

Dependent randomized rounding for clustering and partition systems with knapsack constraints*

David G. Harris

DAVIDGHARRIS29@GMAIL.COM

Department of Computer Science, University of Maryland, College Park, MD 20742

Thomas Pensyl

TPENSYL@BANDWIDTH.COM

Bandwidth, Inc. Raleigh, NC.

Aravind Srinivasan

SRIN@CS.UMD.EDU

Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742

Khoa Trinh

KHOATRINH@GOOGLE.COM

Google, Mountain View, CA 94043

Editor: Maya Gupta

Abstract

Clustering problems are fundamental to unsupervised learning. There is an increased emphasis on *fairness* in machine learning and AI; one representative notion of fairness is that no single group should be over-represented among the cluster-centers. This, and much more general clustering problems, can be formulated with “knapsack” and “partition” constraints. We develop new randomized algorithms targeting such problems, and study two in particular: multi-knapsack median and multi-knapsack center. Our rounding algorithms give new approximation and pseudo-approximation algorithms for these problems.

One key technical tool, which may be of independent interest, is a new tail bound analogous to Feige (2006) for sums of random variables with unbounded variances. Such bounds can be useful in inferring properties of large networks using few samples.

Keywords: Dependent rounding, clustering, fairness

This is an extended version of a paper which appeared in the Proc. 23rd International Conference on Artificial Intelligence and Statistics (AISTATS 2020)

1. Introduction

Clustering is a fundamental technique in unsupervised learning. Our goal is to systematically study clustering with knapsack constraints, particularly in light of fairness. Consider a data-clustering problem, with a set \mathcal{F} of potential cluster-centers with m non-negative cost functions M_1, \dots, M_m , a data-set \mathcal{C} , and a symmetric distance-metric d on $\mathcal{F} \cup \mathcal{C}$. Our goal is to choose a set $\mathcal{S} \subseteq \mathcal{F}$ of cluster-centers to minimize the distances $d(j, \mathcal{S})$, for $j \in \mathcal{C}$, while satisfying the budget constraint $\sum_{i \in \mathcal{S}} M_k(i) \leq 1$ for each k .

The metric used to boil down the values $d(j, \mathcal{S})$ into an objective function $\text{cost}(\mathcal{S})$ is problem-specific. We refer to the $m \times |\mathcal{F}|$ matrix M as the *knapsack-constraint matrix*; we refer to the case $m > 1$ as *multi-knapsack* and the case $m = 1$ as *single-knapsack*. We typically view m as “small”, e.g., as a constant. The RHS value 1 for the budget constraint is just a normalization. Note that it is possible to have $\mathcal{F} = \mathcal{C}$.

As a representative example, which is in fact one of the main motivations behind this work, consider the following scenario. As in classical clustering, we want to select a small number of clusters of small radius for a population \mathcal{C} . In addition, we are given a collection of m groups A_1, \dots, A_m , such that we do not wish to disproportionately select our centers \mathcal{S} from any group. For instance, in healthcare facility location in the face of an epidemic, we may not want to open too many facilities in a geographic region, or with certain types of equipment, or near only some groups of patients, etc. Thus, for each $k = 1, \dots, m$, we have a constraint $|\mathcal{S} \cap A_k| \leq t_k$ where t_k is a target value which is not much larger than the “fair” proportion of A_k compared to the general population. These can be represented as knapsack constraints; due to the normalization we use, the cost functions are given by $M_k(i) = 1/t_k$ if $i \in A_k$ and $M_k(i) = 0$ otherwise.

A common strategy for clustering problems is to first solve a related linear program (LP), leading to a fractional configuration $y \in [0, 1]^{\mathcal{F}}$ over \mathcal{F} satisfying the knapsack constraints. Here, the LP suggests the fractional extent y_i for $i \in \mathcal{F}$ to be chosen as a center. Based on this solution y , we partition \mathcal{F} into groups, which represent sets of cluster-centers “close to” certain items $j \in \mathcal{C}$. Finally, we use some randomized-rounding algorithm to convert the fractional solution y to an integral solution $Y \in \{0, 1\}^{\mathcal{F}}$ representing the chosen solution \mathcal{S} , while ensuring that each group gets a selected center. (This grouping is needed because if some client j has no nearby opened facilities, then $d(j, \mathcal{S})$ may be very large.)

The vector Y should also have other probabilistic properties related to the vector y , for example satisfying $\mathbf{E}[Y] = y$ coordinatewise. Clustering problems pose a particular challenge for randomized rounding because they are fundamentally *non-linear*; the distance from j to its closest center depends on the *joint* behavior of the selected centers. Consequently, the rounding process should also ensure strong independence properties, beyond just bounds on the expected values of individual coordinates of Y .

One obvious rounding method for the LP is for each group to independently choose a single center i , with probability proportional to y_i . This has optimal independence properties, but completely ignores the knapsack constraints. This paper focuses on a new randomized rounding strategy for the LP. We develop two key probabilistic techniques, both of which are quite general and may be of independent interest: (i) a new “Samuels-Feige” type of concentration inequality for *unbounded* random variables; and (ii) a new rounding algorithm in the presence of knapsack constraints plus a single partition constraint.

As we will see, we cannot simultaneously achieve the goals of exactly preserving the knapsack constraints and mimicking the probabilistic guarantees of independent selection. Nevertheless, we obtain significantly stronger guarantees compared to previous algorithms. We summarize these next, and then discuss our new clustering results.

1.1 Dependent rounding and independence

The general problem of randomized rounding while preserving hard combinatorial constraints often goes by the name “dependent rounding”; see, e.g., Byrka et al. (2017, 2010); Charikar and Li (2012). In the most straightforward form of dependent rounding, which we call *cardinality rounding*, we have a fractional solution $x \in [0, 1]^n$ that we wish to round to an integral vector $X \in \{0, 1\}^n$ such that $\mathbf{E}[X] = x$ and the cardinality constraint $\sum_i X_i = \sum_i x_i$ holds with probability one. For example, Charikar and Li (2012) applied cardinality rounding as part of their 3.25-approximation algorithm for k -median.

The cardinality constraint can easily be replaced by a single knapsack constraint Ageev and Sviridenko (2004); Srinivasan (2001). Over the last two decades, increasingly-sophisticated dependent-rounding techniques have been used for optimization problems over various types of polytopes; see, e.g., Ageev and Sviridenko (2004); Srinivasan (2001); Gandhi et al. (2006); Călinescu et al. (2011); Chekuri et al. (2011); Bansal (2019).

Our new rounding algorithm, which we call *Knapsack-Partition Rounding (KPR)*, generalizes this setting in two distinct ways: it allows multiple knapsack constraints, and it allows a partition matroid constraint. Formally, we define a *knapsack-partition system* to be a partition \mathcal{G} over a ground-set U , along with a real-valued $m \times |U|$ matrix M and a fractional vector $y \in [0, 1]^U$ satisfying $y(G) = 1$ for each $G \in \mathcal{G}$.

Our overarching question is: *how well can we approximate independence while preserving the knapsack and partition constraints?* In an ideal scenario, we would like to generate a random vector $Y \in \{0, 1\}^U$ satisfying the following desiderata (which are, to a certain extent, mutually incompatible and unattainable – indeed, our algorithm instead guarantees certain related conditions (E1)–(E6)):

- (D1) $\mathbf{E}[Y_j] = y_j$ for every $j \in U$
- (D2) The random variables Y_j are negatively correlated, in some sense.
- (D3) $Y(G) = 1$ for $G \in \mathcal{G}$.
- (D4) $MY = My$

To explain these further, consider the *independent-selection* rounding strategy. Formally, we define $Y = \text{INDSELECT}(\mathcal{G}, y)$ to be the vector obtained by selecting, independently for each block G , exactly one item j from G and setting $Y_j = 1$, so that each item j is selected with probability y_j . All other (non-selected) items have $Y_j = 0$. This is a valid probability distribution as $y(G) = 1$ and the entries of y are in the range $[0, 1]$. The vector Y satisfies desiderata (D1), (D2), (D3) perfectly. However, it only weakly satisfies (D4): specifically, the value of $M_k Y$ will be a sum of negatively-correlated random variables, which can deviate significantly from its mean $M_k y$.

As we will later discuss in Section 4.2, the cardinality-rounding setting can be viewed as a special case of knapsack-partition constraints. The standard cardinality-rounding algorithm perfectly satisfies (D1) and (D4); (D3) is not meaningful in this case. For (D2), it satisfies a limited but important form of negative correlation known as the *negative cylinder property* Chekuri et al. (2010); Byrka et al. (2017); Srinivasan (2001); Gandhi et al. (2006); namely, for any set $S \subseteq U$ the rounded variables X_i satisfy the conditions

$$\mathbf{E}\left[\prod_{i \in S} X_i\right] \leq \prod_{i \in S} x_i, \quad \text{and} \quad \mathbf{E}\left[\prod_{i \in S} (1 - X_i)\right] \leq \prod_{i \in S} (1 - x_i). \quad (1)$$

Our clustering results will require more general forms of negative correlation. Ideally, we would like for arbitrary disjoint sets $S, T \subseteq U$ to satisfy a similar “near-independence” property:

$$\mathbf{E}\left[\prod_{i \in S} X_i \prod_{i \in T} (1 - X_i)\right] \approx \prod_{i \in S} x_i \prod_{i \in T} (1 - x_i) \quad (2)$$

For example, Byrka et al. (2017) showed this property for the cardinality-rounding algorithm with a random permutation of the input vector, in some parameter ranges. Note that

Eq. (2) cannot be preserved exactly in an integral solution; for example, if $x_1 = x_2 = 1/2$, then any integral solution must either satisfy $\mathbf{E}[X_1(1 - X_2)] \geq 1/2$ or $\mathbf{E}[(1 - X_1)X_2] \geq 1/2$. This is part of the reason why property (D2) and general negative correlation are much more challenging than the negative cylinder property.

To overcome this fundamental barrier, our KPR algorithm terminates with a vector which has a small number t of fractional entries. To provide intuition, let us discuss how our rounding algorithm works in the cardinality-rounding setting. (We emphasize that it can handle much more general scenarios.) The precise sense in which it satisfies Eq. (2) is somewhat technical, but, as one example, we get

$$\mathbf{E}\left[\prod_{i \in S} X_i \prod_{i \in T} (1 - X_i)\right] \leq \left(\prod_{i \in S} x_i \prod_{i \in T} (1 - x_i)\right) + O(1/t) \quad (3)$$

We emphasize that Eq. (3) is only a simplified form of our results. In particular, we will handle cases where $S \cup T$ has a large size but only a few elements are “significant.”

Let us briefly compare KPR with other dependent rounding algorithms. The first main genre of such algorithms is based on variants of the Lovász Local Lemma (e.g., Harris and Srinivasan (2019); Leighton et al. (2001); Srinivasan (2006)). These algorithms have very good independence properties, but also lead to knapsack violations on the order of the “standard deviation.” A second genre of algorithm is based, like KPR, on Brownian motion in the constraint polytope. These algorithms are often targeted to discrepancy minimization, see e.g., Beck and Fiala (1981); Karp et al. (1987); Bansal and Nagarajan (2017); Bansal (2019), where the central goal is to show concentration bounds on linear functions of the variables. This is closely related to pairwise correlation (covariance) among the variables. Our algorithm gets tighter bounds and finer negative correlation properties, including correlations among many variables, by taking advantage of the special properties of the knapsack-partition constraints.

1.2 Additive pseudo-approximation

Knapsack constraints can be intractable to satisfy exactly, and so *pseudo-approximations* (i.e., solutions which only approximately satisfy the knapsack constraints) are often used instead. Many previous algorithms (e.g., Byrka et al. (2017)) have focused on what we refer to as ϵ -*multiplicative pseudo-solutions*: namely that $\sum_{i \in S} M_k(i) \leq 1 + \epsilon$ for each k . This should be distinguished from a true approximation algorithm, which finds a feasible solution whose objective function is within some constant factor of the optimal one.

As we have discussed, our rounding algorithm does not generate a fully-integral vector Y , it only produces a vector \tilde{Y} which is “mostly” integral, that is, $\tilde{Y} \in [0, 1]^U$ has only a small (essentially constant) number of fractional entries. This is critical to overcoming the tradeoff between probabilistic independence and satisfying the knapsack constraints. This naturally leads to an alternative, *additive* form of knapsack pseudo-approximation. We define this formally as follows:

Definition 1 (q -additive pseudo-approximation) *For a single-knapsack constraint (vector of weights w with capacity 1), a set \mathcal{S} is a q -additive pseudo-solution if it has the form $\mathcal{S} = \mathcal{S}_0 \cup \mathcal{S}_1$, where $\sum_{i \in \mathcal{S}_0} w_i \leq 1$ and $|\mathcal{S}_1| \leq q$. (Equivalently, \mathcal{S} satisfies the budget constraint after removing its q highest-weight items.)*

For a multi-knapsack constraint $M = M_1, \dots, M_m$, we say that \mathcal{S} is an q -additive pseudo-solution for M if it is a q -additive solution for each of the m knapsack constraints M_1, \dots, M_m

separately. That is, for each $k = 1, \dots, m$ we have $\mathcal{S} = \mathcal{S}_0^{(k)} \cup \mathcal{S}_1^{(k)}$ where $\sum_{i \in \mathcal{S}_0^{(k)}} M_k(i) \leq 1$ and $|\mathcal{S}_1^{(k)}| \leq q$.

Such additive pseudo-solutions are naturally connected to the *method of alteration* in the probabilistic method, where we delete/alter some items in a random structure to establish a desired property Alon and Spencer (2016). Additive pseudo-approximations have appeared implicitly in prior algorithms, e.g., Li and Svensson (2016); Krishnaswamy et al. (2011). We can summarize some of their advantages here, from both practical and technical points of view.

First, additive pseudo-approximation can be a useful tool to obtain true approximations. Given a q -additive pseudo-solution, we can often “fix” the q additional items in some problem-specific way. If q is small, this may incur only a small overhead in the cost or computational complexity. This strategy was used in the k -median approximation algorithm of Li and Svensson (2016). We use it here for our (true) approximation algorithm for knapsack center. These problems are described in Section 1.3.

