lambda) = \{\hat{\mathbf{c}} + \mathbf{C}\boldsymbol{\xi} : \|\boldsymbol{\xi}\|_2 \leq \lambda\}$ be an ellipsoidal uncertainty set with $\lambda \in \Lambda \subseteq \mathbb{R}_+$. Then, an optimal solution to (CMM) can be found by solving a single robust problem with ellipsoidal uncertainty.

Proof. Using Lemma 2, we find

$$\begin{aligned} \text{val}(\mathbf{x}) &= \int_{\Lambda} w(\lambda) (\hat{\mathbf{c}}\mathbf{x} + \lambda \|\mathbf{C}^t \mathbf{x}\|_2) d\lambda \\ &= \left(\int_{\Lambda} w(\lambda) d\lambda \right) \hat{\mathbf{c}}\mathbf{x} + \left(\int_{\Lambda} \lambda w(\lambda) d\lambda \right) \|\mathbf{C}^t \mathbf{x}\|_2 \end{aligned}$$

To find a minimizer of $\text{val}(\mathbf{x})$, we can solve the min-max robust counterpart of $P(\mathbf{c})$ using an uncertainty set $\mathcal{U}(\lambda')$ with $\lambda' = (\int_{\Lambda} \lambda w(\lambda) d\lambda) / (\int_{\Lambda} w(\lambda) d\lambda)$ by applying Lemma 2 again. \square

Note that $(\int_0^1 \lambda w(\lambda) d\lambda) / (\int_0^1 w(\lambda) d\lambda) = \frac{1}{2}$, if $w(\lambda) = 1$, i.e., the compromise solution simply hedges against the average size of the uncertainty. In general, recall that this formula gives the centroid of the curve defined by w .

The results of this section show that compromise solutions are easy to compute, as the resulting problems have a simple structure. This is due to the linearity of the robust objective value in the uncertainty size λ , which was not the case for min-max regret.

6 Experiments

6.1 Setup

In this section, we present two experiments on compromise solutions to min-max regret problems with variable-sized uncertainty. The first experiment is concerned with the computational effort to find such a solution using the iterative algorithms presented in Section 3.3. In the second experiment we compare

these solutions to the alternatives of using classic regret solutions for various uncertainty set sizes.

Both experiments are conducted on shortest path instances of two types.

The first type consists of complete layered graphs. We parameterize such instances by the number of layers, width and cost types. Each graph consists of a source node, a sink node, and node layers of equal width between. For $N + 1$ layers of width k , an instance has a total of $(N + 1)k + 2$ nodes and $Nk^2 + 2k$ edges. We use $N = 5$ to $N = 55$ in steps of size 5, and $k = 5, 10, 15, 20$. Graph sizes thus vary from 32 nodes and 130 edges to 1,122 nodes and 22,040 edges.

Edges connect all nodes of one layer to the nodes of the next layer. Source and sink are completely connected to the first and last layer, respectively. We considered two types of cost structures to generate $\hat{\mathbf{c}}$. For type A, all costs are chosen uniformly from the interval $[1, 100]$. For type B costs, we generate nominal costs in $[1, 30] \cup [70, 100]$, i.e., they are either low or high.

In total, there are $11 \cdot 4 \cdot 2 = 88$ parameter combinations. For each combination, we generate 20 instances, i.e., a total of 1,760 instances.

The second type consists of graphs with two paths, that are linked by diagonal edges. For some length parameter L , we generated two separate paths from a node s to a node t , each with L nodes between. We then generate diagonal edges in the following way. On one of the two paths, we choose the i th node uniformly randomly. We then connect this node with the j th node on the other path, where $j > i$. The j th node is chosen with probability $\frac{3}{4}(\frac{1}{4})^{j-i-1}$, i.e., long diagonal edges are unlikely (ensuring that j is at most L).

Edges along the two base paths have length chosen uniformly from the interval $[1, 100]$. For diagonal edges, we determine their length by sampling from the same interval $(j - i)$ times, and adding these values, i.e., all paths have the same expected length.

We generate instances with length L from 50 to 850 in steps of 100, and set the number of diagonal edges to be $d \cdot L$ for $d \in \{0.05, 0.10, 0.15\}$. The smallest instances therefore contain 102 nodes and 105 edges; the largest instances contain 1,702 nodes and 1,830 edges. For each parameter combination, we generate 20 instances, i.e., a total of $9 \cdot 3 \cdot 20 = 540$ instances.

