On Tractable Problems of Diversity Optimization

Tesshu Hanaka[1] Yasuaki Kobayashi[2] Kazuhiro Kurita[3] See Woo Lee[2]
Yota Otachi[1]

Abstract: Finding diverse solutions in combinatorial problems recently has received some attention. In this paper we study the following type of problems: given an integer k, the problem asks for k solutions such that the sum of pairwise Hamming distances between these solutions is maximized. We investigate the tractability of the "diverse version" of several classical combinatorial problems, such as finding bases of matroids, arborescences in directed graphs, bipartite matchings, shortest st-paths in directed graphs, and minimum cuts of undirected graphs.

1. Introduction

In many combinatorial problems, we usually seek a *single* solution satisfying some prescribed constraints and/or optimizing given an objective function. However, such a solution may not be adequate for real-world problems since several intricate constraints emerging in real-world problems are overly simplified or even ignored to make those problem amenable. To address this issue, seeking *multiple* solutions is a straightforward but promising approach. One of the best known approaches to do this is *k-best enumeration* [9]. Here, an algorithm is called a k-best enumeration algorithm for some optimization problem if given an integer k, the algorithm finds k feasible solutions $\mathcal{S} = \{S_1, \dots, S_k\}$ such that every feasible solution not in $\mathcal{S}$ is not strictly better than that in $\mathcal{S}$. There are many k-best enumeration algorithms for various optimization problems (see [9] for a survey). One potential drawback of k-best enumeration algorithms is the lack of *diversity* of solutions. Most of k-best enumeration algorithms, such as Lawler's framework [18], recursively generate solutions from a single optimal solution $X = \{x_1, x_2, \dots, x_t\}$ by finding a solution including $\{x_1, \dots, x_{i-1}\}$ and excluding x_i for each $1 \leq i \leq t$. This implies that solutions tend to be similar to each other in nature.

Motivated by this, (explicitly) optimizing diversity of solutions has received considerable attention in the literature. There are many results for finding "diverse" CSP or MIP solutions [6], [15], [19], [22], [23]. According to [3], Michael Fellows proposed *the Diverse X Paradigm*, where X is a placeholder for an optimization problem. Based on this proposal, they studied the parameterized complexity of several diverse versions of combinatorial problems, such as Vertex Cover, Feedback Vertex Set, and d-Hitting Set, and showed that these problems are fixed-parameter tractable parameterized by the solution size plus the number of solutions [3]. Baste et al. [2] also discussed the fixed-parameter tractability of diverse versions of several combinatorial problems on bounded-treewidth graphs.

Before describing our results, we need to define known diversity measures and discuss known results relevant to our results. There are mainly two diversity measures in these theoretical studies. Let U be a finite set. Let $S_1, \dots, S_k$ be (not necessarily disjoint) subsets of U. We define

$$d(S_1, \dots, S_k) = \sum_{1 \leq i < j \leq k} |S_i \triangle S_j|,$$

where $\triangle$ is the symmetric difference of two sets. Also, define

$$d_{\min}(S_1, \dots, S_k) = \min_{1 \leq i < j \leq k} |S_i \triangle S_j|.$$

Fomin et al. [10] showed that the problem of finding two maximum matchings M_1, M_2 maximizing its symmetric difference in bipartite graphs can be solved in polynomial time, whereas it is NP-hard on general graphs and gave an FPT-algorithm with respect to parameter $|M_1 \triangle M_2|$. Also, Fomin et al. [11] gave FPT-algorithms for finding k solutions for several problems related to matroids and matchings such that the weighted symmetric difference between any pair of them is at least d, parameterized by $k + d$ (i.e., the running time of these algorithms is $f(k, d)n^{O(1)}$, where f is some computable function and n is the input size). In particular, they showed that finding k bases of a matroid maximizing the weighted version of $d_{\min}$ is NP-hard even on uniform matroids. Contrary to this hardness result, Hanaka et al. [14] showed that finding k bases of a matroid maximizing d is solvable in polynomial time.[*1]

[1] Nagoya University
[2] Kyoto University
[3] National Institute of Informatics

[*1] In this problem setting, we assume that the independent oracle can be evaluated in polynomial time.

In this paper, we expand the tractability border of the diverse version of classical combinatorial optimization problems.