As another example, there is a common strategy to obtain a multiplicative pseudo-approximation by “guessing” – exhaustively enumerating – the “big” items in an optimal solution, i.e., items with $M(i) > \rho$ for some parameter ρ . Then, a q -additive pseudo-approximation for the residual problem yields a ρq -multiplicative pseudo-approximation to the original problem. This approach can be more efficient than generating the multiplicative pseudo-solution directly. The reverse direction does not hold, in general; there does not seem to be any way to go from multiplicative to additive pseudo-solutions.

To see a practical advantage of additive pseudo-approximation, consider our motivating example concerning fair representation. In this setting, a q -additive pseudo-solution \mathcal{S} leads to a relatively modest violation of the fairness constraint, namely, it has $|\mathcal{S} \cap A_k| \leq t_k + q$ for each population A_k and associated target value t_k . By contrast, an ϵ -multiplicative pseudo-solution \mathcal{S} would give a substantially larger violation, namely $|\mathcal{S} \cap A_k| \leq t_k(1 + \epsilon)$.

1.3 Clustering results

In describing our clustering results, it is convenient to use the language of classical facility location. We refer to \mathcal{C} as a set of *clients*, \mathcal{F} as a set of *facilities*, and we say that $i \in \mathcal{F}$ is *open* if i is placed into the solution set \mathcal{S} . The distance $d(j, \mathcal{S})$ for a client j can be interpreted as the connection cost of j to its nearest open facility.

We study two clustering problems in particular: the *knapsack median* and *knapsack center* problems. In the knapsack median problem, we minimize the *total* connection cost $\text{cost}(\mathcal{S}) = \sum_j d(j, \mathcal{S})$ subject to m knapsack constraints. The knapsack center problem is the same except that the objective is to minimize $\text{cost}(\mathcal{S}) = \max_{j \in \mathcal{C}} d(j, \mathcal{S})$ instead of the sum $\sum_j d(j, \mathcal{S})$.

Knapsack median was first studied by Krishnaswamy et al. Krishnaswamy et al. (2011), who obtained an additive pseudo-approximation with an approximation factor of 16. The current best true approximation factor is 7.08 for the single-knapsack case, due to Krishnaswamy et al. (2018). The special case when all facilities in \mathcal{F} have unit weight and $m = 1$, known as the classical k -median problem, can be approximated to within a factor of $2.675 + \epsilon$ Byrka et al. (2017). Our KPR rounding algorithm gives the following pseudo-approximation results:

Theorem 2 *Let $\gamma, \epsilon \in (0, 1)$. For single-knapsack median, there is a polynomial-time algorithm to obtain an $O(1/\gamma)$ -additive pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (1 + \sqrt{3} + \gamma) \cdot \text{OPT} \leq 2.733 \cdot \text{OPT}$*

and an algorithm with $n^{O(\epsilon^{-1}\gamma^{-1})}$ runtime to obtain an ϵ -multiplicative pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (1 + \sqrt{3} + \gamma) \cdot \text{OPT}$.

We also consider multi-knapsack median, where it is NP-hard to obtain a true approximation. We apply KPR for a key step in an algorithm of Charikar & Li Charikar and Li (2012) to get an additive pseudo-approximation. This can also be leveraged into a multiplicative pseudo-approximation. We summarize these results as follows:

Theorem 3 *Let $\gamma, \epsilon \in (0, 1)$. For multi-knapsack median, there is a polynomial-time algorithm to obtain an $O(\frac{m}{\sqrt{\gamma}})$ -additive pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (3.25 + \gamma) \cdot \text{OPT}$, and an algorithm with $n^{O(\frac{m^2}{\epsilon\sqrt{\gamma}})}$ runtime to obtain an ϵ -multiplicative pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (3.25 + \gamma) \cdot \text{OPT}$.*

By contrast, independent selection in the Charikar-Li algorithm would take $n^{\tilde{O}(m/\epsilon^2)}$ runtime for an ϵ -multiplicative approximation (see Theorem 29). This illustrates how additive pseudo-approximation can be useful for efficient multiplicative pseudo-approximations; in particular, this gives a better dependence on the parameter ϵ (though a worse dependence on parameters m and γ).

The single-knapsack center problem was first studied by Hochbaum & Shmoys in Hochbaum and Shmoys (1986), under the name “weighted k -center”. They gave a 3-approximation algorithm and proved that this is best possible unless $P = NP$; see also Khuller et al. (2000). Our approximation algorithms ensure each client has better bounds on *expected* connection distance, in addition to the usual bound on maximum distance. We summarize this as follows:

Theorem 4 *Let $\gamma \in (0, 1)$. There is an algorithm for single-knapsack center with $n^{\tilde{O}(1/\gamma)}$ runtime which returns a feasible solution \mathcal{S} such that every client $j \in \mathcal{C}$ has*

$$\mathbf{E}[d(j, \mathcal{S})] \leq (1 + 2/e + \gamma) \cdot \text{OPT}, \quad d(j, \mathcal{S}) \leq 3 \cdot \text{OPT} \text{ with probability one.}$$

More recently, Chen et. al. Chen et al. (2016) considered the multi-knapsack center problem. They showed that it is intractable to obtain a true constant-factor approximation, and gave a multiplicative pseudo-approximation with approximation ratio 3. We obtain a number of new pseudo-approximations for this setting.

Theorem 5 *Let $\gamma, \epsilon \in (0, 1)$. For multi-knapsack median, we describe three algorithms to generate different types of pseudo-solutions \mathcal{S} such that every client $j \in \mathcal{C}$ has $\mathbf{E}[d(j, \mathcal{S})] \leq (1 + 2/e + \gamma) \cdot \text{OPT}$ and $d(j, \mathcal{S}) \leq 3 \cdot \text{OPT}$ with probability one.*

- (a) *A polynomial-time algorithm for an $\tilde{O}(m/\sqrt{\gamma})$ -additive pseudo-solution.*
- (b) *An algorithm with $n^{\tilde{O}(m^2/\gamma)}$ runtime for an $\tilde{O}(\sqrt{m})$ -additive pseudo-solution.*
- (c) *An algorithm with $n^{\tilde{O}(m^{3/2}/\epsilon + m^2/\gamma)}$ runtime for an ϵ -multiplicative pseudo-solution.*

Our algorithms thus give finer guarantees: all clients have distance $3 \cdot \text{OPT}$ to an open facility with probability one and also all clients have expected connection cost at most $(1 + 2/e) \cdot \text{OPT} \approx 1.74 \cdot \text{OPT}$. This can be helpful in *flexible* facility location, such as a streaming-service provider periodically reshuffling its service locations. It can also be interpreted as a type of fairness in clustering, where the fairness is in terms of individual users instead of demographic groups. Note that the constant factor $1 + 2/e$ cannot be improved unless $P = NP$, even in the k -supplier setting Harris et al. (2019).

1.4 Notation

In the context of clustering problems, we let $n = |\mathcal{F} \cup \mathcal{C}|$. For a client $j \in \mathcal{C}$ and a real number $x \geq 0$, we define the *facility-ball* $\mathcal{B}(j, x) = \{i \in \mathcal{F} \mid d(i, j) \leq x\}$. For a metric d and a set Y , we write $d(x, Y) = \min_{y \in Y} d(x, y)$.

For a non-negative integer t , we write $[t] = \{1, \dots, t\}$. For a set X and an t -dimensional vector y , we write $y(X) = \sum_{i \in X} y_i$. The *support* of y is defined to be the set of indices i where $y_i \neq 0$.

Given an $m \times n$ budget matrix M , we write $M(i)$ for the m -dimensional vector corresponding to the costs of item i , and we write M_k for the n -dimensional vector, which is a single knapsack-constraint, corresponding to the k^{th} row of M . Likewise, for a set $X \subseteq [n]$, we write $M(X) = \sum_{i \in X} M(i)$. For an n -dimensional vector y , we write $M_k y \leq \vec{1}$ to denote that $M_k y \leq 1$ for all $k = 1, \dots, m$.

Given a vector $Y \in \{0, 1\}^n$ where the support of Y is a q -additive pseudo-solution for a multi-knapsack constraint matrix M , we sometimes say for brevity that Y is a q -additive pseudo-solution for M , i.e. for each k we can modify Y to some Y' by zeroing out at most q entries such that $M_k Y' \leq \vec{1}$.

We use Iverson notation where $[[\phi]]$ is the indicator function for a Boolean predicate ϕ , i.e., $[[\phi]] = 1$ if ϕ is true, and $[[\phi]] = 0$ otherwise. The $\tilde{O}()$ notation is defined as $\tilde{O}(x) = x \cdot \text{polylog}(x)$.

A partition \mathcal{G} of ground set U is a collection of pairwise-disjoint sets G_1, \dots, G_k with $U = G_1 \cup \dots \cup G_k$. We refer to G_i as the *blocks* of the partition. For each $u \in U$, we define $\mathcal{G}(u)$ to be the unique block with $u \in \mathcal{G}(u) \in \mathcal{G}$. For $W \subseteq U$, we define $\mathcal{G}(W) \subseteq \mathcal{G}$ to be the set of blocks involved in W , i.e., $\mathcal{G}(W) = \{\mathcal{G}(w) \mid w \in W\}$. For a set of blocks $\mathcal{D} \subseteq \mathcal{G}$ and a set of items $W \subseteq U$, we define the *restriction of W to \mathcal{D}* , denoted $W \wedge \mathcal{D}$, to be the set of items in W which are also part of a block of \mathcal{D} , i.e., $W \wedge \mathcal{D} = \bigcup_{G \in \mathcal{D}} W \cap G = \{w \in W \mid \mathcal{G}(w) \in \mathcal{D}\}$.

1.5 Organization

In Section 2, we develop a new ‘‘Samuels-Feige’’ type of concentration inequality for unbounded random variables. This result is quite general, and may be of independent interest. In Section 3, we develop and analyze the new KPR dependent rounding algorithm. Section 4 describes a few extensions of this algorithm. Sections 5, 6, 7 describe the applications to knapsack median and knapsack center problems. Section 8 provides further analysis of concentration bounds and near-independence properties of the KPR algorithm.

2. A concentration inequality for additive knapsack pseudo-solutions

The main result of this section is to show an intriguing connection between independent rounding and pseudo-additive solutions. For maximum generality, we state it in terms of a broader class of random variables satisfying a property known as *negative association* (NA) Joag-Dev and Proschan (1983), defined as follows:

Definition 6 (Negatively associated random variables Joag-Dev and Proschan (1983)) *Random variables X_1, \dots, X_n are negatively associated (NA) if for every subset $A \subseteq [n]$, and any pair of non-decreasing functions f_1, f_2 , the random variables $f_1(X_i : i \in A)$ and $f_2(X_j : j \in [n] - A)$ have non-positive covariance. (Here, ‘‘ $f_1(X_i : i \in A)$ ’’ means f_1 applied to the tuple $(X_i : i \in A)$, and similarly for ‘‘ $f_2(X_j : j \in [n] - A)$ ’’.)*

If X_1, \dots, X_n are independent random variables, then they are NA. The class of NA random variables includes other random processes; for example, the load-vector of the urns in balls-and-urns processes Dubhashi and Ranjan (1998).

With this definition, we state our main result:

Theorem 7 *Let X_1, \dots, X_n be negatively-associated, non-negative random variables. Then with probability at least $1 - \delta$, there is a set $W \subseteq [n]$ (which may depend on the values of X_1, \dots, X_n) with $|W| \leq O(\sqrt{n \log \frac{1}{\delta}})$ such that $\sum_{i \in [n]-W} X_i \leq \sum_{i \in [n]} \mathbf{E}[X_i]$.*

Theorem 7 can be rephrased in the language of knapsack constraints:

Corollary 8 *Let X_1, \dots, X_n be negatively-associated, non-negative random variables, and let a_1, \dots, a_n be non-negative coefficients with $\sum_i a_i \mathbf{E}[X_i] = 1$. Then with probability at least $1 - \delta$, the values X_1, \dots, X_n form an $O(\sqrt{n \log \frac{1}{\delta}})$ -additive pseudo-solution to the knapsack constraint $a_1 x_1 + \dots + a_n x_n \leq 1$.*

Notably, this bound does not depend on the variance of the variables X_1, \dots, X_n , which is quite different from conventional concentration bounds such as Chernoff-Hoeffding. We remark that the bound is tight for many values of n and δ . For example, consider a system with n independent Bernoulli(p) variables for any constant $p \in (0, 1)$. In this case, we have $|W| = \max\{(\sum_i X_i) - np, 0\}$, which is on the order of $O(\sqrt{n \log \frac{1}{\delta}})$ with probability at least $1 - \delta$.

Let us briefly summarize the role of this concentration inequality in our overall rounding algorithm. As we have discussed, the KPR rounding algorithm stops with a vector $\tilde{Y} \in [0, 1]^U$ with some t remaining fractional entries, and which satisfies the knapsack constraints exactly. One attractive option to obtain a fully integral vector Y is to apply independent selection to \tilde{Y} , ignoring the knapsack constraints. This may violate the knapsack constraints, but by how much? This is precisely the random process governed by Corollary 8. The integral entries of \tilde{Y} have no effect and so \tilde{Y} effectively has t variables. If we set $q = O(\sqrt{t \log m})$, then Corollary 8 shows that Y is a q -additive pseudo-solution to each knapsack constraint M_k with probability $1 - \frac{1}{2m}$. A union bound over the m knapsack constraints shows that Y is a q -additive pseudo-solution for M with constant probability.

By way of comparison, Feige (2006); Garnett (2020); He et al. (2010); Samuels (1966) give other concentration bounds for sums of nonnegative independent variables without regard to size or variance. Such results are useful for problems such as hypergraph matchings and probabilistic estimation of network parameters Goldreich and Ron (2008); Eden et al. (2019); Alon et al. (2012). We anticipate that Theorem 7 may have broader applications beyond our rounding algorithm.

