

Obtaining Approximately Optimal and Diverse Solutions via Dispersion

Jie Gao^{1[0000-0001-5083-6082]}, Mayank Goswami^{2[0000-0002-2111-3210]}, Karthik C.S.¹, Meng-Tsung Tsai³, Shih-Yu Tsai⁴, and Hao-Tsung Yang⁵

¹ Department of Computer Science, Rutgers University, USA

² Department of Computer Science, Queens College, City University of New York, USA

³ Institute of Information Science, Academia Sinica, Taiwan

⁴ Department of Computer Science, Stony Brook University, USA

⁵ School of Informatics, University of Edinburgh, UK

Abstract. There has been a long-standing interest in computing diverse solutions to optimization problems. In 1995 J. Krarup [29] posed the problem of finding k -edge disjoint Hamiltonian Circuits of minimum total weight, called the peripatetic salesman problem (PSP). Since then researchers have investigated the complexity of finding diverse solutions to spanning trees, paths, vertex covers, matchings, and more. Unlike the PSP that has a constraint on the total weight of the solutions, recent work has involved finding diverse solutions that are all optimal.

However, sometimes the space of exact solutions may be too small to achieve sufficient diversity. Motivated by this, we initiate the study of obtaining sufficiently-diverse, yet approximately-optimal solutions to optimization problems. Formally, given an integer k , an approximation factor c , and an instance I of an optimization problem, we aim to obtain a set of k solutions to I that a) are all c approximately-optimal for I and b) maximize the diversity of the k solutions. Finding such solutions, therefore, requires a better understanding of the global landscape of the optimization function.

Given a metric on the space of solutions, and the diversity measure as the sum of pairwise distances between solutions, we first provide a general reduction to an associated budget-constrained optimization (BCO) problem, where one objective function is to be (maximized) minimized subject to a bound on the second objective function. We then prove that bi-approximations to the BCO can be used to give bi-approximations to the diverse approximately optimal solutions problem.

As applications of our result, we present polynomial time approximation algorithms for several problems such as diverse c -approximate *maximum matchings*, *$s-t$ shortest paths*, *global min-cut*, and *minimum weight bases of a matroid*. The last result gives us *diverse c -approximate minimum spanning trees*, advancing a step towards achieving diverse c -approximate TSP tours.

We also explore the connection to the field of multiobjective optimization and show that the class of problems to which our result applies includes those for which the associated DUALRESTRICT problem defined by Papadimitriou and Yannakakis [37], and recently explored by Herzl et al. [27] can be solved in polynomial time.

Keywords: diversity, minimum spanning tree, maximum matching, shortest path, travelling salesman problem, dispersion problem

1 Introduction

Techniques for optimization problems focus on obtaining optimal solutions to an objective function and have widespread applications ranging from machine learning, operations research, computational biology, networks, to geophysics, economics, and finance. However, in many scenarios, the optimal solution is not only computationally difficult to obtain, but can also render the system built upon its utilization vulnerable to adversarial attacks. Consider a patrolling agent tasked with monitoring n sites in the plane. The most efficient solution (i.e., maximizing the frequency of visiting each of the n sites) would naturally be to patrol along the tour of shortest length⁶ (the solution to TSP - the Traveling Salesman Problem). However, an adversary who wants to avoid the patroller can also compute the shortest TSP tour and can design its actions strategically [41]. Similarly, applications utilizing the minimum spanning tree (MST) on a communication network may be affected if an adversary gains knowledge of the network [14]; systems using solutions to a linear program (LP) would be vulnerable if an adversary gains knowledge of the program's function and constraints.

One way to address the vulnerability is to use a set of approximately optimal solutions and randomize among them. However, this may not help much to mitigate the problem, if these approximate solutions are combinatorially too “similar” to the optimal solution. For example, all points in a sufficiently small neighborhood of the optimal solution on the LP polytope will be approximately optimal, but these solutions are not too much different and the adversaries can still effectively carry out their attacks. Similarly one may use another tree instead of the MST, but if the new tree shares many edges with the MST the same vulnerability persists. Thus k -best enumeration algorithms ([20,25,31,32,35]) developed for a variety of problems fall short in this regard.

One of the oldest known formulations is the Peripatetic Salesman problem (PSP) by Krarup [29], which asks for k -edge disjoint Hamiltonian circuits of minimum total weight in a network. Since then, several researchers have tried to compute diverse solutions for several optimization problems [4,5,18,24]. Most of these works are on graph problems, and diversity usually corresponds to the size of the symmetric difference of the edge sets in the solutions. Crucially, almost all of the aforementioned work demands either every solution individually be optimal, or the set of solutions in totality (as in the case of the PSP) be optimal. Nevertheless, **the space of optimal solutions may be too small to achieve sufficient diversity**, and it may just be singular (unique solution). In addition, for NP-complete problems finding just one optimal solution is already difficult. While there is some research that takes the route of developing FPT algorithms

⁶ We assume without loss of generality that the optimal TSP is combinatorially unique by a slight perturbation of the distances.

for this setting [5,19], to us it seems practical to also consider the relaxation to approximately-optimal solutions.

This motivates the problem of finding a set of diverse and *approximately optimal* solutions, which is the problem considered in this article. The number of solutions k and the desired approximation factor $c > 1$ is provided by the user as input. Working in the larger class gives one more hope of finding diverse solutions, yet every solution has a guarantee on its quality.

1.1 Our Contributions

We develop approximation algorithms for finding k solutions to the given optimization problem: for every solution, the quality is bounded by a user-given approximation ratio $c > 1$ to the optimal solution and the diversity of these k solutions is maximized. Given a metric on the space of solutions to the problem, we consider the diversity measure given by the sum (or average) of pairwise distances between the k solutions. Combining ideas from the well-studied problem on dispersion (which we describe next), we reduce the above problem to a budget constrained optimization (BCO) program.

1.2 Dispersion

Generally speaking, if the optimization problem itself is \mathcal{NP} -hard, finding diverse solutions for that problem is also \mathcal{NP} -hard (see Proposition 1 for more detail). On the other hand, interestingly, even if the original problem is not \mathcal{NP} -hard, finding diverse and approximately optimal solutions can still be \mathcal{NP} -hard. This is due to the connection of the diversity maximization objective with the general family of problems that consider selecting k elements from the given input set with maximum “dispersion”, defined as max-min distance, max-average distance, and so on.

The dispersion problem has a long history, with many variants both in the metric setting and the geometric setting [17,30,40]. For example, finding a subset of size k from an input set of n points in a metric space that maximizes the distance between closest pairs or the sum of distances of the k selected points are both \mathcal{NP} -hard [1,39]. For the max-sum dispersion problem, the best known approximation factor is 2 for general metrics [7,26], although PTAS are available for Euclidean metrics or more generally, metrics of negative type, even with matroid constraints [10,11].

Dispersion in exponentially-sized space We make use of the general framework of the 2-approximation algorithm [8,39] to the max-sum k -dispersion problem, a greedy algorithm where the $i + 1$ th solution is chosen to be the most distant/diverse one from the first i solutions. Notice that in our setting, there is an important additional challenge to understand the space within which the approximate solutions stay. In all of the problems we study, the total number of solutions can be *exponential in the input size*. Thus we need to have a non-trivial way of navigating within this large space and carry furthest insertion

without considering all points in the space. This is where our reduction to budget constrained problem comes in.

Self avoiding dispersion Furthermore, even after implicitly defining the $i + 1$ th furthest point insertion via some optimization problem, one needs to take care that the (furthest, in terms of sum of distances) solution does not turn out to equal one of the previously found i solutions, as this is a requirement for the furthest point insertion algorithm. This is an issue one faces because of the implicit nature of the furthest point procedure in the exponential-sized space of solutions: in the metric k -dispersion problem, it was easy to guarantee distinctness as one only considered the $n - i$ points not yet selected.

1.3 Reduction to Budget Constrained Optimization

Combining with dispersion, we reduce the diversity computational problem to a budget constrained optimization (BCO) problem where the budget is an upper (resp. lower) bound if the quality of solution is described by a minimization (resp. maximization) problem. Recall that the number of solutions k and the approximation factor c is input by the user; a larger c allows for more diversity.

We show how using an (a, b) bi-approximation algorithm for the BCO problem provides a set of $O(a)$ -diverse, bc approximately-optimal solutions to the diversity computational problem (the hidden constant is at most 4). This main reduction is described in Theorem 1.

The main challenge in transferring the bi-approximation results because of a technicality that we describe next. Let $\Omega(c)$ be the space of c approximate solutions. A $(*, b)$ bi-approximation algorithm to the BCO relaxes the budget constraint by a factor b , and hence only promises to return a faraway point in the larger space $\Omega(b \cdot c)$. Thus bi-approximation of BCO do not simply give a furthest point insertion in the space of solutions, and instead *return a point in a larger space*. Nevertheless, we prove in Lemma 4 that in most cases, one loses a factor of at most 4 in the approximation factor for the diversity.

Once the reduction to BCOs is complete, for diverse approximate matchings, spanning trees and shortest paths we exploit the special characteristics of the corresponding BCO to solve it optimally ($a = b = 1$). For other problems such as global min-cut, diverse approximate minimum weight spanning trees, and the more general minimum weight bases of a matroid, we utilize known bi-approximations to the BCO to obtain bi-approximations for the diversity problem. For all problems except diverse (unweighted) spanning trees⁷, our algorithms are the first polynomial time bi-approximations for these problems.

We also connect to the wide literature on multicriteria optimization and show that our result applies to the entire class of problems for which the associated DUALRESTRICT problem (defined by Papadimitriou and Yannakakis [37], and

⁷ While an exact algorithm for diverse unweighted spanning trees is known [24], we give a faster (by a factor $\Omega(n^{1.5}k^{1.5}/\alpha(n, m))$ where $\alpha(\cdot)$ denotes the inverse of the Ackermann function), 2-approximation here.

recently studied by Herzel et al. [27]) has a polynomial time solution. We discuss this in more detail after presenting our reduction.

Layout: The rest of this paper is organized as follows: we survey related work in Section 2, and formulate the problem in Section 3. In Section 4 we mention the connection to dispersion and describe the reduction to the budget constrained optimization problem (Theorem 1). Sections 5-8 describe four applications of our technique to various problems such as diverse approximate matchings, global min-cuts, shortest paths, minimum spanning trees, and minimum weight bases of a matroid. We remark that this list is by no means exhaustive, and we leave finding other interesting optimization problems which are amenable to our approach for future research. We end with open problems in Section 9.