Let $\mathcal{S} \subseteq 2^U$ be a set of solutions. The common goal of our problems is to find a set of k solutions $S_1, S_2, \dots, S_k \in \mathcal{S}$ maximizing $d_w(S_1, \dots. S_k)$, where d_w is the weighted version of the diversity measure d (see Section 2). We show that if $\mathcal{S}$ consists of either (1) the bases of a matroid, (2) the arborescences of a directed graphs, (3) the set of t-matchings of a bipartite graph, or (4) the set of st-paths in a directed graph, then the problem can be solved in polynomial time. The algorithm for (1) is a straightforward extension of the algorithm of [14] to the weighted Hamming distance and that for (3) is an extension of [10] that allows to find more than two diverse bipartite matchings in polynomial time. On the negative side, we show that if $\mathcal{S}$ consists of the set of minimum cut of a graph, then the problem is NP-hard even if the size of a minimum cut is three. To tackle this intractability, we study Diverse Minimum Cuts from the perspective of fixed-parameter tractability. We show that Diverse Minimum Cuts is polynomial-time solvable if the input graph G has a minimum cut of size at most two, fixed-parameter tractable parameterized by k plus the edge-connectivity of G, and W[1]-hard parameterized by k only.

2. Preliminaries

Let $G = (V, E)$ be a (directed) graph. We denote by $V(G)$ and $E(G)$ the sets of vertices and edges of G, respectively. For $X \subseteq V$, the set of edges between X and $V \setminus X$ is denoted by $E_G(X, V \setminus X)$.

Let U be a finite set and let $w : U \to \mathbb{N}_{\geq 0}$. Let $S_1, \dots, S_k$ be (not necessarily disjoint) subsets of U. We define

$$d_w(S_1, \dots, S_k) = \sum_{1 \leq i < j \leq k} w(S_i \triangle S_j),$$

where $w(X) = \sum_{x \in X} w(x)$. This notation extends the diversity measure d defined in the previous section.

Let k be a positive integer and let $\mathcal{S} \subseteq 2^U$ be subsets of U. We expand each element e in U into k copies: Let $U^* = \{e_1, \dots, e_k : e \in U\}$. We define a function $f : U^* \to U$ such that $f(e_i) = e$ for all $e_i \in U^*$. We say that $S^* \subseteq U^*$ is a *k-packing of U^* with respect to $\mathcal{S}$* if S^* can be partitioned into $S_1, \dots, S_k$ such that $\{f(e^*) : e^* \in S_i\} \in \mathcal{S}$ for all $1 \leq i \leq k$.

We consider a weight function $w^* : U^* \to \mathbb{Z}$ such that $w^*(e_i) = w(e) \cdot (k - 2i + 1)$ for $e \in U$ and $1 \leq i \leq k$. Then, we have the following lemma.

Lemma 1. There are $S_1, \dots, S_k \in \mathcal{S}$ with $d_w(S_1, \dots, S_k) \geq t$ if and only if there is a k-packing S^* of U^* with respect to $\mathcal{S}$ such that $w^*(S^*) \geq t$.

Proof. Suppose that there are $S_1, \dots, S_k \in \mathcal{S}$ with $d_w(S_1, \dots, S_k) \geq t$. For each $e \in U$, we denote by $m(e)$ the number of occurrences of e in the collection $\{S_1, \dots, S_k\}$. Then, we have

$$\begin{aligned} d_w(S_1, \dots, S_k) &= \sum_{e \in U} (w(e) \cdot m(e) \cdot (k - m(e))) \\ &= \sum_{e \in U} (w(e) \cdot \sum_{1 \leq i \leq m(e)} (k - 2i + 1)) \\ &= \sum_{e \in U} \sum_{1 \leq i \leq m(e)} w^*(e_i) \\ &= w^*(S^*). \end{aligned}$$

Conversely, assume that there is a k-packing S^* of U^* with respect to $\mathcal{S}$ with $w^*(S^*) \geq t$. We assume moreover that, for each $e \in U$, S^* contains consecutive elements $e_1, \dots, e_m$ for some m, that is, $\{e_1, \dots, e_m\} \subseteq S^*$ and $\{e_{m+1} \dots, e_k\} \cap S^* = \emptyset$. This assumption is legitimate as $w(e_i) > w(e_j)$ for $1 \leq i < j \leq k$. We denote the multiplicity m of e by $m(e)$. Let $\{S_1^*, \dots, S_k^*\}$ be a partition of S^* such that $S_i = \{f(e^*) : e^* \in S_i^*\} \in \mathcal{S}$ for $1 \leq i \leq k$. For each $e \in U$, the contribution of e to $w^*(S^*)$ is indeed $w(e) \cdot \sum_{1 \leq i \leq m(e)} (k - 2i + 1)$. Hence, we have

$$\begin{aligned} w^*(S^*) &= \sum_{e \in U} \sum_{1 \leq i \leq m(e)} w^*(e_i) \\ &= d_w(S_1, \dots, S_k) \end{aligned}$$

as in the "only-if" direction. □

3. Diverse Matroid Bases

As an application of Lemma 1, we consider the following problem.

