Techniques for Generalized Colorful k-Center Problems

Georg Anegg
ETH Zürich, Switzerland

Laura Vargas Koch
ETH Zürich, Switzerland
University of Chile, Santiago, Chile

Rico Zenklusen
ETH Zürich, Switzerland

Abstract

Fair clustering enjoyed a surge of interest recently. One appealing way of integrating fairness aspects into classical clustering problems is by introducing multiple covering constraints. This is a natural generalization of the robust (or outlier) setting, which has been studied extensively and is amenable to a variety of classic algorithmic techniques. In contrast, for the case of multiple covering constraints (the so-called colorful setting), specialized techniques have only been developed recently for k-Center clustering variants, which is also the focus of this paper.

While prior techniques assume covering constraints on the clients, they do not address additional constraints on the facilities, which has been extensively studied in non-colorful settings. In this paper, we present a quite versatile framework to deal with various constraints on the facilities in the colorful setting, by combining ideas from the iterative greedy procedure for Colorful k-Center by Inamdar and Varadarajan with new ingredients. To exemplify our framework, we show how it leads, for a constant number γ of colors, to the first constant-factor approximations for both Colorful Matroid Supplier with respect to a linear matroid and Colorful Knapsack Supplier. In both cases, we readily get an $O(2^\gamma)$-approximation.

Moreover, for Colorful Knapsack Supplier, we show that it is possible to obtain constant approximation guarantees that are independent of the number of colors γ, as long as $\gamma = O(1)$, which is needed to obtain a polynomial running time. More precisely, we obtain a 7-approximation by extending a technique recently introduced by Jia, Sheth, and Svensson for Colorful k-Center.

2012 ACM Subject Classification Theory of computation → Facility location and clustering

Keywords and phrases Approximation Algorithms, Fair Clustering, Colorful k-Center

Digital Object Identifier 10.4230/LIPIcs.ESA.2022.7

Related Version *Full Version*: http://arxiv.org/abs/2207.02609

Funding *Georg Anegg*: Research supported in part by Swiss National Science Foundation grant number 200021_184622.
Rico Zenklusen: Research supported in part by Swiss National Science Foundation grant number 200021_184622. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 817750).

1 Introduction

As more and more decisions are automated, there has been an increasing interest in incorporating fairness aspects in algorithms by design. This applies in particular to clustering problems, where considerable attention has recently been dedicated to developing and studying various models of fair clustering, see, e.g., [8], [3], and [2].

In this paper, we focus on the so-called colorful setting, which was introduced in [3]. In colorful clustering, each client is a member of certain subgroups and every clustering is required to cover at least a given number of clients of each subgroup. This may be considered under various clustering objectives (like k-median and k-mean), though only the k-center case has been studied so far.

Colorful clustering is an appealing notion as it is a natural generalization of the robust (or outlier) setting, where there is only a single group which every client belongs to. Various clustering problems have been studied in depth in the robust setting, see, e.g., [5], [9], and [2].

While the robust setting is amenable to a variety of well-known and basic algorithmic techniques, the only constant-factor approximations for the colorful setting, which imposes multiple covering constraints leading to more balanced clusterings, are based on significantly more sophisticated techniques, tailored specifically to those settings. More precisely, three distinct techniques have been successful at achieving constant-factor approximations in the context of colorful k-center clustering, namely the combinatorial approach of [11], the round-or-cut-based approach of [1], and the iterative greedy reductions of [10].

However, these approaches do not immediately generalize to variants with constraints on the facilities, even for the common Matroid Center or Knapsack Center clustering variants. On the other hand, techniques for the Knapsack and Matroid k-Center problems in the robust setting (see [5] and [9]) do not easily extend to multiple covering constraints.

Thus, prior to this work, no approaches have been known that lead to constant-factor approximations for colorful variants of otherwise well-studied k-center problems like Matroid Center or Knapsack Center. Filling this gap is the goal of this paper.

1.1 Our contributions

Our main contribution is a partitioning procedure which leads to a general reduction of colorful k-center clustering problems with constraints on the facilities to a significantly simpler multi-dimensional covering problem (see Theorem 3). This reduction comes at the cost of a constant factor depending on the number of colors.

It is inspired by recent insights of [10] on decoupling multiple covering constraints and iteratively applying a greedy partitioning procedure of [6]. By taking into account multiple colors at the same time, our framework gives an improved way of dealing with multiple covering constraints while also becoming more versatile. Our framework also extends and simplifies ideas of the approximation algorithm for Robust Matroid Center of [7].

We start by introducing the γ-Colorful $\mathcal{F}$-Supplier problem, which formalizes colorful k-center problems with (down-closed) constraints on the facilities.

▶ **Definition 1** (γ-Colorful $\mathcal{F}$-Supplier problem). *Let $(C \,\dot\cup\, F, d)$ be a finite metric space on a set of clients C and facilities F, let $\mathcal{F} \subseteq 2^F$ be a down-closed family of subsets of F, and let $\gamma \in \mathbb{Z}_{\geq 0}$. Moreover, we are given for each $\ell \in [\gamma]$:*

- *a unary encoded weight/color function $w_\ell : C \to \mathbb{Z}_{\geq 0}$, and*
- *a covering requirement $m_\ell \in \mathbb{Z}_{\geq 0}$.*

The γ-Colorful $\mathcal{F}$-Supplier *problem asks to find the smallest radius r together with a set $S \subseteq \mathcal{F}$ such that $w_\ell(B_C(S, r)) \geq m_\ell$ for all $\ell \in [\gamma]$.*[1]

[1] We use the common notation $w(T) := \sum_{t \in T} w(t)$ for functions $w\colon U \to \mathbb{R}_{\geq 0}$ and $T \subseteq U$, as well as $B(q, r) := \{v \in C \cup F \mid d(q, v) \leq r\}$ for the ball of radius r around point q. Moreover, we use the shorthand $B_U(V, r) := \left\{U \cap \bigcup_{v \in V} B(v, r)\right\}$ for sets $U, V \subseteq C \cup F$.

We note that it is also common to define colorful k-center versions in an unweighted way (thus not using weight functions w_ℓ) by assigning to each client a subset of the γ many colors and requiring that, for each color, m_ℓ many clients of that color are covered. The definition we use clearly captures this case (and can easily be seen to be equivalent). This connection also explains why the weights w_ℓ are assumed to be given in unary encoding.

Following common terminology in the literature, when $\mathcal{F}$ is the family of independent sets of a matroid or feasible sets with respect to a knapsack constraint, we call the problem *γ-Colorful Matroid Supplier* and *γ-Colorful Knapsack Supplier*, respectively.

Our main contribution is a general reduction of γ-Colorful $\mathcal{F}$-Supplier to an auxiliary problem, which we call $\mathcal{F}$-Cover-Promise ($\mathcal{F}$-CP). $\mathcal{F}$-CP, which is formally defined below, is a multi-dimensional cover problem with the added promise that highly structured solutions exist. The promise is key, as the problem without the promise can be thought of as a multi-dimensional max-cover problem.