2 Related Work

Recently there has been a surge of interest in the tractability of finding diverse solutions for a number of combinatorial optimization problems, such as spanning trees, minimum spanning trees, k -paths, shortest paths, k -matchings, etc. [18,19,22,23,24]. Most of the existing work focuses on finding diverse optimal solutions. In cases when finding the optimal solution is NP-complete, several works have focused on developing FPT algorithms [5,19]. Nevertheless, as pointed out in [23], it would be more practical to consider finding a set of diverse “short” paths rather than one set of diverse shortest paths. They show that finding a set of approximately shortest paths with the maximum diversity is NP-hard, but leave the question of developing approximation algorithms open, a question that we answer in our paper for several problems. Similarly the problem of finding diverse maximum matchings was proved to be NP-hard in [18]. We remark that the main difference between our result and previous work is that our algorithms can find a diverse set of c -approximate solutions in polynomial time. If the attained diversity is not sufficient for the application, the user can input a larger c , in hopes of increasing it.

Multicriteria Optimization: In this domain, several optimization functions are given on a space of solutions. Clearly, there may not be a single solution that is the best for all objective functions, and researchers have focused on obtaining Pareto-optimal solutions, which are solutions that are non-dominated by other solutions. Put differently, a solution is Pareto-optimal if no other solution can have a better cost for all criteria. Since exact solutions are hard to find, research has focused on finding ϵ Pareto-optimal solutions, which are a $1+\epsilon$ factor approximations of Pareto-optimal solutions. Papadimitriou and Yannakakis [37] showed that under pretty mild conditions, any multicriteria optimization problem admits an ϵ Pareto-optimal set of fully polynomial cardinality. In terms of being able to *find* such an ϵ Pareto-optimal set, they show that a (FPTAS) PTAS exists for the problem if and only if an associated GAP problem can be solved in (fully) polynomial time. Very recently, Herzel et al. [27] study the class of problems for which a FPTAS or PTAS exists for finding ϵ Pareto-optimal solutions that are *exact* in one of the criteria. Such problems are a subset of the ones characterized

by GAP. Herzel et al. [27] characterize the condition similarly: a FPTAS (PTAS) exists if and only if an associated DUALRESTRICT problem can be solved in (fully) polynomial time. For more details we refer the reader to the survey by Herzel et al. [28].

3 Diversity Computational Problem (DCP)

First, we define some notations. We use the definition of optimization problems given in [3] with additional formalism as introduced in [21].

Definition 1 (Optimization Problem). *An **optimization problem** Π is characterized by the following quadruple of objects $(I_\Pi, \text{Sol}_\Pi, \Delta_\Pi, \text{goal}_\Pi)$, where:*

- I_Π is the set of instances of Π . In particular for every $d \in \mathbb{N}$, $I_\Pi(d)$ is the set of instances of Π of input size at most d (bits);
- Sol_Π is a function that associates to any input instance $x \in I_\Pi$ the set of feasible solutions of x ;
- Δ_Π is the measure function⁸, defined for pairs (x, y) such that $x \in I_\Pi$ and $y \in \text{Sol}_\Pi(x)$. For every such pair (x, y) , $\Delta_\Pi(x, y)$ provides a non-negative integer which is the value of the feasible solution y ;
- $\text{goal}_\Pi \in \{\min, \max\}$ specifies whether Π is a maximization or minimization problem.

We would like to identify a subset of our solution space which are (approximately) optimal with respect to our measure function. To this effect, we define a notion of approximately optimal feasible solution.

Definition 2 (Approximately Optimal Feasible Solution). *Let $\Pi(I_\Pi, \text{Sol}_\Pi, \Delta_\Pi, \text{goal}_\Pi)$ be an optimization problem and let $c \geq 1$. For every $x \in I_\Pi$ and $y \in \text{Sol}_\Pi(x)$ we say that y is a **c -approximate optimal solution of x** if for every $y' \in \text{Sol}_\Pi(x)$ we have $\Delta_\Pi(x, y) \cdot c \geq \Delta_\Pi(x, y')$ if $\text{goal}_\Pi = \max$ and $\Delta_\Pi(x, y) \leq \Delta_\Pi(x, y') \cdot c$ if $\text{goal}_\Pi = \min$.*

Definition 3 (Computational Problem). *Let $\Pi(I_\Pi, \text{Sol}_\Pi, \Delta_\Pi, \text{goal}_\Pi)$ be an optimization problem and let $\lambda : \mathbb{N} \rightarrow \mathbb{N}$. The **computational problem associated with (Π, λ)** is given as input an instance $x \in I_\Pi(d)$ (for some $d \in \mathbb{N}$) and real $c := \lambda(d) \geq 1$ find a c -approximate optimal feasible solution of x .*

Definition 4 (DCP - Diversity Computational Problem). *Let $\Pi(I_\Pi, \text{Sol}_\Pi, \Delta_\Pi, \text{goal}_\Pi)$ be an optimization problem and let $\lambda : \mathbb{N} \rightarrow \mathbb{N}$. Let $\sigma_{\Pi, t}$ be a diversity measure that maps every t feasible solutions of an instance of I_Π to a non-negative real number. The **diversity computational problem associated with $(\Pi, \sigma_{\Pi, t}, k, \lambda)$** is given as input an instance $x \in I_\Pi(d)$ (for some $d \in \mathbb{N}$), an integer $k := k(d)$, and real $c := \lambda(d) \geq 1$, find k -many c -approximate solutions y_1, \dots, y_k to x which maximize the value of $\sigma_{\Pi, k}(x, y_1, \dots, y_k)$.*

⁸ We define the measure function only for feasible solutions of an instance. Indeed if an algorithm solving the optimization problem outputs a non-feasible solution then, the measure just evaluates to -1 in case of maximization problems and ∞ in case of minimization problems.

Proposition 1. *Let $\Pi(I_\Pi, \text{Sol}_\Pi, \Delta_\Pi, \text{goal}_\Pi)$ be an optimization problem and let $\lambda : \mathbb{N} \rightarrow \mathbb{N}$. Let $\sigma_{\Pi,t}$ be a diversity measure that maps every t feasible solutions of an instance of I_Π to a non-negative real number. If the computational problem associated with (Π, λ) is \mathcal{NP} -hard then, the diversity computational problem associated with $(\Pi, \sigma_{\Pi,t}, \lambda)$ is also \mathcal{NP} -hard.*

Therefore the interesting questions arise when we compute problems associated with (Π, λ) which are in \mathcal{P} , or even more when, $(\Pi, \mathbb{1})$ is in \mathcal{P} where $\mathbb{1}$ is the constant function which maps every element of the domain to 1. For the remainder of this paper, we will consider $\lambda(d)$ to be the constant function, and will simply refer to the constant as c .

Finally, we define bicriteria approximations for the diversity computational problem:

Definition 5 ((α, β) bi-approximation for the Diversity Computational Problem). *Consider the diversity computational problem associated with $(\Pi, \sigma_{\Pi,t}, k, c)$, and a given instance $x \in I_\Pi(d)$ (for some $d \in \mathbb{N}$). An algorithm is called an (α, β) bi-approximation for the diversity computational problem if it outputs k feasible solutions y_1, \dots, y_k such that a) y_i is a $\beta \cdot c$ -approximate optimal feasible solution to x for all $1 \leq i \leq k$, and b) for any set y'_1, \dots, y'_k of k -many c -approximate optimal feasible solutions, $\sigma_{\Pi,k}(y_1, \dots, y_k) \cdot \alpha \geq \sigma_{\Pi,k}(y'_1, \dots, y'_k)$. Furthermore, such an algorithm is said to run in polynomial time if the running time is polynomial in d and k .*

4 The Reduction: Enter Dispersion and Biobjective Optimization

As stated in the introduction, our problems are related to the classical dispersion problem in a metric space. Here we state the dispersion problem and show that under the exponential time hypothesis, 2-approximation is actually tight for the k -dispersion problem. We will then use dispersion to reduce the problem of finding diverse, approximately optimal solutions to solving an associated budget constrained optimization problems.

4.1 Dispersion Problem

Definition 6 (k -Dispersion, total distance). *Given a finite set of points P whose pairwise distances satisfy the triangle inequality and an integer $k \geq 2$, find a set $S \subseteq P$ of cardinality k so that $W(S)$ is maximized, where $W(S)$ is the sum of the pairwise distances between points in S .*

The main previous work on the k -dispersion problem relevant to us is [39], where the problem was named as Maximum-Average Facility Dispersion problem with triangle inequality (MAFD-TI). The problems are equivalent as maximizing the average distance between the points also maximizes the sum of pairwise distances between them and vice-versa.

The k -dispersion problem is \mathcal{NP} -hard, but one can find a set S whose $W(S)$ is at least a constant factor of the maximum possible in polynomial time by a greedy procedure [39]. We call the greedy procedure *furthest insertion*. It works as follows. Initially, let P be a singleton set that contains an arbitrary point from the given set. While $|S| < k$, add to S a point $x \notin S$ so that $W(S \cup \{x\}) \geq W(S \cup \{y\})$ for any $y \notin S$. Repeat the greedy addition until S has size k . The final S is a desired solution, which is shown to be a 4-approximation in [39]. It is worth noting that the furthest insertion in [39] initializes S as a furthest pair of points in the given set, and the above change does not worsen the approximation factor. In a later paper [8], the greedy algorithm of choosing an arbitrary initial point is shown to be a 2-approximation, which is a tight bound for this algorithm [7].

Lemma 1 (Furthest Insertion in [8,39]). *The k -dispersion problem can be 2-approximated by the furthest insertion algorithm.*

The running time of the furthest insertion algorithm is polynomial in $|S|$ the size of S , as it performs k iterations, each performing at most $O(k|S|)$ distance computations/lookups. Note that in our case S is the collection of objects of a certain type (matchings, paths, trees, etc.). Hence the size of our metric space is typically exponential in $|V|$ and $|E|$. This adds a new dimension of complexity to the traditional dispersion problems studied.

4.2 Reduction to Budget Constrained Optimization

Recall the definitions of the Diversity Computational Problem (Definition 4) and (a, b) bi-approximations (Definition 5). As the input instance $x \in I_{\Pi}$ will be clear from context, we drop the dependence on x , and assume a fixed input instance to a computational problem. Thus Sol_{Π} will denote the set of feasible solutions, and $\Delta_{\Pi}(y)$ the measure of the feasible solution y .