Definition 1. Given a matroid $\mathcal{M} = (E, \mathcal{I})$ with a weight function $w : E \to \mathbb{N}_{\geq 0}$ and an integer k, Weighted Diverse Matroid Bases asks for k bases $B_1, \dots, B_k$ of $\mathcal{M}$ such that $d_w(B_1, \dots, B_k)$ is maximized.

In [14], they consider a special case of Weighted Diverse Matroid Bases where each element in the ground set E has a unit weight and give a polynomial-time algorithm for it, assuming that the independence oracle $\mathcal{I}$ can be evaluated in polynomial time. This result is obtained by reducing the problem to that of finding disjoint bases of a matroid, which can be solved in polynomial time.

Theorem 1 ([7], [21]). Let $\mathcal{M} = (E, \mathcal{I})$ be a matroid and let $w : E \to \mathbb{Z}$. Suppose that the membership of $\mathcal{I}$ can be checked in polynomial time. Then, the problem of deciding whether there is a set of mutually disjoint k bases $B_1, \dots, B_k$ of $\mathcal{M}$ can be solved in polynomial time. Moreover, if the answer is affirmative, we can find such bases that maximize the total weight (i.e., $\sum_{1 \leq i \leq k} \sum_{e \in B_k} w(e)$) in polynomial time.

By applying Lemma 1, we have a polynomial-time algorithm for Weighted Diverse Matroid Bases as well.

Theorem 2. Weighted Diverse Matroid Bases can be solved in polynomial time.

Proof. The proof is almost analogous to that in [14]. Let $\mathcal{M} = (E, \mathcal{I})$ be a matroid and let $e \in E$. Define $\mathcal{J} = \mathcal{I} \cup \{(F \setminus \{e\}) \cup \{e'\} : F \in \mathcal{I} \wedge e \in F\}$. Then, $(E \cup \{e'\}, \mathcal{J})$

is also a matroid [14], [20]. We define k copies $e_1, e_2, \ldots, e_k$ for each $e \in E$ and $E^* = \{e_1, \ldots, e_k : e \in E\}$. Then, the pair $\mathcal{M}^* = (E^*, \mathcal{I}^*)$ is a matroid if $\mathcal{I}^*$ consists of all sets $F \subseteq 2^{E^*}$ such that F contains at most one copy of $e_1, \ldots, e_k$ for each $e \in E$ and $\bigcup_{e_i \in F} f(e_i) \in \mathcal{I}$, where $f(e_i) = e$ for $e \in E$ and $1 \leq i \leq k$.

To find a set of k bases $B_1, \ldots, B_k$ of $\mathcal{M}$ maximizing $d_w(B_1, \ldots, B_k)$, by Lemma 1, it suffices to find a maximum weight k-packing with respect to the base family of $\mathcal{M}^*$ under a weight function w^* with $w^*(e_i) = w(e) \cdot (k - 2i + 1)$ for $e \in E$ and $1 \leq i \leq k$, which can be solved in polynomial time by Theorem 1. □

4. Diverse Arborescences

Theorem 2 allows us to find diverse spanning trees in undirected graphs as the set of spanning trees of a graph forms the set of bases of a graphic matroid. In this section, we develop a polynomial-time algorithm for a directed version of this problem. Let $G = (V, E)$ be a directed graph and let $r \in V$. We say that a subgraph T of G is an *arborescence* (with root r) if for every $v \in V$, there is exactly one directed path from r to v in T. In other words, an arborescence is a spanning subgraph of G in which each vertex except r has in-degree one and its underlying undirected graph is a tree. In this section, we consider the following problem.

Definition 2. Given an arc-weighted directed graph $D = (V, A)$ with weight function $w : A \to \mathbb{N}_{\geq 0}$, $r \in V$, and an integer k, WEIGHTED DIVERSE ARBORESCENCES asks for k arborescences $T_1, \ldots, T_k$ of D with root r such that $d_w(E(T_1), \ldots, E(T_k))$ is maximized.