▶ **Definition 2** ($\mathcal{F}$-Cover-Promise ($\mathcal{F}$-CP)). *In the $\mathcal{F}$-Cover-Promise problem ($\mathcal{F}$-CP), we are given a set family $\mathcal{H} \subseteq 2^{\mathcal{U}}$ over a finite universe $\mathcal{U}$, a family $\mathcal{F} \subseteq 2^{\mathcal{H}}$ of feasible subsets of $\mathcal{H}$, and γ many unary encoded weight functions $w_1, \dots, w_\gamma : \mathcal{U} \to \mathbb{R}_{\geq 0}$ each with a requirement m_ℓ (for $\ell \in [\gamma]$). The task is to find a feasible family of sets $S \in \mathcal{F}$ such that*

$$w_\ell\left(\bigcup_{H\in S} H\right) \geq m_\ell \qquad \forall\, \ell \in [\gamma] \ .$$

The promise is that there exists a family $S \subseteq \mathcal{F}$ and a way to pick for each $H \in S$ a single representative $u_H \in H$ such that

$$w_\ell\left(\{u_H \colon H \in S\}\right) \geq m_\ell \qquad \forall\, \ell \in [\gamma] \ .$$

In words, the promise is that there is a solution that picks a family of sets and the requirements can be fulfilled by only using a single representative u_H in each set. However, the solution we are allowed to build is such that the weight of all elements covered by our sets are counted instead of just a single representative per set.

We are now ready to state our main reduction theorem, which, as we discuss later, readily leads, for a constant number of colors γ, to the first constant-factor approximations for γ-Colorful Matroid Supplier for linear matroids and γ-Colorful Knapsack Supplier. Our reduction to $\mathcal{F}$-CP comes at the cost of an $O(2^\gamma)$-factor in the approximation guarantee.

▶ **Theorem 3.** *For any family of down-closed set systems, we have that if $\mathcal{F}$-CP can be solved efficiently for any $\mathcal{F}$ in that family, then there is an $O(2^\gamma)$-approximation algorithm for γ-Colorful $\mathcal{F}$-Supplier for any $\mathcal{F}$ in the family.*[2]

While the dependence of the approximation factor on γ may be undesirable, the algorithmic barriers for prior approaches remain even when $\gamma = 2$ and, for hardness reasons, we do not expect approximation algorithms to exist at all when γ grows too quickly. In particular, [1] showed that even a simple version of colorful clustering, where any k centers can be chosen, does not admit an $O(1)$-approximation algorithm when $\gamma = \omega(\log |C \,\dot{\cup}\, F|)$ under the Exponential Time Hypothesis. Thus, in what follows, we restrict ourselves to $\gamma = O(1)$.

[2] When talking about the same set system $\mathcal{F}$ both in the context of $\mathcal{F}$-CP and γ-Colorful $\mathcal{F}$-Supplier, we consider $\mathcal{F}$ to be the same set system in both settings even if the ground sets are different, as long as there is a one-to-one relation between the ground sets mapping sets of one system to sets of the other one and vice versa.

We now discuss implications of Theorem 3 to γ-Colorful Matroid Supplier for linear matroids and γ-Colorful Knapsack Supplier. When $\mathcal{F}$ is the family of independent sets of a linear matroid, we show how $\mathcal{F}$-CP can be solved with techniques relying on an efficient randomized procedure for the *Exact Weight Basis* (XWB) problem for linear matroids.[3] Linear matroids include as special cases many other well-known matroid classes, including uniform matroids, and more generally partition and laminar matroids, graphic matroids, transversal matroids, gammoids, and regular matroids.

▶ **Theorem 4.** *For $\gamma = O(1)$ and $\mathcal{F}$ being the independent sets of a linear matroid,* $\mathcal{F}$-CP *can be solved efficiently by a randomized algorithm. Hence (by Theorem 3), there is a randomized $O(2^\gamma)$-approximation algorithm for γ-Colorful Matroid Supplier for linear matroids.*

The restriction to linear matroids and the fact that the algorithm is randomized are not artifacts of our framework. Indeed, by an observation in [11], rephrased for matroids below, we do not only have that XWB implies results for γ-Colorful Matroid Supplier (which will follow from our reduction), but also a reverse implication. More precisely, even for 2-Colorful Matroid Supplier, deciding whether there is a solution of radius zero requires being able to solve XWB on that matroid. However, it is unknown whether XWB can be solved efficiently on general matroids, and the only technique known for XWB on linear matroids is inherently randomized [4]. (Derandomization is a long-standing open question in this context.)

▶ **Lemma 5** (based on [11]). *If there is an efficient algorithm for deciding whether 2-Colorful Matroid Supplier with respect to a given class of matroids admits a solution of radius zero, then XWB can be solved efficiently on the same class of matroids.*

Note that if we cannot decide the existence of a radius zero solution, then no approximation algorithm with any finite approximation guarantee can exist.

For the case where $\mathcal{F}$ are the feasible sets for a knapsack problem, one can use standard dynamic programming techniques to see that $\mathcal{F}$-CP can be solved efficiently, which readily leads to a $O(2^\gamma)$-approximation for γ-Colorful Knapsack Supplier.

Whereas our reduction given by Theorem 3 is broadly applicable and readily leads to first constant-factor approximations for γ-Colorful $\mathcal{F}$-Supplier problems, it remains open whether and in which settings a dependence of the approximation factor on the number of colors is necessary. We make first progress toward this question for γ-Colorful Knapsack Supplier, where we show how techniques from [11] can be modified and extended to give a 7-approximation (independent of the number of colors).

▶ **Theorem 6.** *For $\gamma = O(1)$, there is a 7-approximation algorithm for γ-Colorful Knapsack Supplier.*

Our technical contribution here lies in handling the knapsack constraint in this approach – modifying the algorithm of [11] to the supplier setting and to weighted instances is straightforward. In fact, their algorithm can be seen to give a 3-approximation even for γ-Colorful k-Supplier, which is tight in light of a hardness result in [6], namely that it is NP-hard to approximate Robust k-center with forbidden centers to within $3 - \epsilon$. This remains the strongest hardness result even for γ-Colorful $\mathcal{F}$-Supplier problems.

[3] In XWB, one is given a matroid on a ground set with unary encoded weights and a target weight; the goal is to find a basis of the matroid of weight equal to the target weight. The technique in [4] to solve XWB for linear matroids needs an explicit linear representation of the linear matroid. We make the common assumption that this is the case whenever we make a statement about linear matroids.

1.2 Organization of this paper

Our main reduction, Theorem 3, is based on what we call (L, r)-partitions, which is a way to judiciously partition the clients into parts that we want to cover together. We introduce (L, r)-partitions in Section 2 and show how the existence of certain strong (L, r)-partitions implies Theorem 3. In Section 3, we show how our reduction framework can be used to obtain first constant-factor approximations for γ-Colorful Matroid Supplier for linear matroids (thus showing Theorem 4) and γ-Colorful Knapsack Supplier. Finally, in Section 4 we prove existence of strong (L, r)-partitions. The proof of Lemma 5 as well as our 7-approximation for γ-Colorful Knapsack Supplier, i.e., the proof of Theorem 6, can be found in the extended version of this paper.

2 Reducing to $\mathcal{F}$-CP through (L, r)-partitions

Consider a γ-Colorful $\mathcal{F}$-Supplier problem on a metric space $(X = (C \dot\cup F), d)$ with weights $w_\ell \colon C \to \mathbb{Z}_{\geq 0}$ for $\ell \in [\gamma]$ and covering requirements $m_\ell \in \mathbb{Z}_{\geq 0}$ for $\ell \in [\gamma]$. An (L, r)-partition is a partition of the clients into parts of small diameter each of which we consider in our analysis to be either fully covered or not covered at all. The key property of an (L, r)-partition is that, if our instance admits a radius-r solution, then there is a radius-$(L + 1)r$ solution where we allow each center to cover only *a single part* of the partition. It is the existence of such highly structured solutions that we exploit to design $O(1)$-approximation algorithms.