Diversity and similarity measures from metrics Let $d : \text{Sol}_{\Pi} \times \text{Sol}_{\Pi} \rightarrow \mathbb{R}^+$ be a metric on the space of feasible solutions. When such a metric is available, we will consider the diversity function $\sigma_{\Pi, t} : \text{Sol}_{\Pi} \times \dots \times \text{Sol}_{\Pi} \rightarrow \mathbb{R}^+$ that assigns the diversity measure $\sum_{i,j} d(y_i, y_j)$ to a t -tuple of feasible solutions (y_1, \dots, y_t) . Also, given such a metric d , define D to be the diameter of Sol_{Π} under d , i.e., $D = \max_{y, y' \in \text{Sol}_{\Pi}} d(y, y')$. In many cases, we will be interested in the *similarity* measure $s_{\Pi, t}$ defined by $s_{\Pi, t}(y_1, \dots, y_t) = \sum_{i,j} (D - d(y_i, y_j))$. The examples the reader should keep in mind are graph objects such as spanning trees, matchings, shortest paths, Hamiltonian circuits, etc., such that $d(y, y')$ denotes the Hamming distance, a.k.a. size of the symmetric difference of the edge sets of y and y' , and s denotes the size of their intersection.

In the remainder of the paper we consider the above total distance (resp. similarity) diversity measures $\sigma_{\Pi, t}$ arising from the metric d (resp. similarity measure s), and we will parameterize the problem by d (resp. s) instead.

Definition 7 (Budget Constrained Optimization). *Given an instance of a computational problem Π , a $c \geq 1$, and a set y_1, \dots, y_i of feasible solutions in Sol_{Π} , define the **metric budget constrained optimization problem** $BCO(\Pi, (y_1, \dots, y_i), c, d)$ as follows:*

- If $\text{goal}_\pi = \min$, define $\Delta^* := \min_{y \in \text{Sol}_\Pi} \Delta_\Pi(y)$. Then $\text{BCO}(\Pi, (y_1, \dots, y_i), c, d)$ is the problem

$$\begin{aligned} \max \quad & f_d(y) := \sum_{j=1}^i d(y, y_j) \\ \text{s.t.} \quad & \Delta_\Pi(y) \leq c \cdot \Delta^* \\ & y \in \text{Sol}_\Pi \setminus \{y_1, \dots, y_i\} \end{aligned} \tag{1}$$

- If $\text{goal}_\pi = \max$, define $\Delta^* := \max_{y \in \text{Sol}_\Pi} \Delta_\Pi(y)$. Then $\text{BCO}(\Pi, (y_1, \dots, y_i), c, d)$ is the problem

$$\begin{aligned} \max \quad & f_d(y) := \sum_{j=1}^i d(y, y_j) \\ \text{s.t.} \quad & \Delta_\Pi(y) \cdot c \geq \Delta^* \\ & y \in \text{Sol}_\Pi \setminus \{y_1, \dots, y_i\} \end{aligned} \tag{2}$$

- Given a similarity measure s , define the **similarity budget constrained optimization problem** $\text{BCO}(\Pi, (y_1, \dots, y_i), c, s)$ with the same constraint set as above (depending on goal_π), but with the objective function changed to $g_s(y) := \min \sum_{j=1}^i s(y, y_j)$ instead of $f_d(y) = \max \sum_{j=1}^i d(y, y_j)$.

Definition 8 (Bi-approximation to BCO). An algorithm for an associated BCO is called an (a, b) bi-approximation algorithm if for any $1 \leq i \leq k$, it outputs a solution y such that the following holds.

- If $\text{goal}_\Pi = \min$ and the associated BCO is $\text{BCO}(\Pi, (y_1, \dots, y_i), c, d)$, then
 - a) $y \in \text{Sol}_\Pi \setminus \{y_1, \dots, y_i\}$,
 - b) $\Delta_\Pi(y) \leq b \cdot c \cdot \Delta^*$,
 - c) for all y' satisfying the constraints of $\text{BCO}(\Pi, (y_1, \dots, y_i), c, d)$, $f_d(y) \cdot a \geq f_d(y')$.
- If $\text{goal}_\Pi = \max$ and the associated BCO is $\text{BCO}(\Pi, (y_1, \dots, y_i), c, d)$, then
 - a) $y \in \text{Sol}_\Pi \setminus \{y_1, \dots, y_i\}$,
 - b) $\Delta_\Pi(y) \cdot b \cdot c \geq \Delta^*$,
 - c) for all y' satisfying the constraints of $\text{BCO}(\Pi, (y_1, \dots, y_i), c, d)$, $f_d(y) \cdot a \geq f_d(y')$.
- If $\text{goal}_\Pi = \min$ and the associated BCO is $\text{BCO}(\Pi, (y_1, \dots, y_i), c, s)$, then
 - a) $y \in \text{Sol}_\Pi \setminus \{y_1, \dots, y_i\}$,
 - b) $\Delta_\Pi(y) \leq b \cdot c \cdot \Delta^*$,
 - c) for all y' satisfying the constraints of $\text{BCO}(\Pi, (y_1, \dots, y_i), c, s)$, $g_s(y) \leq g_s(y') \cdot a$.
- If $\text{goal}_\Pi = \max$ and the associated BCO is $\text{BCO}(\Pi, (y_1, \dots, y_i), c, s)$, then
 - a) $y \in \text{Sol}_\Pi \setminus \{y_1, \dots, y_i\}$,
 - b) $\Delta_\Pi(y) \cdot b \cdot c \geq \Delta^*$,
 - c) for all y' satisfying the constraints of $\text{BCO}(\Pi, (y_1, \dots, y_i), c, s)$, $g_s(y) \leq g_s(y') \cdot a$.

Remark: Minimization and maximization are essentially equivalent (by changing the sign), and so optimally solving one solves the other. The reason why we continue to treat them separately is because obtaining an approximation to minimizing total similarity $g_s(y) := \sum_{j=1}^i s(y, y_j)$ is not equivalent to an approximation to maximizing total distance $f_d(y) := \sum_{j=1}^i d(y, y_j)$ – in fact, these functions are the “opposite” of each other, as $f_d(y) = D_i - g_s(y)$.

We are now ready to state our main theorem.

Theorem 1 (Reduction of DCP to BCO). *Consider an input (Π, k, d, c) to the diversity computational problem (DCP).*

- For metric BCO,
 1. An $(a, 1)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, d)$ can be used to give a $(2a, 1)$ approximation to the DCP, and
 2. An (a, b) bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, d)$ can be used to give a $(4a, b)$ approximation to the DCP.
- For similarity BCO,
 3. A $(1, 1)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, s)$ can be used to give a $(2, 1)$ approximation to the DCP,
 4. A $(1, b)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, s)$ can be used to give $(4, b)$ approximation to the DCP,
 5. A $(1 + \epsilon, 1)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, s)$ can be used to give $(4, 1)$ approximation to the DCP, under the condition that the average pairwise distance in the optimal solution to the DCP is at least $D \frac{4\epsilon}{1+2\epsilon}$.

In all of the above, the overhead for obtaining a bi-approximation for the DCP, given a bi-approximation for BCO problem, is $O(k)$.

The detailed proof of Theorem 1 is in Appendix 10. A few remarks are in order:

- The above theorem provides a recipe for solving the diversity computational problem for any given optimization problem. As long as either the metric or the similarity budget constrained optimization problems can be solved or approximated in polynomial time, one has an analogous result for the DCP.
- In the remainder of this paper we will see several application that follow from the above 5 “types” of bi-approximations available. These include DCP for Maximum Matching and Global Min-Cut (Type 1), DCP for shortest path (Type 3), DCP for minimum weight bases of a matroid, minimum spanning trees (Types 4 and 5).
- Whenever either a or b (or both) is set to be $1 + \epsilon$, we call a bi-approximation for the BCO problem an FPTAS if the running time is polynomial in $1/\epsilon$ in addition to being polynomial in d and i . Otherwise we call it a PTAS.

Relation to Multicriteria Optimization: Observe that for similarity BCOs, we need either a or b to be 1. This class of biobjective problems that have a PTAS that is exact in one of the criteria is a special case of the multicriteria problems that have a PTAS that is exact in one of the criteria. Herzel et al. [27] showed that this class is exactly the class of problems for which the DUALRESTRICT version of the problem, posed by Diakonikolas and Yannakakis [15], can be solved in polynomial time. These are also the class of problems having a polynomial-time computable approximate ϵ -Pareto set that is exact in one objective. This equivalence means that our theorem is applicable to this entire class of problems.

4.3 Relaxed BCOs and Self-Avoidance

Before we delve into our applications, we describe another challenge in directly applying results from multicriteria optimization literature. For a BCO, the second constraint demands that $y \in \text{Sol}_H \setminus \{y_1, \dots, y_i\}$. Intuitively y is the farthest point to the set of already discovered solutions $\{y_1, \dots, y_i\}$, and because it is defined implicitly, without the second constraint y may equal one of the y_j ($1 \leq j \leq i$). Consider an alternate BCO, which we call BCO^r where the constraint is relaxed to $y \in \text{Sol}_H$. For many graph problems, solving BCO^r combined with the approach by Lawler [31] gives a solution to the original BCO. This is extremely useful because most of the literature on multicriteria optimization concerns optimization of the relaxed type of problems BCO^r , and one can borrow results derived before without worrying about the second constraint. We remark that for other problems, k -best enumeration algorithms (see [20, 25, 31, 32, 35] for examples) may be useful to switch from the BCO to its relaxed version. *Thus any algorithm for BCO^r can be used, modulo the self-avoiding constraint (to be handled using Lawler's approach), to give a polynomial time algorithm for the Diversity Computational Problem with the same guarantees as in Theorem 1.* We provide examples of the approach by Lawler in subsequent sections where we consider specific problems.

5 Application 1: Diverse Spanning Trees

In this section, we discuss the diverse spanning trees problem, which is the diversity computational problem for spanning trees with Hamming distance function as the diversity measure. Let $G = (V, E)$ be an undirected graph. The problem aims to output a set S of k spanning trees T_1, \dots, T_k of G such that the sum of the pairwise distances $\sum_{i,j \in S} d(T_i, T_j)$ is maximized, where d is the Hamming distance between the edge sets of the trees. While this problem actually has an exact algorithm running in time $O((kn)^{2.5}m)$ [24], we get a faster approximation algorithm.

Theorem 2. *Given a simple graph $G = (V, E)$, there exists an $O(kn\alpha(n, m))$ -time algorithm, where $\alpha(\cdot)$ is the inverse of the Ackermann function, that generates k spanning trees T_1, \dots, T_k , such that the sum of all pairwise Hamming distances is at least half of an optimal set of k diverse spanning trees.*

We prove the above theorem in Section 11 in the appendix, by developing an exact $(1, 1)$ polynomial time subroutine for the associated BCO problem.