2.1 Formal proof of Theorem 7

We begin by recalling some useful properties of NA random variables; see Joag-Dev and Proschan (1983); Dubhashi and Ranjan (1998) for further details.

Proposition 9 *Let X_1, \dots, X_n be NA random variables.*

(Q1) *If f_1, \dots, f_n are univariate non-decreasing functions, then $f_1(X_1), \dots, f_n(X_n)$ are NA as well.*

(Q2) If X_1, \dots, X_n are bounded in the range $[0, 1]$, then the Chernoff-Hoeffding bounds (for both the upper- and lower-tails of sums) apply to them as they do to independent random variables.

The proof of Theorem 7 has two parts. First, we prove it under the assumption that the random variables X_1, \dots, X_n have continuous cumulative density functions (CDF's); we then use a “smoothing” argument to extend it to arbitrary distributions.

Let us define $\lambda = \log \frac{1}{\delta}$. If $n < c\lambda$ for any chosen constant c , then the result will hold trivially by taking $W = [n]$. Thus, we assume that $n > c\lambda$ for any needed constant c in the proof.

Part I. Suppose that X_1, \dots, X_n have continuous CDF's. By rescaling, we assume without loss of generality that $\sum_i \mathbf{E}[X_i] = 1$. Since the CDF of the X_i variables is continuous and the X_i variables are non-negative, there is a real number $\alpha \geq 0$ such that

$$\sum_{i \in [n]} \Pr(X_i > \alpha) = 10\lambda \quad (4)$$

We will take W to be the set of all indices $i \in [n]$ with $X_i > \alpha$, thus $\mathbf{E}[|W|] = 10\lambda$. We need to show that $|W| \leq O(\lambda)$ and $\sum_{i \in [n]-W} X_i \leq 1$. Let us define the random variables $Y_i = \min(\alpha, X_i)$. Noting that $Y_i = X_i$ for $i \notin W$ while $Y_i = \alpha$ for $i \in W$, we see that the following equation holds with probability one:

$$\sum_{i \in [n]-W} X_i = \sum_{i \in [n]} (Y_i - \alpha \mathbb{1}[i \in W]) = -\alpha|W| + \sum_{i \in [n]} Y_i \quad (5)$$

Let \mathcal{E}_1 denote the event that $5\lambda \leq |W| \leq 20\lambda$ and let \mathcal{E}_2 denote the event that $\sum_i Y_i \leq 1 + 5\alpha\lambda$. By Eq. (5), when \mathcal{E}_1 and \mathcal{E}_2 both hold, we have $\sum_{i \in [n]-W} X_i \leq 1$ and $|W| \leq O(\lambda)$, as is desired. It remains to compute the probabilities of \mathcal{E}_1 and \mathcal{E}_2 .

First, since the function mapping x to $\mathbb{1}[x > \alpha]$ is non-decreasing, the indicator variables for $i \in W$ remain negatively associated by (Q1). By (Q2) and Eq. (4) we therefore have

$$\Pr(\mathcal{E}_1) = \Pr(5\lambda \leq |W| \leq 20\lambda) \geq 1 - 2e^{-\Omega(\lambda)} \geq 1 - 2e^{-\Omega(\sqrt{n \log(1/\delta)})} \geq 1 - \delta/2$$

using our assumption that $n > c\lambda$ for any sufficiently large constant c .

Next, since $Y_i \leq X_i$, we have $\mathbf{E}[\sum_i Y_i] \leq 1$. The random variables Y_i are all bounded in the range $[0, \alpha]$. Also since the function mapping x to $\min(x, \alpha)$ is non-decreasing, by (Q1) the random variables Y_i remain negatively associated. By (Q2) we can apply Hoeffding's bound, and thus

$$\Pr(\neg \mathcal{E}_2) = \Pr\left(\sum_{i \in [n]} Y_i > 1 + 5\alpha\lambda\right) \leq e^{-\frac{-2(5\alpha\lambda)^2}{n\alpha^2}} = e^{-50\lambda^2/n} = e^{-50 \log(1/\delta)} \leq \delta/2$$

Thus, we have $\Pr(\mathcal{E}_1 \cap \mathcal{E}_2) \geq 1 - \delta$, and in this case we have $\sum_{i \in [n]-W} X_i \leq 1$ for $|W| \leq O(\lambda)$.

Part II. Now consider random variables X_1, \dots, X_n (with no assumption on their CDF's). Let Y_1, \dots, Y_n be independent random variables which are uniform in the range $[0, \frac{\epsilon}{n}]$, and let $Z_i = X_i + Y_i$. The random variables Z_i are clearly non-negative, and they have continuous CDF's. For any fixed values of Y_1, \dots, Y_n , the random variables Z_1, \dots, Z_n are NA by (Q1); this remains true after integrating over Y_1, \dots, Y_n .

So as we showed in Part I of the proof, with probability at least $1 - \delta$ we have

$$\min_{\substack{W \subseteq [n] \\ |W| \leq r}} \sum_{i \in [n]-W} Z_i \leq \sum_{i \in [n]} \mathbf{E}[Z_i]$$

for some parameter $r = O(\sqrt{n\lambda})$. Since $\mathbf{E}[Z_i] = \mathbf{E}[X_i] + \epsilon n/2$ and $X_i \leq Z_i$, this implies that

$$\Pr\left(\min_{\substack{W \subseteq [n] \\ |W| \leq r}} \sum_{i \in [n]-W} X_i \leq \sum_{i \in [n]} \mathbf{E}[X_i] + \epsilon/2\right) \geq 1 - \delta$$

Since this holds for every $\epsilon > 0$, this implies:

$$\Pr\left(\min_{\substack{W \subseteq [n] \\ |W| \leq r}} \sum_{i \in [n]-W} X_i \leq \sum_{i \in [n]} \mathbf{E}[X_i]\right) \geq 1 - \delta$$

This concludes the proof. We remark that our result is algorithmically “local” in that it discards each i if and only if $X_i > \alpha$ for a certain threshold value α ; this automatically ensures that the discarded set is small with good probability.

3. Knapsack-partition systems and rounding

The KPR rounding algorithm takes as input the partition \mathcal{G} over ground set U , the constraint matrix M , the fractional vector y satisfying $y(G) = 1$ for all $G \in \mathcal{G}$, and an integer parameter t . It returns a mostly rounded vector \tilde{Y} .

For $G \in \mathcal{G}$, we define the following functions to count the number of fractional entries:

$$T_G(y) = \max\{0, |\{i \in G \mid y_i \in (0, 1)\}| - 1\}, \quad \text{and} \quad T(y) = \sum_{G \in \mathcal{G}} T_G(y)$$

We then define the algorithm as follows:

Algorithm 1 KPR(\mathcal{G}, M, y, t)

- 1: **for** each block $G \in \mathcal{G}$ **do**
 - 2: Execute an unbiased walk to generate random vector $y' \in [0, 1]^U$ with $\mathbf{E}[y'] = y$, and such that y' is an extreme point of the polytope: $\{My' = My, y'(G) = 1, y'_j = y_j \text{ for } j \notin G\}$.
 - 3: Update $y \leftarrow y'$
 - 4: **while** $T(y) > t$ **do**
 - 5: Form a set $\mathcal{J} \subseteq \mathcal{G}$, wherein each $G \in \mathcal{G}$ goes into \mathcal{J} independently with probability $p = 3m/T(y)$.
 - 6: **if** $\sum_{G \in \mathcal{J}} T_G(y) \geq m + 1$ **then**
 - 7: Choose $\delta \in \mathbb{R}^U$ such that
 - $M\delta = 0$, $y + \delta \in [0, 1]^U$, and $y - \delta \in [0, 1]^U$
 - There is at least one index i with $y_i \in (0, 1)$ such that $y_i + \delta_i \in \{0, 1\}$ or $y_i - \delta_i \in \{0, 1\}$.
 - $\delta_j = 0$ if $\mathcal{G}(j) \notin \mathcal{J}$
 - $\delta(G) = 0$ for all $G \in \mathcal{G}$.
 - 8: With probability $1/2$, update $y \leftarrow y + \delta$; else, update $y \leftarrow y - \delta$
 - 9: **return** y
-

The loop at lines 1–3 is a preprocessing step consisting of straightforward dependent rounding within each block G . We write $y' = \text{INTRABLOCKREDUCE}(y)$ for the vector obtained at the termination of the loop. After this step, the algorithm repeatedly applies a more-complicated rounding

process which modifies multiple blocks. We write $y' = \text{KPR-ITERATION}(y)$ to denote a single iteration of the loop at lines 4–8. Thus, the overall algorithm is equivalent to the following:

Algorithm 2 $\text{KPR}(\mathcal{G}, M, y, t)$, summarized

```

1:  $y \leftarrow \text{INTRABLOCKREDUCE}(y)$ 
2: while  $T(y) > t$  do
3:   update  $y \leftarrow \text{KPR-ITERATION}(y)$ 
4: return  $y$ 
    
```

The KPR algorithm *requires throughout* that $t > 12m$; this assumption will not be stated explicitly again. Because of this condition, the probability p in line 5 is at most $3m/t \leq 1/4$. (Note that it is likely impossible to obtain fewer than m fractional entries, while still respecting the knapsack constraints.)

Also, although budget matrices for clustering problems are usually assumed to be non-negative, we do not require this for KPR. The entries of the matrix M can be arbitrary real numbers.

3.1 KPR algorithm: formal results

As we have discussed, desiderata (D1)–(D4) cannot be exactly satisfied. To describe the negative correlation properties of KPR, we use a potential function $Q(W, x)$ defined for a set $W \subseteq U$ and vector $x \in [0, 1]^U$ as follows:

$$Q(W, x) = \prod_{G \in \mathcal{G}} (1 - x(W \cap G)).$$

We will show that the vector $\tilde{Y} = \text{KPR}(\mathcal{G}, M, y, t)$ satisfies the following constraints:

- (E1) For all $W \subseteq U$, $\mathbf{E}[Q(W, \tilde{Y})]$ is “not much more than” $Q(W, y)$;
- (E2) Every $j \in U$ has $\mathbf{E}[\tilde{Y}_j] = y_j$;
- (E3) $\tilde{Y}(G) = 1$ for $G \in \mathcal{G}$;
- (E4) $M\tilde{Y} = My$;
- (E5) At most $2t$ entries of \tilde{Y} are fractional.
- (E6) For each block $G \in \mathcal{G}$, at most $m + 1$ entries of \tilde{Y} are fractional.

(E1) is intentionally vague, as the relationship between $\mathbf{E}[Q(W, \tilde{Y})]$ and $Q(W, y)$ is quite complex. Our main result covers a setting needed for a number of our clustering algorithms, where there is a relatively small set \mathcal{D} of blocks G which have $y(G \cap W)$ close to one. Formally, we show the following:

Theorem 10 *Let $\mathcal{D} \subseteq \mathcal{G}$ with $|\mathcal{D}| = d$ and let $\tilde{Y} = \text{KPR}(\mathcal{G}, M, y, t)$ with $t \geq 5000m(d + 1)$. Then for any set $W \subseteq U$, there holds*

$$\mathbf{E}[Q(W, \tilde{Y})] \leq Q(W, y) + Q(W \wedge \mathcal{D}, y)(e^{O((d+1)^2 m^2/t)} - 1)$$

Recall that we define $W \wedge \mathcal{D} = \bigcup_{G \in \mathcal{D}} W \cap G$. Theorem 10 is complex and hard to use directly. We derive a number of simplified results, such as the following three estimates:

Theorem 11 *Let $\tilde{Y} = \text{KPR}(\mathcal{G}, M, y, t)$ and let $W \subseteq U$.*

- (a) *For $\mathcal{D} \subseteq \mathcal{G}$ with $t > 5000m^2(|\mathcal{D}| + 1)^2$ there holds $\mathbf{E}[Q(W, \tilde{Y})] \leq Q(W, y) + O((|\mathcal{D}| + 1)^2 m^2 / t) \cdot Q(W \wedge \mathcal{D}, y)$.*
- (b) *For $t > 12m$, there holds $\mathbf{E}[Q(W, \tilde{Y})] \leq Q(W, y) + O(m^2/t)$.*
- (c) *For $t > 10000md$ where $d = |\mathcal{G}(W)|$, there holds $\mathbf{E}[Q(W, \tilde{Y})] \leq Q(W, y)e^{O(m^2 d^2 / t)}$.*

There are two key steps in KPR to ensure property (E1). First, the modification vector δ in line 4 is always bounded by the current value of y . This ensures that the typical change in the value of $Q(W, y)$ is proportional to the current value of $Q(W, y)$. Second, the set \mathcal{J} , which determines the entries of y to modify, is randomly selected. This spreads out the (inevitable) correlation among the entries of y .

Before we show Theorem 10, let us explain the role played by the potential function Q . Observe that for $Y = \text{INDSELECT}(\mathcal{G}, y)$ we have $Y(W) = 0$ if and only if $Q(W, Y) = 1$ and so

$$\Pr(Y(W) = 0) = \mathbf{E}[Q(W, Y)] = Q(W, y);$$

thus, $Q(W, \tilde{Y})$ is a smoothed measure of whether the KPR output \tilde{Y} satisfies $\tilde{Y}(W) = 0$.

Why might one be interested in upper-bounding terms of the form $\Pr(Y(W) = 0)$? We have briefly touched on this, but let us spell out in greater detail how such bounds arise in clustering algorithms, such as our algorithms for knapsack median and knapsack center. The simplest versions of these algorithms first cluster the facilities in some greedy manner; these are the blocks G of the partition. They then open a facility suitably at random from each block. Any given client j will first check if some “nearby” facility gets opened; if not, then it must use a “backup” facility which, however, is farther away. The bad event of no opened “nearby” facility corresponds to $Y(W) = 0$, where Y is the indicator vector for which facilities are open and W is the set of nearby facilities.

In more advanced algorithms, multiple facilities may be opened from a cluster, or the clusters may have even more complex interactions. These cases can also be interpreted as knapsack-partition systems, and again the distance for a client j can be recast in terms of events of the form $Y(W) = 0$.

