The Minimum Cost Query Problem on Matroids with Uncertainty Areas

Arturo I. Merino 🔟

Dept. of Mathematical Engineering and CMM, Universidad de Chile & UMI-CNRS 2807,
Santiago, Chile
amerino@dim.uchile.cl

José A. Soto[1] 🔟

Dept. of Mathematical Engineering and CMM, Universidad de Chile & UMI-CNRS 2807,
Santiago, Chile
www.dim.uchile.cl/~jsoto
jsoto@dim.uchile.cl

─── **Abstract** ───

We study the minimum weight basis problem on matroid when elements' weights are uncertain. For each element we only know a set of possible values (an uncertainty area) that contains its real weight. In some cases there exist bases that are uniformly optimal, that is, they are minimum weight bases for every possible weight function obeying the uncertainty areas. In other cases, computing such a basis is not possible unless we perform some queries for the exact value of some elements.

Our main result is a polynomial time algorithm for the following problem. Given a matroid with uncertainty areas and a query cost function on its elements, find the set of elements of minimum total cost that we need to simultaneously query such that, no matter their revelation, the resulting instance admits a uniformly optimal base. We also provide combinatorial characterizations of all uniformly optimal bases, when one exists; and of all sets of queries that can be performed so that after revealing the corresponding weights the resulting instance admits a uniformly optimal base.
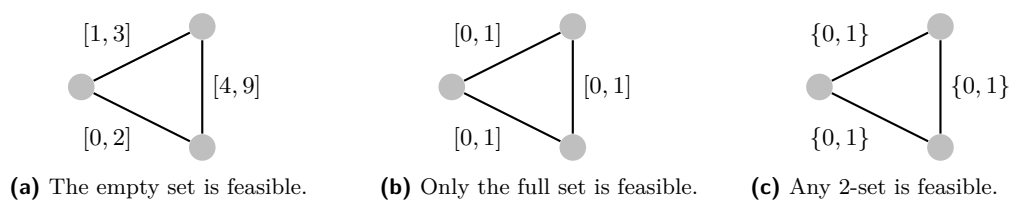
1 Introduction

We study fundamental combinatorial optimization problems on weighted structures where the numerical data is uncertain but it can be queried at a cost. We focus on the problem of finding a minimum weight base of a matroid under uncertainty, a problem that includes finding the smallest k elements of a list and the minimum spanning tree (MST) problem. In our setting, for every element e of the matroid we know a set $\mathcal{A}(e)$, called uncertainty area, of possible values that contains its real weight w_e. We can reveal this real weight by paying some query cost c_e. We assume that the queries are done in a non-adaptive way, or equivalently, that all the elements queried reveal their values at the same time. A set of elements F is a *feasible query* if for every possible revelation of the weights of F it is possible to compute a minimum weight base T of the resulting instance. The task of the *Minimum Cost Query Problem on Matroids* is to determine a minimum-cost feasible query.

─────────────

[1] Corresponding author

(a) The empty set is feasible.   **(b)** Only the full set is feasible.   **(c)** Any 2-set is feasible.

Figure 1 Feasible queries.

To better illustrate this, consider the problem of computing an MST of a triangle graph in three possible situations as shown in Figure 1. In the first situation, the edges with areas $[0, 2]$ and $[1, 3]$ always form an MST, so we don't need to query any element. In this case we say that the matroid (the graph) admits a *uniformly optimal basis*, that is, a basis (a spanning tree) having minimum weight for every possible realization of the elements' weights. In the second situation, all edges have uncertainty area $[0, 1]$ and the only feasible query is the entire set of edges. For if we only query two edges, we could be in a situation where both have weight $1/2$. With that information we cannot compute an MST: if the unqueried edge e had weight 0, then e must be in the MST. However, if it had weight 1, then e cannot be in an MST. In the last situation, all uncertainty areas are finite: they have two elements $\{0, 1\}$. Here, every set of size two is feasible. Indeed, if both elements reveal a weight of 0, then they form an MST. Otherwise, the set obtained by deleting any edge with weight 1 is an MST.

Paper outline and results

In Section 2 we give formal definitions and show a simple but strong result: the uniformly optimal bases (UOB) of an uncertainty matroid form the bases of a second matroid. In Section 3 we extend the classic MST red and blue rules to uncertainty matroids, introducing the concept of *colored* elements, and study their properties. We then show that an uncertainty matroid admits a UOB if and only if all its elements are colored, and use this to give a combinatorial description of the matroid of uniformly optimal bases. We provide polynomial time algorithms for testing the existence of a UOB and for finding one if it exists. In Section 4 we study the minimum cost feasible query problem in detail. By using our coloring framework we construct a partition of the elements into groups that characterize minimal feasible queries. We show that every minimal feasible query set is formed by taking the first group (denoted as the core) completely and by deleting exactly one element from each other group. Our main result is an algorithm to find this partition that we use to fully solve the minimum cost feasible query problem on matroids. Unlike related work on the MST, the uncertainty areas in our setting can be arbitrary sets of real numbers, and not just intervals. Our algorithms only assume access to an independence oracle for the matroid, and a very mild type of access to the uncertainty areas. At the end of our study we show that the interval uncertainty area case is special, as there is a unique minimal size feasible query. We also relate the solutions for the MST with $\{0, 1\}$-areas case with the 2-connected components of the graph.

Related Work

Traditional research in optimization with uncertain data mostly focuses on finding solutions whose value is good, either in the worst case (robust optimization) or in some probabilistic sense (stochastic optimization), without gaining new information about the uncertain data. Our problem contributes to the *query model setting*, a different approach that has gained some strength in the last years. In this model one assumes that the algorithm can learn the

exact value of an uncertain input data by paying some query cost in order to solve a certain problem P (e.g., determining the MST of a graph). The aim is to minimize the cost of the queries while guaranteeing that an exact/approximate solution can be computed. Work in this area (see [3] for a survey) falls into three main categories:

Adaptive online.   In this setting, the algorithm can query elements one by one, using the information revealed until a certain step to guide the decision of the next element to query. Even though algorithms for this version can be analyzed with a traditional worst-case approach (e.g., minimizing the depth of the decision tree associated to the algorithm's strategy), most of the work in this setting prefers to measure performance in terms of competitive analysis, comparing the number of queries an algorithm makes with the minimum number of queries that an adversary algorithm (that knows the real values beforehand) would make in order to verify that an answer is correct. Probably the first one to consider this model is Kahan [9] who provide optimal online algorithms to compute *all the elements* achieving the maximum, median and minimum gap of a list of closed intervals. Feder et al. [7] devise optimal online competitive algorithms to determine the numerical *value* of the median, and more generally, the K-th top element of the list within some prescribed tolerance. Bruce et al. [1] introduce a general method, called the *witness algorithm*, for adaptive online problems with open intervals and singletons uncertainty areas. They apply this method to geometric problems such as finding all maximum points in the plane from a family with uncertain coordinates. Erlebach et al. [5] studied the MST problem under two types of uncertainty, the edge uncertainty one (which is the same as ours) and the vertex uncertainty setting, in which the graph is complete, the vertices are points in the plane whose coordinates are uncertain, and the weight of an edge is the distance between its endpoints. They get 2 and 4 competitive algorithms, respectively, for both types of uncertainty, under open intervals and singletons areas, which is optimal for deterministic algorithms. The algorithm for edge uncertainty, denoted as U-RED, was later extended by Erlebach, Hoffmann and Kammer [4] to the minimum weight base problem on matroids achieving an optimal 2 competitive algorithm. Megow, Meißner and Skutella [13] show that by using randomization one can do better, lowering the competitive ratio down to $1 + 1/\sqrt{2}$. They also studied the non-uniform cost case. Gupta et al. [8] studied variants where queries return refined estimates of the areas, instead of a single value.