6 Application 2: Diverse Approximate Shortest Paths

Given a graph $G = (V, E)$, non-negative edge weights $w(e)$, two vertices s and t , and a factor $c > 1$, the diversity computational problem asks to output k many st paths, such that the weight of each path is within a factor c of the weight of the shortest st path, and subject to this constraint, the total pairwise distance

between the paths is maximized. Here the distance between two paths is again the Hamming distance, or size of symmetric difference of their edge sets.

In [23], it is shown that finding k shortest paths with the maximum diversity (i.e. the average Hamming distance between solutions) can be solved in polynomial time, but finding k “short” paths with the maximum diversity is NP-hard. In contrast, in what follows, we will show that finding k “short” paths with constant approximate diversity is polynomial-time solvable.

We will show that the associated budget constrained optimization problem for this is of Type 3 in Theorem 1. In other words, we will show that the BCO can be solved exactly. This will result in a $(2, 1)$ approximation algorithm for the diversity computational problem.

Hence, we need an algorithm that implements: given a set S of c -st-shortest paths, find a c -st-shortest path $P \notin S$ so that $W(S \cup \{P\})$ is maximum among all $W(S \cup \{P'\})$ for c -st-shortest path $P' \notin S$. Here, $W(S')$ is the sum of all pairwise Hamming distances between two elements in S' . This is a special case of the bicriteria shortest paths, for which there is an algorithm in [36]. In our case, **one of the two weight functions is an integral function with range bounded in $[0, k]$** . Hence, it can be solved more efficiently than the solution in [36], which can be summarized as following.

Lemma 2 (Exact solution to the relaxed BCO^r problem). *Given a real $c \geq 1$ and a directed simple graph $G = (V \cup \{s, t\}, E)$ associated with two weight functions on edges $\omega : E \rightarrow \mathbb{R}^+$ and $f : E \rightarrow \{0, 1, \dots, r\}$, there is an $O(r|V|^3)$ -time algorithm to output an st-path P^* so that $\sum_{e \in E(P^*)} f(e)$ is minimized while retaining $\sum_{e \in E(P^*)} \omega(e) \leq c \sum_{e \in E(P)} \omega(e)$ for all st-paths P .*

Please see Section 12 in appendix for the detailed proof of Lemma 2.

Self-Avoiding Constraint We now turn to solving the associated (non-relaxed) BCO problem, by generalizing the above lemma to Corollary 1 (proof in appendix Section 12.). Thus [corollary 1](#) will help us avoid the situation that a furthest insertion returns a path that is already picked by some previous furthest insertion.

Corollary 1 (Exact solution to the BCO problem). *Given a real $c \geq 1$, a directed simple graph $G = (V \cup \{s, t\}, E)$ associated with two weight functions on edges $\omega : E \rightarrow \mathbb{R}^+$, $f : E \rightarrow \{0, 1, \dots, r\}$, and two disjoint subsets of edges $E_{in}, E_{ex} \subseteq E$ so that all edges in E_{in} together form a directed simple path P_{prefix} starting from node s , there exists an $O(r|V|^3)$ -time algorithm to output an c -st-shortest path P^* under ω so that $\sum_{e \in E(P^*)} f(e)$ is minimum among all the c -st-shortest paths P that contain P_{prefix} as a prefix and contain no edges from E_{ex} , if such an c -st-shortest path exists.*

We are ready to state our main result (proof in appendix Section 12) for the diverse c -st-shortest paths.

Theorem 3 ((2, 1) bi-approximation to the Diversity Problem on Shortest Paths). *For any directed simple graph $G = (V \cup \{s, t\}, E)$, given a $c > 1$ and a $k \in \mathbb{N}$, there exists an $O(k^3|V|^4)$ -time algorithm that, if G contains at least*

k distinct c -st-shortest paths, computes a set S of k distinct c -st-shortest paths so that the sum of all pairwise Hamming distances between two paths in S is at least one half of the maximum possible; otherwise, reports “Non-existent.”

7 Application 3: Diverse Approximate Maximum Matchings, and Global Min-cut

Consider the diversity computational problem for computing k many c maximum matchings for undirected graphs. In [18], the authors present an algorithm, among others, to find a pair of maximum matchings for bipartite graphs whose Hamming distance is maximized. In contrast, our result can be used to find $k \geq 2$ approximate maximum matchings for any graph whose diversity (i.e. the average Hamming distance) approximates the largest possible by a factor of 2.

We show that this problem can be restated into the budgeted matching problem [6]. As noted in [6], though the budgeted matching is in general \mathcal{NP} -hard, if both the weight and cost functions are integral and have a range bounded by a polynomial in $|V|$, then it can be solved in polynomial time with a good probability by a reduction to the exact perfect matching problem [9,34]. The exact running time for such a case is not stated explicitly in [6]. We combine the algorithm in [6] and the approach by Lawler [31] to prove:

Theorem 4. *There exists a $O(k^4|V|^7 \log^3 k|V|)$ time, $(2, 1)$ bi-approximation to the diversity computational problem for c -maximum matchings, with failure probability $1/|V|^{\Omega(1)}$.*

The details of our algorithm can be found in appendix Section 13.

DCP for Global Min-Cuts: Next, consider the diversity computational problem for computing k many c -approximate global min-cuts: given a graph G and a positive weight function w on its edges, a c -approximate min-cut is a cut C whose cut-edge set $E(C)$ satisfies $\sum_{e \in E(C)} w(e) \leq c \sum_{e \in E(C')} w(e)$ for any other cut C' . Given i cuts, we define the (integral) cost of an edge as the number of cuts in which it appears as a cut edge. Consider the BCO with cost minimization in the objective function (as the cost of a cut is now inversely proportional to its sum of distances from the found cuts) and constraint with upper bound (the weight of the cut should be at most c times that of a global min weight cut). In [2] the authors provide a polynomial-time algorithm for this problem, implying that the BCO can be solved exactly in polynomial time. This gives us a $(2, 1)$ bi-approximation to the diversity computational problem for c -approximate global minimum cuts. We remark that one may be able to exploit integrality of our cost function to obtain a faster algorithm than that in [2].

8 Application 4: Diverse Minimum Weight Matroid Bases and Minimum Spanning Trees

One of the original ways to attack the peripatetic salesman problem (Krarup [29]) was to study the k edge-disjoint spanning trees problem [13]. Note that the

existence of such trees is not guaranteed, and one can use our results in Section 5 to maximize diversity of the k trees found.

However, for an application to the TSP problem, cost conditions must be taken into account. Here we study the diverse computational problem (DCP) on minimum spanning trees: Given a weighted undirected graph $G = (V, E)$ with nonnegative weights $w(e)$, $c > 1$ and a $k \in \mathbb{N}$, return k spanning trees of G such that each spanning tree is a c -approximate minimum spanning tree, and subject to this, the diversity of the k trees is maximized. Here again the diversity of a set of trees is the sum of pairwise distances between them, and the distance between two trees is the size of their symmetric difference.

Our results in this section generalize to the problem of finding k diverse bases of a matroid such that every basis in the solution set is a c approximate minimum-weight basis. The DCP on MSTs is a special case of this problem. However, in order to not introduce extra notation and definitions here, we will describe our method for minimum spanning trees. We will then briefly sketch how to extend the algorithm to the general matroid case.

Starting with $T_1 = MST(G)$ (a minimum spanning tree on G , computable in polynomial time), assume we have obtained i trees T_1, \dots, T_i , all of which are c -approximate minimum spanning trees. Assign to each edge a length $\ell(e)$ which equals $|j : 1 \leq j \leq i, e \in T_j|$.

Claim. Given T_1, \dots, T_i , finding T_{i+1} that maximizes $\sum_{j=1}^i d(T, T_j)$ is equivalent to finding T that minimizes $\sum_{e \in T} \ell(e)$.

Proof. An explicit calculation reveals that $\sum_{e \in T} \ell(e) = (n-1)i - \sum_{j=1}^i d(T, T_j)$.

Consider now the associated *similarity budget constrained optimization problem*

$$\begin{aligned} \min \quad & \sum_{e \in T} \ell(e) \\ \text{s.t.} \quad & w(T) \leq c \cdot w(MST(G)) \\ & T \in \mathbf{Sol}_\Pi \setminus \{T_1, \dots, T_i\} \end{aligned} \tag{3}$$

Here \mathbf{Sol}_Π is just the set of spanning trees on G . We will handle the self-avoiding constraints in a similar fashion as in Section 5. For the moment consider the relaxed BCO^r where the last constraint is simply $T \in \mathbf{Sol}_\Pi$. This is a budget constrained MST with two weights. This problem has been considered by Ravi and Goemans [38], who termed it the CMST problem. They provide a $(1, 2)$ bi-approximation that runs in near-linear time, and a $(1, 1 + \epsilon)$ bi-approximation that runs in polynomial time⁹. Also, they show that the $(1, 1 + \epsilon)$ bi-approximation can be used as a subroutine to compute a $(1 + \epsilon, 1)$ bi-approximation in pseudopolynomial time.

Applying their results and observing that we are in cases 4 and 5 of Theorem 1, we get

⁹ The latter is a PTAS, not an FPTAS.

Theorem 5 (DCP for Minimum Spanning Trees). *There exists a*

- *polynomial (in n, m and k) time algorithm that outputs a $(4, 2)$ bi-approximation to the DCP problem for MSTs.*
- *polynomial (in n, m and k) and exponential in $1/\epsilon$ time algorithm that outputs a $(4, 1 + \epsilon)$ bi-approximation to the DCP problem for MSTs.*
- *pseudopolynomial time algorithm that outputs a $(4, 1)$ bi-approximation to the DCP problem for MSTs, as long as the average distance between the trees in the optimal solution to the k DCP on c -approximate minimum spanning trees does not exceed $\frac{4\epsilon(n-1)}{1+2\epsilon}$.*

Extension to Matroids: It is stated in the paper by Ravi and Goemans [38] that the same result holds if one replaces the set of spanning trees by the bases of any matroid. It is straightforward to show that the analog of Claim 8 hold in the matroid setting too. With a bit of work, one can also generalize the approach of Lawler [31] to avoid self-intersection (the bases found so far), and thus all the techniques generalize to the matroid setting. In all of this, we assume an independence oracle for the matroid, as is standard. In [19], it is shown that, given integers k, d , finding k perfect matchings so that every pair of the found matchings have Hamming distance at least d is NP-hard. This hardness result also applies to finding weighted diverse bases and weighted diverse common independent sets.