A crucial property of (L, r)-partitions is that they neither depend on $\mathcal{F}$ nor the covering requirements m_ℓ, but only on the metric space and the weight functions, which we call a γ-colorful space for convenience.

▶ **Definition 7** (γ-colorful space (X, d, w)). *A γ-colorful space $(X = C \dot\cup F, d, w)$ consists of*
1. *a metric space (X, d), and*
2. *color functions $w_\ell \colon C \to \mathbb{R}_{\geq 0}$ for $\ell \in [\gamma]$.*

We assume for convenience that the supports of the color functions, i.e., $\mathrm{supp}(w_\ell)$ for $\ell \in [\gamma]$, are pairwise disjoint. One can reduce to this case without loss of generality by co-locating copies of clients. We are now ready to formally define the notion of (L, r)-partition.

▶ **Definition 8** ((L, r)-partition). *Let $(X = C \dot\cup F, d, w)$ be a γ-colorful space and $r, L \in \mathbb{R}_{\geq 0}$. A partition $\mathcal{P} \subseteq 2^C$ is an* (L, r)-partition *if*
1. $\mathrm{diam}(A) := \max_{u,v \in A} d(u, v) \leq L \cdot r \quad \forall A \in \mathcal{P}$, *and*
2. *for any $Z \subseteq F$, there exists a subfamily $\mathcal{A} \subseteq \mathcal{P}$ and injection $h : \mathcal{A} \to Z$ such that*
   a. *$d(A, h(A)) \leq r$,[4] and*
   b. *$w_\ell \left(\bigcup_{A \in \mathcal{A}} A\right) \geq w_\ell (B_C(Z, r)) \quad \forall \ell \in [\gamma]$.*

To connect (L, r)-partitions to colorful clustering problems, think of $Z \in \mathcal{F}$ as centers of a γ-Colorful $\mathcal{F}$-Supplier problem that satisfy the covering requirements with radius r. The definition of an (L, r)-partition $\mathcal{P}$ then implies that there is a subset $\mathcal{A} \subseteq \mathcal{P}$ of the parts such that (i) for each $A \in \mathcal{A}$ there exists an element $h(A) \in Z$ such that any client in A has distance at most $(L + 1) \cdot r$ from $h(A)$, which follows from property 1 and 2a of the definition, and (ii) the clients in $\mathcal{A}$ cover as much as $B_C(Z, r)$ in each color. Thus, the set of facilities $h(\mathcal{A})$ satisfies the covering requirements with respect to the radius $(L + 1) \cdot r$, and,

[4] For any set $V \subseteq F \dot\cup C$ and $x \in F \dot\cup C$, we use the shorthand $d(V, x) := \min\{d(v, x) \colon v \in V\}$.

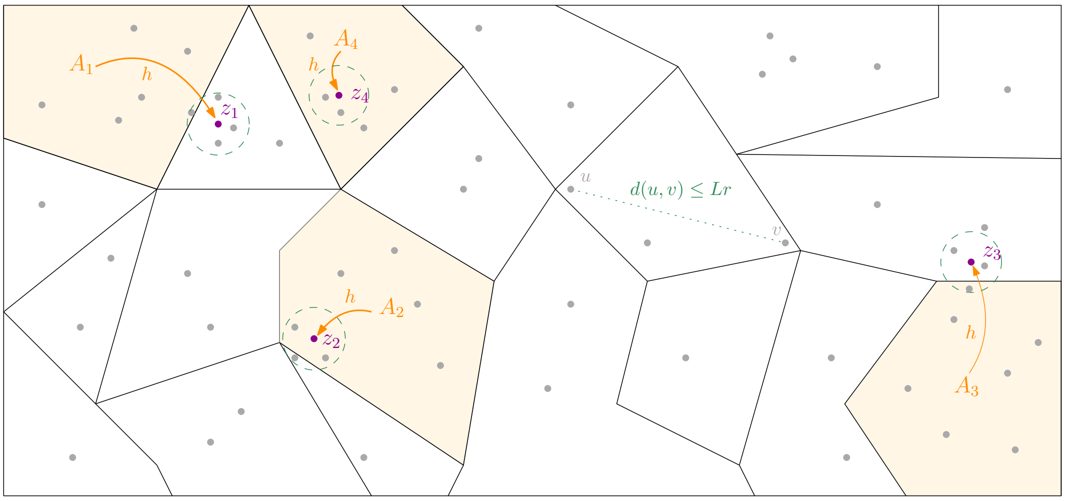

Figure 1 Illustration of an (L, r)-partition of a 1-colorful space (where all points have unit weight). For $Z = \{z_i \mid i \in [4]\}$, the mapping h maps A_i to z_i for $i \in [4]$. Note that $\cup_{i\in[4]} A_i$ contains at least as many points than $\cup_{i\in[4]} B(r, z_i)$ and that $d(z_i, A_i) \leq r$ for $i \in [4]$. Furthermore, the largest distance between any two points in a set A_i is bounded by Lr.

furthermore, $h(\mathcal{A})$ is feasible because $h(\mathcal{A}) \subseteq Z$ and $\mathcal{F}$ is down-closed. In short, $h(\mathcal{A})$ is an $(L+1)$-approximate solution to the γ-Colorful $\mathcal{F}$-Supplier problem. Hence, to obtain an $(L+1)$-approximation, the problem reduces to deciding which of the parts of $\mathcal{P}$ to cover. A key simplification we gain from this connection is that the client sets in $\mathcal{P}$ are non-overlapping because $\mathcal{P}$ is a partition, which we will heavily exploit later to design our algorithms.

The key structural result of our work is to show that (L, r)-partitions with constant L (for a fixed γ) exist and can also be constructed efficiently, which is summarized below.

▶ **Lemma 9.** *For every γ-colorful space (X, d, w) and $r \in \mathbb{R}_{\geq 0}$, one can construct in polynomial time a $(10(2^\gamma - 1), r)$-partition.*[5]

We defer the proof of Lemma 9 to Section 4, and first show how it implies our main reduction theorem, Theorem 3, and how this reduction readily leads to $O(1)$-approximations for γ-Colorful Matroid Supplier for linear matroids and γ-Colorful Knapsack Supplier.

Proof of Theorem 3. Consider an instance of the γ-Colorful $\mathcal{F}$-Supplier Problem on a γ-colorful space (X, d, w). We can guess the radius r of an optimal solution to the problem. This can be achieved by considering all pairwise distances between facilities F and clients C, repeating the steps below for each guess and only considering the best output (and discarding outputs where the procedure fails). Hence, assume that r is the optimal radius from now on.

By Lemma 9, we can efficiently construct an (L, r)-partition $\mathcal{P}$ of (X, d, w) for $L = 10(2^\gamma - 1) = O(2^\gamma)$. Consider the $\mathcal{F}$-CP instance with universe $\mathcal{U} := \mathcal{P}$, family of sets

$$\mathcal{H} := \{H_f : f \in F\} \quad \text{, where}$$
$$H_f := \{A \in \mathcal{P} \text{ with } d(A, f) \leq r\} \quad \forall\, f \in F \ .$$

[5] As we highlight later, a more careful analysis of our approach allows for a slight improvement in the constant factor, leading to the construction of $(8 \cdot 2^\gamma - 10, r)$-partitions. However, in the interest of simplicity, we present a simpler analysis that shows the bound claimed in the lemma.