Verification.   The verification problem is the one the offline adversary of the previous setting has to solve. That is, given both the uncertainty areas and a family of assumed real values, to determine the minimum number of queries one has to make so that, if the values obtained from the queries and the assumed values coincide, then no more queries are needed in order to obtain an optimal solution. Charalambous and Hoffman [2] show that the verification problem for maximal points in the plane is NP-hard for uncertainty areas of size at most 2. Erlebach and Hoffmann [3] show that the verification problem of MST with (open interval and singleton) uncertainty in the edges is in P, while that for vertex uncertainty is NP-hard.

Non-adaptive online.   This setting encompasses our work and is sometimes called the *offline* problem. In it, an algorithm must determine a set F of queries to perform simultaneously, in order to have enough information to solve the problem. The only conceptual difference between the non-adaptive online problem and the verification one, is that in the latter, the algorithm can make use of the real values of the elements to guide the decision of which queries to perform, while in the former, that information is not available. Feder et al. [7] provide optimal algorithms for finding the K-th top element of a list up to additive tolerance.

Later, Feder et al. [6] consider the problem of finding the shortest *s-t* path on a DAG with closed-intervals uncertainty on the edges. They show that determining the length of the shortest *s-t* path within a given additive error is neither in NP nor in co-NP unless NP=co-NP, and provide exact algorithms for some special cases. To the best of our knowledge, the MST and, more generally, the matroid case have not been considered before this work.

Further related work.   A common approach to deal with closed-interval uncertainty for a problem, without querying extra information, is to find a solution that minimizes the maximal regret, that is, the difference between the (real) weight of the chosen solution and the weight of the best solution that could have been picked had the true weights been known. Zero maximal regret bases and uniformly optimal bases of a matroid coincide. In the context of the MST problem, Yaman et al. [16] characterize trees with zero regret when they exists. In [10], Kaspersky and Zielinsky give a 2-approximation algorithm for the minmax regret problem on general matroids with interval data, and in [11] they give algorithms to find zero maximal regret bases. It is worth noting that all these results assume interval areas on the elements, while we allow for arbitrary uncertainty areas.

2   Preliminaries – Uniformly Optimal Bases

We assume familiarity with basic concepts in matroid theory such as bases, independent sets, span, circuit, cocircuits and duality, contraction, deletion and matroid connectivity. For an introduction and specific results, we refer to Oxley's book [14]. However, most of this paper can be understood by having the graphic matroid of a connected graph in mind. This is a matroid whose elements are the edges. The bases, independent sets, circuits and cocircuits, are the spanning trees, forests, cycles and minimal edge cut-sets, respectively. An element e is in the span (cospan) of a set F if there is a circuit (cocircuit) in $F \cup \{e\}$ containing e. We also use the standard notation $X + e$ and $X - e$ to denote $X \cup \{e\}$ and $X \setminus \{e\}$.

▶ **Definition 1** (Uncertainty Matroid). *An uncertainty matroid is a pair* $(\mathcal{M}, \mathcal{A})$ *where* $\mathcal{M} = (E, \mathcal{I})$ *is a matroid and* $\mathcal{A} \colon E \to 2^{\mathbb{R}} \setminus \{\emptyset\}$ *is a function mapping each element* $e \in E$ *in the ground set to a nonempty set* $\mathcal{A}(e)$ *of real numbers denoted the uncertainty area of* e. *We denote* $\inf \mathcal{A}(e)$ *by* L_e *and* $\sup \mathcal{A}(e)$ *by* U_e.

If $\mathcal{A}(e)$ is a singleton, we say that the element e is *certain*, otherwise, we say it is *uncertain*. If all the elements are certain, then we can identify $\mathcal{A}$ with the associated weight function $w \colon E \to \mathbb{R}$, such that $\mathcal{A}(e) = \{w_e\}$ so that $(\mathcal{M}, w)$ becomes a weighted matroid.

▶ Remark 2. For our algorithms, we assume access to matroid $\mathcal{M}$ via an independence oracle. We also need a very mild access to the uncertainty areas. More precisely, we assume that both L_e and U_e are known for every element, and that for every pair of (not necessarily distinct) elements e and f we can test if $(L_e, U_e) \cap \mathcal{A}(f)$ is empty in constant time. A polynomial time algorithm will, therefore, only use a polynomial number of calls to the independence oracle and to the uncertainty areas. Furthermore, for simplicity we will assume that all the infima and suprema L_e and U_e are finite. Otherwise we can apply a suitable strictly increasing function mapping the reals to a bounded set (such as arctan), and work in its image instead.

Intuitively, an uncertainty matroid models the situation in which we do not know the actual weight of every element e in the matroid, but we do know a set $\mathcal{A}(e)$ containing it. We can learn the actual weight of e by querying it. In this work we are concerned with non-adaptive (simultaneous) queries. The new uncertainty area function obtained after querying a subset of elements will be called a *revelation* of $\mathcal{A}$. The formal definition is below.

▶ **Definition 3** (Revelations and realizations). *Let $X \subseteq E$. A revelation of X in $(\mathcal{M}, \mathcal{A})$ is a function $\mathcal{B}: E \to 2^{\mathbb{R}} \setminus \{\emptyset\}$ such that:*

(i) *$\forall e \in X$, $\mathcal{B}(e) = \{b_e\} \subseteq \mathcal{A}(e)$ is a singleton, and*

(ii) *$\forall e \in E \setminus X$, $\mathcal{B}(e) = \mathcal{A}(e)$.*

In particular, $(\mathcal{M}, \mathcal{B})$ is also an uncertainty matroid. A realization is a revelation of the entire ground set E. The collection of all revelations of X in $(\mathcal{M}, \mathcal{A})$ is denoted by $\mathcal{R}(X, \mathcal{A})$.

Suppose we want to compute a minimum weight basis of a matroid, but we only know uncertainty areas for its elements. In certain situations (e.g., Figure 1 (a)), we can find sets that are optimal bases for every realization, we call them uniformly optimal bases.