The classic min-max regret shortest path problem on instances of both types is known to be NP-hard, see [Cha16]. We investigate both types, as we expect the nominal solution to show a different performance: For layered graphs, the nominal solution is also optimal for $\mathcal{U}(1)$, as for every path there also exists a disjoint path. Therefore, the regret of a path P with respect to $\mathcal{U}(1)$ is $\sum_{e \in P} 2\hat{c}_e$. For the second type of instances, a good solution with respect to min-max regret can be expected to intersect with as many other paths as possible. We can therefore expect the nominal solution to be different to the optimal solution of $\mathcal{U}(1)$.

All experiments were conducted using one core of a computer with an Intel Xeon E5-2670 processor, running at 2.60 GHz with 20MB cache, with Ubuntu 12.04 and Cplex v.12.6.

6.2 Experiment 1: Computational Effort

6.2.1 Layered Graphs

We solve the compromise approach to variable-sized uncertainty for each instance using the algorithms described in Section 3.3 and record the computation times. Average computation times in seconds are presented in Table 1. In each column we average over all instances for which this parameter is fixed; i.e., in column "width 5" we show the results over all 440 instances that have a width of 5, broken down into classes of different length. The results indicate that computation times are still reasonable given the complexity of the problem, and mostly depend on the size of the instance (width parameter) and the density of the graph, while the cost structure has no significant impact on computation times.

	Width				Costs	
	5	10	15	20	A	B
5	0.05	0.13	0.21	0.40	0.20	0.20
10	0.22	0.49	0.89	1.43	0.75	0.77
15	0.47	0.99	1.78	2.98	1.47	1.64
20	1.17	2.31	3.61	6.78	3.45	3.49
25	1.99	4.17	7.53	11.47	5.97	6.61
30	3.77	7.86	13.13	21.14	11.50	11.45
35	6.05	11.08	19.87	35.51	18.28	17.97
40	9.46	21.85	35.37	49.58	28.64	29.48
45	13.77	29.64	56.30	85.48	47.23	45.37
50	21.58	46.37	67.88	141.14	66.33	72.16
55	26.61	69.56	125.95	193.30	105.94	101.76

Table 1: Average computation times to solve (CMMR) in seconds.

We present more details in Tables 2 and 3, where the number of iterations (i.e., how often was the relaxation of (CMMR) solved in Line 2 of Algorithm 2) and the size of $\bar{\Lambda}$ at the end of the algorithm are presented, respectively.

We find that the average number of iterations is stable and small, with around two iterations on average (the maximum number of iterations is three). This value seems largely independent of the problem size. For the number of generated change points $|\bar{\Lambda}|$, however, this is different. It increases with the number of layers, but it decreases with the width of the graph. Recall that the regret of a solution \mathbf{x} is roughly determined by the number of edges a regret path \mathbf{y} has in common. With increasing width, regret paths are less likely to use the same edges, which explains why the size of $\bar{\Lambda}(\mathbf{x})$ decreases. As before, we find that the cost structure does not have a significant impact on the performance of the solution algorithm.

6.2.2 Two-Path Graphs

The Tables 4, 5 and 6 correspond to the Tables 1, 2 and 3 from the last experiment, respectively. Computation times are sensitive to the parameter d , i.e., the number of diagonal edges. For small values of d , the computational complexity

	Width				Costs	
	5	10	15	20	A	B
5	1.98	1.98	1.73	1.80	1.89	1.85
10	2.02	1.98	1.98	2.00	1.98	2.01
15	2.08	2.02	1.98	1.98	2.01	2.01
20	2.08	2.08	2.00	2.08	2.06	2.05
25	2.00	2.02	2.02	2.05	2.04	2.01
30	2.15	2.10	2.08	2.05	2.11	2.08
35	2.10	2.02	2.10	2.05	2.06	2.08
40	2.10	2.15	2.10	2.02	2.11	2.08
45	2.12	2.15	2.12	2.05	2.10	2.12
50	2.17	2.08	2.05	2.15	2.05	2.17
55	2.10	2.10	2.02	2.12	2.10	2.08

Table 2: Average numbers of iterations.