The family of feasible subsets of $\mathcal{H}$ is the same as $\mathcal{F}$ when identifying H_f with the element f. To make this relation explicit, if we denote by $\mathcal{F}_{\mathcal{H}}$ the family of feasible subsets, then some subset of $\mathcal{H}$, say $\{H_f \colon f \in I\}$ where $I \subseteq F$, is in $\mathcal{F}_{\mathcal{H}}$ if and only if $I \in \mathcal{F}$. Moreover, the weights and coverage thresholds are inherited from those of the given γ-Colorful $\mathcal{F}$-Supplier problem; formally, for $\ell \in [\gamma]$, the ℓ-th weight of $A \in \mathcal{U}$ is given by $w_\ell(A)$.

To make sure that this indeed leads to an $\mathcal{F}$-CP problem, we have to verify that the promise holds. Thus, let $Z \subseteq F$ be a solution to the given γ-Colorful $\mathcal{F}$-Supplier problem for radius r, which exists because we assume that r was guessed correctly. As $\mathcal{P}$ is an (L, r)-partition of (X, d, w), there is a subfamily $\mathcal{A} \subseteq \mathcal{P}$ and injection $h \colon \mathcal{A} \to Z$ satisfying property 2 of Definition 8. We claim that a solution fulfilling the promise is given by choosing

$$S = \{H_f \colon f \in h(\mathcal{A})\} \in \mathcal{F}_{\mathcal{H}} \enspace ,$$

and setting as representative element $u_{H_f} \in H_f$ the element $u_{H_f} = A_f$, where $A_f = h^{-1}(f)$. Note that because $h(\mathcal{A}) \subseteq Z \in \mathcal{F}$ and $\mathcal{F}$ is down-closed, we indeed have $S \in \mathcal{F}_{\mathcal{H}}$. Furthermore, because the injection h satisfies $d(A_f, h(A_f)) \leq r$, we have $u_{H_f} \in H_f$, as desired. Moreover,

$$\sum_{f \in h(\mathcal{A})} w_\ell\left(u_{H_f}\right) = \sum_{f \in h(\mathcal{A})} w_\ell\left(A_f\right) = \sum_{A \in \mathcal{A}} w_\ell(A) \geq w_\ell(B_C(Z, r)) \geq m_\ell \quad \forall \ell \in [\gamma] \enspace ,$$

where the first inequality follows because $\mathcal{A}$ fulfills the second property of Definition 8, and the last inequality is a consequence of Z being centers that are a radius-r solution to the given γ-Colorful $\mathcal{F}$-Supplier problem. Hence, the promised solution exists.

Thus, we can compute an $\mathcal{F}$-CP solution $S_{\mathcal{H}} \subseteq \mathcal{F}_{\mathcal{H}}$, which can be written as $S_{\mathcal{H}} := \{H_f \colon f \in S\}$ for some $S \in \mathcal{F}$. We claim that S is a solution to the given γ-Colorful $\mathcal{F}$-Supplier problem with radius $(L+1) \cdot r$, which finishes the proof. This follows from the fact that $S_{\mathcal{H}}$ is an $\mathcal{F}$-CP solution, and that, for any $f \in F$, each client in $\bigcup_{A \in H_f} A$ has distance at most $(L+1) \cdot r$ from f because $\mathcal{P}$ is an (L, r)-partition. Hence, the clustering solution with centers S and radius $(L+1) \cdot r$ covers all clients in

$$\bigcup_{f \in S} \bigcup_{A \in H_f} A \enspace ,$$

and the w_ℓ-weight (for any $\ell \in [\gamma]$) that it covers is at least

$$w_\ell\left(\bigcup_{f \in S} \bigcup_{A \in H_f} A\right) = \sum_{A \in \bigcup_{f \in S} H_f} w_\ell(A) \geq m_\ell \enspace ,$$

where the equality uses that the ground set $\mathcal{U} = \mathcal{P}$ consists of sets A that are disjoint, and the inequality holds because $S_{\mathcal{H}} = \{H_f \colon f \in S\}$ is a solution to $\mathcal{F}$-CP. Thus, all coverage requirements are fulfilled by the clustering with centers S and radius $(L+1) \cdot r$, as desired. ◀

3 Applications of our reduction framework

We now discuss implications of our reduction framework, Theorem 3, to γ-Colorful Matroid Supplier for linear matroids and γ-Colorful Knapsack Supplier.

3.1 γ-Colorful Matroid Supplier

To apply our reduction framework to γ-Colorful Matroid Supplier for linear matroids, we have to solve $\mathcal{F}$-CP when $\mathcal{F}$ are the independent sets of a linear matroid. We show how this problem can be reduced to XWB in a suitably defined matroid. More precisely, we use a reduction to the *Exact Weight Independent Set* (XWI) problem for matroids. This problem is identical to XWB except that an *independent set* with the desired target weight needs to be returned, instead of a basis. However, XWI easily reduces to XWB on linear matroids, by adding zero weight copies of the elements.

This reduction relies on Rado matroids, which is a way to construct a matroid from another one (see, e.g., [13, Section 8.2]).[6] It relies on the notation of a *system of representatives*, where, for a finite universe $\mathcal{U}$ and a set system $S \subseteq 2^{\mathcal{U}}$, a *system of representatives of* S is any set $\{u_H\}_{H\in S}$ with $u_H \in H$ for $H \in S$. In words, a system of representatives is obtained by replacing each set in S by an element in that set (its representative). (Note that an element can be chosen more than once as a representative, but, as defined above, only appears once in the system of representatives.)

▶ **Definition 10** (Rado matroid). *Let $\mathcal{U}$ be a finite universe, $\mathcal{H} \subseteq 2^{\mathcal{U}}$ be some set system, and let $M = (\mathcal{H}, \mathcal{I})$ be a matroid. The Rado matroid $(\mathcal{U}, \overline{\mathcal{I}})$ induced by $(\mathcal{U}, \mathcal{H}, M)$ is a matroid on the ground set $\mathcal{U}$ with independent sets*

$$\{U \subseteq \mathcal{U} \colon U \text{ is a system of representatives for some } I \in \mathcal{I}\} \ .$$

A proof that a Rado matroid is indeed a matroid can be found, e.g., in [13, Section 8.2]. We will reduce $\mathcal{F}$-CP to XWI on a Rado matroid obtained from a linear matroid. For this, we need that also the Rado matroid we obtain is linear and, moreover, that an explicit linear representation of it can be found efficiently, which is the case due to a result from [12].

▶ **Lemma 11** (see Theorem 3 of [12]). *For a set family $\mathcal{H} \subseteq 2^{\mathcal{U}}$ and a linear matroid $M = (\mathcal{H}, \mathcal{I})$, the Rado matroid $\overline{M} = (\mathcal{U}, \overline{\mathcal{I}})$ induced by $(\mathcal{U}, \mathcal{H}, M)$ is a linear matroid. Moreover, given a linear representation of M, one can find a linear representation of $\overline{M}$ in time polynomial in $|\mathcal{H}|$, $|\mathcal{U}|$, and the size of the linear representation of M.*

We are now ready to show that $\mathcal{F}$-CP can be solved efficiently for linear matroids, which implies Theorem 4.