▶ **Definition 4** (Uniformly optimal bases). *A set $T \subseteq E$ is a uniformly optimal basis of $(\mathcal{M}, \mathcal{A})$ (or simply, an $\mathcal{A}$-basis) if for every realization w, T is a minimum weight basis (a w-basis).*

Recall that a nonempty family of sets form the bases of a matroid if and only if they satisfy the strong basis exchange axiom. Our first basic result is the following:

▶ **Lemma 5.** *Let $\mathfrak{B}$ be the collection of all uniformly optimal bases of $(\mathcal{M}, \mathcal{A})$. Suppose $\mathfrak{B} \neq \emptyset$ and let T_1 and T_2 be two sets in $\mathfrak{B}$. If e is an element in $T_1 \setminus T_2$ then:*

(i) *e is certain, and*

(ii) *strong basis exchange holds, i.e., there is an element f in $T_2 \setminus T_1$ such that both $T_1 - e + f$ and $T_2 - f + e$ are in $\mathfrak{B}$.*

In particular, if $\mathfrak{B} \neq \emptyset$, then $\mathfrak{B}$ is the set of bases of a matroid, that we denote by $\mathtt{mat}(\mathcal{M}, \mathcal{A})$

Proof. Let $e \in T_1 \setminus T_2$. By strong basis exchange of $\mathcal{M}$, there is an element $f \in T_2 \setminus T_1$ such that both $T_1 - e + f$ and $T_2 + e - f$ are bases of $\mathcal{M}$. Assume first by contradiction that e is uncertain, then there is a realization $w \in \mathcal{R}(E, \mathcal{A})$ such that $w_e \neq w_f$. If $w_e > w_f$ then $w(T_1 - e + f) < w(T_1)$ contradicting the fact that T_1 is uniformly optimal. On the other hand, if $w_e < w_f$ then $w(T_2 + e - f) < w(T_2)$ contradicting that T_2 is uniformly optimal. We conclude not only that e is certain, but also that in every realization $w_f = w_e$. In particular, for every realization w, $w(T_1 - e + f) = w(T_1) = w(T_2) = w(T_2 - f + e)$, i.e. both $T_1 - e + f$ and $T_2 - f + e$ are uniformly optimal bases. ◀

3 Blue and red rules for uncertainty matroids

By Lemma 5, we conclude that if an uncertain element e is in some uniformly optimal basis then it is in *every* uniformly optimal basis. Our next task is to characterize the set of uncertain elements that are in every uniformly optimal basis. Now it is useful to remember the classic blue and red rules for computing an MST. An edge of a weighted graph is called blue if it is in at least one MST, and it is called red if it is outside at least one MST. Virtually every algorithm for the MST work in steps: if it detects a red edge from the current graph, it deletes it from the graph, and if it detects a blue edge then it adds it to the solution and contracts it. The following definitions extend the coloring notion to uncertainty matroids.

▶ **Definition 6** (Blue and red elements). *An element $e \in E$ is blue (resp. red) if for every realization $w \in \mathcal{R}(E, \mathcal{A})$ there exists a w-basis T such that $e \in T$ (resp. $e \notin T$). We say that e is colored if e is red or blue (or both at the same time), otherwise we say that e is uncolored.*

Note that e is blue (resp. red) on $(\mathcal{M}, \mathcal{A})$ if and only if e is blue (resp. red) on $(\mathcal{M}, w)$ for every realization $w \in \mathcal{R}(E, \mathcal{A})$. Standard matroid arguments show that for any such w, if an element e is any heaviest element of a circuit, then it avoids some w-basis (i.e., it is red on $(\mathcal{M}, w)$), and if it is the unique heaviest element, then it cannot be in a w-basis (i.e., it is

not blue on $(\mathcal{M}, w)$). Blue elements in graphs / matroids with interval uncertainty areas have been studied before under the name of strong edges [16] or necessarily optimal elements [11]. We start with a basic result on colored elements, they are preserved under revelations.

▶ **Lemma 7.** *Let $X \subseteq E$ and $\mathcal{B} \in \mathcal{R}(X, \mathcal{A})$. If e is blue (resp. red) in $(\mathcal{M}, \mathcal{A})$, then it is blue (resp. red) in $(\mathcal{M}, \mathcal{B})$.*

Proof. Suppose e is blue in $(\mathcal{M}, \mathcal{A})$ and consider any $w \in \mathcal{R}(E, \mathcal{B})$. Since $\mathcal{R}(E, \mathcal{B}) \subseteq \mathcal{R}(E, \mathcal{A})$ we get that there exists some w-basis T such that $e \in T$, implying that e is blue in $(\mathcal{M}, \mathcal{B})$. The proof works analogously if e is red in $(\mathcal{M}, \mathcal{A})$. ◀

We can characterize blue and red elements of an uncertainty matroid using its span and cospan functions[2] of a set X in the matroid $\mathcal{M}$ (resp., in the dual matroid $\mathcal{M}^*$). The span and cospan of a set can be computed using a polynomial number of calls to the independence oracle.

▶ **Lemma 8.** *Let $F(e) = \{f \in E - e : L_f < U_e\}$ and $F^*(e) = \{f \in E - e : L_e < U_f\}$.*
   (i) *$e \in E$ is blue if and only if $e \notin \operatorname{span} F(e)$.*
   (ii) *$e \in E$ is red if and only if $e \notin \operatorname{cospan} F^*(e)$.*
In particular, we can test in polynomial time if any element is blue, red, both or none.

▶ Remark 9. Proofs involving blue and red elements follow the same main ideas. Most of the time we only consider the blue case, since the red case follows by *duality arguments*. More precisely, consider the dual matroid $\mathcal{M}^*$ with inverted uncertainty area function $-\mathcal{A}(e) = \{-x : x \in \mathcal{A}(e)\}$. Since the bases of $\mathcal{M}^*$ are complements of the bases of $\mathcal{M}$ we also get that the uniformly optimal bases of $(\mathcal{M}, \mathcal{A})$ are the complements of those in $(\mathcal{M}^*, -\mathcal{A})$, that the red elements in $(\mathcal{M}, \mathcal{A})$ are the blue elements of $(\mathcal{M}^*, -\mathcal{A})$ and vice versa.

Proof of Lemma 8. We only prove (i), since (ii) follows by duality arguments using that $F^*(e)$ is the analogue of $F(e)$ for $(\mathcal{M}^*, -\mathcal{A})$.

Let e be a blue element and let $K = \min\limits_{f \in F(e)}(U_e - L_f) > 0$. For each $f \in E - e$ choose $\varepsilon_f \in [0, K/2)$ such that $L_f + \varepsilon_f \in \mathcal{A}(f)$. In a similar way, choose $\varepsilon_e \in [0, K/2)$ such that $U_e - \varepsilon_e \in \mathcal{A}(e)$. Consider the realization $w \in \mathcal{R}(E, \mathcal{A})$ given by $w_f = \begin{cases} U_e - \varepsilon_e & \text{if } f = e, \\ L_f + \varepsilon_f & \text{if } f \neq e. \end{cases}$

By construction, for every $f \in F(e)$, $w_f < w_e$. Suppose now that $e \in \operatorname{span} F(e)$. Then, there exists a circuit $C \subseteq F(e) + e$ such that e is the unique heaviest (with respect to w) element of C. This implies that e is outside every w-basis, which contradicts that e was blue.

Now let e be an element outside $\operatorname{span} F(e)$ and suppose that e is not blue. Then, there exists a realization $w \in \mathcal{R}(E, \mathcal{A})$ such that e is not in any w-basis. Choose T to be any w-basis and let C be the fundamental circuit[3] of $T + e$. As $e \notin \operatorname{span} F(e)$ we have that $C - e \nsubseteq F(e)$. Select any $f \in (C - e) \setminus F(e)$, then: $w_f \geq L_f \geq U_e \geq w_e$. We conclude that $T - f + e$ is a w-basis that contains e, which is a contradiction. ◀

Even though certain elements can be blue and red at the same time (e.g., in a circuit in which every element has the same weight every element has both colors), this is not possible for uncertain ones as shown by the next lemma whose proof is deferred to the full version.