9 Conclusion

We obtained bi-approximation algorithms for the diversity computational problems associated with approximate minimum-weight bases of a matroid, shortest paths, matchings, min-cut and spanning trees. Our general reduction lemma can be applied to any computational problem for which the DUALRESTRICT version defined by Papadimitriou and Yannakakis [37] can be solved in polynomial time. One may also exploit special integrality properties of our BCO to get faster algorithms than those for general multicriteria problems. There are (at least) three open questions:

- Are there other interesting examples of problems to which our reduction lemma applies?
- It is noteworthy that while Types 3, 4 and 5 in Theorem 1 contain the class for which DUALRESTRICT can be solved efficiently, the case of $a = b = 1 + \epsilon$ missing in our theorem would include the larger GAP class studied by Papadimitriou and Yannakakis [37]. Our current techniques to prove Theorem 1 do not allow us to translate such bi-approximations of BCOs into those for DCPs (at least not without demanding certain stronger condition than in Type 5).
- Can the diversity computational problem for approximate TSP tours or Max-TSP tours be solved in polynomial time?

References

1. Abbar, S., Amer-Yahia, S., Indyk, P., Mahabadi, S., R. Varadarajan, K.: Diverse near neighbor problem. Proceedings of the Annual Symposium on Computational Geometry (06 2013) [3](#)
2. Armon, A., Zwick, U.: Multicriteria global minimum cuts. *Algorithmica* **46**(1), 15–26 (2006) [13](#)
3. Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., Protasi, M.: Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties. Springer-Verlag Berlin Heidelberg (1999) [6](#)
4. Baste, J., Fellows, M.R., Jaffke, L., Masařík, T., de Oliveira Oliveira, M., Philip, G., Rosamond, F.A.: Diversity of solutions: An exploration through the lens of fixed-parameter tractability theory. *Artificial Intelligence* **303**, 103644 (2022) [2](#)
5. Baste, J., Jaffke, L., Masařík, T., Philip, G., Rote, G.: Fpt algorithms for diverse collections of hitting sets. *Algorithms* **12**(12), 254 (2019) [2](#), [3](#), [5](#)
6. Berger, A., Bonifaci, V., Grandoni, F., Schäfer, G.: Budgeted matching and budgeted matroid intersection via the gasoline puzzle. *Math. Program.* **128**(1-2), 355–372 (2011) [13](#), [11](#)
7. Birnbaum, B., Goldman, K.: An improved analysis for a greedy remote-clique algorithm using factor-revealing LPs. *Algorithmica* **55**, 42–59 (2009) [3](#), [8](#), [2](#)
8. Borodin, A., Jain, A., Lee, H.C., Ye, Y.: Max-sum diversification, monotone submodular functions, and dynamic updates. *ACM Trans. Algorithms* **13**(3), 41:1–41:25 (2017) [3](#), [8](#)
9. Camerini, P.M., Galbiati, G., Maffioli, F.: Random pseudo-polynomial algorithms for exact matroid problems. *J. Algorithms* **13**(2), 258–273 (1992) [13](#), [11](#)
10. Cevallos, A., Eisenbrand, F., Zenklusen, R.: Max-sum diversity via convex programming. arXiv preprint arXiv:1511.07077 (2015) [3](#)
11. Cevallos, A., Eisenbrand, F., Zenklusen, R.: Local search for max-sum diversification. In: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 130–142. SIAM (2017) [3](#)
12. Chazelle, B.: A minimum spanning tree algorithm with Inverse-Ackermann type complexity. *J. ACM* **47**(6), 1028–1047 (2000) [6](#)
13. Clausen, J., Hansen, L.A.: Finding k edge-disjoint spanning trees of minimum total weight in a network: an application of matroid theory. In: Combinatorial Optimization II, pp. 88–101. Springer (1980) [13](#)
14. Commander, C.W., Pardalos, P.M., Ryabchenko, V., Uryasev, S., Zrazhevsky, G.: The wireless network jamming problem. *Journal of Combinatorial Optimization* **14**(4), 481–498 (2007) [2](#)
15. Diakonikolas, I., Yannakakis, M.: Small approximate pareto sets for biobjective shortest paths and other problems. *SIAM Journal on Computing* **39**(4), 1340–1371 (2010) [10](#)
16. Dijkstra, E.W.: A note on two problems in connexion with graphs. *Numerische mathematik* **1**(1), 269–271 (1959) [8](#)
17. Erkut, E.: The discrete p -dispersion problem. *European Journal of Operational Research* **46**(1), 48–60 (1990). [https://doi.org/https://doi.org/10.1016/0377-2217\(90\)90297-O](https://doi.org/https://doi.org/10.1016/0377-2217(90)90297-O), <https://www.sciencedirect.com/science/article/pii/0377221790902970> [3](#)
18. Fomin, F.V., Golovach, P.A., Jaffke, L., Philip, G., Sagunov, D.: Diverse pairs of matchings. arXiv preprint arXiv:2009.04567 (2020) [2](#), [5](#), [13](#)

19. Fomin, F.V., Golovach, P.A., Panolan, F., Philip, G., Saurabh, S.: Diverse collections in matroids and graphs. arXiv preprint arXiv:2101.04633 (2021) [3](#), [5](#), [15](#)
20. Gabow, H.N.: Two algorithms for generating weighted spanning trees in order. *SIAM Journal on Computing* **6**(1), 139–150 (1977) [2](#), [11](#)
21. Goldenberg, E., Karthik C. S.: Hardness amplification of optimization problems. In: Vidick, T. (ed.) 11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA. LIPIcs, vol. 151, pp. 1:1–1:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). <https://doi.org/10.4230/LIPIcs.ITCS.2020.1>, <https://doi.org/10.4230/LIPIcs.ITCS.2020.1> [6](#)
22. Hanaka, T., Kiyomi, M., Kobayashi, Y., Kobayashi, Y., Kurita, K., Otachi, Y.: A framework to design approximation algorithms for finding diverse solutions in combinatorial problems. *CoRR* [abs/2201.08940](#) (2022) [5](#)
23. Hanaka, T., Kobayashi, Y., Kurita, K., Lee, S.W., Otachi, Y.: Computing diverse shortest paths efficiently: A theoretical and experimental study. *CoRR* [abs/2112.05403](#) (2021) [5](#), [12](#)
24. Hanaka, T., Kobayashi, Y., Kurita, K., Otachi, Y.: Finding diverse trees, paths, and more. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 3778–3786 (2021) [2](#), [4](#), [5](#), [11](#)
25. Hara, S., Maehara, T.: Enumerate lasso solutions for feature selection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 31 (2017) [2](#), [11](#)
26. Hassin, R., Rubinstein, S., Tamir, A.: Approximation algorithms for maximum dispersion. *Oper. Res. Lett.* **21**(3), 133–137 (Oct 1997) [3](#)
27. Herzel, A., Bazgan, C., Ruzika, S., Thielen, C., Vanderpooten, D.: One-exact approximate pareto sets. *Journal of Global Optimization* **80**(1), 87–115 (2021) [1](#), [5](#), [6](#), [10](#)
28. Herzel, A., Ruzika, S., Thielen, C.: Approximation methods for multiobjective optimization problems: A survey. *INFORMS Journal on Computing* **33**(4), 1284–1299 (2021) [6](#)
29. Krarup, J.: The peripatetic salesman and some related unsolved problems. In: Combinatorial programming: methods and applications, pp. 173–178. Springer (1995) [1](#), [2](#), [13](#)
30. Kuby, M.: Programming models for facility dispersion: The p -dispersion and maxisum dispersion problems. *Geographical Analysis* **19**(4), 315–329 (Oct 1987). <https://doi.org/10.1111/j.1538-4632.1987.tb00133.x> [3](#)
31. Lawler, E.L.: 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) [2](#), [11](#), [13](#), [15](#), [5](#), [9](#), [12](#)
32. Lindgren, E.M., Dimakis, A.G., Klivans, A.: Exact map inference by avoiding fractional vertices. In: International Conference on Machine Learning. pp. 2120–2129. PMLR (2017) [2](#), [11](#)
33. Micali, S., Vazirani, V.V.: An $O(\sqrt{|V|} \cdot |E|)$ algorithm for finding maximum matching in general graphs. In: 21st Annual Symposium on Foundations of Computer Science. pp. 17–27 (1980) [12](#)
34. Mulmuley, K., Vazirani, U.V., Vazirani, V.V.: Matching is as easy as matrix inversion. *Comb.* **7**(1), 105–113 (1987) [13](#), [11](#)
35. Murty, K.G.: An algorithm for ranking all the assignments in order of increasing cost. *Operations research* **16**(3), 682–687 (1968) [2](#), [11](#)
36. Namorado Climaco, J.C., Queirós Vieira Martins, E.: A bicriterion shortest path algorithm. *European Journal of Operational Research* **11**(4), 399–404 (1982) [12](#)

37. Papadimitriou, C.H., Yannakakis, M.: On the approximability of trade-offs and optimal access of web sources. In: Proceedings 41st annual symposium on foundations of computer science. pp. 86–92. IEEE (2000) [1](#), [4](#), [5](#), [15](#)
38. Ravi, R., Goemans, M.X.: The constrained minimum spanning tree problem. In: Scandinavian Workshop on Algorithm Theory. pp. 66–75. Springer (1996) [14](#), [15](#)
39. Ravi, S.S., Rosenkrantz, D.J., Tayi, G.K.: Heuristic and special case algorithms for dispersion problems. *Operations Research* **42**(2), 299–310 (1994) [3](#), [7](#), [8](#)
40. Wang, D., Kuo, Y.S.: A study on two geometric location problems. *Information Processing Letters* **28**(6), 281–286 (1988). [https://doi.org/https://doi.org/10.1016/0020-0190\(88\)90174-3](https://doi.org/https://doi.org/10.1016/0020-0190(88)90174-3), <https://www.sciencedirect.com/science/article/pii/0020019088901743> [3](#)
41. Yang, H.T., Tsai, S.Y., Liu, K.S., Lin, S., Gao, J.: Patrol scheduling against adversaries with varying attack durations. In: Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems. pp. 1179–1188 (2019) [2](#)

10 Proof of Reduction Theorem

Theorem 1 (Reduction of DCP to BCO). *Consider an input (Π, k, d, c) to the diversity computational problem (DCP).*

- For metric BCO,
 1. An $(a, 1)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, d)$ can be used to give a $(2a, 1)$ approximation to the DCP, and
 2. An (a, b) bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, d)$ can be used to give a $(4a, b)$ approximation to the DCP.
- For similarity BCO,
 3. A $(1, 1)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, s)$ can be used to give a $(2, 1)$ approximation to the DCP,
 4. A $(1, b)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, s)$ can be used to give $(4, b)$ approximation to the DCP,
 5. A $(1 + \epsilon, 1)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, s)$ can be used to give $(4, 1)$ approximation to the DCP, under the condition that the average pairwise distance in the optimal solution to the DCP is at least $D \frac{4\epsilon}{1+2\epsilon}$.