	Width				Costs	
	5	10	15	20	A	B
5	3.05	2.95	2.17	2.27	2.59	2.64
10	5.05	3.75	3.33	3.10	3.71	3.90
15	6.72	4.95	4.33	4.10	4.92	5.12
20	9.00	6.53	5.03	5.20	6.40	6.47
25	10.47	7.70	6.65	5.60	7.42	7.79
30	11.68	9.43	7.28	6.70	9.00	8.54
35	13.10	9.32	7.62	7.10	8.99	9.59
40	14.95	11.05	9.35	7.78	10.95	10.61
45	16.10	11.35	10.35	8.53	11.85	11.31
50	18.73	12.57	10.05	9.47	12.56	12.85
55	19.77	14.40	11.43	9.62	13.68	13.94

Table 3: Average size of $\bar{\Lambda}$ at the end of the algorithm.

of problem (CMMR) scales well with the length of the graph; however, for larger values of d , the problem becomes intractable.

	d		
	0.05	0.10	0.15
50	0.04	0.05	0.08
150	0.13	0.29	0.67
250	0.22	0.79	2.23
350	0.48	2.07	11.76
450	0.78	5.37	28.22
550	1.33	12.01	57.44
650	2.06	19.65	165.17
750	3.01	36.70	488.51
850	3.84	73.42	3186.18

Table 4: Average computation times to solve (CMMR) in seconds.

While the number of iterations is relatively small overall, as in the last experiment, the size of $\bar{\Lambda}$ increases with d , which makes the master problems larger and more difficult to solve.

	d		
	0.05	0.10	0.15
50	2.00	2.05	2.00
150	2.05	2.15	2.20
250	2.05	2.10	2.30
350	2.10	2.20	2.45
450	2.10	2.35	2.45
550	2.20	2.30	2.40
650	2.15	2.30	2.55
750	2.20	2.35	2.50
850	2.05	2.50	2.35

Table 5: Average numbers of iterations.

6.3 Experiment 2: Comparison of Solutions

6.3.1 Layered Graphs

In our second experiment, we compare the compromise solution to the nominal solution (which is also the min-max regret solution with respect to the uncertainty sets $\mathcal{U}(0)$ and $\mathcal{U}(1)$), and to the min-max regret solutions with respect to $\mathcal{U}(0.3)$, $\mathcal{U}(0.5)$ and $\mathcal{U}(0.7)$.

To compare solutions, we calculate the regret of the compromise solution for values of λ in $[0, 1]$. We take this regret as the baseline. For all other solutions, we also calculate the regret depending on λ , and compute the difference to the baseline. We then compute the average differences for fixed λ over all instances of the same size. The resulting average differences are shown in Figure 5 for four

	d		
	0.05	0.10	0.15
50	2.75	3.55	3.95
150	4.20	6.45	8.90
250	5.15	9.15	13.15
350	6.95	12.25	18.05
450	8.05	16.25	21.40
550	9.95	18.45	27.80
650	11.20	20.20	30.25
750	12.35	22.45	32.70
850	13.05	27.00	39.00

Table 6: Average size of $\bar{\Lambda}$ at the end of the algorithm.

instance sizes. To set the differences in perspective, the average regret ranges from $\mathcal{U}(0)$ to $\mathcal{U}(1)$ of the compromise solutions are shown in the captions.

By construction, a min-max regret solution with respect to $\mathcal{U}(\bar{\lambda})$ has the smallest regret for this $\bar{\lambda}$. Generally, all presented solutions have higher regret than the nominal solution for small and for large values of λ , and perform better in between. By construction, the compromise solution has the smallest integral under the shown curve. It can be seen that it presents an interesting alternative to the other solutions by having a relatively small regret for small and large values of λ , but also a relatively good performance in between.

6.3.2 Two-Path Graphs

We generate the same plots as in Section 6.3.1 using the two-path instances. Recall that in this case, the nominal solution is not necessarily an optimal solution with respect to $\mathcal{U}(1)$. We therefore include an additional line for $\mathcal{U}(1)$ in Figure 6.

It can be seen that the nominal solution performs different to the last experiment; the regret increases with λ in a rate that part of the line needed to be cut off from the plot for better readability. The solution to $\mathcal{U}(0.5)$ performs very close to the compromise solution overall. Additionally, the scale of the plots show that differences in regret are much larger than in the previous experiment. Overall, it can be seen that using a robust solution plays a more significant role than in the previous experiment, as the nominal solution shows poor performance. The solutions that hedge against large uncertainty sets ($\mathcal{U}(0.7)$ and $\mathcal{U}(1.0)$) are relatively expensive for small uncertainty sets and vice versa. The compromise solution (as $\mathcal{U}(0.5)$, in this case) presents a reasonable trade-off over all uncertainty sizes.