---

[2] We recall that $\operatorname{span}(X)$ for $X \subseteq E$ is the unique maximal set $U \supseteq X$ with the same rank as X and $\operatorname{cospan}(X)$ is its span in the dual matroid. The span and cospan of a set in a matroid are related by the expression $\operatorname{cospan}(X) = X \cup \{e \in E : e \notin \operatorname{span}[(E - e) \setminus X]\}$

[3] Recall that if T is a basis and e an element, the fundamental circuit C of $T + e$ is the only circuit of $T + e$. For any $f \in C$, $T - f + e$ is a basis.

▶ **Lemma 10.** *If e is red and blue, then e is certain.*

The definition of blue and red will be useful even if e is certain, or if the uncertainty matroid has no uniformly optimal basis. The following lemma, whose proof is in the full version, highlights the utility of this definition.

▶ **Lemma 11.** *Let $(\mathcal{M}, \mathcal{A})$ be an uncertainty matroid such that an $\mathcal{A}$-basis exists.*
  (i) *If e is blue (resp. red) then there exists an $\mathcal{A}$-basis T such that $e \in T$ (resp. $e \notin T$).*
  (ii) *Let e be an uncertain element. Then e is blue (resp. red) if and only if e is inside (resp. outside) every $\mathcal{A}$-basis.*

The previous lemma shows that if the set of uniformly optimal bases is nonempty then all uncertain elements are colored either red or blue. Furthermore the uncertain elements contained in every $\mathcal{A}$-basis are exactly the blue uncertain elements. The next theorem, which is the main result of this section, shows that a converse also holds.

▶ **Theorem 12.** *An $\mathcal{A}$-basis exists if and only if every uncertain element is colored.*

In order to prove this theorem we need the following simple lemmas, whose proofs are deferred to the full version. Basically, they show that contracting blue elements and/or deleting red elements preserves structure and colors.

▶ **Lemma 13.** *Let $e \in E$ be a colored element (certain or uncertain).*
  (i) *If e is blue: T is an $(\mathcal{M}/e, \mathcal{A}|_{E-e})$-basis if and only if $T + e$ is an $(\mathcal{M}, \mathcal{A})$-basis.*
  (ii) *If e is red: T is an $(\mathcal{M} \setminus e, \mathcal{A}|_{E-e})$-basis if and only if T is an $(\mathcal{M}, \mathcal{A})$-basis.*

▶ **Lemma 14.** *Let $e \in E$ be a colored uncertain element, and $f \in E - e$.*
  (i) *Suppose e is blue. If f is blue (resp. red) in $(\mathcal{M}, \mathcal{A})$, then f is blue (resp. red) in $(\mathcal{M}/e, \mathcal{A}|_{E-e})$.*
  (ii) *Suppose e is red. If f is blue (resp. red) in $(\mathcal{M}, \mathcal{A})$, then f is blue (resp. red) in $(\mathcal{M} \setminus e, \mathcal{A}|_{E-e})$.*

Proof of Theorem 12. We only need to prove the converse and we proceed by induction on the number of uncertain elements k. If $k = 0$ an $\mathcal{A}$-basis is simply a basis of minimum weight, which clearly exists. Suppose that $k > 0$, and let $e \in E$ be any uncertain element. If e is blue, we have that every uncertain element of $E - e$ is colored in $(\mathcal{M}/e, \mathcal{A}|_{E-e})$, as colors were preserved. By inductive hypothesis, we have an $(\mathcal{M}/e, \mathcal{A}|_{E-e})$-basis T and by Lemma 13 we have that $T + e$ is an $(\mathcal{M}, \mathcal{A})$-basis. If e is red, one can proceed similarly but deleting instead. ◀

▶ Remark 15. The previous theorem gives an algorithmic way to test if $(\mathcal{M}, \mathcal{A})$ admits an $\mathcal{A}$-basis: we simply check if every element is colored, using Lemma 8. Let us now consider the problem of finding one such base. It is worth noting that algorithms for this task are available for closed interval and open interval uncertainty areas. For closed interval uncertainty areas, one can find an $\mathcal{A}$-basis by using Kaspersky and Zielinsky's approach [10] for finding a zero maximal regret basis. For open intervals (or singletons), one can simply run the 2-competitive U-RED algorithm by Erlebach, Hoffmann and Kammer [4] for the online adaptive variant: if this algorithm does not perform any query, then it outputs a uniformly optimal basis. Otherwise, the algorithm finds a witness set, i.e., a set for which at least one element must be queried in order to find a solution. In what follows we provide a new algorithm that finds $\mathcal{A}$-bases for arbitrary uncertainty areas (not just intervals).

If uniformly optimal bases exist, then the elements are partitioned into blue uncertain, red uncertain and certain elements. After contraction of the set B of all blue uncertain elements and deletion of the set R of all red uncertain elements, we are left with a matroid that only has certain elements, we call such weighted matroid the *certain weighted matroid* $(\mathcal{M}^c, w^c)$, where $\mathcal{M}^c = \mathcal{M}/B \setminus R$, and $w^c = \mathcal{A}|_{E \setminus (R \cup B)}$. The following theorem shows that every uniformly optimal basis arises by extending some optimal basis on the certain weighted matroid.

▶ **Theorem 16.** *Let $(\mathcal{M}, \mathcal{A})$ be an uncertainty matroid for which an $\mathcal{A}$-basis exists and $(\mathcal{M}^c, w^c)$ its certain weighted matroid. Then T is an $\mathcal{A}$-basis if and only if:*

(i) *T contains every blue uncertain element,*

(ii) *T avoids each red uncertain element, and*

(iii) *The certain elements of T form a minimum weight basis of $(\mathcal{M}^c, w^c)$.*

Proof. $\mathcal{A}$-bases always contain every blue uncertain element and avoid each red uncertain element by Lemma 11. By Lemma 14, we can delete each red uncertain element while preserving colors. More so, from Lemma 13 we conclude that T is uniformly optimal after these deletions. A similar argument allows us to now contract each blue uncertain element while preserving colors in each contraction. We are only left with the certain elements of T and by Lemma 13 we conclude they form a minimum weight basis of $(\mathcal{M}^c, w^c)$.

We show the converse by induction on the number of uncertain elements k of E. If $k = 0$, using (iii) we get that T is a minimum weight basis of $(\mathcal{M}^c, w^c)$. Noting that $\mathcal{M} = \mathcal{M}^c$ and $\mathcal{A} = w^c$ we conclude that T is an $(\mathcal{M}, \mathcal{A})$-basis.

If $k > 0$ select any uncertain element $e \in E$. If e is blue, we have from (i) that $e \in T$. As colors are preserved when contracting e, it follows by inductive hypothesis that $T - e$ is an $(\mathcal{M}/e, \mathcal{A}|_{E-e})$-basis, and using Lemma 13 we conclude that T is an $(\mathcal{M}, \mathcal{A})$-basis. If e is red, then $e \notin T$ by (ii). We now proceed as before but deleting e instead. ◀

Theorems 12 and 16 allow for algorithmic implementation. We can decide if an $\mathcal{A}$-basis exists by checking if all elements are colored. If every element is colored, we contract every blue uncertain element, delete each red uncertain element, compute a minimum weight basis of the certain weighted matroid and output the optimal certain basis along with every blue uncertain element. We summarize this result and discuss it in more detail in the full version. We also use the previous theorem to characterize the matroid $\mathtt{mat}(\mathcal{M}, \mathcal{A})$.