Theorem 3. WEIGHTED DIVERSE ARBORESCENCES can be solved in polynomial time.

The proof of Theorem 3 is almost analogous to that in Theorem 2. We define a directed graph D^* with vertex set V from $D = (V, A)$ such that for $e = (u, v) \in A$, we add k parallel arcs $e_1, \ldots, e_k$ directed from u to v to D'. Then, we set $w^*(e_i) = w(e) \cdot (k - 2i + 1)$ for $e \in A$ and $1 \leq i \leq k$. By Lemma 1, it is sufficient to find a maximum weight k-packing of the arc set of D^* with respect to the family of arborescences of D^*, which can be found in polynomial time by the following result.

Theorem 4 ([8]). Given an arc-weighted directed multigraph $D = (V, A)$ with weight function $w : A \to \mathbb{Z}$, $r \in V$, and an integer k, the problem of finding k arc disjoint arborescences $T_1, \ldots, T_k$ with root r maximizing $w(\bigcup_{1 \leq i \leq k} E(T_i))$ is solved in strongly polynomial time.

5. Diverse Bipartite Matchings

A *matching* of a graph $G = (V, E)$ is a set $M \subseteq E$ of edges such that no two edges share their end vertices. In this section, we consider the following problem.

Definition 3. Let $G = (A \cup B, E, w)$ be an edge-weighted bipartite graph with $w : E \to \mathbb{N}_{\geq 0}$, where A and B are color classes of G. Let k, p be positive integers. We denote by $\mathcal{M}$ the collection of all matchings of G with cardinality exactly p. DIVERSE BIPARTITE MATCHINGS asks for k matchings $M_1, \ldots, M_k \in \mathcal{M}$ maximizing $d_w(M_1, \ldots, M_k)$.

As mentioned in Section 1, the problem of finding two edge-joint perfect matchings in a general graph is known to be NP-complete [16]. In this section, we design a polynomial-time algorithm for DIVERSE BIPARTITE MATCHINGS by applying Lemma 1.

We construct a bipartite multigraph G^* from G by replacing each edge $e = \{a, b\} \in E$ with k parallel edges $e_1, \ldots, e_k$. We set $w^*(e_i) = w(e) \cdot (k - 2i + 1)$ for each $e \in E$ and $1 \leq i \leq k$. By Lemma 1, it suffices to show that there is a polynomial-time algorithm that computes a maximum weight k-packing of E^* with respect to $\mathcal{M}$, where $\mathcal{M}$ is the collection of matchings M of G with cardinality exactly p. This problem can be solved in polynomial time by reducing to the minimum cost flow problem as follows.

Let $G^* = (A \cup B, E^*)$ be bipartite and let $w^* : E^* \to \mathbb{Z}$. We construct a directed acyclic graph from G^* by orienting each edge $\{a, b\}$ of G^* directed from a to b, where $a \in A$ and $b \in B$. Each arc (a, b) in G^* has capacity one and cost $-w^*(\{a, b\})$. We also add a source vertex s, sink vertex t, and then arcs (s, a) for each $a \in A$ and (b, t) for each $b \in B$. The arcs incident to the source or the sink have capacity k and have weight zero. Now, we set the flow requirement from s to t to kp. Thanks to the integral theorem of the minimum cost flow problem, we can find, in polynomial time, a maximum weight subgraph H^* of G^* such that H^* has exactly kp edges and each vertex has degree at most k. From this subgraph H^*, we need to construct a maximum weight k-packing of E^* with respect to $\mathcal{M}$. The following lemma ensures that it is always possible.

Lemma 2. Let H^* be a bipartite graph with kp edges. Suppose the maximum degree of a vertex in H^* is at most k. Then, the edges of H^* can be partitioned into k matchings of cardinality exactly p. Moreover, such a partition can be computed in polynomial time from H^*.