7 Conclusion

Classic robust optimization approaches assume that the uncertainty set \mathcal{U} is part of the input, i.e., it is produced using some expert knowledge in a previous step. If the modeler has access to a large set of data, it is possible to follow recently developed data-driven approaches to design a suitable set \mathcal{U} . In our

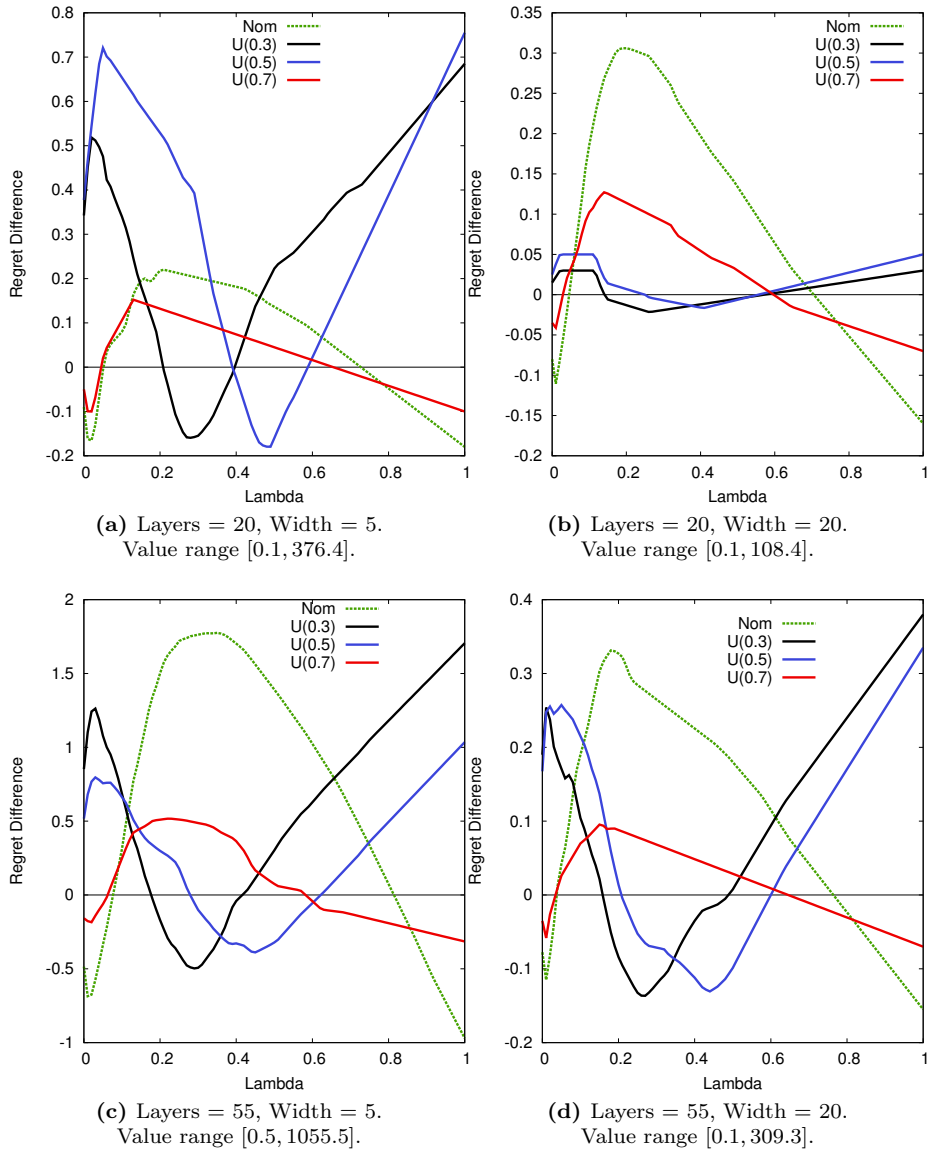


Figure 5: Difference in regret compared to nominal solution depending on λ .