▶ **Corollary 17.** *There is an algorithm that finds an $\mathcal{A}$-basis or decides that none exists in polynomial time.*

▶ **Corollary 18.** *If $(\mathcal{M}, \mathcal{A})$ admits an $\mathcal{A}$-basis, then the matroid $\mathtt{mat}(\mathcal{M}, \mathcal{A})$ of all $\mathcal{A}$-bases is a sum of minors of $\mathcal{M}$. In particular, if $\mathcal{M}$ belongs to some minor closed class of matroids (e.g., graphic, linear, gammoid) then so does $\mathtt{mat}(\mathcal{M}, \mathcal{A})$.*

4 Feasible Queries

▶ **Definition 19.** *A set $F \subseteq E$ is a feasible query (or simply feasible) if no matter its revelation it guarantees the existence of a uniformly optimal basis. That is, $\forall \mathcal{B} \in \mathcal{R}(F, \mathcal{A})$ there exists some $\mathcal{B}$-basis.*

Any superset of a feasible query is also feasible. Using this, one can show that feasible sets for a given uncertainty area function are also feasible for any revelation of a subset.

▶ **Lemma 20.** *Let $F \subseteq E$ be feasible in $(\mathcal{M}, \mathcal{A})$ and $X \subseteq E$. If $\mathcal{B} \in \mathcal{R}(X, \mathcal{A})$, then F is feasible for $(\mathcal{M}, \mathcal{B})$.*

Proof. Since F is feasible in $(\mathcal{M}, \mathcal{A})$, so is $X \cup F$. Consider any revelation $\mathcal{B}' \in \mathcal{R}(X, \mathcal{B})$. Since $\mathcal{B}' \in \mathcal{R}(X \cup F, \mathcal{A})$ and $X \cup F$ is feasible, we conclude that there exists a $\mathcal{B}'$-basis. ◀

Since revealing elements that are certain does not yield extra information, we also get that all minimal (for inclusion) queries only contain uncertain elements. A simple, yet strong result is that it never pays off to query a colored element. Due to space consideration, we defer the proof of the next lemma to the full version.

▶ **Lemma 21.** *Let $X \subseteq E$ a feasible query. If $e \in X$ is colored, then $X - e$ is a feasible query. In particular, if X is a feasible query minimal for inclusion then X consists only of uncertain uncolored elements.*

An exciting application of this lemma is that the set of all uncolored uncertain elements is always a feasible query set. To see this, start with all the uncertain elements (which are a feasible query) and repeatedly remove colored elements while applying Lemma 21.

By Theorem 12, a set F is feasible only if after its revelation all uncertain elements are colored. So, consider any uncertain element e before any revelation. Intuitively, its color depends on the possible relative position between its real weight w_e and the real weight of elements that could potentially span it (or cospan it). In particular, it is not hard too see that the *color* of e is unaffected if we reveal an element f whose uncertainty area is too low (say $U_f \leq L_e$), because the real value of f will be in every case at most that of e. A similar situation happens if the uncertainty area is too high (say $U_e \leq L_f$). A complicated thing occurs if $\mathcal{A}(f)$ intersects (L_e, U_e), because after revealing f we may still don't know the relative positions of w_e and w_f until we reveal e. Finally, if f is in none of the previous situation, then after revealing it we will know for sure the relative position of w_e, w_f even without revealing e. The previous discussion motivates the following definitions.

▶ **Definition 22.** *For each $e \in E$ define the sets $\mathtt{low}(e)$, $\mathtt{mid}(e)$, $\mathtt{high}(e)$ and $\mathtt{both}(e)$ by:*

$$\mathtt{low}(e) = \{f \in E - e : U_f \leq L_e\}, \qquad \mathtt{high}(e) = \{f \in E - e : U_e \leq L_f\},$$
$$\mathtt{mid}(e) = \{f \in E - e : \mathcal{A}(f) \cap (L_e, U_e) \neq \emptyset\},$$
$$\mathtt{both}(e) = \{f \in E \setminus \mathtt{mid}(e) - e : \mathcal{A}(f) \cap (-\infty, L_e] \neq \emptyset \ \wedge \ \mathcal{A}(f) \cap [U_e, \infty) \neq \emptyset\}.$$

Note that that for $e \in E$, $F(e) = (E - e) \setminus \mathtt{high}(e)$ and $F^*(e) = (E - e) \setminus \mathtt{low}(e)$. Furthermore if e is uncertain (i.e., $L_e < U_e$) then $E - e$ is partitioned into the sets $\mathtt{low}(e), \mathtt{high}(e), \mathtt{mid}(e)$ and $\mathtt{both}(e)$.

▶ **Definition 23.** *For each $e \in E$ denote $\mathcal{M} / \mathtt{low}(e) \setminus \mathtt{high}(e)$ by $\mathcal{M}'_e$*

▶ **Remark 24.** In this section we talk about different revelations simultaneously (for instance, $\mathcal{A}$ and $\mathcal{B}$). We differentiate the objects that arise this way by using superscript denoting these dependencies. For example $F^{\mathcal{A}}(e)$ denotes $F(e)$ with respect to the areas given by $\mathcal{A}$.

The next technical lemma formalizes the idea that the elements that influence the color of an uncertain element e are those in $\mathtt{mid}(e)$ and those in $\mathtt{both}(e)$ that are not queried.

▶ **Lemma 25.** *Let $X \subseteq E$ and $\mathcal{B} \in \mathcal{R}(E \setminus X, \mathcal{A})$ a revelation of its complement. If $e \in X$ is uncertain and uncolored in $\mathcal{B}$, then there exists a circuit C in $\mathcal{M}'_e$ such that $e \in C$ and $(C - e) \cap [\mathtt{mid}^{\mathcal{A}}(e) \cup (X \cap \mathtt{both}^{\mathcal{A}}(e))] \neq \emptyset$.*

Proof. Consider the revelation $\widetilde{\mathcal{B}} \in \mathcal{R}(\mathtt{both}(e) \setminus X, \mathcal{A})$ such that $\widetilde{\mathcal{B}}(f) = \mathcal{B}(f)(= \{B_f\})$ if $f \in \mathtt{both}(e) \setminus X$ and $\widetilde{\mathcal{B}}(f) = \mathcal{A}(f)$ otherwise and note that $\mathcal{B} \in \mathcal{R}(E \setminus X, \widetilde{\mathcal{B}})$. By Lemma 7, since e is uncolored in $(\mathcal{M}, \mathcal{B})$, we get that e is also uncolored in $(\mathcal{M}, \widetilde{\mathcal{B}})$. Define the auxiliary sets $Y = \{f \in \mathtt{both}^{\mathcal{A}}(e) \setminus X : B_f \leq L_e\}$ and $\overline{Y} = \{f \in \mathtt{both}^{\mathcal{A}}(e) \setminus X : B_f > L_e\}$. Note that $\mathtt{low}^{\widetilde{\mathcal{B}}}(e) = \mathtt{low}^{\mathcal{A}}(e) \cup Y; \mathtt{high}^{\widetilde{\mathcal{B}}}(e) = \mathtt{high}^{\mathcal{A}}(e) \cup \overline{Y}; \mathtt{mid}^{\widetilde{\mathcal{B}}}(e) = \mathtt{mid}^{\mathcal{A}}(e)$; and $\mathtt{both}^{\widetilde{\mathcal{B}}}(e) = X \cap \mathtt{both}^{\mathcal{A}}(e)$. Since e is uncertain and it is not blue nor red in $(\mathcal{M}, \widetilde{B})$ we get: $e \in \mathrm{span}[\mathtt{low}^{\widetilde{\mathcal{B}}}(e) \cup \mathtt{mid}^{\widetilde{\mathcal{B}}}(e) \cup \mathtt{both}^{\widetilde{\mathcal{B}}}(e)] = \mathrm{span}[\mathtt{low}^{\mathcal{A}}(e) \cup \mathtt{mid}^{\mathcal{A}}(e) \cup Y \cup (X \cap \mathtt{both}^{\mathcal{A}}(e))]$, and $e \in \mathrm{cospan}[\mathtt{high}^{\widetilde{\mathcal{B}}}(e) \cup \mathtt{mid}^{\widetilde{\mathcal{B}}}(e) \cup \mathtt{both}^{\widetilde{\mathcal{B}}}(e)] = \mathrm{cospan}[\mathtt{high}^{\mathcal{A}}(e) \cup \mathtt{mid}^{\mathcal{A}}(e) \cup \overline{Y} \cup (X \cap \mathtt{both}^{\mathcal{A}}(e))]$. Using that $e \in \mathrm{cospan}[Q]$ implies $e \notin \mathrm{span}[(E - e) \setminus Q]$ for any $Q \subseteq E$, we conclude:

$$e \in \mathrm{span}[\mathtt{low}^{\mathcal{A}}(e) \cup \mathtt{mid}^{\mathcal{A}}(e) \cup Y \cup (X \cap \mathtt{both}^{\mathcal{A}}(e))] \setminus \mathrm{span}[\mathtt{low}^{\mathcal{A}}(e) \cup Y]. \tag{1}$$

From (1), $e \in \mathrm{span}[\mathtt{low}^{\mathcal{A}}(e) \cup \mathtt{mid}^{\mathcal{A}}(e) \cup Y \cup (X \cap \mathtt{both}^{\mathcal{A}}(e))] \setminus \mathtt{low}^{\mathcal{A}}(e) = \mathrm{span}_{M'_e}[\mathtt{mid}^{\mathcal{A}}(e) \cup Y \cup (X \cap \mathtt{both}^{\mathcal{A}}(e))]$, where the equality follows from properties of contraction and deletion. Therefore there is a circuit C in $\mathcal{M}'_e$ such that $e \in C$ and $C - e \subseteq \mathtt{mid}^{\mathcal{A}}(e) \cup Y \cup X \cap \mathtt{both}^{\mathcal{A}}(e)$. If $(C - e) \cap [\mathtt{mid}^{\mathcal{A}}(e) \cup (X \cap \mathtt{both}^{\mathcal{A}}(e))] = \emptyset$, we would have that $C - e \subseteq Y$, implying that $e \in \mathrm{span}_{\mathcal{M}'_e} Y = \mathrm{span}[\mathtt{low}^{\mathcal{A}}(e) \cup Y] \setminus \mathtt{low}^{\mathcal{A}}(e)$ which contradicts (1). ◀

The previous lemma is useful to characterize the sets that intersect every feasible set.

▶ **Definition 26.** *A set $X \subseteq E$ is a witness set if it intersects every feasible set.*

Witness sets have been studied before in the context of online adaptive algorithms for MST and matroids. Since the definition of feasible is slightly different in that settings (they are feasible for the verification problem, in which one knows the real values a priori), these witness sets are also different from ours.

▶ **Lemma 27.** *Let $X \subseteq E$. The following statements are equivalent:*
  (i) *X is a witness set.*
  (ii) *There exists an uncertain element $e \in X$ and a circuit C in $\mathcal{M}'_e$ such that $e \in C$ and $C \cap [\mathit{mid}(e) \cup (X \cap \mathit{both}(e))] \neq \emptyset$.*

Proof. Let X be a witness set. Since $E \setminus X$ isn't feasible there is a revelation $\mathcal{B} \in \mathcal{R}(E \setminus X, \mathcal{A})$ such that there is no $(\mathcal{M}, \mathcal{B})$-basis and by Theorem 12 we must have an element e that is uncertain and uncolored in $(\mathcal{M}, \mathcal{B})$. Note that $e \in X$ as every element in $E \setminus X$ is certain in $(\mathcal{M}, \mathcal{B})$. We conclude by using Lemma 25 on e.

For the converse, set $Y = [\mathtt{both}(e) \setminus X] \cap C$, $\overline{Y} = [\mathtt{both}(e) \setminus X] \setminus C$ and note that $e \in \mathrm{span}_{\mathcal{M}'_e}(C - e) = \mathrm{span}_{\mathcal{M}'_e}[(C \cap \mathtt{mid}(e)) \cup (C \cap (\mathtt{both}(e) \setminus X)) \cup (C \cap \mathtt{both}(e) \cap X)] \subseteq \mathrm{span}_{\mathcal{M}'_e}[\mathtt{mid}(e) \cup Y \cup (X \cap \mathtt{both}(e))]$.

If $e \in \mathrm{span}_{\mathcal{M}'_e}(Y)$ there would be a circuit D in $\mathcal{M}'_e$ such that $D - e \subseteq Y \subseteq C$, but since $C \cap [\mathtt{mid}(e) \cup (X \cap \mathtt{both}(e))] \neq \emptyset$ we get $D \subsetneq C$ which contradicts the minimality of C as circuit. Then, $e \in \mathrm{span}_{\mathcal{M}'_e}[\mathtt{mid}(e) \cup Y \cup (X \cap \mathtt{both}(e))] \setminus \mathrm{span}_{\mathcal{M}'_e} Y$, which is included in $\mathrm{span}[\mathtt{low}(e) \cup \mathtt{mid}(e) \cup Y \cup (X \cap \mathtt{both}(e))] \setminus \mathrm{span}[\mathtt{low}(e) \cup Y]$.

As $e \notin \mathtt{high}(e) \cup \mathtt{mid}(e) \cup \overline{Y} \cup (X \cap \mathtt{both}(e))$ we conclude that $e \in \mathrm{span}[\mathtt{low}(e) \cup \mathtt{mid}(e) \cup Y \cup (X \cap \mathtt{both}(e))]$ and $e \in \mathrm{cospan}[\mathtt{high}(e) \cup \mathtt{mid}(e) \cup \overline{Y} \cup (X \cap \mathtt{both}(e))]$.