Section 8 shows some further upper bounds on $\mathbf{E}[Q(W, \tilde{Y})]$. These do not follow from Theorem 10, and are not directly useful for our clustering algorithms. For simplicity, we do not attempt to optimize the constant factors here or elsewhere in the analysis.

We will begin by showing some easier properties of this algorithm, including that it is well-defined and terminates in polynomial time. The proof of (E1), which is much harder, comes next.

3.2 Simple properties and convergence of KPR

Proposition 12 *The vector $y' = \text{INTRABLOCKREDUCE}(y)$ satisfies properties (E2), (E3), (E4), (E6). Furthermore, for any $W \subseteq U$ we have $\mathbf{E}[Q(W, y')] = Q(W, y)$.*

Proof The polytope in line 2 of KPR has $m + 1$ constraints among entries y_j for $j \in G$, so an extreme point has at most $m + 1$ fractional entries. The polytope conditions preserve properties (E3) and (E4), and preserve (E2) since the walk is unbiased. Finally, the change in y during each

iteration is confined to block G . Since $Q(W, y)$ is linear function of y_j for $j \in G$, (E2) ensures that $\mathbf{E}[Q(W, y)]$ does not change. \blacksquare

Proposition 13 *The vector δ in line 7 of KPR exists and can be found efficiently.*

Proof In forming the vector δ , there is one degree of freedom for each entry $i \in G$ with $y_i \in (0, 1)$ and $G \in \mathcal{J}$. We further have m linear constraints (from the matrix M) and $|\mathcal{J}|$ linear constraints (from the condition $\delta(G) = 0$ for each $G \in \mathcal{J}$). This gives a total of $\sum_{G \in \mathcal{J}} T_G(y) - m$ degrees of freedom.

So the linear system has a non-zero solution vector v as long as $\sum_{G \in \mathcal{J}} T_G(y) \geq m + 1$, which is precisely the condition at line 6 of KPR. Now let $a \in \mathbb{R}$ be maximal such that $y + av \in [0, 1]^U$ and $y - av \in [0, 1]^U$. One may verify that $a < \infty$ and that $\delta = a\gamma$ achieves the claimed result. \blacksquare

Proposition 14 *In each iteration $y' \leftarrow \text{KPR-ITERATION}(y)$ of KPR, there is a probability of at least 0.24 that y' has at least one more integral coordinate than y .*

Proof By Proposition 12, the vector y after INTRABLOCKREDUCE has $T_G(y) \leq m$ for each $G \in \mathcal{G}$. We also have $T(y) > t$, as otherwise KPR would have terminated.

For each $G \in \mathcal{G}$, define $Z_G = \mathbb{I}[G \in \mathcal{J}]T_G(y)$ and define $Z = \sum_G Z_G$. If $Z \geq m + 1$, then there is at least 1/2 probability of producing at least one new rounded entry in y' . Also, Z is a sum of independent random variables with mean $pT(y) = 3m$; furthermore, each Z_G is bounded in the range $[1, m]$. A simple analysis with Chernoff's bound shows that $\Pr(Z \geq m + 1) \geq 0.48$, which gives the claimed result. \blacksquare

We will later need the following stronger version of Proposition 14 where we condition on a given iteration not touching a given subset of blocks.

Proposition 15 *Consider an iteration $y' \leftarrow \text{KPR-ITERATION}(y)$ of KPR. For any set of blocks $\mathcal{D} \subseteq \mathcal{G}$ with $t \geq 5000m(|\mathcal{D}| + 1)$, there is a probability of at least 1/10 that $\mathcal{D} \cap \mathcal{J} = \emptyset$ and $T(y') < T(y)$.*

Proof Let $d = |\mathcal{D}|$. Letting \mathcal{E} denote the event that $\mathcal{D} \cap \mathcal{J} = \emptyset$, we have

$$\Pr(\mathcal{E}) = (1 - 3m/T(y))^d \geq (1 - 3m/t)^d \geq \left(1 - \frac{3m}{5000m(d+1)}\right)^d \geq 0.99$$

Next, as in Proposition 14, let us define $Z = \sum_{G \in \mathcal{J}} \mathbb{I}[G \in \mathcal{J}]T_G(y)$. If $Z \geq m + 1$, then with probability 1/2 there will be at least one rounded variable. Furthermore, conditional on event \mathcal{E} , here Z is a sum of independent random variables in the range $[1, m]$ and with mean $\mu = \sum_{G \in \mathcal{G} - \mathcal{D}} pT_G(y)$. Since $T(y) \geq t$ and $T_G(y) \leq m$ for all blocks G , we have $\mu \geq (3m/t) \cdot (t - dm) \geq \frac{3m(5000m(d+1) - dm)}{5000md} \geq 2.99m$. So, by Chernoff's bound, $\Pr(Z < m + 1 \mid \mathcal{E}) \leq e^{-(2.99)^2(0.34)^2/2} \leq 0.60$. Overall, the desired event happens with probability at least $0.99 \cdot (1 - 0.60) \cdot 1/2 \geq 1/10$. \blacksquare

Proposition 16 *The output \tilde{Y} of $\text{KPR}(\mathcal{G}, M, y, t)$ satisfies properties (E2) — (E6).*

Proof By Proposition 12, the conditions (E2), (E3), (E4), (E6) hold after INTRABLOCKREDUCE . Each application of KPR-ITERATION updates $y \leftarrow y \pm \delta$. Since $\delta(G) = 0$ and $M\delta = 0$, properties (E3), (E4) are preserved. The expected change in y is $\frac{1}{2}\delta + \frac{1}{2}(-\delta) = 0$, so property (E2) is preserved. For property (E5), note that KPR terminates when $T(y) \leq t$, in which case $\tilde{Y} = y$ has at most $2t$ fractional entries. \blacksquare

Proposition 17 *KPR runs in expected polynomial time.*

Proof Clearly INTRABLOCKREDUCE runs in polynomial time. By Proposition 14, each iteration of KPR-ITERATION has a probability $\Omega(1)$ of causing a new entry to become integral, in which case $T(y)$ decreases by at least one. This implies that the expected number of iterations is $O(|U|)$, and by Proposition 13, each iteration can be implemented in polynomial time. \blacksquare

3.3 Property (E1): Proof of Theorem 10

Before the formal analysis, let us provide an overview. Consider the evolution of the random variable $Q(W, y)$ for some set W . There may be small increase in the expected value of $Q(W, y)$ in each iteration of KPR , but this is compensated by steady decrease in $T(y)$. To measure this, we will define a potential function Φ which depends on $Q(W, y)$ and $T(y)$; we will show that $\mathbf{E}[\Phi(y)]$ does not increase in any iteration of KPR . Furthermore, when $T(y) \leq t$ at the end of the process, we have $\Phi(y) = Q(W, y)$. Consequently, the final value $\mathbf{E}[Q(W, \tilde{Y})] = \mathbf{E}[\Phi(\tilde{Y})]$ is at most the initial value $\Phi(y)$.

We first show a useful result on how the potential function Q changes during a single iteration of KPR .

Lemma 18 *Suppose we are in the middle of executing KPR with state vector y , and let y' be state vector at the next iteration. If we condition on the vector y , then for each $W \subseteq U$ there holds*

$$\mathbf{E}[Q(W, y') \mid y] \leq Q(W, y) \cosh\left(\frac{6m}{T(y)} \sum_{G \in \mathcal{G}} y(G \cap W)\right)$$

Proof Define $S = Q(W, y)$ and $S' = Q(W, y')$. All probability calculations here are conditioned on y . Let us first condition as well on the random variable δ . Note $y' = y + \delta$ or $y' = y - \delta$, each with probability $1/2$. If we define $b_G = y(G \cap W)$ and $\mu_G = \delta(G \cap W)$ for each $G \in \mathcal{G}$, we therefore get:

$$\begin{aligned} \mathbf{E}[S' \mid \delta] &= 1/2 \prod_G (1 - b_G - \mu_G) + 1/2 \prod_G (1 - b_G + \mu_G) \\ &= 1/2 \sum_{X \subseteq \mathcal{G}} \prod_{G \in X} (-\mu_G) \prod_{G \in U-X} (1 - b_G) + 1/2 \sum_{X \subseteq \mathcal{G}} \prod_{G \in X} (\mu_G) \prod_{G \in U-X} (1 - b_G) \\ &= \sum_{\substack{X \subseteq \mathcal{G} \\ |X| \text{ even}}} \prod_{G \in X} \mu_G \prod_{i \in U-X} (1 - b_G) = S \sum_{\substack{X \subseteq \mathcal{G} \\ |X| \text{ even}}} \prod_{G \in X} \frac{\mu_G}{1 - b_G} \leq S \sum_{\substack{X \subseteq \mathcal{G} \\ |X| \text{ even}}} \prod_{G \in X} \frac{|\mu_G|}{1 - b_G} \end{aligned}$$

For each block G define $R_G = 2[[G \in \mathcal{J}]]b_G(1 - b_G)$. We claim that $|\mu_G| \leq R_G$ with probability one. For, if $G \notin \mathcal{J}$, then $\mu_G = 0 = R_G$. If $G \in \mathcal{J}$, then necessarily $\sum_{j \in G \cap W} (y_j + \delta_j)$ is bounded in $[0, 1]$, and hence $|\mu_G| \leq \min(b_G, 1 - b_G) \leq 2b_G(1 - b_G) = R_G$. Therefore,

$$\mathbf{E}[S' \mid \delta] \leq S \sum_{\substack{X \subseteq \mathcal{G} \\ |X| \text{ even}}} \prod_{G \in X} \frac{|\mu_G|}{1 - b_G} \leq S \sum_{\substack{X \subseteq \mathcal{G} \\ |X| \text{ even}}} \prod_{G \in X} \frac{R_G}{1 - b_G}$$

The random variables R_G are independent, and each has mean $\mathbf{E}[R_G] = (3m/T(y)) \cdot 2b_G(1 - b_G)$. Now integrate over random variables δ and \mathcal{J} to obtain:

$$\begin{aligned} \mathbf{E}[S'] &\leq S \sum_{\substack{X \subseteq \mathcal{G} \\ |X| \text{ even}}} \prod_{G \in X} \frac{\mathbf{E}[R_G]}{1 - b_G} \leq S \sum_{v=0}^{\infty} \frac{1}{(2v)!} \left(\sum_G \frac{\mathbf{E}[R_G]}{1 - b_G} \right)^{2v} = S \cosh \left(\sum_G \frac{\mathbf{E}[R_G]}{1 - b_G} \right) \\ &\leq S \cosh \left(\sum_G \frac{(3m/T(y)) \cdot 2b_G(1 - b_G)}{1 - b_G} \right) = S \cosh \left(\frac{6m \sum_G b_G}{T(y)} \right) \end{aligned}$$

■

We are now ready to show Theorem 10; Theorem 11 will follow as an immediate corollary.

Proof of Theorem 10 Let us fix $\mathcal{D} \subseteq \mathcal{G}$ and $W \subseteq U$ with $|D| = d$ and $t \geq 5000m(d + 1)$. We define parameter $a = 8000m^2(d + 1)^2$ and we define potential function Φ for a vector $x \in [0, 1]^U$ by:

$$\Phi(x) = \begin{cases} Q(W, x) + Q(W \wedge \mathcal{D}, x)(e^{a(\frac{1}{t} - \frac{1}{T(x)})} - 1) & \text{if } T(x) \geq t \\ Q(W, x) & \text{if } T(x) < t \end{cases}$$

The key to the proof is to show that, if we are in the middle of executing KPR, with state vector y , and we let y' be the state vector at the next iteration, then there holds

$$\mathbf{E}[\Phi(y') \mid y] \leq \Phi(y) \tag{6}$$

To show Eq. (6), let us fix the state y . We may assume that $T(y) > t$, as otherwise the algorithm is done and $y = y'$. We define a number of parameters as follows:

$$\begin{aligned} S_0 &= Q(W \wedge \mathcal{D}, y) & S'_0 &= Q(W \wedge \mathcal{D}, y') \\ S_1 &= Q(W \wedge (\mathcal{G} - \mathcal{D}), y) & S'_1 &= Q(W \wedge (\mathcal{G} - \mathcal{D}), y') \\ S &= Q(W, y) = S_0 S_1 & S' &= Q(W, y') = S'_0 S'_1 \\ T &= T(y) & T' &= T(y') \\ \beta &= e^{a(1/t - 1/T)} & \beta' &= \max\{1, e^{a(1/t - 1/T')}\} \end{aligned}$$

With this notation, we observe that

$$\Phi(y) = S + S_0(\beta - 1), \quad \Phi(y') = S' + S'_0(\beta' - 1).$$

Now let \mathcal{E}_1 denote the event that $T' < T$; when \mathcal{E}_1 occurs, then we have $\beta' \leq e^{a(1/t - 1/(T-1))}$ since $T > t$. Our condition on t implies that $T \geq t \geq 2\sqrt{a}$; as we show in Proposition 49, we thus have

$$\beta' \leq \beta - \frac{a[[\mathcal{E}_1]]\beta}{2T^2}.$$

So $\Phi(y') \leq S' + S'_0(\beta - 1 - \frac{a[[\mathcal{E}_1]]^\beta}{2T^2})$. Taking expectations gives

$$\mathbf{E}[\Phi(y')] - \Phi(y) \leq \mathbf{E}[S'] + \mathbf{E}[S'_0(\beta - 1)] - \mathbf{E}\left[S'_0 \frac{a[[\mathcal{E}_1]]^\beta}{2T^2}\right] - S - S_0(\beta - 1)$$

which we rearrange as:

$$\mathbf{E}[\Phi(y')] - \Phi(y) \leq \mathbf{E}[S' - S] + (\beta - 1) \mathbf{E}[S'_0 - S_0] - \frac{a\beta \mathbf{E}[S'_0 | \mathcal{E}_1] \Pr(\mathcal{E}_1)}{2T^2} \quad (7)$$