Proof. It is known that every bipartite graph of maximum degree at most k has a proper edge-coloring with k color, and such an edge-coloring can be computed in polynomial time [5], [12]. We can assume that each color is used at least once by recoloring an edge whose color is used at least twice. Then, the edges of H^* can be decomposed into k non-empty matchings $M_1, \ldots, M_k$. If $|M_1| = \cdots = |M_k| = p$, we are done. Suppose that there is a pair of matchings M_i and M_j with $|M_i| > p$ and $|M_j| < p$. The union of M_i and M_j induces a subgraph of H^* of maximum degree at most two. As $|M_i| > |M_j|$, the subgraph contains an augmenting path $P = (v_1, \ldots, v_t)$ with $\{v_\ell, v_{\ell+1}\} \in M_i$ for odd ℓ and $\{v_\ell, v_{\ell+1}\} \in M_j$ for even ℓ. Let $M'_i = (M_i \setminus E(P)) \cup (E(P) \cap M_j)$ and $M'_j = (M_j \setminus E(P)) \cup (E(P) \cap M_i)$. Then, we have matchings M'_i and M'_j with $|M'_i| = |M_i| - 1$ and $|M'_j| = |M_j| + 1$. By repeating this argument, we have a desired set of matchings. □

Therefore, $E(H^*) = M_1 \cup \cdots \cup M_k$ is a maximum weight k-packing of $E(H^*)$ with respect to $\mathcal{M}$, which implies the following theorem.

Theorem 5. There is a polynomial-time algorithm that, given a bipartite graph G and positive integers k, p, computes (not necessarily edge-disjoint) k matchings $M_1, \ldots, M_k$ with cardinality p such that $d_w(M_1, \ldots, M_k)$ is maximized.

6. Diverse Shortest *st*-Paths

This section is devoted to solving the following problem.

Definition 4. Let $G = (V, E)$ be a directed graph with specified vertices $s, t \in V$. Let $\ell : E \to \mathbb{N}_{\geq 0}$ be a length function on edges. Let $\mathcal{P}$ be the set of all shortest paths from s to t in (G, ℓ). Given an integer k and a weight function $w : E \to \mathbb{N}_{\geq 0}$, Diverse Shortest st-Paths asks for k paths $P_1, \ldots, P_k \in \mathcal{P}$ such that $d_w(E(P_1), \ldots, E(P_k))$ is maximized.

Theorem 6. Diverse Shortest st-Paths can be solved in polynomial time.

We first compute the shortest distance label $\mathrm{dist} : V \to \mathbb{N}_{\geq 0}$ from s in polynomial time. For each edge $e = (u, v) \in E$ with $\mathrm{dist}(u) \neq \mathrm{dist}(v) + \ell(e)$, we remove it from G. We also remove vertices that are not reachable from s in the removed graph. This does not change the optimal solutions since every path in $\mathcal{P}$ does not include such vertices and edges. Then, the obtained graph, denoted by $G' = (V', E')$, has no directed cycles, and every path from s to t belongs to $\mathcal{P}$. From this directed acyclic graph G', we construct an weighted directed multigraph G^* by replacing each edge $e = (u, v)$ with k copies $e_1, \ldots, e_k$ and setting $w^*(e_i)$ to $w(e)(k - 2i + 1)$ for $e \in E'$ and $1 \leq i \leq k$. By Lemma 1, it is sufficient to find a maximum weight k-packing of E^* with respect to $\mathcal{P}$.

Lemma 3. Let G^* and $w : E^* \to \mathbb{N}_{\geq 0}$ be as above. Then, there is a polynomial-time algorithm that finds a maximum weight k-packing of E^* with respect to $\mathcal{P}$.

Proof. We reduce the k-packing problem to the minimum-cost flow problem, which can be solved in a polynomial time. The source and the sink vertices of G^* are defined as s and t, respectively. For each $e \in E^*$, we assign the capacity value of 1 (to prevent edge sharing) and the cost value of $-w^*(e)$. The flow requirement is set to k. Then we can find a flow $f : E^* \to \mathbb{R}_{\geq 0}$ maximizing $\sum_{e \in E^*} f(e) \cdot w^*(e)$ in polynomial time. Moreover, it is well known that f is integral, that is, $f(e) \in \mathbb{N}_{\geq 0}$ for every $e \in E^*$, as the capacity is integral. Since the all of the edges in E^* has a capacity of 1, f can be decomposed into k edge-disjoint st-paths $P_1, \ldots, P_k$, which implies that the maximum weight k-packing with respect to $\mathcal{P}$ can be found in polynomial time as well. □

7. Diverse Minimum Cuts

In previous sections, we design polynomial-time algorithms for several diverse version of well-known combinatorial problems. In this section, we discuss the diverse version of the minimum cut problem.