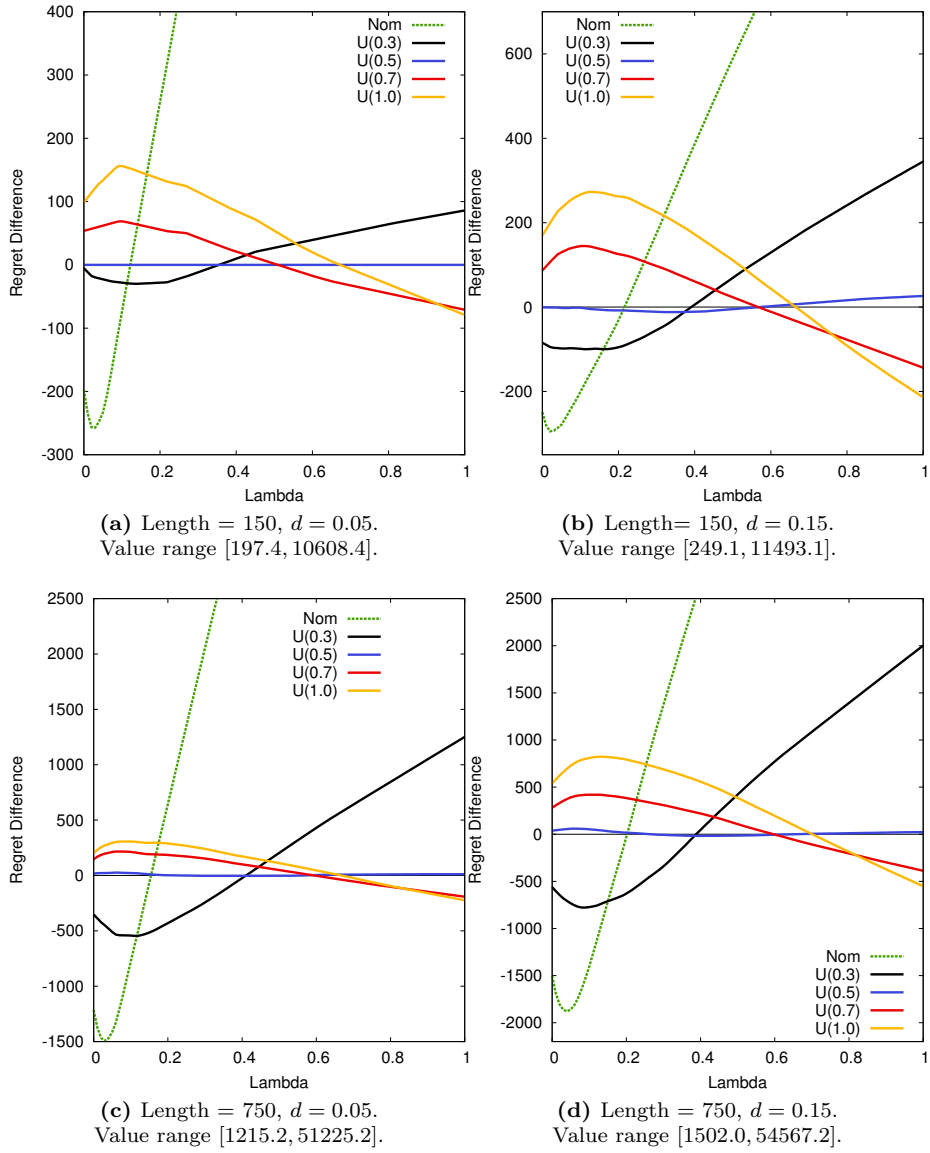


Figure 6: Difference in regret compared to nominal solution depending on λ .

approach, we remove the necessity of defining \mathcal{U} by using a single nominal scenario, and considering all uncertainty sets generated by scaling this scenario using a single parameter. Following a goal programming perspective, the aim of the compromise approach is to find a single solution that performs well on average in the robust sense over all possible uncertainty set sizes.

For min-max combinatorial problems, we showed that our approach can be reduced to solving a classic robust problem of particular size. The setting is more involved for min-max regret problems, where the regret objective is a piecewise linear function in the uncertainty size. We presented a general solution algorithm for this problem, which is based on a reduced master problem, and the iterative solution of subproblems of nominal structure.

For specific problems, positive and negative complexity results were demonstrated. The compromise selection problem can be solved in polynomial time. Solutions to the compromise minimum spanning tree problem can be evaluated in polynomial time, but it is NP-hard to find an optimal solution. For compromise shortest path problems, the same results hold in case of layered graphs; however, it is still an open problem if polynomial time algorithms exist for special classes such as series-parallel graphs.

In computational experiments we highlighted the value of our approach in comparison with different min-max regret solutions, and showed that computation times can be within few minutes for instances with up to 22,000 edges.