Let us consider the terms in Eq. (7) in turn. For the last term, define \mathcal{E}_2 to be the event that the set \mathcal{J} formed at line 5 of KPR satisfies $\mathcal{D} \cap \mathcal{J} = \emptyset$. We can estimate $\mathbf{E}[S'_0 | \mathcal{E}_1] \Pr(\mathcal{E}_1) \geq \mathbf{E}[S'_0 | \mathcal{E}_1 \cap \mathcal{E}_2] \Pr(\mathcal{E}_1 \cap \mathcal{E}_2)$. When \mathcal{E}_1 and \mathcal{E}_2 occur, then $S'_0 = S_0$ as none of the entries in $W \wedge \mathcal{D}$ are modified. By Proposition 15, $\Pr(\mathcal{E}_1 \cap \mathcal{E}_2) \geq 1/10$ so overall $\mathbf{E}[S'_0 | \mathcal{E}_1] \Pr(\mathcal{E}_1) \geq S_0/10$. Substituting the value of a , we get

$$\frac{a\beta \mathbf{E}[S'_0 | \mathcal{E}_1] \Pr(\mathcal{E}_1)}{2T^2} \geq 400m^2(d+1)^2\beta S_0/T^2 \quad (8)$$

Our next step is to estimate the term $\mathbf{E}[S'_0 - S_0]$. By Lemma 18 applied to the set $W \wedge \mathcal{D}$, we have

$$\mathbf{E}[S'_0] \leq S_0 \cosh\left(6m \sum_{G \in \mathcal{D}} y(G \cap W)/T\right) \leq S_0 \cosh(6md/T)$$

By our assumption that $T \geq t \geq 5000m(d+1)$, this is at most $S_0(1 + (6md/T)^2)$, and so

$$\mathbf{E}[S'_0] - S_0 \leq 36S_0(md/T)^2 \leq 36S_0m^2(d+1)^2/T^2 \quad (9)$$

Finally, we turn to estimating $\mathbf{E}[S' - S]$. By Lemma 18 applied to the set W , we have $\mathbf{E}[S'] \leq S \cosh(6m \sum_G y(G \cap W)/T)$. Since $0 \leq y(G \cap W) \leq 1$ and $x \leq -\ln(1-x)$ for all $x \in [0, 1]$, we have:

$$\mathbf{E}[S'] \leq S \cosh\left(\frac{6m}{T} \left(\sum_{G \in \mathcal{D}} 1 - \sum_{G \in \mathcal{G}-\mathcal{D}} \ln(1 - y(G \cap W))\right)\right) = S_0 S_1 \cosh(6m(d - \ln S_1)/T)$$

As we show in Proposition 48, this implies that

$$\mathbf{E}[S'] - S \leq S_0 \left(S_1 \cosh(6m(d - \ln S_1)/T) - S_1\right) \leq S_0(6m(d+1)/T)^2 \quad (10)$$

Substituting the estimates of Eqs. (8), (9), (10) into Eq. (7), we see that

$$\begin{aligned} \mathbf{E}[\Phi(y')] - \Phi(y) &\leq 36S_0m^2(d+1)^2/T^2 + 36(\beta - 1)S_0m^2(d+1)^2/T^2 - 400m^2(d+1)^2S_0/T^2 \\ &= S_0m^2(d+1)^2(36 + 36(\beta - 1) - 400\beta)/T^2 = -364S_0m^2(d+1)^2/T^2 \end{aligned}$$

which is non-positive; this shows Eq. (6) as desired.

To complete the proof of Theorem 10, suppose we execute KPR with input vector y . Let y' be the vector after INTRABLOCKREDUCE and let \tilde{Y} be the output vector. Since $\mathbf{E}[Q(X, y')] = \mathbf{E}[Q(X, y)]$ for all sets $X \subseteq U$, we have

$$\mathbf{E}[\Phi(y')] \leq \mathbf{E}[Q(W, y')] + (e^{a/t} - 1) \mathbf{E}[Q(W \wedge \mathcal{D}, y')] = Q(W, y) + (e^{a/t} - 1)Q(W \wedge \mathcal{D}, y)$$

By Eq. (6), and induction on all iterations of KPR, the output \tilde{Y} satisfies $\mathbf{E}[\Phi(\tilde{Y})] \leq \mathbf{E}[\Phi(y')]$. At the termination of KPR, we have $T(\tilde{Y}) \leq t$ and so $\Phi(\tilde{Y}) = Q(W, \tilde{Y})$. Putting these inequalities together, we have shown that

$$\mathbf{E}[Q(W, \tilde{Y})] \leq Q(W, y) + Q(W \wedge \mathcal{D}, y)(e^{a/t} - 1).$$

■

Proof [Proof of Theorem 11] For part (a), observe that when $t \geq 5000m^2(d+1)^2$ we have $e^{O(m^2(d+1)^2/t)} - 1 = O(m^2(d+1)^2/t)$. Part (b) follows from part (a) with $\mathcal{D} = \emptyset$; note that if $t < 10000m^2$ then the bound holds vacuously since $Q(W, \tilde{Y}) \leq 1$ with probability one. Part (c) follows Theorem 10 with $\mathcal{D} = \mathcal{G}(W)$; if if $d = 0$, then $W = \emptyset$ so the bound holds vacuously. ■

4. Variants of KPR

We summarize here some simpler ways to use KPR, which will occur in a number of algorithmic scenarios.

4.1 KPR followed by independent selection

One natural rounding strategy for a knapsack-partition problem is to execute KPR up to some stage t , and then finish by independent rounding. We define this formally as the algorithm FULLKPR:

Algorithm 3 FULLKPR(\mathcal{G}, M, y, t)

- 1: $\tilde{Y} \leftarrow \text{KPR}(\mathcal{G}, M, y, t)$
 - 2: $Y \leftarrow \text{INDSELECT}(\mathcal{G}, \tilde{Y})$
 - 3: **return** Y
-

The resulting vector $Y \in \{0, 1\}^U$ is fully integral; it will not exactly satisfy the knapsack constraints, but it will be relatively close (depending on the value of t). Since independent selection does not change the expectation of $Q(W, y)$, all the analysis for KPR carries over immediately to FULLKPR.

Theorem 19 *Let $Y = \text{FULLKPR}(\mathcal{G}, M, y, t)$ with $t > 12m$. Suppose that the constraint matrix M is non-negative and satisfies $My \leq \vec{1}$. Then with probability at least $1 - \delta$, the vector Y is a q -additive pseudo-solution to M where $q = O(\sqrt{t \log \frac{m}{\delta}})$. This probability bound holds even after conditioning on the fixed vector \tilde{Y} .*

Proof Let us fix $\tilde{Y} = \text{KPR}(\mathcal{G}, M, y, t)$, and let \mathcal{G}' denote the set of blocks where \tilde{Y} has a fractional entry. By Proposition 16, we have $M\tilde{Y} = My \leq \vec{1}$, and $|\mathcal{G}'| \leq 2t$.

Consider some row k of the constraint matrix. Define the random variable $Z_G = \sum_{j \in G} M_k(j) Y_j$ for each block $G \in \mathcal{G}$, and note that $M_k Y = \sum_{G \in \mathcal{G}} Z_G$. By Corollary 8, with probability of $1 - \delta/m$ there is a subset of blocks $\mathcal{G}'_k \subseteq \mathcal{G}'$ with $|\mathcal{G}'_k| \leq q = O(\sqrt{t \log \frac{m}{\delta}})$, such that $\sum_{G \in \mathcal{G}' - \mathcal{G}'_k} Z_G \leq \sum_{G \in \mathcal{G}'} \mathbf{E}[Z_G]$. The vector Y has one non-zero entry per block, and if we

zero out the entries in the blocks of \tilde{G}'_k , the resulting vector Y' has $M_k Y' \leq \sum_{G \in \mathcal{G} - \mathcal{G}'_k} Z_G \leq \sum_{G \in \mathcal{G}} \mathbf{E}[Z_G] = M_k \tilde{Y} = M_k y$. Thus Y is a q -additive pseudo-solution for M_k with probability at least $1 - \delta/m$. Now take a union bound over all m rows. ■

Proposition 20 *Let $\tilde{Y} = \text{FULLKPR}(\mathcal{G}, M, y, t)$ and let $W \subseteq U$.*

1. *For $t > 12m$, there holds $\mathbf{E}[\prod_{j \in W} Y_j] \leq O(m^2/t) + \prod_{j \in W} y_j$.*
2. *For $t > 10000m|W|$, there holds $\mathbf{E}[\prod_{j \in W} Y_j] \leq e^{O(m^2|W|^2/t)} \prod_{j \in W} y_j$.*

Proof We assume that the elements of W all come from distinct blocks, as otherwise the LHS is zero and this holds immediately. Consider the set $W' = \{j \in U - W : \mathcal{G}(j) \in \mathcal{G}(W)\}$. We have $|\mathcal{G}(W')| \leq |\mathcal{G}(W)| = |W|$ and $\mathbf{E}[\prod_{j \in W} Y_j] = Q(W, Y)$. Finally, note that $\mathbf{E}[Q(W, Y)] = \mathbf{E}[Q(W, \tilde{Y})]$ where \tilde{Y} is the vector at line 1 of FULLKPR. Now apply Theorem 11. ■

4.2 Dependent rounding for knapsack constraints

Given a vector $x \in [0, 1]^U$ and a multi-knapsack constraint M with $Mx \leq \vec{1}$, the problem of *knapsack rounding* is to produce an integral vector $X \in \{0, 1\}^U$ which (as closely as possible) satisfies the knapsack constraint $MX \leq \vec{1}$ and has probabilistic properties related to vector x such as $\mathbf{E}[X] = x$ coordinatewise. Note that this includes cardinality rounding as a special case, with $m = 1$ and $M(i) = 1/r$ for all $i \in U$.

We can interpret knapsack rounding as a special case of KPR. To do so, we extend the set U to a larger ground-set \bar{U} ; for each item $i \in U$, we have a corresponding “dummy” item \bar{i} . We then form a vector $y \in [0, 1]^{\bar{U}}$ by setting $y_i = x_i$ and $y_{\bar{i}} = 1 - x_i$ for each $i \in U$, and we define a block $G_i = \{i, \bar{i}\}$. We lift the knapsack constraints M to \bar{U} by setting $M(\bar{i}) = 0$ for all i .

We can then run KPR on this resulting knapsack-partition instance, and return the fractional vector $\tilde{X} \in [0, 1]^{\bar{U}}$ defined as $\tilde{X}_i = \tilde{Y}_i$ for $i \in U$. We let $\tilde{X} = \text{KNAPROUND}(x, M, t)$ be the result of this process. Note that $M\tilde{X} = M\tilde{Y} = My$, and that \tilde{X} has at most $2t$ fractional entries. We can state a particularly crisp form of our near negative-correlation bounds in this setting:

Proposition 21 *Let $\tilde{X} = \text{KNAPROUND}(x, M, t)$, and let S, T be disjoint subsets of U . Let $d = |S \cup T|$.*

1. *For $t > 12m$ there holds $\mathbf{E}[\prod_{i \in S} \tilde{X}_i \prod_{i \in T} (1 - \tilde{X}_i)] \leq O(m^2/t) + \prod_{i \in S} x_i \prod_{i \in T} (1 - x_i)$*
2. *For $t > 10000md$ there holds $\mathbf{E}[\prod_{i \in S} \tilde{X}_i \prod_{i \in T} (1 - \tilde{X}_i)] \leq e^{O(m^2 d^2/t)} \cdot \prod_{i \in S} x_i \prod_{i \in T} (1 - x_i)$*

Proof Define $W \subseteq \bar{U}$ by $W = \{\bar{i} \mid i \in S\} \cup \{i \mid i \in T\}$. Clearly $|W| = d$. Also, we have $Q(W, y) = \prod_{i \in S} x_i \prod_{i \in T} (1 - x_i)$ and $Q(W, \tilde{Y}) = \prod_{i \in S} \tilde{X}_i \prod_{i \in T} (1 - \tilde{X}_i)$. Now apply Theorem 11. ■

This technique of creating “dummy elements” (in this case, the elements \bar{i}), as indicators for not selecting items, will appear in a number of constructions.

5. Pseudo-approximation algorithm for single-knapsack median

We now describe a $(1 + \sqrt{3} + \gamma)$ -pseudo-approximation algorithm for single-knapsack median. This is inspired by an approximation algorithm of Li & Svensson Li and Svensson (2016) for k -median. The idea is to solve a relaxation called a “bi-point solution”, and then round it to an additive pseudo-solution. This can also be used to obtain a multiplicative pseudo-approximation. The k -median algorithm of Li and Svensson (2016) has an additional postprocessing step to correct it to a true solution; however, this step does not seem to work for knapsack median.

Recall that we define the cost for a facility set $\mathcal{S} \subseteq \mathcal{F}$ by $\text{cost}(\mathcal{S}) = \sum_{j \in \mathcal{C}} d(j, \mathcal{S})$ and the weight by $M(\mathcal{S}) = \sum_{i \in \mathcal{S}} M(i)$. We define OPT to be the minimum value of $\text{cost}(\mathcal{S})$ over feasible sets \mathcal{S} . By a straightforward adaptation of Williamson and Shmoys (2011); Jain et al. (2002) to the knapsack setting, we get the following result:

Theorem 22 *There is a polynomial-time algorithm to compute two sets $\mathcal{F}_1, \mathcal{F}_2 \subseteq \mathcal{F}$ and a parameter $b \in [0, 1]$ satisfying the following properties:*

- $M(\mathcal{F}_1) \leq 1 \leq M(\mathcal{F}_2)$,
- $(1 - b) \cdot M(\mathcal{F}_1) + b \cdot M(\mathcal{F}_2) \leq 1$
- $(1 - b) \cdot \text{cost}(\mathcal{F}_1) + b \cdot \text{cost}(\mathcal{F}_2) \leq 2 \cdot \text{OPT}$.

The sets $\mathcal{F}_1, \mathcal{F}_2$ are called the *bi-point solution*. For $i \in \mathcal{F}_2$ we define $\sigma(i)$ to be the closest facility of \mathcal{F}_1 . Following Li and Svensson (2016), we define $\text{Star}(i)$ for each facility $i \in \mathcal{F}_1$ to be the set of facilities $k \in \mathcal{F}_2$ with $\sigma(k) = i$, that is, $\text{Star}(i) = \sigma^{-1}(i)$. The intent is that for each $i \in \mathcal{F}_1$, with probability $1 - b$ we open i and with the complementary probability b we open *all* the facilities of $\text{Star}(i)$. In order to preserve the knapsack constraints, there are a few exceptional cases where we open both i and some subset of $\text{Star}(i)$.