Definition 5. Let $G = (V, E)$ be an edge-weighted graph with weight function $w : E \to \mathbb{N}_{\geq 0}$, and let k be an integer. Diverse Minimum Cut asks for k minimum edge cuts $C_1, \ldots, C_k \subseteq E$ such that $d_w(C_1, \ldots, C_k)$ is maximized.

In contrast to results in previous sections, Diverse Minimum Cut is intractable. Let $\lambda(G)$ be the size of a minimum cut of G.

Theorem 7. Diverse Minimum Cut is NP-complete even if $\lambda(G) = 3$.

The problem obviously belongs to NP. The NP-hardness is shown by performing a polynomial-time reduction from the maximum independent set problem on cubic graphs, which is known to be NP-complete [13]. For a graph H, we denote by $\alpha(H)$ the maximum size of an independent set of H. Let H be a graph in which every vertex has degree exactly three. Let H' be the graph obtained from H by sudividing each edge twice, that is, each edge is replaced by a path of three edges. The set of vertices in H' that do not appear in H is denoted by D. The following folklore lemma ensures that the value of α increases exactly by m.

Lemma 4 (folklore). Let m be the number of edges in H. Then, $\alpha(H') = \alpha(H) + m$.

We construct a graph G from H' by adding a new vertex v^* and connecting v^* and a vertex in D. Note that the degree of v^* in H' is more than three.

Lemma 5. G has k edge-disjoint cuts of size three if and only if H' has an independent set of size k.

Proof. Let $G = (V, E)$. Suppose first that H' has an independent set S of size k. Since every vertex in S appears also in G, we can construct a cut of the form $C_i = E_G(\{v_i\}, V \setminus \{v\})$ for each $v_i \in S$. As S is an independent set of G, these k cuts are edge disjoint. Moreover, these cuts have exactly three edges since every vertex in S has degree three in G.

Conversely, suppose G has k edge-disjoint cuts $C_1, C_2, \ldots, C_k \subseteq E$ with $|C_i| = 3$ for $1 \leq i \leq k$. It is sufficient to prove that each of these cuts forms $C_i = E_G(\{v\}, V \setminus \{v\})$ for some $v \in V \setminus \{v^*\}$. Let $C_i = E_G(X, V \setminus X)$ for some $X \subseteq V$. Without loss of generality, we assume that $v^* \in V \setminus X$. In the following, we show that X contains exactly one vertex. Since every vertex in D is adjacent to v^*, X contains at most three vertices of D. Suppose first that $|X \cap D| = 3$. Since every vertex of D has a neighbor in D, $V \setminus X$ has a vertex in D that has a neighbor in $X \cap D$. However, as every vertex of $X \cap D$ is adjacent to v^*, there are at least four edges between X and $V \setminus X$, contradicting to the fact that $|C_i| = 3$.

Suppose next that $|X \cap D| = 2$. Let $u, v \in X \cap D$ be distinct. If u is not adjacent to v, there are two vertices u' and v' in $(V \setminus X) \cap D$ that are adjacent to u and v, respectively. This implies C_i contains four edges $\{u,u'\},\{v,v'\},\{u,v^*\},\{v,v^*\}$, yielding a contradiction. Thus, u is adjacent to v. Let u' and v' be the vertices in $V \setminus (D \cup \{v^*\})$ that are adjacent to u and v, respectively. Observe that at least one of u' and v', say u', belongs to X as otherwise there are four edges $(\{u,u'\},\{v,v'\},\{u,v^*\},\{v,v^*\})$ between X and $V \setminus X$. Since $|X \cap D| = 2$ and u' has three neighbors in D, the other two neighbors of u' belongs to $V \setminus D$, which ensures at least four edges between X and $V \setminus X$.

Suppose that $|X \cap D| = 1$. Let $u \in X \cap D$. In this case, we show that $X = \{u\}$. To see this, consider the neighbors of u. Since $v^* \in V \setminus X$ and $|X \cap D| = 1$, at least two neighbors of u, which are v^* and a vertex in D, belong to $V \setminus X$. If the other neighbor v is in X, then by the assumption that $|X \cap D| = 1$, the two neighbors of v other than u belong to $V \setminus X$, which implies there are at least four edges between X and $V \setminus X$. Thus, all the neighbors of u belong to $V \setminus X$. Since G is connected, all the vertices except for u belong to $V \setminus X$ as well. Thus, we have $C_i = E_G(\{u\}, V \setminus \{u\})$.