▶ **Lemma 12.** $\mathcal{F}$-CP *can be solved efficiently when $\mathcal{F}$ is the family of independent sets of a linear matroid.*

Proof. We recall that we are given an $\mathcal{F}$-CP instance, which defines a set system $\mathcal{H} \subseteq 2^{\mathcal{U}}$ over a finite universe $\mathcal{U}$, and a family $\mathcal{F} \subseteq 2^{\mathcal{H}}$ such that $M = (\mathcal{H}, \mathcal{F})$ is a linear matroid. Let $\overline{M} = (\mathcal{U}, \overline{\mathcal{I}})$ be the Rado matroid induced by $(\mathcal{U}, \mathcal{H}, M)$. $\overline{M}$ is a linear matroid by Lemma 11 and we can obtain a linear representation of $\overline{M}$ in polynomial time. The promise of $\mathcal{F}$-CP implies the existence of an independent set T of $\overline{M}$ satisfying the covering requirements, i.e.,

$$w_\ell(T) \geq m_\ell \qquad \forall \ell \in [\gamma] \ . \tag{1}$$

To solve $\mathcal{F}$-CP, we guess, for each color $\ell \in [\gamma]$, the weight $\lambda_\ell \coloneqq w_\ell(T)$ that T covers. Note that λ_ℓ is at most $W_\ell \coloneqq w_\ell(\mathcal{U})$, which, due to the unary encoding of w_ℓ, is polynomially bounded in the input. Hence, the guessing of the λ_ℓ, for $\ell \in [\gamma]$, can be performed in time $\prod_{\ell\in[\gamma]} W_\ell$, which is polynomially bounded because $\gamma = O(1)$.

We now determine an independent set $\widetilde{T}$ in $\overline{M}$ with $w_\ell(\widetilde{T}) = w_\ell(T)$ for each $\ell \in [\gamma]$. This can be achieved by encoding all ℓ many (unary encoded) weight functions w_ℓ for $\ell \in [\gamma]$ into a single one $\overline{w}$ and then solving an appropriate XWI problem with respect to $\overline{w}$. More precisely,

[6] This construction of Rado matroids is also called the *induction of a matroid by a bipartite graph.*

for an element $u \in \mathcal{U}$, we obtain a new single weight $\overline{w}(u)$ whose first $\lceil \log_2(|W_1|+1)\rceil$ bits represent the weight $w_1(u)$, the next $\lceil \log_2(|W_2|+1)\rceil$ bits the weight $w_2(u)$, and so on. Because $\gamma = O(1)$ and all w_ℓ have unary encoding, this leads to combined weights $\overline{w}$ whose unary encoding is polynomially bounded. Analogously, we encode the guessed weights λ_ℓ for $\ell \in [\gamma]$ into a single one $\overline{\lambda}$. We now solve XWI on $\overline{M}$ with weights $\overline{w}$ and target weight $\overline{\lambda}$. As $\overline{M}$ is linear, this is possible by a randomized algorithm in time pseudo-polynomial in the total weight [4]. Moreover, because the weights are unary encoded in our setting, this implies a polynomial running time as desired.

Let $\widetilde{T}$ be a solution of this XWI problem, which must exist for the correct guess of the λ_ℓ because of the promised solution T. $\widetilde{T}$ being independent in $\overline{M}$ implies that it is a system of representatives for some independent set $S \in \mathcal{F}$ of M. Such a set S can be found through matroid intersection. More precisely, it is known that the minimal (inclusion-wise) sets $I \subseteq \mathcal{H}$ such that $\widetilde{T}$ is a system of representatives for I form the basis of a matroid $\widetilde{M}$, for which an efficient independence oracle can be obtained. (See [13, Section 7.3].) Hence, the desired set S can be obtained by finding a basis of $\widetilde{M}$ that is independent in M, which can be computed through matroid intersection algorithms. The set S is the solution of $\mathcal{F}$-CP that we return. Because $\widetilde{T} \subseteq \bigcup_{H\in S} H$, the set S fulfills the covering requirements due to (1). ◀

3.2 γ-Colorful Knapsack Supplier

To showcase the versatility of our reduction, we now show how it implies an $O(2^\gamma)$-approximation for γ-Colorful Knapsack Supplier, by discussing an efficient way to solve $\mathcal{F}$-CP when $\mathcal{F}$ are the feasible solutions to a knapsack constraint. Even though there is a stronger (and more sophisticated) approximation result for this problem (as stated in Theorem 6), this application is a nice example of how one can readily obtain constant-factor approximations through our reduction technique combined with known methods; in this case, by solving $\mathcal{F}$-CP through a standard dynamic programming approach.

▶ **Lemma 13.** *Let $\mathcal{F}$ be the feasible sets of a knapsack constraint, i.e., $\mathcal{F} = \{S \subseteq \mathcal{H} \colon \kappa(S) \leq K\}$ for some $\kappa\colon \mathcal{H} \to \mathbb{R}_{\geq 0}$ and budget $K \in \mathbb{R}_{\geq 0}$. Then $\mathcal{F}$-CP can be solved efficiently.*

Proof. Recall that the $\mathcal{F}$-CP problem to be solved defines a family $\mathcal{H} \subseteq 2^{\mathcal{U}}$ over a finite universe $\mathcal{U}$, and a family $\mathcal{F} \subseteq 2^{\mathcal{H}}$, which is defined by a knapsack constraint, i.e., $\mathcal{F} = \{S \subseteq \mathcal{H} \colon \kappa(S) \leq K\}$. We define the following weight function on $\mathcal{U}$:

$$\eta(u) \coloneqq \min\{\kappa(H) \colon H \in \mathcal{H} \text{ with } u \in H\} \enspace .$$

In words, $\eta(u)$ corresponds to the cost of the cheapest set in $\mathcal{H}$ that covers u. Consider the following binary program, which can be solved efficiently by standard dynamic programming techniques due to the unary encoding of the weights w_ℓ for $\ell \in [\gamma]$ (see, e.g., [1] for details):

$$\begin{array}{rrcll} \min & \sum_{u\in\mathcal{U}} \eta(u)\cdot z(u) & & & \\ & \sum_{u\in\mathcal{U}} w_\ell(u)\cdot z(u) & \geq & m_\ell & \forall \ell \in [\gamma] \\ & z & \in & \{0,1\}^{\mathcal{U}} \enspace . & \end{array}$$

We compute an optimal solution z^* to the above binary program. Let $Q \coloneqq \{u \in \mathcal{U} \colon z^*(u) = 1\}$. For each $u \in Q$, let $H_u \in \mathcal{H}$ be a set of minimum cost that contains u; hence, $\kappa(H_u) = \eta(u)$. We claim that $\{H_u \colon u \in Q\}$ is a solution to $\mathcal{F}$-CP. Because

z^* fulfills the constraints of the binary program, we have that $\{H_u : u \in Q\}$ fulfills the covering requirements. It remains to show that it fulfills the knapsack constraint, i.e., its cost is at most K. This reduces to show that the optimal value of the binary program is at most K. We claim that this holds because of the promise of $\mathcal{F}$-CP. Indeed, the promise guarantees that there is $S \subseteq \mathcal{F}$ and a system of representatives u_H for $H \in S$ such that $w_\ell(\{u_H : H \in S\}) \geq m_\ell$ for $\ell \in [\gamma]$. Hence, setting $z_{u_H} = 1$ for all $H \in S$, and setting all other coordinates of $z \in \{0,1\}^{\mathcal{U}}$ to zero, is a solution to the binary program which has objective value at most $\kappa(S) \leq K$. ◀

4 Existence and construction of strong (L, r)-partitions

We now prove our key structural result, Lemma 9, which guarantees the existence and efficient constructability of $(O(2^\gamma), r)$-partitions for γ-colorful spaces. Our proof proceeds by induction on γ. The base case, i.e., $\gamma = 0$, holds because the family $\{\{c\} : c \in C\}$ is a $(0, r)$-grouping on every 0-colorful space $(C \,\dot\cup\, F, d, w)$. The key step is extending an (L, r)-partition of a $(\gamma - 1)$-colorful space to a suitable partition of a γ-colorful space.