In all of the above, the overhead for obtaining a bi-approximation for the DCP, given a bi-approximation for BCO problem, is $O(k)$.

Given a $BCO(\Pi, (y_1, \dots, y_i), c, d)$, denote by $\Omega(c, (y_1, \dots, y_i))$ its feasible set. When $b = 1$, the returned solution y lies in $\Omega(c, (y_1, \dots, y_i))$. On the other hand when $b > 1$, the returned solution lies in the larger space $\Omega(bc, (y_1, \dots, y_i))$. Our proof will handle these cases separately. We will first state two lemmas concerning the greedy heuristic of farthest point insertion, the first of which we suspect is folklore, but we could not find a reference. The second lemma is the heart of the argument, and is (as far as we know) novel.

Lemma 3. *Let A be a metric space, and consider the k -dispersion problem in A as in Definition 6. Consider an oracle O that, given a set $P_i := \{p_1, \dots, p_i\}$ of i points in A , returns a point $p \in A \setminus P_i$ such that $\sum_{j=1}^i d(p, p_i) \cdot a \geq \sum_{j=1}^i d(p', p_i)$ for all $p' \in A \setminus P_i$ for some constant $a > 1$. Then in k calls to O one can obtain a $2a$ -approximate solution to the k -dispersion problem on A . In other words, the total pairwise distance between the k points of the solution obtained via O is at most that of the optimal solution to the k -dispersion divided by $2a$.*

Lemma 4. *Assume A is a metric space, and $B \subseteq A$. Suppose there is an oracle O that, given a set of i points $P_i := \{p_1, \dots, p_i\}$ in A , outputs a point $p \in A \setminus P_i$ such that $\sum_{j=1}^i d(p_j, p) \geq \max_{p' \in B} \sum_{j=1}^i d(p_j, p')$. Then the oracle can be used to give a set of k points in A whose diversity is at least $1/4$ of an optimal solution to the k -dispersion problem in B .*

Assuming the above lemmas, we now prove the theorem.

- Type 1: An $(a, 1)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, d)$ returns a point in $\Omega(c, (y_1, \dots, y_i))$. Therefore we only apply Lemma 3 and we are done.
- Type 2: An (a, b) bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, d)$ returns a point in $\Omega(bc, (y_1, \dots, y_i))$, that is an a -approximation to the farthest point in $\Omega(c, (y_1, \dots, y_i))$. Applying Lemma 3 and Lemma 4 together, we get the claimed result.
- Type 3: A $(1, 1)$ bi-approximation to the similarity $BCO(\Pi, (y_1, \dots, y_i), c, s)$ returns the farthest point in $\Omega(c, (y_1, \dots, y_i))$, and so this case follows from the 2-approximation guarantee of the farthest insertion heuristic, without applying either lemma.
- Type 4: A $(1, b)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, d)$ returns the farthest point in $\Omega(bc, (y_1, \dots, y_i))$. In this case, we only apply Lemma 4 to get the claimed result.

We prove the remaining case of Type 5 separately. A $(1+\epsilon, 1)$ bi-approximation to $BCO(\Pi, (y_1, \dots, y_i), c, s)$ returns a point in $\Omega(c, (y_1, \dots, y_i))$. However, *this point is not the farthest point, nor necessarily an approximation of it*. This is because the guarantee is only on the total similarity, and not the total distance. In other words, a $1 + \epsilon$ approximation to $g_s(y) := \sum_{j=1}^i s(y, y_j)$ is not a $1 + \epsilon$ approximation to $f_d(y) := \sum_{j=1}^i d(y, y_j)$ – in fact, these functions are the “opposite” of each other, as $f_d(y) = Di - g_s(y)$.

Let y^* be the point in $\Omega(c, (y_1, \dots, y_i))$ that maximizes $f_d(y)$, i.e., the farthest point. The $(1 + \epsilon, 1)$ bi-approximation algorithm returns a point y such that $g_s(y) \leq a \cdot g_s(y^*)$. Consider the condition that the average pairwise distance in the optimal solution OPT to the DCP is at least $D \frac{4\epsilon}{1+2\epsilon}$. We refer the reader now to Lemma 1 in [7], which proves that during farthest point insertion, at step $i+1$ in the algorithm when $P_i = \{p_1, \dots, p_i\}$ have been selected, there is a point $p \notin P_i$ whose average distance is at least $OPT/2$. While this lemma seems like it can only be applied to point sets P_i formed during an iterated farthest insertion, a careful analysis of Lemmas 3,4,5 and 6 in [7] used to prove it reveals that the argument does not demand that this be the case. That is, for any set P_i of i points, there always exist a $p \notin P_i$ whose average distance is at least $OPT/2$.

This implies the existence of a point $y' \in \Omega(c, (y_1, \dots, y_i))$ such that the average distance of y' to the y_i s is at least $OPT/2 \geq D \frac{2\epsilon}{1+2\epsilon} = D(1 - \frac{1}{1+2\epsilon})$. Since y^* is the farthest point, the same condition holds for y^* .

This implies that the total distance $f_d(y^*) \geq Di(1 - \frac{1}{1+2\epsilon})$, implying that $g_s(y^*) = Di - f_d(y^*) \leq Di/(1 + 2\epsilon)$. We then have the following string of inequalities

$$\begin{aligned}
g_s(y^*) &\leq \frac{Di}{1+2\epsilon} \\
\iff Di &\geq (1+2\epsilon)g_s(y^*) \\
\iff \frac{Di}{2} &\geq \left(\frac{1}{2} + \epsilon\right)g_s(y^*) \\
\iff Di - (1+\epsilon)g_s(y^*) &\geq \frac{1}{2}(D - g_s(y^*))
\end{aligned}$$

But we know that the returned point y satisfies $g_s(y) \leq (1+\epsilon)g_s(y^*)$, implying that $f_d(y) = Di - g_s(y) \geq \frac{1}{2}(D - g_s(y^*)) = f_d(y^*)/2$. Hence y is a 2-approximation to the farthest point, and we can apply Lemma 3 to complete the proof.

10.1 Proofs of Lemmas 3 and 4

Proof of Lemma 3: We will prove the statement of the lemma for the average distance, which will imply the claim for the total distance. Let OPT be the average distance between solutions in the optimal solution to the k -dispersion problem. By Lemma 1 in [7], at any step i , there is always a solution whose distance to the already selected set P_i is at least $OPT/2$, and hence the same holds for the farthest solution. This implies that the point p returned by the algorithm has a distance at least $OPT/2a$. A simple inductive argument now finishes the proof: assume the lemma holds until stage i , i.e., the average pairwise distance between points in P_i is at least $OPT/2a$. When $p_{i+1} = p$ is chosen, its average distance to the i points in P_i is at least $OPT/2a$ by the above reasoning, and so the P_{i+1} has an average pairwise distance of at least $OPT/2a$.

Proof of Lemma 4 For a point $p \in A$, denote by $n_B(p)$ its closest point in B . By definition, if $p \in B$ then $n_B(p) = p$.

Assume that the oracle O in the hypothesis of the lemma exists. Then given the point set P_i , consider the point set $Q_i = \{n_B(p_j) : 1 \leq j \leq i\}$. As mentioned before, a careful analysis of Lemmas 3-6 in [7] reveals that there always exists a point $p^* \in B$ whose average distance to the points in Q_i is at least $OPT(B)/2$, where $OPT(B)$ denotes the optimal average pairwise distance between points in the optimal solution to the k dispersion problem on B .

Consider the three points $p_i, n_B(p_i)$ and p^* . By the triangle inequality, $d(p_i, p^*) \geq d(p^*, n_B(p_i)) - d(p_i, n_B(p_i))$. However, by definition, $d(p_i, n_B(p_i)) \leq d(p_i, p^*)$. This implies that $d(p_i, p^*) \geq d(p^*, n_B(p_i)) - d(p_i, p^*)$, implying that $d(p_i, p^*) \geq d(p^*, n_B(p_i))/2$. Summing over all i , we get that the total (and hence the average) distance of p^* to points P_i is at least half of its total distance to points in Q_i . Since the average distance of p^* to points in Q_i is at least $OPT(B)/2$, we get that the average distance of p^* to points in P_i is at least $OPT(B)/4$.

However, the oracle O returns a point $p \in A$ whose average distance to P_i is at least that of p^* to points in P_i , which as argued above is at least $OPT(B)/4$.

By the same inductive argument as in Lemma 3, we get that the average distance between points returned by k calls to O is at least $OPT(B)/4$, proving that it gives a 4-approximation.

11 Diverse Spanning Trees

In this section, we discuss the diverse spanning trees problem, which is the diversity computational problem for spanning trees with Hamming distance function as the diversity measure. Let $G = (V, E)$ be an undirected graph. The problem aims to output a set S of k spanning trees T_1, \dots, T_k of G such that the sum of the pairwise distances $\sum_{i,j \in S} d(T_i, T_j)$ is maximized, where d is the Hamming distance between the edge sets of the trees. While this problem actually has an exact algorithm running in time $O((kn)^{2.5}m)$ [24], we get a faster approximation algorithm.

Theorem 6. *Given a simple graph $G = (V, E)$, there exists an $O(knma(n, m))$ -time algorithm (here $\alpha(\cdot)$ denotes the inverse of the Ackermann function) that generates k spanning trees T_1, \dots, T_k , such that the sum of all pairwise Hamming distances is at least half of an optimal set of k diverse spanning trees.*

Using our reduction Theorem 1, we observe that since one is only required to output spanning trees, the associated budget constrained problem has no inequalities. Having obtained i trees T_1, \dots, T_i (starting with an arbitrary spanning tree T_1), the BCO problem looks like

$$\begin{aligned} \max \quad & f_d(y) := \sum_{j=1}^i d(T, T_j) \\ \text{s.t.} \quad & T \in \text{Sol}_\Pi \setminus \{T_1, \dots, T_i\} \end{aligned} \tag{4}$$

Note that the relaxed budget constrained problem then simply asks to maximize $\sum_{j=1}^i d(T, T_j)$. We first show how to solve this problem exactly, and then adapt the approach by Lawler [31] to handle the self-avoiding constraint. The algorithm to maximize $\sum_{j=1}^i d(T, T_j)$ is very simple: give each edge e a weight $w(e) = \sum_{j=1}^i \mathbb{1}(e \in T_j)$ and compute the minimum spanning tree T with respect to these edge weights.