Finally, suppose that $X \cap D = \emptyset$. In this case, at least one vertex of $V \setminus (D \cup \{v^*\})$ is included in X. Let $u \in X \setminus (D \cup \{v^*\})$. Since $X \cap D = \emptyset$, every neighbor of u belongs to $V \setminus D$. Similarly to the previous case, we have $X = \{u\}$, which completes the proof. □

Note that the proof of Lemma 5 also proves that G has no cut of size at most two. Therefore, by Lemmas 4 and 5, Theorem 7 follows.

When $\lambda(G) = 1$, then Diverse Minimum Cut is trivially solvable in polynomial time: We can select (not necessarily edge-disjoint) k bridges maximizing d_w. If $\lambda(G) = 2$, the problem in fact is solvable in polynomial time as well.

Theorem 8. Diverse Minimum Cut is polynomial time solvable if $\lambda(G) \leq 2$.

We reduce the problem to that of finding a subgraph of prescribed size with maximizing the sum of convex functions on their degrees of vertices.

Theorem 9 ([1]). Given an undirected graph H, an integer k, and convex functions $f_v : \mathbb{N}_{\geq 0} \to \mathbb{Q}$ for $v \in V(H)$, the problem of finding k-edge subgraph H' of H maximizing $\sum_{v \in V(H)} f_v(d_{H'}(v))$ is solvable in polynomial time, where $d_{H'}(v)$ is the degree of v in H'.

We first enumerate all minimum cuts of G in polynomial time. Then, we construct a graph H whose vertex set corresponds to E, and the edge set of H is defined as follows. For each pair $e, f \in E$, we add k parallel edges between e and f to H if $\{e, f\}$ is a cut of G. Obviously, the graph H can be constructed in polynomial time. For each $e \in E$, we let $f_e(i) := w(e) \cdot i \cdot (k-i)$ for $0 \leq i \leq k$ and $f_e(i) = \infty$ for $i > k$. Clearly, the function f_e is convex. Let $C_1, \dots, C_k \subseteq E$ be k cuts of G. For each e, we denote by $m(e)$ the number of occurrences of e among $C_1, \dots, C_k$. Since each edge in E contributes $w(e) \cdot m(e) \cdot (k - m(e))$ to $d_w(C_1, \dots, C_k)$, we immediately have the following lemma.

Lemma 6. H has a subgraph H' of k edges such that $\sum_{v \in V(H)} f_e(d_{H'}(e)) \geq t$ if and only if there are k edge cuts $C_1, \dots, C_k \subseteq E$ of G with $|C_i| = 2$ for $1 \leq i \leq k$ such that $d_w(C_1, \dots, C_k) \geq t$.

By Lemma 6 and Theorem 9, Diverse Minimum Cut can be solved in polynomial time, provided $\lambda(G) \leq 2$.

We also obtain the following fixed-parameter tractability and intractability results, which are omitted in this paper.

Theorem 10. There is an $2^{O(k+\lambda(G))} n^{O(1)}$-time algorithm for Diverse Minimum Cuts, where n is the number of vertices in G.

In other words, Diverse Minimum Cuts is fixed-parameter tractable parameterized by $k + \lambda(G)$. The proof of this theorem is made by combining algorithm for enumerating all minumum cuts in graphs [17] and a general FPT result due to [14]. In contrast to this tractability, we can show that the problem is unlikely to be fixed-parameter tractable when parameterized by k only.

Theorem 11. Diverse Minimum Cuts is W[1]-hard parameterized by k.

Similarly to Theorem 7, we can reduce Independent Set on d-regular graphs, which is known to be W[1]-hard [4], to Diverse Minimum Cuts.

Acknowledgments

This work was partially supported by JSPS Kakenhi Grant Numbers JP18K11168, JP18K11169, JP18H04091, JP19J10761 , JP19K21537, JP20K19742, and JP20H05793.

References

[1] Nicola Apollonio and András Sebö. Minconvex factors of prescribed size in graphs. *SIAM J. Discret. Math.*, 23(3):1297–1310, 2009.

[2] Julien Baste, Michael R. Fellows, Lars Jaffke, Tomás Masarík, Mateus de Oliveira Oliveira, Geevarghese Philip, and Frances A. Rosamond. Diversity of solutions: An exploration through the lens of fixed-parameter tractability theory. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 1119–1125. ijcai.org, 2020.

[3] Julien Baste, Lars Jaffke, Tomás Masarík, Geevarghese Philip, and Günter Rote. FPT algorithms for diverse collections of hitting sets. *Algorithms*, 12(12):254, 2019.