To this end, we extend ideas on the greedy algorithm of [6], which was originally introduced to deal with a single color k-center problem. More precisely, to augment a partition of a $(\gamma - 1)$-colorful space, we apply a greedy subroutine on the points of color γ. A careful construction and analysis (which takes into account the earlier colors) then shows that this yields a $(2L + 10, r)$-partition of the γ-colorful space. Our refined charging scheme improves on a decoupled analysis of [10] (which gives an $O(5^\gamma)$ approximation algorithm for γ-Colorful k-Center).

The lemma below formalizes the induction step.

▶ **Lemma 14.** *Given a (L, r)-partition for a $(\gamma - 1)$-colorful space, then one can efficiently construct a $(2L + 10, r)$-partition for any γ-colorful space obtained by adding one color to the $(\gamma - 1)$-colorful space.*

Proof. Let $(C \,\dot\cup\, F, d, w)$ be a γ-colorful space, and let $\widehat{w} = (w_1, \ldots, w_{\gamma-1})$ be the first $\gamma - 1$ colors. (Hence, we omitted the last color.) Let $C_\gamma := \operatorname{supp}(w_\gamma)$ and $C_{<\gamma} := C \setminus C_\gamma$, and let $\mathcal{P}$ be a (L, r)-partition of the $(\gamma - 1)$-colorful space $(C_{<\gamma} \,\dot\cup\, F, d, \widehat{w})$. Note that we assumed that the supports of the weights w_ℓ are disjoint. Hence, $w_\ell(C_\gamma) = 0$ for $\ell \in [\gamma - 1]$. Moreover, without loss of generality, we assume that for every client $c \in C$, there is a facility $f \in F$ with $d(f, c) \leq r$. All clients not fulfilling this condition can be deleted from the instance without changing the statement as they can never be covered by any radius-r solution. Indeed, a partition of the clients of this purged instance can simply be extended to a partition of all clients by adding the deleted clients as singleton sets to the partition.

We now prove that Algorithm 1 returns an $(\overline{L}, r)$-partition $\overline{\mathcal{P}}$ of $(C \,\dot\cup\, F, d, w)$, where $\overline{L} := 2L + 10$. Algorithm 1 goes through all facilities in a well-chosen order and iteratively builds new parts consisting of parts in $\mathcal{P}$ together with a subset of C_γ. (See Figure 2 for an illustration of this procedure.)

First, observe that $\overline{\mathcal{P}}$ is a partition. It clearly covers all clients as no client is farther than distance r away from its nearest facility, and we consider all facilities. Moreover, the sets in $\overline{\mathcal{P}}$ are disjoint by construction. Now, observe that any $\overline{A}_i \in \overline{\mathcal{P}}$ has small diameter, because

$$\operatorname{diam}\left(\overline{A}_i\right) \leq 2 \cdot \max_{c \in \overline{A}_i} d\left(g_i, c\right) \leq 10r + 2Lr \ ,$$

where the second inequality holds because $d(g_i, c) \leq 5r + Lr$ for any $c \in \overline{A}_i$ due to the following. Consider $c \in \overline{A}_i$. If $c \in C_\gamma$, then we even have $d(g_i, c) \leq 3r$. Otherwise, let

Algorithm 1 GreedyPartitioning($C, F, d, \mathcal{P}, w_\gamma$).

for $i = 1$ **to** $|F|$ **do**

$$g_i = \operatorname*{argmax}_{f \in F \setminus \{g_1,\dots,g_{i-1}\}} w_\gamma\left(B_C(f,r) \setminus \bigcup_{t=1}^{i-1} \overline{A}_t\right);$$

$$\overline{A}_i \leftarrow \left(B_{C_\gamma}(g_i, 3r) \cup \bigcup_{\substack{A \in \mathcal{P} \text{ with} \\ d(g_i,A) \le 5r}} A\right) \setminus \bigcup_{t=1}^{i-1} \overline{A}_t;$$

end

return $\overline{\mathcal{P}} \coloneqq \{\overline{A}_i : i \in [|F|]\}$;

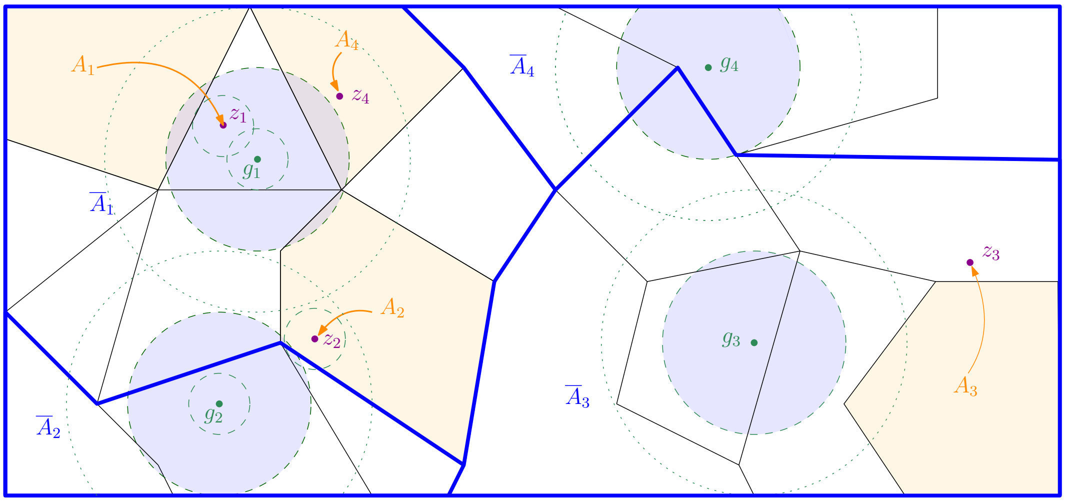

Figure 2 Visualization of an (L, r)-partition and of Algorithm 2. The black polygons depict an (L, r)-partition $\mathcal{P}$ of the clients $C_{<\gamma}$. The blue polygons shows how the clients $C_{<\gamma}$ are partitioned by $\overline{\mathcal{P}}$. Moreover, the blue $3r$-balls around g_i illustrate which clients of C_γ get assigned to the part $\overline{A}_i \in \overline{\mathcal{P}}$. The dashed circles have radius r and $3r$ respectively, while the dotted circles have radius $5r$. We assume $Z = \{z_i \mid i \in [4]\}$ is given and we construct the respective $\overline{\mathcal{A}}$ and $\overline{h}$, given $\mathcal{A}$ and h. We have $\mathcal{A} = \{A_i \mid i \in [4]\}$ (the orange areas) and $\overline{\mathcal{A}} = \{\overline{A}_i \mid i \in [3]\}$. Moreover $h(A_i) = z_i$ for $i \in [4]$ (depicted by an orange arrow), while $\overline{h}(\overline{A}_i) = z_i$ for $i \in [3]$.