Lemma 5. *The minimum spanning tree T with respect to the edge weights w satisfies $\sum_{j=1}^i d(T_j, T) \geq \sum_{j=1}^i d(T_j, T')$ for any spanning tree T' .*

The proof of Lemma 5 is at the end of this section. If $T \neq T_j$ for all $1 \leq j \leq i$, that is, the new spanning tree is different from all previous trees, then we are done and Theorem 2 is proved using Lemma 5.

Self-Avoiding Constraint However, this is not guaranteed as our measure is the sum-of-distances and not the minimum distance. Note that for furthest point insertion, we need the point that is furthest from the current set, but does not belong to the current set. This is an issue we face because of the implicit nature of the furthest point procedure in the exponential-sized space of spanning trees: in the metric k -dispersion problem, it was easy to guarantee distinctness as one only considered the $n - i$ points not yet selected.

We now show how to guarantee that the new tree is distinct. In case that $T = T_j$ for some $1 \leq j \leq i$, we use the approach by Lawler [31]. We will

obtain $i + 1$ distinct spanning trees $\tilde{T}_1, \dots, \tilde{T}_{i+1}$, at least one of which must then be distinct from T_1, \dots, T_i , which will be chosen as T_{i+1} . Set $\tilde{T}_1 = T_j$. For every edge $e \in \tilde{T}_1$, we find the minimum spanning tree T_e (with respect to the same edge weights $w(e)$ as before) after deleting e from the graph. Among the collection $\{T_e\}_{e \in \tilde{T}_1}$ of trees thus obtained, we find the one whose sum of distances from T_1, \dots, T_i is as large as possible, and set it to be \tilde{T}_2 . Note that $\tilde{T}_2 \neq \tilde{T}_1$. If $\tilde{T}_2 \neq T_j$ for all $1 \leq j \leq i$, then we are done; else, we repeat the above procedure to obtain the collection $\{T_e\}_{e \in \tilde{T}_2}$ and set \tilde{T}_3 as the best one in $\{T_e\}_{e \in \tilde{T}_1} \cup \{T_e\}_{e \in \tilde{T}_2} \setminus \{\tilde{T}_1, \tilde{T}_2\}$, after which we are either done or we repeat to get \tilde{T}_4 , and so on. We stop after obtaining \tilde{T}_{i+1} , and select the $\tilde{T}_{(.)}$ that is distinct from all of T_1, \dots, T_i . The proof of Lemma 5 implies that the T_{i+1} thus obtained is the furthest tree from the T_1, \dots, T_i among all trees $T \notin \{T_1, \dots, T_i\}$, and we have accomplished furthest point insertion in the space of spanning trees. In total, the above needs to compute $O(kn)$ instances of the minimum spanning tree, so the running time is $O(knma(n, m))$ [12].

11.1 Proof of Lemma 5

Proof. Given input spanning trees T_1, \dots, T_i , suppose T^A is the tree returned by the algorithm in Section 4.1, and T^* is an optimal spanning tree. Let E_j be the subset of edges with weight j . That is, $E_j = \{e \in E \mid w(e) = j\}$ for all $0 \leq j \leq i$ and partition E as $E = E_0 \cup E_1 \cup \dots \cup E_i$.

Let $c_j = |T^A \cap E_j|$ and $\beta_j = |T^* \cap E_j|$. To prove that T^A is optimal, we will show that $\sum_{j=1}^i d(T^A, T_j) \geq \sum_{j=1}^i d(T^*, T_j)$.

Define $A = \sum_{j=1}^i d(T^A, T_j)$ and $B = \sum_{j=1}^i d(T^*, T_j)$. Plugging α_j and β_j into the sum of pairwise distances from our tree and the optimal tree to constructed trees, we get

$$A = \sum_{j=1}^i d(T^A, T_j) = (n-1)i - \sum_{j=1}^i j\alpha_j$$

and

$$B = \sum_{j=1}^i d(T^*, T_j) = (n-1)i - \sum_{j=1}^i j\beta_j.$$

We will show that $A - B = \sum_{j=0}^i j(\beta_j - \alpha_j) \geq 0$. We first express $A - B$ as the dot product of two vectors

$$A - B = \langle 0, 1, \dots, i \rangle \langle \beta_0 - \alpha_0, \beta_1 - \alpha_1, \dots, \beta_i - \alpha_i \rangle^t,$$

where $\langle \cdot \rangle^t$ denotes transpose, and makes a column vector out of a row vector. We develop some more notation:

$$E^j := E_0 \cup E_1 \cup \dots \cup E_j$$

$$\alpha^j := \alpha_0 + \alpha_1 + \dots + \alpha_j$$

$$\begin{aligned}
\beta^j &:= \beta_0 + \beta_1 + \cdots + \beta_j \\
c^j &:= \langle 0, 1, \dots, j \rangle \\
\alpha_v^j &:= \langle \alpha_0, \alpha_1, \dots, \alpha_j \rangle^t \\
\beta_v^j &:= \langle \beta_0, \beta_1, \dots, \beta_j \rangle^t
\end{aligned}$$

With this we have that $A - B = c^i(\beta_v^i - \alpha_v^i)$.

We will show $c^i\beta_v^i \geq c^i\alpha_v^i$ by induction on i , which will prove that $A \geq B$, completing the proof.

We first claim that $\alpha^j \geq \beta^j$ for all $0 \leq j \leq i$. To prove this, fix a j and let C_1 be the number of connected components (including any isolated vertices) of $T^A \cap E^j$ and C_2 be the number of connected components of $T^* \cap E^j$. Clearly, $C_1 \leq C_2$ since for Kruskal's algorithm for minimum spanning trees is the greedy algorithm that adds as many edges from E^j into the solution as possible without creating a cycle, and the addition of an edge decreases the number of connected components by one. Since, $C_1 = n - \alpha^j$ and $C_2 = n - \beta^j$, $C_1 \leq C_2$ implies that $\alpha^j \geq \beta^j$.

Now, we continue our proof that $c^i\beta_v^i \geq c^i\alpha_v^i$ by induction on i . Observe $\|\alpha_v^i\|_1 = \|\beta_v^i\|_1 = n' := n - 1$, as both spanning trees contain exactly $n - 1$ edges.

Base Case For the base case of $i = 1$, $c^1 = \langle 0, 1 \rangle$, $\alpha_v^1 = \langle \alpha_0, n' - \alpha_0 \rangle$, and $\beta_v^1 = \langle \beta_0, n' - \beta_0 \rangle$. From the previously mentioned property, we know $\alpha_0 = \alpha^0 \geq \beta^0 = \beta_0$. Hence, $c^1\beta_v^1 = n' - \beta_0 \geq n - \alpha_0 = c^1\alpha_v^1$.

The induction hypothesis assumes that $c^i\beta_v^i \geq c^i\alpha_v^i$ for $i = K - 1$, i.e., the statement is true for two vectors of length $K - 1$. Now, consider when $i = K$. Since $\|\alpha_v^{K-1}\|_1 = \alpha^{K-1} \geq \|\beta_v^{K-1}\|_1 = \beta^{K-1}$, there are two cases,

Case 1: $\|\beta_v^{K-1}\|_1 = \|\alpha_v^{K-1}\|_1$ Then since $\|\beta_v^K\|_1 = \|\alpha_v^K\|_1$, we get that $\alpha_K = \beta_K$. Using the induction hypothesis, we get

$$c^K\beta_v^K = c^{K-1}\beta_v^{K-1} + K\beta_K \geq c^{K-1}\alpha_v^{K-1} + K\alpha_K = c^K\alpha_v^K.$$

Case 2) $\|\alpha_v^{K-1}\|_1 > \|\beta_v^{K-1}\|_1$. Then $\|\alpha_v^K\|_1 = \|\beta_v^K\|_1$ yields $\alpha_K < \beta_K$. Let $d = \beta_K - \alpha_K = \|\alpha_v^{K-1}\|_1 - \|\beta_v^{K-1}\|_1$. Then, $c^K\beta_v^K - c^K\alpha_v^K = c^{K-1}(\beta_v^{K-1} - \alpha_v^{K-1}) + Kd$. We claim that this quantity is positive. This is because the minimum value of $c^{K-1}(\beta_v^{K-1} - \alpha_v^{K-1})$, subject to the conditions that a) $\|\alpha_v^{K-1}\|_1 - \|\beta_v^{K-1}\|_1 = d > 0$ and b) $\|\alpha_v^K\|_1 = \|\beta_v^K\|_1$, occurs when $\beta_j = \alpha_j$ for all $0 \leq j \leq K - 2$, $\beta_{K-1} = \alpha_{K-1} - d$, and $\beta_K = \alpha_K + d$. For this setting, the value of $c^{K-1}(\beta_v^{K-1} - \alpha_v^{K-1}) + Kd$ equals $-(K-1)d + Kd = d$, which is positive, completing the proof.

12 DCP on Approximate Shortest Paths

Lemma 2 (Exact solution to the relaxed BCO^r problem). *Given a real $c \geq 1$ and a directed simple graph $G = (V \cup \{s, t\}, E)$ associated with two weight functions on edges $\omega : E \rightarrow \mathbb{R}^+$ and $f : E \rightarrow \{0, 1, \dots, r\}$, there is an $O(r|V|^3)$ -time algorithm to output an st -path P^* so that $\sum_{e \in E(P^*)} f(e)$ is minimized while retaining $\sum_{e \in E(P^*)} \omega(e) \leq c \sum_{e \in E(P)} \omega(e)$ for all st -paths P .*

Proof. Let $D(.,.)$ be a $|V| \times r(|V| - 1)$ table with the entry values defined as the following:

$$D(v, c) = \min \left\{ \sum_{e \in E(P)} \omega(e) : P \text{ is an } sv\text{-path with } \sum_{e \in E(P)} f(e) \leq c \right\}.$$

If the RHS takes the minimum from an empty set, then we define $D(v, c) = \infty$. By definition, P^* has $\sum_{e \in E(P^*)} f(e) = c^*$ so that

$$c^* = \min\{c \in [0, r(|V| - 1)] : D(t, c) \leq c \cdot \text{dis}_\omega(s, t)\}$$

where $\text{dis}_\omega(s, t)$ denotes the distance under ω from s to t . By backtracking from the $D(t, c^*)$, one can construct P^* from the table D in $O(r|V|)$ time. Hence, our goal is to fill out the entries in D using $O(r|V|^3)$ time.

To compute $D(v, 0)$ for all $v \in V$, it suffices to run Dijkstra's algorithm [16] on the graph obtained from G with the removal of the edges whose $f(e) > 0$. This step takes $O(|V|^2)$ time.