The full details of our rounding algorithm are spelled out in Algorithm 4. Here, t is an integer parameter to be specified; we assume throughout that $t \geq 20000$. Note that the modified constraint matrix M' can have negative entries, but this does not cause a problem for executing KPR.

Algorithm 4 ROUNDSTARS(t)

- 1: Define vectors $y, M' \in [0, 1]^{\mathcal{F}_1}$ by setting $y_i = 1 - b, M'(i) = M(i) - M(\text{Star}(i))$ for each $i \in \mathcal{F}_1$.
 - 2: $\tilde{Y} \leftarrow \text{KNAPROUND}(y, M', t)$
 - 3: Define the vector $z \in [0, 1]^{\mathcal{F}_2}$, by setting $z_i = 1 - \tilde{Y}_{\sigma(i)}$.
 - 4: $\tilde{Z} \leftarrow \text{KNAPROUND}(z, M, t)$
 - 5: **return** $\mathcal{S} = \{i \in \mathcal{F}_1 \mid \tilde{Y}_i > 0\} \cup \{i \in \mathcal{F}_2 \mid \tilde{Z}_i > 0\}$.
-

Proposition 23 *The solution \mathcal{S} is a $4t$ -additive pseudo-solution with probability one.*

Proof Let \mathcal{S}'_1 denote the facilities $i \in \mathcal{F}_1$ with $\tilde{Y}_i \in (0, 1)$ and let \mathcal{S}'_2 denote the set of facilities $i \in \mathcal{F}_2$ with $\tilde{Z}_i \in (0, 1)$, and let $\mathcal{S}' = \mathcal{S}'_1 \cup \mathcal{S}'_2$. Since \tilde{Y} and \tilde{Z} have at most $2t$ fractional entries, we have $|\mathcal{S}'| \leq 4t$. We now claim $M(\mathcal{S} - \mathcal{S}') \leq 1$, which shows that \mathcal{S} is a $4t$ -additive pseudo-solution. We have:

$$M(\mathcal{S} - \mathcal{S}') = \sum_{i \in \mathcal{F}_1} [[\tilde{Y}_i = 1]]M(i) + \sum_{i \in \mathcal{F}_2} [[\tilde{Z}_i = 1]]M(i) \leq \sum_{i \in \mathcal{F}_1} M(i)\tilde{Y}_i + \sum_{i \in \mathcal{F}_2} M(i)\tilde{Z}_i.$$

By Property (E4) of KPR, this equals

$$\sum_{i \in \mathcal{F}_1} M(i) \tilde{Y}_i + \sum_{i \in \mathcal{F}_2} M(i) z_i = \sum_{i \in \mathcal{F}_1} M(i) \tilde{Y}_i + \sum_{i \in \mathcal{F}_2} M(i) \cdot (1 - \tilde{Y}_{\sigma(i)}).$$

Here, $\sum_{i \in \mathcal{F}_2} M(i) \cdot (1 - \tilde{Y}_{\sigma(i)})$ contributes $M(\text{Star}(i))(1 - \tilde{Y}_i)$ for each $i \in \mathcal{F}_1$. Thus, the sum is at most

$$\sum_{i \in \mathcal{F}_1} M(i) \tilde{Y}_i + (1 - \tilde{Y}_i) M(\text{Star}(i)) = \sum_{i \in \mathcal{F}_1} M(\text{Star}(i)) + M'(i) \tilde{Y}_i$$

By Property (E4) of (KPR), this equals $\sum_{i \in \mathcal{F}_1} M(\text{Star}(i)) + M'(i) y_i$, which can be simplified as

$$\sum_{i \in \mathcal{F}_1} M(\text{Star}(i)) + (1 - b)(M(i) - M(\text{Star}(i))) = (1 - b)M(\mathcal{F}_1) + bM(\mathcal{F}_2)$$

which is at most 1 by the properties of the bi-point solution. ■

Proposition 24 For any facilities $i_1 \in \mathcal{F}_1, i_2 \in \mathcal{F}_2$, Algorithm 4 yields

$$\Pr(i_2 \notin \mathcal{S}) \leq 1 - b, \quad \Pr(i_1 \notin \mathcal{S} \wedge i_2 \notin \mathcal{S}) \leq b(1 - b)(1 + O(1/t))$$

Proof Let $k = \sigma(i_2)$. Conditioned on the vector \tilde{Y} we have

$$\Pr(i_2 \notin \mathcal{S} \mid \tilde{Y}) = \Pr(\tilde{Z}_{i_2} = 0 \mid z) \leq \mathbf{E}[1 - \tilde{Z}_{i_2} \mid z] = 1 - z_{i_2} = \tilde{Y}_k \quad (11)$$

Integrating over z shows that $\Pr(i_2 \notin \mathcal{S}) \leq \mathbf{E}[\tilde{Y}_k] = y_k = 1 - b$, as claimed.

For the second result, note that in order to have $i_1 \notin \mathcal{S}$ and $i_2 \notin \mathcal{S}$, we must have $\tilde{Y}_{i_1} = \tilde{Z}_{i_2} = 0$ and $i_1 \neq k$. So by Eq. (11) we have

$$\Pr(i_1 \notin \mathcal{S} \wedge i_2 \notin \mathcal{S} \mid \tilde{Y}) = [[\tilde{Y}_{i_1} = 0]] \Pr(i_2 \notin \mathcal{S} \mid \tilde{Y}) \leq [[\tilde{Y}_{i_1} = 0]] \tilde{Y}_k \leq (1 - \tilde{Y}_{i_1}) \tilde{Y}_k.$$

Since $i_1 \neq k$, and we are assuming $t \geq 20000$, Proposition 21 gives:

$$\mathbf{E}[\tilde{Y}_{i_1} (1 - \tilde{Y}_k)] \leq y_{i_1} (1 - y_k) e^{O(m^2(1+1)^2/t)} = b(1 - b)(1 + O(1/t))$$
■

Proposition 25 Each client $j \in \mathcal{C}$ has expected cost

$$\mathbf{E}[d(j, \mathcal{S})] \leq (1 - b)d(j, \mathcal{F}_1) + bd(j, \mathcal{F}_2) + 2d(j, \mathcal{F}_2)b(1 - b)(1 + O(1/t))$$

Proof Let i_1, i_2 denote the closest facilities in $\mathcal{F}_1, \mathcal{F}_2$ to j , and let $k = \sigma(i_2)$. We also define $d_1 = d(j, \mathcal{F}_1) = d(j, i_1)$ and $d_2 = d(j, \mathcal{F}_2) = d(j, i_2)$.

If $i_2 \in \mathcal{S}$ then $d(j, \mathcal{S}) \leq d_2$ and likewise if $i_1 \in \mathcal{S}$ then $d(j, \mathcal{S}) \leq d_1$. If neither holds, then necessarily facility k is open. Since k is the closest facility in \mathcal{F}_1 , we must have $d(i_2, k) \leq d(i_2, i_1)$. We then have

$$d(j, \mathcal{S}) \leq d(j, k) \leq d(j, i_2) + d(i_2, k) \leq d(j, i_2) + d(i_2, i_1) \leq d(j, i_2) + d(i_2, j) + d(j, i_1) = 2d_2 + d_1$$

Putting these together gives

$$\begin{aligned} \mathbf{E}[d(j, \mathcal{S})] &\leq \Pr(i_2 \in \mathcal{S})d_2 + \Pr(i_1 \in \mathcal{S} \wedge i_2 \notin \mathcal{S})d_1 + \Pr(i_1 \notin \mathcal{S} \wedge i_2 \notin \mathcal{S})(2d_2 + d_1) \\ &= d_2 + (d_1 - d_2)\Pr(i_2 \notin \mathcal{S}) + 2\Pr(i_1 \notin \mathcal{S} \wedge i_2 \notin \mathcal{S})d_2 \end{aligned}$$

Proposition 24 gives $\Pr(i_2 \notin \mathcal{S}) \leq 1 - b$ and $\Pr(i_1 \notin \mathcal{S} \wedge i_2 \notin \mathcal{S}) \leq b(1 - b)(1 + O(1/t))$. ■

We are now ready to obtain our bi-factor approximation algorithm.

Theorem 26 *There is an algorithm with $\text{poly}(n/\gamma)$ runtime to obtain an $O(1/\gamma)$ -additive pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (1 + \sqrt{3} + \gamma) \cdot \text{OPT}$.*

Proof We will use Algorithm 4, and output either the solution \mathcal{S} it returns or the feasible solution \mathcal{F}_1 (whichever has least cost). By Proposition 23, the solution \mathcal{S} is a $4t$ -additive pseudo-solution.

Define $D_1 = \sum_j d(j, \mathcal{F}_1)$ and $D_2 = \sum_j d(j, \mathcal{F}_2)$. Applying Proposition 25 and summing over all clients $j \in \mathcal{C}$, we see that \mathcal{S} satisfies

$$\mathbf{E}[\text{cost}(\mathcal{S})] \leq (1 - b)D_1 + bD_2 + 2b(1 - b)D_2 \cdot (1 + O(1/t))$$

Since the cost of \mathcal{A} is the minimum of the cost of \mathcal{F}_1 and \mathcal{S} , this implies that

$$\mathbf{E}[\text{cost}(\mathcal{A})] \leq (1 + O(1/t)) \cdot \min\{D_1, (1 - b)D_1 + bD_2 + 2b(1 - b)D_2\}$$

It can be routinely verified that¹

$$\min\{D_1, (1 - b)D_1 + bD_2 + 2b(1 - b)D_2\} \leq \frac{1 + \sqrt{3}}{2}((1 - b)D_1 + bD_2)$$

Since $\mathcal{F}_1, \mathcal{F}_2$ is a bi-point solution, we have $(1 - b)D_1 + bD_2 \leq 2 \cdot \text{OPT}$. Therefore,

$$\mathbf{E}[\text{cost}(\mathcal{A})] \leq (1 + O(1/t)) \cdot 2 \cdot \text{OPT} \cdot \frac{1 + \sqrt{3}}{2} = (1 + O(1/t)) \cdot (1 + \sqrt{3}) \cdot \text{OPT}$$

If we set $t = \Omega(1/\gamma)$, then after an expected $O(1/\gamma)$ repetitions of this process, we obtain a solution of cost at most $(1 + O(\gamma)) \cdot \text{OPT}$. ■

We can leverage this to obtain a multiplicative pseudo-approximation.

Theorem 27 *There is an algorithm with $n^{O(\epsilon^{-1}\gamma^{-1})}$ runtime to obtain an ϵ -multiplicative pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (1 + \sqrt{3} + \gamma) \cdot \text{OPT}$.*

Proof Call a facility i *big* if $M(i) \geq \rho = \Theta(\epsilon\gamma)$. The solution may have at most $1/\rho$ big facilities, which we can guess in $n^{O(1/\rho)}$ time. Now construct a residual instance where all other big facilities, aside from the ones guessed to be in our solution, are removed. Apply Theorem 26 to this residual instance; after rescaling, the resulting solution \mathcal{S} satisfies $\text{cost}(\mathcal{S}) \leq (1 + \sqrt{3} + \gamma) \cdot \text{OPT}$ and $M(\mathcal{S}) \leq 1 + O(\rho/\gamma)$, since each facility now has weight at most ρ . Our choice of ρ ensures that $M(\mathcal{S}) \leq 1 + \epsilon$ and gives a runtime of $n^{O(1/\rho)} \cdot (n/\gamma)^{O(1)} n^{O(\epsilon^{-1}\gamma^{-1})}$. ■

1. For example, this inequality can be encoded in the first-order theory of real-closed fields, in terms of indeterminates b, D_1, D_2 . This theory is decidable, so it can be checked that it holds for all values $b \in [0, 1], D_1 \geq 0, D_2 \geq 0$.

6. Pseudo-approximation algorithm for multi-knapsack median

In this section, we give a 3.25-pseudo-approximation algorithm for multi-knapsack median. This is based on applying KPR for a key rounding step in the 3.25-approximation algorithm of Charikar & Li Charikar and Li (2012) for k -median. Although this is not the best approximation ratio for k -median, the main benefits of that algorithm is its good approximation ratio as a function of the “obvious” LP relaxation defined as follows:

$$\begin{aligned} & \text{minimize } \sum_{j \in \mathcal{C}} \sum_{i \in \mathcal{F}} x_{i,j} d(i, j) \\ & \text{subject to } \sum_{i \in \mathcal{F}} x_{i,j} = 1 \quad \forall j \in \mathcal{C} \\ & \quad 0 \leq x_{i,j} \leq y_i \leq 1 \quad \forall i \in \mathcal{C}, j \in \mathcal{F} \\ & \quad My \leq \vec{1} \end{aligned}$$

Here, $x_{i,j}$ represents fractionally how client j is matched to facility i , and y_i is an indicator that facility i is open. The final constraint specifies that each of the m knapsack constraints is satisfied. For any client $j \in \mathcal{C}$, let $r_j = \sum_i x_{i,j} d(i, j)$ denote the fractional connection cost of j . By standard facility-splitting methods (see e.g., Swamy (2004)), we can ensure that $y_i > 0$ and $x_{i,j} \in \{0, y_i\}$ for all i, j .