$A \in \mathcal{P}$ be the set in the partition $\mathcal{P}$ containing c. Note that $c \in \overline{A}_i$ implies $A \subseteq \overline{A}_i$. Hence, $d(g_i, c) \le d(g_i, A) + \max\{d(b, c) : b \in A\} \le 5r + Lr$, where we use $d(g_i, A) \le 5r$, because $A \subseteq \overline{A}_i$, and $\operatorname{diam}(A) \le Lr$, which holds because $\mathcal{P}$ is an (L, r)-partition. Thus, property 1 of the definition of an $(2L + 10, r)$-partition (Definition 8) is fulfilled for $\overline{\mathcal{P}}$.

It remains to show that property 2 holds for a given selection Z. To this end, we use that $\mathcal{P}$ is an (L, r)-partition, which implies that there is a subfamily $\mathcal{A} \subseteq \mathcal{P}$ and a corresponding injection $h : \mathcal{A} \to Z$ fulfilling property 2 of Definition 8 for the $(\gamma - 1)$-colorful space $(C_{<\gamma} \,\dot\cup\, F, d, \widehat{w})$. In the following we construct $\overline{\mathcal{A}} \subseteq \overline{\mathcal{P}}$ and $\overline{h} : \overline{\mathcal{A}} \to Z$ such that property 2 of Definition 8 is satisfied for $\overline{\mathcal{A}}$ and $\overline{h}$. At the same time when constructing $\overline{\mathcal{A}}$, we employ a careful charging argument that makes sure that $w_\gamma\big(\bigcup_{\overline{A} \in \overline{\mathcal{A}}} \overline{A}\big) \ge w_\gamma(B_C(Z, r))$, i.e., that the constructed $\overline{\mathcal{A}}$ covers at least as much as Z of color γ. For the remaining colors, we show that the new selection $\overline{\mathcal{A}}$ includes all of $\mathcal{A}$; formally, we show that for each $A \in \mathcal{A}$, there is an $\overline{A} \in \overline{\mathcal{A}}$ such that $A \subseteq \overline{A}$. This, as well as $d(\overline{A}, \overline{h}(\overline{A})) \le r$ for all $\overline{A} \in \overline{\mathcal{P}}$ and injectivity of $\overline{h}$, are proved later.

For $i \in [|F|]$, we define

$$U_i := C \setminus \bigcup_{t=1}^{i-1} \overline{A}_t$$

to be the clients that are "uncovered" at step i. By the way Algorithm 1 selects g_i in each iteration $i \in [|F|]$, we have

$$w_\gamma \left(B_C\left(g_i, r\right) \cap U_i\right) \geq w_\gamma \left(B_C(f, r) \cap U_i\right) \quad \forall i \in [|F|] \text{ and } f \in F \ ,$$

which we call the *greediness property.*

We now describe the construction of $\overline{\mathcal{A}}$ and the charging scheme in detail. We successively add sets $\overline{A}_i \in \overline{\mathcal{P}}$ to $\overline{\mathcal{A}}$, where the sets $\overline{A}_i$ are considered in increasing order of their index. When adding a set $\overline{A}_i$ to $\overline{\mathcal{A}}$, we also perform two further steps: (i) we identify an element $f \in Z$ and set $\overline{h}(\overline{A}_i) = f$, and (ii) we mark f as assigned to make sure that we never assign it again in the future (as $\overline{h}$ needs to be an injection). For convenience, for $i \in [|F|]$ and $f \in Z$, we write $\textsc{Assign}(i, f)$ for performing these steps, i.e., adding $\overline{A}_i$ to $\overline{\mathcal{A}}$, setting $\overline{h}(\overline{A}_i)$ to f, and marking f as assigned.

The charging argument charges the coverage of color γ of $B_C(Z, r)$ against the γ-coverage in $\bigcup_{\overline{A} \in \overline{\mathcal{A}}} \overline{A}$. Whenever we charge a set $Q \subseteq B_C(Z, r)$ against some subset $W \subseteq \bigcup_{\overline{A} \in \overline{\mathcal{A}}} \overline{A}$, we make sure that $w_\gamma(Q) \leq w_\gamma(W)$. Algorithm 2 shows our procedure to construct both $\overline{\mathcal{A}}$ and the desired injection $\overline{h}\colon \overline{\mathcal{A}} \to Z$ together with the charging argument. (See also Figure 2.)

Algorithm 2 Construction of $\overline{\mathcal{A}}$ and injection $\overline{h}\colon \overline{\mathcal{A}} \to Z$ together with charging argument.

Mark all facilities in Z as unassigned.
for $i = 1$ **to** $|F|$ **do**

- **Rule 1** If there is an unassigned $f \in Z$ with $B_C(f, r) \cap B_C(g_i, r) \cap U_i \neq \emptyset$: $\textsc{Assign}(i, f)$.
- **Rule 2** Else if there is an unassigned $f \in Z$ with $B_C(f, r) \cap B_C(g_i, 3r) \cap U_i \neq \emptyset$: $\textsc{Assign}(i, f)$ and charge $B_C(f, r) \cap U_i$ against $B_C(g_i, r) \cap U_i$.
- **Rule 3** Else if there is an $A \in \mathcal{A}$ such that $A \subseteq \overline{A}_i$: $\textsc{Assign}(i, h(A))$ and charge $B_C(h(A), r) \cap U_i$ against $B_C(g_i, r) \cap U_i$.

If $\textsc{Assign}$ was called, charge against themselves all points in $\overline{A}_i$ that have not been charged yet.

end

We start by showing that $\overline{h}$ is an injection. Suppose f is assigned using Rule 1 or 2. Then f was not assigned so far as we only assign unassigned facilities. Now suppose $h(A) = \overline{h}(\overline{A}_i)$ is assigned using Rule 3. We claim that $h(A)$ is not assigned so far. Assume by the sake of deriving a contradiction that it was assigned in a previous iteration $j < i$. It cannot have been assigned by Rule 3, since h is injective. So assume it is was assigned by Rule 1 or 2. Hence, g_j satisfies $B_C(g_j, 3r) \cap B_C(h(A), r) \cap U_i \neq \emptyset$. This implies that $d(g_j, h(A)) \leq 4r$ and thus $A \subseteq \overline{A}_j$, which contradicts $A \subseteq \overline{A}_i$.

Moreover, $\overline{A}_i$ fulfills property 2a of a $(2L + 10, r)$-partition because of the following. Let $f \in Z$ and $\overline{A}_i := \overline{h}^{-1}(f)$, and we have to show that $d(\overline{A}_i, f) \leq r$. Because $h(\overline{A}_i) = f$, we called at some point during Algorithm 2 the procedure $\textsc{Assign}(i, f)$. In both Rule 1 and Rule 2 we have $B_C(f, r) \cap B_C(g_i, 3r) \cap U_i \neq \emptyset$, which implies that $\overline{A}_i$ contains a client in $B_C(f, r)$, as desired. If $\textsc{Assign}(i, f)$ was called in Rule 3, then we have $h^{-1}(f) \subseteq \overline{A}_i$, which implies $d(\overline{A}_i, f) \leq d(h^{-1}(f), f) \leq r$ by the fact hat $\mathcal{P}$ is an (L, r)-partition.