References

- [ABV09] H. Aissi, C. Bazgan, and D. Vanderpooten. Min-max and min-max regret versions of combinatorial optimization problems: A survey. *European Journal of Operational Research*, 197(2):427–438, 2009.
- [AL04] I. Averbakh and V. Lebedev. Interval data minmax regret network optimization problems. *Discrete Applied Mathematics*, 138(3):289–301, 2004.
- [BB09] D. Bertsimas and D. B. Brown. Constructing uncertainty sets for robust linear optimization. *Operations research*, 57(6):1483–1495, 2009.
- [BBC11] D. Bertsimas, D. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.
- [BGK13] D. Bertsimas, V. Gupta, and N. Kallus. Data-driven robust optimization. *arXiv preprint arXiv:1401.0212*, 2013.
- [BTdHdW⁺13] A. Ben-Tal, D. den Hertog, A. de Waegenaere, B. Melenberg, and G. Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.
- [BTGN09] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, Princeton and Oxford, 2009.

- [BTN98] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of operations research*, 23(4):769–805, 1998.
- [BTN99] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations Research Letters*, 25(1):1–13, 1999.
- [BTN00] A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical programming*, 88(3):411–424, 2000.
- [BTN02] A. Ben-Tal and A. Nemirovski. Robust optimization–methodology and applications. *Mathematical Programming*, 92(3):453–480, 2002.
- [Car83] P. J. Carstensen. *The complexity of some problems in parametric linear and combinatorial programming*. University of Michigan, 1983.
- [CG16] A. Chassein and M. Goerigk. Performance analysis in robust optimization. In *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, pages 145–170. Springer, 2016.
- [CG17] A. Chassein and M. Goerigk. Minmax regret combinatorial optimization problems with ellipsoidal uncertainty sets. *European Journal of Operational Research*, 258(1):58–69, 2017.
- [CG18] A. Chassein and M. Goerigk. Variable-sized uncertainty and inverse problems in robust optimization. *European Journal of Operational Research*, 264(1):17–28, 2018.
- [Cha16] A. Chassein. *Robust Optimization: Complexity and Solution Methods*. PhD thesis, TU Kaiserslautern, 2016.
- [Con04] E. Conde. An improved algorithm for selecting p items with uncertain returns according to the minmax-regret criterion. *Mathematical Programming*, 100(2):345–353, 2004.
- [DP88] D. Dubois and H. Prade. Possibility theory: An approach to the computerized processing of information, 1988.
- [Ehr06] M. Ehrgott. *Multicriteria optimization*. Springer Science & Business Media, 2006.
- [GMT14] V. Gabrel, C. Murat, and A. Thiele. Recent advances in robust optimization: An overview. *European journal of operational research*, 235(3):471–483, 2014.
- [GS10] J. Goh and M. Sim. Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1):902–917, 2010.
- [GS16] M. Goerigk and A. Schöbel. Algorithm engineering in robust optimization. In L. Kliemann and P. Sanders, editors, *Algorithm Engineering: Selected Results and Surveys*, volume 9220, pages 245–279. Springer, 2016.

- [GYdH15] B. L. Gorissen, İ. Yanıkoğlu, and D. den Hertog. A practical guide to robust optimization. *Omega*, 53:124–137, 2015.
- [KZ06] A. Kasperski and P. Zieliński. An approximation algorithm for interval data minmax regret combinatorial optimization problems. *Information Processing Letters*, 97(5):177–180, 2006.
- [KZ10] A. Kasperski and P. Zieliński. Minmax regret approach and optimality evaluation in combinatorial optimization problems with interval and fuzzy weights. *European Journal of Operational Research*, 200(3):680–687, 2010.
- [KZ11] A. Kasperski and P. Zieliński. On the approximability of robust spanning tree problems. *Theoretical Computer Science*, 412(4):365–374, 2011.
- [KZ16] A. Kasperski and P. Zieliński. Robust discrete optimization under discrete and interval uncertainty: A survey. In *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, pages 113–143. Springer, 2016.
- [MJC15] Andrew Mastin, Patrick Jaillet, and Sang Chin. Randomized minmax regret for combinatorial optimization under uncertainty. In *International Symposium on Algorithms and Computation*, pages 491–501. Springer, 2015.
- [WKS14] W. Wiesemann, D. Kuhn, and M. Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.
- [WOBD13] K. Witting, S. Ober-Blöbaum, and M. Dellnitz. A variational approach to define robustness for parametric multiobjective optimization problems. *Journal of Global Optimization*, 57(2):331–345, 2013.
- [YdH12] İ. Yanıkoğlu and D. den Hertog. Safe approximations of ambiguous chance constraints using historical data. *INFORMS Journal on Computing*, 25(4):666–681, 2012.
- [YKP01] H. Yaman, O. E. Karaşan, and M. Ç. Pınar. The robust spanning tree problem with interval data. *Operations Research Letters*, 29(1):31–40, 2001.