6.1 The Charikar-Li algorithm

Given an LP solution, the Charikar-Li algorithm has two phases, which we briefly summarize here. Please see Charikar and Li (2012) for further details. The *bundling* phase can be divided into four stages:

1. A client set $\mathcal{C}' \subseteq \mathcal{C}$ is chosen, such that the clients $j, j' \in \mathcal{C}'$ are relatively far apart from each other. For $j \in \mathcal{C}$ let $\sigma(j)$ be the closest client in \mathcal{C}' .
2. For each $j \in \mathcal{C}'$, we define a set $\mathcal{U}_j \subseteq \mathcal{F}$ which are the facilities “claimed by” j . The sets \mathcal{U}_j are called “bundles”; they are disjoint and have $1/2 \leq y(\mathcal{U}_j) \leq 1$. We define $R_j = \frac{1}{2} d(j, \mathcal{C}' - j)$ for $j \in \mathcal{C}'$.
3. The clients in \mathcal{C}' are paired up, giving a partition of \mathcal{C}' into cardinality-two sets.² We refer to this as the *matching* \mathcal{M} of \mathcal{C}' .
4. We also define $\mathcal{U}_0 = \mathcal{F} - \bigcup_{j \in \mathcal{C}'} \mathcal{U}_j$; these are the “unbundled” facilities.

In the *selection* phase, each pair $\{j, j'\} \in \mathcal{M}'$ selects either one or two facilities to open from $\mathcal{U}_j \cup \mathcal{U}_{j'}$, and in addition some unbundled facilities are opened.

The simplest strategy for this, which is a baseline for more advanced algorithms, is *independent selection*. Here, we first choose a set of “open clients” in \mathcal{C}' , wherein for each pair $\{j, j'\} \in \mathcal{M}$, we open j or j' with the following probabilities: (i) open j alone with probability $1 - y(\mathcal{U}_{j'})$, (ii) open

2. If $|\mathcal{C}'|$ is odd, then \mathcal{C}' has one remaining unmatched client. The Charikar & Li algorithm has some additional steps to handle this case. We can avoid these exceptional steps by adding an additional dummy client and dummy zero-cost facility, with distance zero to each other and distance ∞ to all other clients. This allows us to assume without loss of generality that $|\mathcal{C}'|$ is even.

j' alone with probability $1 - y(\mathcal{U}_j)$, or (iii) open both j, j' with probability $y(\mathcal{U}_j) + y(\mathcal{U}_{j'}) - 1$. Then, for each open client j , we open exactly one facility i in \mathcal{U}_j , wherein i is chosen with probability proportional to y_i . Also, for each facility $i \in \mathcal{U}_0$, we open i independently with probability y_i .

Charikar & Li shows a number of powerful bounds and properties for this process.

Proposition 28 (Charikar and Li (2012)) *Under independent selection, the following bounds hold:*

(B1) $\Pr(i \in \mathcal{S}) \leq y_i$ for any facility i .

(B2) The events $[[i \in \mathcal{S}]]$, for $i \in \mathcal{F}$, are cylindrically negatively correlated.

(B3) For any client $j \in \mathcal{C}$ we have $\mathbf{E}[d(j, \mathcal{S})] \leq 3.25r_j$.

(B4) For any client $j \in \mathcal{C}$ we have $r_{\sigma(j)} \leq r_j$ and $d(j, \sigma(j)) \leq 4r_j$

(B5) For any client $j \in \mathcal{C}'$, it holds that $d(j, \mathcal{S}) \leq \beta R_j$ with probability one, where β is some constant.

(B6) For any client $j \in \mathcal{C}'$, there holds $r_j \leq R_j$ and $\mathcal{B}(j, R_j) \subseteq \mathcal{U}_j$

This leads to a simple multiplicative pseudo-approximation algorithm for multi-knapsack median.

Theorem 29 *Let $\gamma, \epsilon \in (0, 1)$. There is an algorithm with $n^{O(m \log(m/\gamma)/\epsilon^2)}$ runtime to obtain an ϵ -multiplicative pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (3.25 + \gamma) \cdot \text{OPT}$.*

Proof [Proof (Sketch)] Say a facility $i \in \mathcal{F}$ is big if $M_k(i) \geq \rho = \frac{\epsilon^2}{10 \log(m/\gamma)}$ for any $k \in [m]$. We can guess the big facilities in an optimal solution in $n^{O(m/\rho)}$ time. Next, solve the LP and run the Charikar-Li rounding on the resulting residual instance. By properties (B1), (B2), each $\sum_i M_k(i) [[i \in \mathcal{S}]]$ is a sum of negatively-correlated random variables with mean $\sum_i M_k(i) y_i \leq 1$, and each is bounded in the range $[0, \rho]$. By Chernoff's bound, the probability that it exceeds $1 + \epsilon$ is at most $e^{-\epsilon^2/(3\rho)} = O(\gamma/m)$. ■

6.2 Facility selection as a knapsack-partition system

In the selection phase, each pair $e \in \mathcal{M}$ needs to select one or two facilities, according to a certain probability distribution. We also need to decide the status of the unbundled facilities. We can encode these requirements in terms of a knapsack-partition system. Here, the ground-set U will be a polynomial-size subset of the power set $2^{\overline{\mathcal{F}}}$, where $\overline{\mathcal{F}}$ is \mathcal{F} plus some dummy items. That is, each element of U is itself a set containing (possibly dummy) facilities.

For this construction, for each pair $e = \{j, j'\} \in \mathcal{M}$, we create a block $G_e \subseteq 2^{\overline{\mathcal{F}}}$ defined as

$$G_e = \left\{ \{i\} \mid i \in \mathcal{U}_j \cup \mathcal{U}_{j'} \right\} \cup \left\{ \{i, i'\}, \mid i \in \mathcal{U}_j, i' \in \mathcal{U}_{j'} \right\}$$

For each unbundled facility $i \in \mathcal{U}_0$, we create a block $G_i \subseteq 2^{\overline{\mathcal{F}}}$ defined as

$$G_i = \left\{ \{i\}, \{\bar{i}\} \right\}$$

where \bar{i} is a dummy item (which indicates that facility i is not to be opened).

The ground set is $U = \bigcup_{e \in \mathcal{M}} G_e \cup \bigcup_{i \in \mathcal{U}_0} G_i$. Since the sets \mathcal{U}_j are disjoint, the blocks G_e and G_i form a partition of U , which we denote by \mathcal{G} . Correspondingly, we also define a vector $z \in [0, 1]^U$ as follows. For each $e = \{j, j'\} \in \mathcal{M}$, and every pair of facilities $i \in \mathcal{U}_j, i' \in \mathcal{U}_{j'}$, we set:

$$\begin{aligned} z_{\{i\}} &= (1 - y(\mathcal{U}_{j'})) \frac{y_i}{y(\mathcal{U}_j)}, & z_{\{i'\}} &= (1 - y(\mathcal{U}_j)) \frac{y_{i'}}{y(\mathcal{U}_{j'})} \\ z_{\{i, i'\}} &= (y(\mathcal{U}_j) + y(\mathcal{U}_{j'}) - 1) \frac{y_i y_{i'}}{y(\mathcal{U}_j) y(\mathcal{U}_{j'})} \end{aligned}$$

Likewise, for each $i \in \mathcal{U}_0$, we define

$$z_{\{i\}} = y_i, \quad z_{\{\bar{i}\}} = 1 - y_i$$

Finally, we extend the knapsack constraints to U , by setting $M_k(\bar{i}) = 0$ for each dummy item i and setting $M_k(V) = \sum_{i \in V} M_k(i)$ for $V \subseteq 2^{\bar{\mathcal{F}}}$. It can be seen that that z is a fractional solution to the partition system; also, by the way we have extended M to U , we have $Mz = My$.

Given any facility set $W \subseteq \mathcal{F}$, we define a corresponding set $W^* \subseteq U$ by $W^* = \{V \in U \mid V \cap W \neq \emptyset\}$. The following observation summarizes how the knapsack partition system is connected to the Charikar-Li selection phase and independent selection.

Observation 30 *Given an integral vector $Z \in \{0, 1\}^U$, we can generate a corresponding facility set \mathcal{S} by opening all (non-dummy) facilities in all sets V with $Z_V = 1$, i.e., $\mathcal{S} = \bigcup_{V: Z_V=1} V \cap \mathcal{F}$. In this case, we have the following:*

- *There holds $M(\mathcal{S}) \leq MZ$.*
- *For any set of facilities $W \subseteq \mathcal{F}$, there holds $\mathcal{S} \cap W \neq \emptyset$ if and only $Q(W^*, Z) = 0$.*
- *If Z is generated as $Z = \text{INDSELECT}(\mathcal{G}, z)$, then the resulting solution set \mathcal{S} has the same probability distribution as in independent selection for the Charikar-Li algorithm.*

We omit the proofs since they follow immediately from definitions. As one important consequence of Observation 30, we can write the expected distance for a client j in terms of the potential function Q for the resulting knapsack-partition system:

Proposition 31 *If \mathcal{S} is generated as in Observation 30, then any client j has*

$$d(j, \mathcal{S}) = \int_{u=0}^{4r_j + \beta R_{\sigma(j)}} Q(\mathcal{B}(j, u)^*, Z) du.$$

Proof Let $s = 4r_j + \beta R_{\sigma(j)}$. By properties (B4) and (B5), we have $d(j, \mathcal{S}) \leq d(j, \sigma(j)) + d(\sigma(j), \mathcal{S}) \leq s$ with probability one. For any $u \geq 0$, we have $\mathcal{S} \cap \mathcal{B}(j, u) = \emptyset$ if and only if $Q(\mathcal{B}(j, u)^*, Z) = 1$, and so

$$d(j, \mathcal{S}) = \int_{u=0}^s [[d(j, \mathcal{S}) > u]] du = \int_{u=0}^s [[\mathcal{S} \cap \mathcal{B}(j, u) = \emptyset]] du = \int_{u=0}^s Q(\mathcal{B}(j, u)^*, Z) du$$

■

6.3 KPR selection strategy

Our approximation algorithm, summarized as Algorithm 5, uses FULLKPR rounding instead of independent selection in the Charikar-Li algorithm.

Algorithm 5 MULTIKNAPSACKMEDIANROUND(t)

- 1: Let x, y be the solution to the LP.
 - 2: Run the Charikar-Li bundling phase, resulting in fractional vector $z \in [0, 1]^U$ and partition \mathcal{G} .
 - 3: Let $Z = \text{FULLKPR}(\mathcal{G}, M, z, t)$.
 - 4: **return** $\mathcal{S} = \bigcup_{V: Z_V=1} V \cap \mathcal{F}$.
-

Our main rounding result is that Algorithm 5 has a similar probability of opening a facility in any given set W compared to independent selection.

Lemma 32 *Let $t \geq 20000m^2$. Then for any client $j \in \mathcal{C}'$ and any set $W \subseteq \mathcal{F}$, Algorithm 5 satisfies*

$$\mathbf{E}[Q(W^*, Z)] \leq Q(W^*, z) + O(m^2/t) \cdot \left(\frac{r_j}{R_j} + \sum_{i \in \mathcal{U}_j - W} x_{i,j} \right)$$

Proof Let $e = \{j, j'\} \in \mathcal{M}$ be the pair in the matching corresponding to j . We apply Theorem 11(a) with respect to $\mathcal{D} = \{G_e\}$. Since $|\mathcal{D}| = 1$, this gives

$$\mathbf{E}[Q(W^*, Z)] \leq Q(W^*, z) + O(m^2/t) \cdot Q(W^* \wedge \mathcal{D}, z)$$

To finish the proof, we need to show that

$$Q(W^* \wedge \mathcal{D}, z) \leq r_j/R_j + \sum_{i \in \mathcal{U}_j - W} x_{i,j}. \quad (12)$$

For this, we calculate:

$$\begin{aligned} Q(W^* \wedge \mathcal{D}, z) &= 1 - \sum_{\substack{V \cap W \neq \emptyset \\ V \in G_e}} z_V \leq 1 - \sum_{i \in W \cap \mathcal{U}_j} z_{\{i\}} - \sum_{i \in W \cap \mathcal{U}_j, i' \in \mathcal{U}_{j'}} z_{\{i, i'\}} \\ &= 1 - \sum_{i \in \mathcal{U}_j \cap W} \frac{(1 - y(\mathcal{U}_{j'}))y_i}{y(\mathcal{U}_j)} - \sum_{i \in W \cap \mathcal{U}_j, i' \in \mathcal{U}_{j'}} \frac{(y(\mathcal{U}_j) + y(\mathcal{U}_{j'}) - 1)y_i y_{i'}}{y(\mathcal{U}_j)y(\mathcal{U}_{j'})} \\ &= 1 - y(\mathcal{U}_j \cap W) \end{aligned}$$

The facility-splitting step ensures that $y_i = x_{i,j}$ for any $i \in \mathcal{U}_j$. So $1 - y(W \cap \mathcal{U}_j) = 1 - \sum_{i \in W \cap \mathcal{U}_j} x_{i,j} = 1 - \sum_{i \in \mathcal{U}_j} x_{i,j} + \sum_{i \in \mathcal{U}_j - W} x_{i,j}$. Finally, to bound the term $1 - \sum_{i \in \mathcal{U}_j} x_{i,j}$, we use property (B6):

$$1 - \sum_{i \in \mathcal{U}_j} x_{i,j} \leq 1 - \sum_{i \in \mathcal{B}(j, R_j)} x_{i,j} = \sum_{i: d(i,j) > R_j} x_{i,j} \leq \sum_i \frac{d(i,j)}{R_j} x_{i,j} = \frac{r_j}{R_j}$$

This shows Eq. (12) and hence shows the claimed result. ■

This, in turn, allows us to bound expected distances for the rounding algorithm.