Choose two revelations $w^+, w^- \in \mathcal{R}(E, \mathcal{A})$ as follows:

$$w_f^+ \in \begin{cases} (L_e, U_e) \cap \mathcal{A}(f) & \text{if } f \in \mathtt{mid}(e), \\ (-\infty, L_e] \cap \mathcal{A}(f) & \text{if } f \in Y \text{ or} \\ & \quad f \in X \cap \mathtt{both}(e), \\ [U_e, \infty) \cap \mathcal{A}(f) & \text{if } f \in \overline{Y}, \\ \mathcal{A}(e) \cap (L_e, U_e] & \text{if } f = e. \end{cases} \qquad w_f^- \in \begin{cases} [U_e, \infty) \cap \mathcal{A}(f) & \text{if } f \in X \cap \mathtt{both}(e), \\ \mathcal{A}(e) \cap [L_e, U_e) & \text{if } f = e, \\ \{w_f^+\} & \text{otherwise.} \end{cases}$$

Note that w^+ and w^- only differ on $X \cap \mathtt{both}(e)$ and e. As $e \in \mathrm{span}[\mathtt{low}(e) \cup \mathtt{mid}(e) \cup Y \cup (X \cap \mathtt{both}(e))]$ it is the unique heaviest element in a circuit in $(\mathcal{M}, w^+)$, therefore it is in no $(\mathcal{M}, w^+)$-basis. Similarly, since $e \in \mathrm{cospan}[\mathtt{high}(e) \cup \mathtt{mid}(e) \cup \overline{Y} \cup (X \cap \mathtt{both}(e))]$ it is the unique lightest element in a cocircuit in $(\mathcal{M}, w^-)$, hence it is in every $(\mathcal{M}, w^-)$-basis. To conclude suppose that there is a feasible query set F such that $X \cap F = \emptyset$, we then pick the revelation $\mathcal{B} \in \mathcal{R}(F, \mathcal{A})$ such that $\mathcal{B}(f) = \{w_f^+\} = \{w_f^-\}$ if $f \in F$, and $\mathcal{B}(f) = \mathcal{A}(f)$ otherwise.

Select any $\mathcal{B}$-basis T. As $w^+, w^- \in \mathcal{R}(E, \mathcal{B})$, T is both a w^+-basis and a w^--basis. By the previous paragraph, this implies that $e \notin T$ and $e \in T$ which is a contradiction. ◄

▶ **Lemma 28.** *The minimal feasible queries are the sets intersecting every witness set.*

Proof. Recall that a *clutter* is a family of sets such that no one is contained in another. The *blocker* of a clutter is the clutter of all minimal sets that intersect the first one. Thus, the minimal witness sets are the blocker of the minimal feasible queries. A basic result in packing and covering theory (see e.g., [15, Theorem 77.1]) states that the blocker of the blocker of a clutter is again the original clutter. The lemma follows from this fact. ◄

As we see below, minimal witness sets cannot be large: they can have at most 2 elements.

▶ **Corollary 29.** *Let X be a witness set such that $|X| \geq 2$. Then, there exists distinct $e, f \in X$ such that $\{e, f\}$ is a witness set.*

Proof. If $X \cap F \neq \emptyset$ for every feasible query set F then, by Lemma 27, there exists an uncertain $e \in X$, a circuit C in $\mathcal{M}'_e$ such that $e \in C$ and $C \cap [\mathtt{mid}(e) \cup (X \cap \mathtt{both}(e))] \neq \emptyset$.

If $C \cap (X \cap \mathtt{both}(e)) = \emptyset$, then $C \cap \mathtt{mid}(e) \neq \emptyset$ and by Lemma 27 we have that $\{e\} \cap F \neq \emptyset$ for every F feasible query set. Picking any $f \in X - e$ we conclude that $\{e, f\} \cap F \neq \emptyset$ for every feasible query set F. If $C \cap (X \cap \mathtt{both}(e)) \neq \emptyset$, select any $g \in C \cap (X \cap \mathtt{both}(e))$. Once again, Lemma 27 lets us conclude that $\{e, g\} \cap F \neq \emptyset$ for every F feasible query set. ◄

▶ **Definition 30.** *Let $\mathtt{core} = \{e \in E \colon \{e\}$ is a witness set$\}$ and $\overline{\mathtt{core}} = E \setminus \mathtt{core}$. Define the graph $G^{wit} = (\overline{\mathtt{core}}, E^{wit})$ where $ef \in E^{wit}$ if $\{e, f\}$ is a witness set.*

By Lemma 28, minimal feasible sets are exactly those sets containing all elements in $\mathtt{core}$ together with a vertex cover of G^{wit}. The following clean characterization of $\mathtt{core}$ follows directly from Lemma 27.

▶ **Lemma 31.** *Let e be an uncertain element. $e \in \mathtt{core}$ if and only if there is a circuit C in $\mathcal{M}'_e$ such that $e \in C$ and $C \cap \mathtt{mid}(e) \neq \emptyset$.*

We can turn the previous lemma into an algorithm that computes $\mathtt{core}$. In order to do this we compute the connected component[4] of $\mathcal{M}'_e$ that contains e and check if it has non-empty intersection with $\mathtt{mid}(e)$. We compute connected components with an algorithm due to Krogdhal [12] that takes polynomial time. Therefore, the previous procedure also takes polynomial time. In what follows we show that G^{wit} has a very nice structure.

▶ **Lemma 32.**
  (i) *Let $e, f \in \overline{\mathtt{core}}$ distinct. $ef \in E^{wit}$ if and only if $\mathcal{A}(e) = \mathcal{A}(f) = \{L_e, U_e\}$ with $L_e < U_e$ and there is a circuit C of $\mathcal{M}'_e$ such that $e, f \in C$.*
  (ii) *The connected components of the graph G^{wit} are cliques.*

---

[4] Recall that e is connected to f in a matroid if and only if there is a circuit that contains e and f. A connected component is an equivalence class of this equivalence relation (see, eg. [14, Section 4.1]).

Proof.

(i) Suppose $ef \in E^{\mathrm{wit}}$. Since $\{e, f\}$ is a witness set, $E \setminus \{e, f\}$ is not feasible. Let $\mathcal{B} \in \mathcal{R}(E \setminus \{e, f\}, \mathcal{A})$ be a revelation without $\mathcal{B}$-basis. Since $f \notin \mathtt{core}$, we have that $E - f$ is feasible in $(\mathcal{M}, \mathcal{A})$ and by Lemma 20 it is also feasible in $(\mathcal{M}, \mathcal{B})$. If e was colored in $(\mathcal{M}, \mathcal{B})$ then, by Lemma 21 we would conclude that $(E - f) - e = E \setminus \{e, f\}$ is feasible in $(\mathcal{M}, \mathcal{B})$. But all elements in $E \setminus \{e, f\}$ are already certain in $(\mathcal{M}, \mathcal{B})$, from which we deduce that $\emptyset$ is feasible in $(\mathcal{M}, \mathcal{B})$, which is a contradiction.

We conclude that e is uncolored in $(\mathcal{M}, \mathcal{B})$. From this, we can use Lemma 25 to obtain a circuit C in $\mathcal{M}'_e$ such that $e \in C$ and $C \cap [\mathtt{mid}^{\mathcal{A}}(e) \cup (\{e, f\} \cap \mathtt{both}^{\mathcal{A}}(e))] \neq \emptyset$. Since $e \notin \mathtt{core}$, we have by Lemma 31 that $C \cap \mathtt{mid}(e) = \emptyset$. Then $C \cap \{e, f\} \cap \mathtt{both}(e) \neq \emptyset$, consequently $e, f \in C$ and $f \in \mathtt{both}(e)$. We can now use the same argument for f, concluding that $e \in \mathtt{both}(f)$. The only way for $e \in \mathtt{both}(f)$ and $f \in \mathtt{both}(e)$ to occur at the same time is that $\mathcal{A}(e) = \mathcal{A}(f) = \{L_e, U_e\}$. Finally, since witness sets do not contain elements that are certain (this follows since minimal feasible sets only have uncertain elements, hence by removing certain elements from a witness set it would still intersect every feasible set) we must have $L_e < U_e$.

We now prove the converse. As e is an uncertain element such that there is a circuit C in $\mathcal{M}'_e$ and $f \in C \cap \{e, f\} \cap \mathtt{both}(e)$ by Lemma 27 we conclude that $ef \in E^{\mathrm{wit}}$.