[4] Leizhen Cai. Parameterized complexity of cardinality constrained optimization problems. *Comput. J.*, 51(1):102–121, 2008.

[5] Richard Cole and John E. Hopcroft. On edge coloring bipartite graphs. *SIAM J. Comput.*, 11(3):540–546, 1982.

[6] Emilie Danna and David L. Woodruff. How to select a small set of diverse solutions to mixed integer programming problems. *Oper. Res. Lett.*, 37(4):255–260, 2009.

[7] Jack Edmonds. Matroid partition. *Mathematics of the Decision Sciences*, 11:335–345, 1968.

[8] Jack Edmonds. Some well-solved problems in combinatorial optimization. In B. Roy, editor, *Combinatorial Programming: Methods and Applications*, pages 285–301, Dordrecht, 1975. Springer Netherlands.

[9] David Eppstein. k-best enumeration. In *Encyclopedia of Algorithms*, pages 1003–1006. 2016.

[10] Fedor V. Fomin, Petr A. Golovach, Lars Jaffke, Geevarghese Philip, and Danil Sagunov. Diverse pairs of matchings. In *31st International Symposium on Algorithms and Computation, ISAAC 2020, December 14-18, 2020, Hong Kong, China (Virtual Conference)*, pages 26:1–26:12, 2020.

[11] Fedor V. Fomin, Petr A. Golovach, Fahad Panolan, Gee-

varghese Philip, and Saket Saurabh. Diverse Collections in Matroids and Graphs. In Markus Bläser and Benjamin Monmege, editors, *38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021)*, volume 187 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 31:1–31:14, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[12] Harold N. Gabow and Oded Kariv. Algorithms for edge coloring bipartite graphs and multigraphs. *SIAM J. Comput.*, 11(1):117–129, 1982.

[13] M. R. Garey, David S. Johnson, and Larry J. Stockmeyer. Some simplified np-complete problems. In Robert L. Constable, Robert W. Ritchie, Jack W. Carlyle, and Michael A. Harrison, editors, *Proceedings of the 6th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1974, Seattle, Washington, USA*, pages 47–63. ACM, 1974.

[14] Tesshu Hanaka, Yasuaki Kobayashi, Kazuhiro Kurita, and Yota Otachi. Finding diverse trees, paths, and more. *CoRR*, abs/2009.03687, 2020.

[15] Emmanuel Hebrard, Brahim Hnich, Barry O'Sullivan, and Toby Walsh. Finding diverse and similar solutions in constraint programming. In *Proceedings of the 20th National Conference on Artificial Intelligence - Volume 1*, AAAI'05, page 372–377. AAAI Press, 2005.

[16] Ian Holyer. The np-completeness of edge-coloring. *SIAM J. Comput.*, 10(4):718–720, 1981.

[17] A. V. Karzanov and E. A. Timofeev. Efficient algorithm for finding all minimal edge cuts of a nonoriented graph. *Cybernetics*, 22:156–162, 1986.

[18] Eugene L. Lawler. A procedure for computing the k best solutions to discrete optimization problems and its application to the shortest path problem. *Management Science*, 18(7):401–405, 1972.

[19] Alexander Nadel. Generating diverse solutions in SAT. In Karem A. Sakallah and Laurent Simon, editors, *Theory and Applications of Satisfiability Testing - SAT 2011 - 14th International Conference, SAT 2011, Ann Arbor, MI, USA, June 19-22, 2011. Proceedings*, volume 6695 of *Lecture Notes in Computer Science*, pages 287–301. Springer, 2011.

[20] Hiroshi Nagamochi, Dao-Zhi Zeng, Naohisa Kabutoya, and Toshihide Ibaraki. Complexity of the minimum base game on matroids. *Math. Oper. Res.*, 22(1):146–164, 1997.

[21] Crispin St. John Alvah Nash-Williams. An application of matroids to graph theory. In *Theory of Graphs; Proceedings of an International Symposium in Rome 1966*, pages 263–265. Gordon and Breach, New York, 1967.

[22] Thierry Petit and Andrew C. Trapp. Finding diverse solutions of high quality to constraint optimization problems. In Qiang Yang and Michael J. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 260–267. AAAI Press, 2015.

[23] Thierry Petit and Andrew C. Trapp. Enriching solutions to combinatorial problems via solution engineering. *INFORMS J. Comput.*, 31(3):429–444, 2019.