It remains to show that $\overline{\mathcal{A}}$ fulfills property 2b of an $(2L+10, r)$-partition. We first consider the last color (color γ) and show $w_\gamma(\bigcup_{\overline{A}\in\overline{\mathcal{A}}} \overline{A}) \geq w_\gamma(B_C(Z,r))$. To this end, observe that the charging indeed charges clients in $B_C(Z,r)$ against clients in $\bigcup_{\overline{A}\in\overline{\mathcal{A}}} \overline{A}$. We allow for charging a client in $B_C(Z,r)$ against more than one client in $\bigcup_{\overline{A}\in\overline{\mathcal{A}}} \overline{A}$. However, no client in $\bigcup_{\overline{A}\in\overline{\mathcal{A}}} \overline{A}$ gets charged against more than once because in iteration i we only charge against clients in $\overline{A}_i$, and the sets $\overline{\mathcal{A}} = \{\overline{A}_1, \ldots, \overline{A}_{|F|}\}$ form a partition of C. Also note that we always charge clients of $B_C(Z,r)$ against clients of $\bigcup_{\overline{A}\in\overline{\mathcal{A}}} \overline{A}$ of at least the same w_γ-weight. This is true whenever charging happens in Rule 2 or Rule 3, because of the greediness property, and holds trivially for all other charging operations, which only charge clients against themselves. To conclude that $w_\gamma(\bigcup_{\overline{A}\in\overline{\mathcal{A}}} \overline{A}) \geq w_\gamma(B_C(Z,r))$, it remains to observe that all of $B_C(Z,r)$ gets charged against something.

To this end, fix a facility $f \in Z$. Consider an iteration j of Algorithm 2 such that $B_C(g_j, 3r) \cap U_j$ intersects $B_C(f,r)$. We claim that for each such iteration, either $\textsc{Assign}(j,f)$ is called, or $B_C(f,r) \cap B_C(g_j,3r) \cap U_j$ is charged. To prove the claim, suppose f is not assigned in iteration j. By Algorithm 2, either Rule 1 or Rule 2 must have applied in this iteration j, as f satisfies the condition of Rule 2. Thus $\textsc{Assign}$ was called on j and all points in $\overline{A}_j$ have been charged. Now suppose the first case applies, i.e., $\textsc{Assign}(j,f)$ is called for some j. Then all of $B_C(f,r) \cap U_j$ is charged (and $B_C(f,r) \setminus U_j$ is already charged by the second case). If the first case never applies, then all of $B_C(f,r)$ is charged by the second case since $U_{|F|}$ is empty. Hence, all of $B_C(f,r)$ is charged, as desired.

To see that property 2b of Definition 8 is fulfilled also for all colors $\ell \in [\gamma-1]$, observe that Rule 3 makes sure that any component that was in $\mathcal{A}$ will still be selected in $\overline{\mathcal{A}}$. Thus, $w_\ell(\overline{\mathcal{A}}) \geq w_\ell(B_C(Z,r))$ for all colors $\ell \in [\gamma]$.

It remains to show that $d(\overline{A}_i, \overline{h}(\overline{A}_i)) \leq r$. If Rule 1 or Rule 2 is applied, this is satisfied as there is a client $c \in B_C(g_i,3r) \cap B_C(\overline{h}(\overline{A}_i), r) \cap U_i$; because $c \in \overline{A}_i$ by construction, we have $d(\overline{h}(\overline{A}_i), \overline{A}_i) \leq d(\overline{h}(\overline{A}_i), c) \leq r$. If Rule 3 is applied for $A \subseteq \overline{A}_i$, we also have $d(\overline{h}(\overline{A}_i), \overline{A}_i) \leq d(\overline{h}(\overline{A}_i), A) = d(h(A), A) \leq r$, where the last inequality follows from $\mathcal{P}$ being an (L,r)-partition. ◀

Lemma 9 now follows readily from Lemma 14.

Proof of Lemma 9. The proof follows by induction on γ. For the induction start, consider $\gamma = 0$. The set $\{\{c\} : c \in C\}$ is a $(0,r)$-partition on every 0-colorful space $(C \mathbin{\dot\cup} F, d, w)$. The induction step is given by Lemma 14. Note that $2\left(10(2^{\gamma-1}-1)\right) + 10 = 10(2^\gamma - 1)$. The running time is clearly $O(\mathsf{poly}(|X|,\gamma))$ as every step in the induction takes time $O(\mathsf{poly}(|X|,\gamma))$.[7] ◀

References

1 G. Anegg, H. Angelidakis, A. Kurpisz, and R. Zenklusen. A technique for obtaining true approximations for k-center with covering constraints. *Mathematical Programming*, 2021. doi:10.1007/s10107-021-01645-y.

2 T. Bajpai, D. Chakrabarty, C. Chekuri, and M. Negahbani. Revisiting priority k-center: Fairness and outliers. In N. Bansal, E. Merelli, and J. Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 21:1–21:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.21.

[7] As briefly mentioned earlier, one can obtain slightly better constants, leading to existence and conductibility of $(8 \cdot 2^\gamma - 10, r)$-partitions. This can be achieved by using $\gamma = 1$ as base case, for which our techniques can be shown to imply that there are $(6,r)$-partitions. In the interest of simplicity, we use the slightly weaker bound in Lemma 9.

3. S. Bandyapadhyay, T. Inamdar, S. Pai, and K. R. Varadarajan. A constant approximation for colorful k-center. In M. A. Bender, O. Svensson, and G. Herman, editors, *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, volume 144 of *LIPIcs*, pages 12:1–12:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ESA.2019.12.
4. P.M. Camerini, G. Galbiati, and F. Maffioli. Random pseudo-polynomial algorithms for exact matroid problems. *Journal of Algorithms*, 13(2):258–273, 1992. doi:10.1016/0196-6774(92)90018-8.
5. D. Chakrabarty and M. Negahbani. Generalized center problems with outliers. *ACM Trans. Algorithms*, 15(3):41:1–41:14, 2019. doi:10.1145/3338513.
6. M. Charikar, S. Khuller, D. M. Mount, and G. Narasimhan. Algorithms for facility location problems with outliers. In S. Rao Kosaraju, editor, *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, January 7-9, 2001, Washington, DC, USA*, pages 642–651. ACM/SIAM, 2001. URL: http://dl.acm.org/citation.cfm?id=365411.365555.
7. D. Z. Chen, J. Li, H. Liang, and H. Wang. Matroid and knapsack center problems. *Algorithmica*, 75(1):27–52, 2016. doi:10.1007/s00453-015-0010-1.
8. F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5029–5037, 2017. URL: https://proceedings.neurips.cc/paper/2017/hash/978fce5bcc4eccc88ad48ce3914124a2-Abstract.html.
9. D. G. Harris, T. W. Pensyl, A. Srinivasan, and K. Trinh. A lottery model for center-type problems with outliers. *ACM Trans. Algorithms*, 15(3):36:1–36:25, 2019. doi:10.1145/3311953.
10. T. Inamdar and K. Varadarajan. Non-uniform k-center and greedy clustering, 2021. arXiv:2111.06362.
11. X. Jia, K. Sheth, and O. Svensson. Fair colorful k-center clustering. *Mathematical Programming*, 2021. doi:10.1007/s10107-021-01674-7.
12. M. J. Piff and D. J. A. Welsh. On the vector representation of matroids. *Journal of The London Mathematical Society-second Series*, pages 284–288, 1970.
13. D.J.A. Welsh. *Matroid theory.* Courier Corporation, 2010.