(ii) We only need to show that whenever $ef, fg \in E^{\mathrm{wit}}$ we also have $eg \in E^{\mathrm{wit}}$. Suppose that $ef, fg \in E^{\mathrm{wit}}$. By the previous item we have $\mathcal{A}(e) = \mathcal{A}(f) = \mathcal{A}(g) = \{L_e, U_e\}$, in particular $\mathcal{M}'_e = \mathcal{M}'_f = \mathcal{M}'_g \doteq \mathcal{M}'$. The previous item also allows us to conclude that there are two circuits C^1, C^2 in $\mathcal{M}'$ such that $e, f \in C^1$ and $f, g \in C^2$. Then e, f and g are in the same matroid connected component in $\mathcal{M}'$. Therefore there is a circuit C^3 in $\mathcal{M}'$ such that $e, g \in C^3$ and using the previous item we conclude that $eg \in E^{\mathrm{wit}}$. ◀

We can test if $ef \in E^{\mathrm{wit}}$ similarly to how we computed $\mathtt{core}$: we start by considering elements with areas of size two (by checking if $(L_e, U_e) \cap \mathcal{A}(e) = \emptyset$ for every element e). If e and f have the same two-element uncertainty area, we check if they belong to the same connected component in $\mathcal{M}'_e = \mathcal{M}'_f$ using Krogdhal's algorithm [12].

▶ **Theorem 33.** $X \subseteq E$ *is a minimal feasible query if and only if* $\mathtt{core} \subseteq X$ *and* X *intersects all but one element in each connected component of* G^{wit}.

Proof. By Lemma 28 and the definition of G^{wit} the minimal feasible queries X satisfies that $\mathtt{core} \subseteq X$ and $X \cap \overline{\mathtt{core}}$ is a minimal vertex cover of G^{wit}. Since every connected component is a clique, the minimal vertex covers of G^{wit} are exactly those sets containing all but one element in each connected component. ◀

▶ **Corollary 34.** *Let* $c : E \to \mathbb{R}$ *be any cost function. We can compute a minimum-cost feasible query in polynomial time.*

Proof. Computing $\mathtt{core}$ and G^{wit} can be done in polynomial time and polynomial number of calls to the independence oracle of $\mathcal{M}$ or to minors of $\mathcal{M}$ (for example, $\mathcal{M}'_e$), since the oracle of independence of any minor of $\mathcal{M}$ can also be implemented using a polynomial number of calls to the oracle of $\mathcal{M}$. One can compute a minimum-cost minimal size feasible query F by simply returning a set containing $\mathtt{core}$ and all but the most expensive element from each connected component of G^{wit}. If we allow negative costs, then the minimum-cost feasible query is F together with all the negative cost elements outside F. An efficient implementation is discussed in the full version of this paper. ◀

We finish our study with two special cases of the last theorem.

▶ **Corollary 35.** *Suppose that $\mathcal{A}(e)$ is an interval (of any type: open, closed, semiopen, trivial) for every element $e \in E$. Then the set S of all uncolored uncertain elements is the only minimum-size feasible query.*

▶ **Corollary 36.** *Let G be a connected graph such that for every $e \in E(G)$, $\mathcal{A}(e) = \{0, 1\}$. A set $F \subseteq E(G)$ is a minimal feasible query for the MST problem if and only if F contains all but one edge from each 2-connected component of G.*

References

1 Richard Bruce, Michael Hoffmann, Danny Krizanc, and Rajeev Raman. Efficient Update Strategies for Geometric Computing with Uncertainty. *Theory Comput. Syst.*, 38(4):411–423, 2005. `doi:10.1007/s00224-004-1180-4`.

2 George Charalambous and Michael Hoffmann. Verification Problem of Maximal Points under Uncertainty. In Thierry Lecroq and Laurent Mouchard, editors, *Combinatorial Algorithms - 24th International Workshop, IWOCA 2013, Rouen, France, July 10-12, 2013, Revised Selected Papers*, volume 8288 of *Lecture Notes in Computer Science*, pages 94–105. Springer, 2013. `doi:10.1007/978-3-642-45278-9_9`.

3 Thomas Erlebach and Michael Hoffmann. Query-competitive algorithms for computing with uncertainty. *Bulletin of EATCS*, 2(116), 2015.

4 Thomas Erlebach, Michael Hoffmann, and Frank Kammer. Query-competitive algorithms for cheapest set problems under uncertainty. *Theor. Comput. Sci.*, 613:51–64, 2016. `doi:10.1016/j.tcs.2015.11.025`.

5 Thomas Erlebach, Michael Hoffmann, Danny Krizanc, Matús Mihal'ák, and Rajeev Raman. Computing Minimum Spanning Trees with Uncertainty. In Susanne Albers and Pascal Weil, editors, *25th International Symposium on Theoretical Aspects of Computer Science*, volume 1 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 277–288, Dagstuhl, Germany, February 2008. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. `doi:10.4230/LIPIcs.STACS.2008.1358`.

6 Tomás Feder, Rajeev Motwani, Liadan O'Callaghan, Chris Olston, and Rina Panigrahy. Computing shortest paths with uncertainty. *J. Algorithms*, 62(1):1–18, 2007. `doi:10.1016/j.jalgor.2004.07.005`.

7 Tomás Feder, Rajeev Motwani, Rina Panigrahy, Chris Olston, and Jennifer Widom. Computing the Median with Uncertainty. *SIAM J. Comput.*, 32(2):538–547, 2003. `doi:10.1137/S0097539701395668`.

8 Manoj Gupta, Yogish Sabharwal, and Sandeep Sen. The Update Complexity of Selection and Related Problems. *Theory Comput. Syst.*, 59(1):112–132, 2016. `doi:10.1007/s00224-015-9664-y`.

9 Simon Kahan. A Model for Data in Motion. In *Proceedings of the Twenty-third Annual ACM Symposium on Theory of Computing*, STOC '91, pages 265–277, New York, NY, USA, 1991. ACM. `doi:10.1145/103418.103449`.

10 Adam Kasperski and Pawel Zielinski. An approximation algorithm for interval data minmax regret combinatorial optimization problems. *Inf. Process. Lett.*, 97(5):177–180, 2006. `doi:10.1016/j.ipl.2005.11.001`.

11 Adam Kasperski and Pawel Zielinski. On combinatorial optimization problems on matroids with uncertain weights. *European Journal of Operational Research*, 177(2):851–864, 2007. `doi:10.1016/j.ejor.2005.12.033`.

12 Stein Krogdahl. The dependence graph for bases in matroids. *Discrete Mathematics*, 19(1):47–59, 1977. `doi:10.1016/0012-365X(77)90118-2`.

13 Nicole Megow, Julie Meißner, and Martin Skutella. Randomization Helps Computing a Minimum Spanning Tree under Uncertainty. *SIAM Journal on Computing*, 46(4):1217–1240, January 2017. `doi:10.1137/16M1088375`.

14 James G. Oxley. *Matroid theory*. Oxford University Press, 1992.

15 Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer Science & Business Media, 2003.

16 Hande Yaman, Oya Ekin Karasan, and Mustafa Ç. Pinar. The robust spanning tree problem with interval data. *Oper. Res. Lett.*, 29(1):31–40, 2001. `doi:10.1016/S0167-6377(01)00078-5`.