Given $D(v, c')$ for all $v \in V, c' \in [0, c - 1]$, to compute $D(v, c)$ for all $v \in V$ there are two cases to discuss. We say an sv -path P *admits* $D(v, c')$ if P has $\sum_{e \in E(P)} \omega(e) = D(v, c')$ and $\sum_{e \in E(P)} f(e) \leq c'$.

Case 1. Some sv -path (s, \dots, x, v) with $f((x, v)) > 0$ admits $D(v, c)$.

Case 2. Some sw -path $P = (s, \dots, y, w)$ with $f((y, w)) > 0$ admits $D(w, c)$, and P together with some edges e whose $f(e) = 0$ admits $D(v, c)$.

For Case 1, it suffices to check $D(x, c - f((x, v)))$ for all the neighbors of v with $f((x, v)) > 0$. This step takes $O(|V|)$ time. For Case 2, observe that $D(w, c)$ can be determined by Case 1. It suffices to initialize $D(v, c)$ for all $v \in V$ as the values obtained from Case 1 and update $D(v, c)$ if there is a path Q from some node w to v so that every edge $e \in E(Q)$ has $f(e) = 0$ and $D(w, c) + \sum_{e \in E(Q)} \omega(e) < D(v, c)$. This can be implemented in $O(|V|^2)$ time for each c by running Dijkstra's algorithm on the graph $G = (V, E \cup E_1 \setminus E_2)$ where E_1 is a set of new directed edges (s, v) for all $v \in V$ with $f((s, v)) = 0$ and $\omega((s, v))$ as the initial $D(v, c)$ obtained from Case 1, and E_2 is the set of all the edges $e \in E$ that have $f(e) > 0$. It is worth noting that, though we do not explicitly check whether the found paths contain cycles or not (i.e. not simple paths), the paths found to admit $D(v, c)$ must be simple since all shortest paths in a positive weight graph are simple.

To sum up, an optimal solution P^* can be found in $O(r|V|^3)$ time.

Corollary 1 (Exact solution to the BCO problem). *Given a real $c \geq 1$, a directed simple graph $G = (V \cup \{s, t\}, E)$ associated with two weight functions on edges $\omega : E \rightarrow \mathbb{R}^+$, $f : E \rightarrow \{0, 1, \dots, r\}$, and two disjoint subsets of edges $E_{in}, E_{ex} \subseteq E$ so that all edges in E_{in} together form a directed simple path P_{prefix} starting from node s , there exists an $O(r|V|^3)$ -time algorithm to output an c -st-shortest path P^* under ω so that $\sum_{e \in E(P^*)} f(e)$ is minimum among all the c -st-shortest paths P that contain P_{prefix} as a prefix and contain no edges from E_{ex} , if such an c -st-shortest path exists.*

Proof. Let $P_{\text{prefix}} = (s, \dots, v)$. This suffices to find a feasible vt -path $P_{\text{suffix}} = (v, \dots, t)$ on the graph obtained from G with the removal of all the nodes in P_{prefix} except v and with the removal of all edges in E_{ex} and those incident to the deleted nodes. This can be computed by Lemma 2.

Theorem 3 ((2,1) bi-approximation to the Diversity Problem on Shortest Paths). *For any directed simple graph $G = (V \cup \{s, t\}, E)$, given a $c > 1$ and a $k \in \mathbb{N}$, there exists an $O(k^3|V|^4)$ -time algorithm that, if G contains at least k distinct c -st-shortest paths, computes a set S of k distinct c -st-shortest paths so that the sum of all pairwise Hamming distances between two paths in S is at least one half of the maximum possible; otherwise, reports “Non-existent.”*

Proof. Let P_1 be an arbitrary 1-st-shortest path (i.e. a shortest path from s to t), which can be computed in $O(|V|^2)$ time by Dijkstra’s algorithm. We perform the furthest insertion mentioned in Lemma 1 to obtain the c -st-shortest paths P_2, P_3, \dots, P_k . The collection $\{P_i : i \in [1, k]\}$ is a desired solution.

To perform the i -th furthest insertion for $i \in [1, k - 1]$, we need an algorithm for the problem that, given $\{P_j : 1 \leq j \leq i\}$, find an c -st-shortest path P_{i+1} so that the following conditions both hold:

1. $\sum_{e \in E(P_{i+1})} f_i(e)$ is minimum among all c -st-shortest paths where $f_i(e)$ for $e \in E$ is defined as the number of times that e appears in P_j for all $j \in [1, i]$.
Hence, f_i maps E to $\{0, 1, \dots, i\}$.
2. $P_{i+1} \neq P_j$ for all $j \in [1, i]$.

To find an c -st-shortest path that satisfies the first condition, it is an instance of the bicriteria shortest path problem with two weight functions ω and f_i . By Lemma 2, this can be done in $O(i|V|^3)$ time.

However, the c -st-shortest path found by minimizing $\sum_e f_i(e)$ is not necessarily different from those found in the previous furthest insertions. To remedy, one can find the c -st-shortest paths whose $\sum_e f_i(e)$ are the j -th smallest for all $j \in [1, i + 1]$. Some of the $i + 1$ c -st-shortest paths can meet the second condition. This step can be implemented by the standard approach due to Lawler [31]. That is, suppose in time T one can compute P_{i+1} with an additional constraint that some given edge set $E_{i,in}$ (all edges in $E_{i,in}$ together form a directed path starting from node s) must be contained in $E(P_{i+1})$ and some given edge set $E_{i,ex}$ must not be contained in $E(P_{i+1})$, then enumerating the $i + 1$ smallest candidates for P_{i+1} can be done in $O(i|V|T)$ time. By Corollary 1, $T = O(i|V|^3)$.

Thus, the i -th furthest insertion can be done in $O(i^2|V|^4)$ time. In total, our approach needs $O(k^3|V|^4)$ time as desired.

13 DCP on Approximate Maximum Matchings

We show that the diverse c -maximum matchings can be 2-approximated by a randomized polynomial-time algorithm with high probability.

Define the distance between two matchings as the Hamming distance between the edge sets. The associated relaxed BCO^r is a maximization problem with lower bounded constraints. It can be restated as: given a non-empty set S of c -maximum matchings, find a c -maximum matching $M \notin S$ so that $W(S \cup \{M\})$ is maximum among all $W(S \cup \{M'\})$ for c -maximum matching $M' \notin S$ (here $W(X)$ is the sum of pairwise distances between all matchings in X). This problem can be restated into the budgeted matching problem [6]. As noted in [6], though the budgeted matching is in general \mathcal{NP} -hard, if both the weight and cost functions are integral and have a range bounded by a polynomial in $|V|$, then it can be solved in polynomial time with a good probability by a reduction to the exact perfect matching problem [9,34]. The exact running time for such a case is not stated explicitly in [6], and we analyze it as below.

Lemma 6 (Restricted Budgeted Matching Problem). *Given an undirected simple graph $G = (V, E)$ and a cost function $c : E \rightarrow \{0, 1, \dots, r\}$ on edges, there exists an $O(r^2|V|^6 \log^2 r|V|)$ -time randomized algorithm that can find a matching of smallest cost among all c -maximum matchings with some constant success probability where the cost of a matching is the sum of the costs of all edges in the matching.*

Proof. Suppose that, for all $x \in [0, X]$ with $X = O(|V|)$, for all $y \in [0, Y]$ with $Y = O(r|V|)$, we can decide whether G contains a matching M that consists of x edges and has cost y . Then, the desired matching M^* can be found.

As the reduction in [6], we construct an edge-weighted graph $H = (U, F)$ so that

- $U = V \cup \{z_1, z_2, \dots, z_{|V|}\}$ where $z_i \notin V$ for all $i \in [1, |V|]$,
- $F = E \cup \{\{z_i, z_j\} : i \neq j \in [1, |V|]\} \cup \{\{x, z_i\} : x \in V, i \in [1, |V|]\}$, and
- for every edge $e \in F$, if e is also in E , then it has weight $\Gamma + c(e)$ for some sufficiently large integer Γ to be determined later; otherwise, it has weight 0.

By setting $\Gamma \geq r\lceil|V|/2\rceil + 1$, G has a matching M that consists of x edges and has cost y if and only if H has a perfect matching of weight $x\Gamma + y$. This is an instance of the exact perfect matching problem [9,34], which can be solved by a randomized algorithm in $O(r|V|^4 \log r|V|)$ time with some constant success probability. Summing over all choices of x and y , the running time is $O(r^2|V|^6 \log^2 r|V|)$.

We are ready to prove our main result for the diverse c -maximum matchings.

Theorem 7. *Given an undirected simple graph $G = (V, E)$ and an integer k , there exists an $O(k^4|V|^7 \log^3 k|V|)$ -time randomized algorithm that computes a set S of k distinct c -maximum matchings so that the sum of all pairwise Hamming distances between two matchings in S is at least half of the maximum possible with failure probability $1/|V|^{\Omega(1)}$, if G contains k distinct c -maximum matchings.*

Proof. Let M_1 be an arbitrary maximum matching, which can be computed in $O(|V|^{2.5})$ time [33]. For each $i \in [1, k-1]$, given $S_i = \{M_j : j \in [1, i]\}$, find a c -maximum matching $M_{i+1} \notin S_i$ so that $W(S_i \cup M_{i+1})$ is maximum among all $W(S_i \cup M')$ for c -maximum matching $M' \notin S_i$. By Lemma 6, set the cost function $c_i(e)$ for every $e \in E$ as the number of times that edge e appears in M_j for all $j \in [1, i]$ and apply the algorithm for Lemma 6 to solve the instance $(G = (V, E), c_i, c)$. This returns an c -maximum matching M^\dagger so that $W(S_i \cup M^\dagger)$ is maximum among all $W(S_i \cup M')$ that M' is an c -maximum matching and may or may not be contained in S_i . Thus we need to guarantee self-avoidance.

To remedy, we enumerate the best $i+1$ candidates for M_{i+1} . Some of them is not contained in S_i . This enumeration can be done by running the algorithm for Lemma 6 $O(i|V|)$ times, each of which preselects some edges to be included in and some to be excluded from M_{i+1} , as the Lawler's approach [31]. For each $i \in [1, k-1]$, this step takes $O(k^3|V|^7 \log^3 k|V|)$ time and may fail with probability $(k|V|)^{-\Omega(1)}$. The failure probability is $(k|V|)^{-\Omega(1)}$ rather than a constant because we can run the algorithm that may err $O(\log k|V|)$ times.

Summing up the running time for the $k-1$ furthest insertions, the total running time is $O(k^4|V|^7 \log^3 k|V|)$, and the failure probability is $1/|V|^{\Omega(1)}$ by the Union bound.