Proposition 33 *If $t \geq 20000m^2$, then for any client j the set \mathcal{S} returned by Algorithm 5 satisfies $\mathbf{E}[d(j, \mathcal{S})] \leq D_j + O(r_j m^2/t)$, where D_j is the expected distance of j under independent selection.*

Proof For $u \geq 0$, define $C_u = \mathcal{B}(j, u)^*$, and let $j' = \sigma(j)$. By Proposition 31, the expected distance of j is given by:

$$\mathbf{E}[d(j, \mathcal{S})] = \int_{u=0}^s \mathbf{E}[Q(C_u, Z)] du$$

where $s = 4r_j + \beta R_{j'}$. Using Lemma 32 with respect to client j' and facility set $W = \mathcal{B}(j, u)$ gives:

$$\begin{aligned} \int_{u=0}^s \mathbf{E}[Q(C_u, Z)] du &\leq \int_{u=0}^s \left(Q(C_u, z) + O(m^2/t) \left(\frac{r_{j'}}{R_{j'}} + \sum_{i \in \mathcal{U}_{j'} - \mathcal{B}(j, u)} x_{i, j'} \right) \right) du \\ &= \int_{u=0}^s Q(C_u, z) du + O(m^2/t) \left(\frac{sr_{j'}}{R_{j'}} + \int_{u=0}^s \sum_{i \in \mathcal{U}_{j'} - \mathcal{B}(j, u)} x_{i, j'} du \right) \end{aligned}$$

Now note that, under independent selection, we have $\mathbf{E}[Q(C_u, Z)] = Q(C_u, z)$ for each u . Thus, by Proposition 31, the first term here is precisely D_j . To estimate the last term, we interchange summation and use the triangle inequality to get:

$$\begin{aligned} \int_{u=0}^s \sum_{i \in \mathcal{U}_{j'} - \mathcal{B}(j, u)} x_{i, j'} du &= \sum_{i \in \mathcal{U}_{j'}} x_{i, j'} \int_{u=0}^s [[d(i, j) > u]] du = \sum_{i \in \mathcal{U}_{j'}} x_{i, j'} \min(s, d(i, j)) \\ &\leq \sum_{i \in \mathcal{U}_{j'}} x_{i, j'} (d(j, j') + d(i, j')) = d(j, j') \sum_{i \in \mathcal{U}_{j'}} x_{i, j'} + \sum_{i \in \mathcal{U}_{j'}} d(i, j') x_{i, j'} \leq d(j, j') + r_{j'} \end{aligned}$$

Therefore

$$\mathbf{E}[d(j, \mathcal{S})] \leq D_j + O(m^2/t) \left(\frac{sr_{j'}}{R_{j'}} + d(j, j') + r_{j'} \right)$$

Now, since $r_{j'} \leq r_j$ and $r_{j'} \leq R_{j'}$, we have $sr_{j'}/R_{j'} = 4r_j(r_{j'}/R_{j'}) + r_{j'} \cdot \beta R_{j'}/R_{j'} \leq O(r_j)$. Also, by (B4), we have $d(j, j') + r_{j'} \leq O(r_j)$. Overall, we see that

$$\frac{sr_{j'}}{R_{j'}} + d(j, j') + r_{j'} \leq O(r_j),$$

which concludes the proof. ■

From property (B3), this immediately shows the following:

Corollary 34 *If $t \geq 20000m^2$, then Algorithm 5 ensures $\mathbf{E}[d(j, \mathcal{S})] \leq (3.25 + O(m^2/t))r_j$ for every client $j \in \mathcal{C}$, and furthermore $\mathbf{E}[\text{cost}(\mathcal{S})] \leq (3.25 + O(m^2/t)) \cdot \text{OPT}$.*

We now show Theorem 3, restated here for convenience.

Theorem 3 *Consider a multi-knapsack median instance with m constraints, and let $\epsilon, \gamma \in (0, 1)$. There is an algorithm with $\text{poly}(n/\gamma)$ runtime to obtain an $O(\frac{m}{\sqrt{\gamma}})$ -additive pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (3.25 + \gamma) \cdot \text{OPT}$, and an algorithm with $n^{\tilde{O}(m^2 \epsilon^{-1} \gamma^{-1/2})}$ runtime to obtain an ϵ -multiplicative pseudo-solution \mathcal{S} with $\text{cost}(\mathcal{S}) \leq (3.25 + \gamma) \cdot \text{OPT}$.*

Proof Let us define $V = \text{cost}(\mathcal{S})/\text{OPT}$; Proposition 28 shows that $\mathbf{E}[V] \leq (3.25 + O(m^2/t))$ for $t \geq 20000m^2$. We will choose $t = m^2/\gamma$ (which satisfies $t \geq 20000m^2$ for γ sufficiently small).

Let \mathcal{E} be the event that Z is an q -additive pseudo-solution for the given value q . By Theorem 19, we have $\Pr(\mathcal{E}) \geq 1 - \gamma$. Since each element in U contains at most two facilities, the solution \mathcal{S} is a $2q$ -additive pseudo-solution to the original knapsack whenever \mathcal{E} holds. Also, we have $\mathbf{E}[V \mid \mathcal{E}] \leq 3.25 + O(\gamma)$. So, after an expected $O(1/\gamma)$ repetitions we find a set \mathcal{S} such that \mathcal{E} holds and such that $V \leq 3.25 + O(\gamma)$. Overall, we get a runtime of $O(1/\gamma) \cdot n^{O(1)}$; the result follows by rescaling γ and simplifying.

For the second result, say a facility i is big if $M_k(i) > \rho = \Theta\left(\frac{\epsilon\sqrt{\gamma}}{m\sqrt{\log(m/\gamma)}}\right)$ for any $k \in [m]$.

We can guess the big facilities in an optimal solution in $n^{O(m/\rho)}$ time. We then construct a residual instance where all other big facilities are discarded, and apply the additive pseudo-approximation to it. The overall runtime is $(n/\gamma)^{O(1)} \cdot n^{O(m/\rho)}$; with simplification of parameters, this gives the claimed runtime. \blacksquare

7. The knapsack center problem

We now analyze the knapsack center problems, proving Theorems 4 and 5. There are only $\binom{n}{2}$ possible values for the optimal radius $R = \text{OPT}$, and so we can guess this value in $O(n^2)$ time. To simplify the notation for this section, let us suppose that we have guessed R and rescaled to have $R = 1$.

Given some arbitrary $\gamma > 0$, our goal is to find a distribution Ω over solution sets \mathcal{S} , such that every client $j \in \mathcal{C}$ has

$$\mathbf{E}_{\mathcal{S} \sim \Omega}[d(j, \mathcal{S})] \leq 1 + 2/e + \gamma, \quad d(j, \mathcal{S}) \leq 3 \text{ with probability one} \quad (13)$$

We refer to the distribution Ω as a γ -fair solution. Let us define a *weighting function* a to be a map $a : \mathcal{C} \rightarrow [0, 1]$ with $\sum_{j \in \mathcal{C}} a_j = 1$. In such a distribution Ω , every weighting function a would have a corresponding solution set \mathcal{S} satisfying

$$\sum_{j \in \mathcal{C}} a_j d(j, \mathcal{S}) \leq 1 + 2/e + \gamma, \quad \max_{j \in \mathcal{C}} d(j, \mathcal{S}) \leq 3 \quad (14)$$

By LP duality, the converse also holds. Furthermore, the multiplicative weights update (MWU) method makes this efficient: if we have an efficient algorithm \mathcal{A} which takes as input a weighting function a and returns a solution set \mathcal{S} satisfying Eq. (14), then it can be converted into an efficient randomized algorithm which returns a $O(\gamma)$ -fair solution. We summarize this as the following Algorithm 6:

Algorithm 6 KNAPSACKCENTERMWU

- 1: Initialize the vector $a_j^{(1)} = 1$ for all clients $j \in \mathcal{C}$.
 - 2: **for** $k = 1, \dots, v$ **do**
 - 3: Use algorithm \mathcal{A} with the weighting function $\frac{a^{(k)}}{a^{(k)}(\mathcal{C})}$ to get solution set \mathcal{S}_k .
 - 4: **for each** $j \in \mathcal{C}$ **do** update $a_j^{(k+1)} = e^{\epsilon d(j, \mathcal{S}_k)} a_j^{(k)}$
 - 5: Return \mathcal{S}_K for K uniformly chosen from $[v]$.
-

Lemma 36 *Suppose that, for every weighting function a , the algorithm \mathcal{A} runs in time T and generates a solution \mathcal{S} satisfying Eq. (14) from some family \mathfrak{S} . Then with appropriate parameters v, ϵ , Algorithm 6 runs in $\text{poly}(n/\gamma) \cdot T$ time and outputs an $O(\gamma)$ -fair solution \mathcal{S} in \mathfrak{S} .*

Proof Let $\beta = 1 + 2/e + \gamma$. The output \mathcal{S}_K is clearly in \mathfrak{S} , since $\mathcal{S}_1, \dots, \mathcal{S}_k$ are all in \mathfrak{S} . Let us define $\Phi_k = \sum_{j \in \mathcal{C}} a_j^{(k)}$. Since $d(j, \mathcal{S}_k) \leq 3$, for sufficiently small ϵ we have

$$\Phi_{k+1} = \sum_{j \in \mathcal{C}} a_j^{(k)} e^{\epsilon d(j, \mathcal{S}_k)} \leq \sum_{j \in \mathcal{C}} a_j^{(k)} (1 + \epsilon d(j, \mathcal{S}_k) + \epsilon^2 d(j, \mathcal{S}_k)^2) \leq a^{(k)}(\mathcal{C}) + (\epsilon + 3\epsilon^2)\beta \sum_j a_j^{(k)} d(j, \mathcal{S})$$

The algorithm \mathcal{A} ensures that $\sum_j \frac{a_j^{(k)}}{a^{(k)}(\mathcal{C})} d(j, \mathcal{S}) \leq \beta$ and therefore

$$\Phi_{k+1} \leq a^{(k)}(\mathcal{C}) + (\epsilon + 3\epsilon^2)\beta a^{(k)}(\mathcal{C}) = (1 + (\epsilon + 3\epsilon^2)\beta)\Phi_k$$

Since $\Phi_1 \leq n$, this implies that $\Phi_{v+1} \leq (1 + (\epsilon + 3\epsilon^2)\beta)^v n \leq e^{v(\epsilon + 3\epsilon^2)\beta} n$.

Now consider some client j at the end of this process. We have $e^{\epsilon \sum_{k=1}^v d(j, \mathcal{S}_k)} = a_j^{(v+1)} \leq \Phi_{v+1} \leq e^{v(\epsilon + 3\epsilon^2)(1 + 2/e + \gamma)} n$. Taking logarithms and simplifying, this shows $\sum_{k=1}^v d(j, \mathcal{S}_k) \leq v\beta(1 + 3\epsilon) + \log n$. Therefore, with $\epsilon = \gamma$ and $v = \frac{\log n}{\gamma^2}$, the output \mathcal{S}_K has $\mathbf{E}[d(j, \mathcal{S}_K)] = \frac{1}{v} \sum_{j=1}^k d(j, \mathcal{S}_k) \leq \beta + O(\gamma)$. \blacksquare

7.1 The knapsack center LP

We use the following LP consisting of points (x, y) satisfying constraints (C1) – (C4):

- (C1) $\sum_{i \in \mathcal{B}(j, 1)} x_{ij} = 1$ for all $j \in \mathcal{C}$ (all clients should get connected to some open facility),
- (C2) $x_{ij} \leq y_i$ for all $i, j \in \mathcal{C}$ (client j can only connect to facility i if it is open),
- (C3) $My \leq \vec{1}$ (the m knapsack constraints),
- (C4) $0 \leq x_{ij}, y_i \leq 1$ for all $i, j \in \mathcal{C}$.

By splitting facilities we may enforce an additional constraint:

- (C5) For all $i \in \mathcal{F}, j \in \mathcal{C}$, we have $x_{ij} \in \{0, y_i\}$,

We say a facility i is *integral* if $y_i \in \{0, 1\}$; else it is *fractional*. For $j \in \mathcal{C}$ we define $F_j := \{i \in \mathcal{F} : x_{ij} > 0\}$ and similarly for $i \in \mathcal{F}$ we define $H_i = \{j \in \mathcal{C} : x_{ij} > 0\}$. We may form a subset $\mathcal{C}' \subseteq \mathcal{C}$, such that the sets $F_j, F_{j'}$ for $j, j' \in \mathcal{C}'$ are pairwise disjoint, and such that \mathcal{C}' is maximal with this property. We also define $F_0 = \mathcal{F} - \bigcup_{j \in \mathcal{C}'} F_j$.

We turn this into a knapsack-partition instance as follows. For each $j \in \mathcal{C}'$, we define a block G_j to be simply F_j . For each $i \in F_0$, we create a dummy item \bar{i} with $y_{\bar{i}} = 1 - y_i$ and $M(\bar{i}) = 0$, and we create a block $G_i = \{i, \bar{i}\}$. (If we select \bar{i} , it simply means that we do not choose to open facility i .) By Property (C5), we have $y(F_j) = 1$ for all j , and so y satisfies the partition constraints.

The following result shows how the knapsack center instance relates to this knapsack-partition system, in particular to the potential function Q .

Proposition 37 *Let $j \in \mathcal{C}$ be any client. If (x, y) is a fractional LP solution, then $Q(F_j, y) \leq 1/e$. If Y is an integral solution to the partition system, and we open the (non-dummy) facilities in the support of Y , then $d(j, \mathcal{S}) \leq 1 + 2Q(F_j, Y)$.*

Proof For the first claim, since $y(F_j) = 1$ and the blocks G are pairwise disjoint, we have

$$Q(F_j, y) = \prod_{G \in \mathcal{G}} (1 - y(G \cap F_j)) \leq \prod_{G \in \mathcal{G}} e^{-y(G \cap F_j)} = e^{-y(F_j)} = 1/e$$

For the second claim, if some $i \in F_j$ is opened (i.e., if $Q(F_j, Y) = 0$), then $d(j, \mathcal{S}) \leq 1$. Also, by maximality of \mathcal{C}' , there must exist some facility $k \in \mathcal{C}'$ with $F_j \cap F_k \neq \emptyset$ (possibly $k = j$). There will be some facility opened in F_k , and so $d(j, \mathcal{S}) \leq d(j, i) + d(i, k) + d(k, \mathcal{S}) \leq 3$ with probability one. \blacksquare

7.2 Removing dense facilities

Given some fixed weighting function a , we first need a preprocessing step to ensure that no facility i serves a large (weighted) fraction of the clients.

Proposition 38 *For any $\delta > 0$ and weighting function a , there is an algorithm with $n^{O(1/\delta)}$ runtime, which returns a fractional LP solution (x, y) such that every fractional facility $i \in \mathcal{F}$ has $a(H_i) \leq \delta$. A solution (x, y) with this property is called δ -sparse with respect to a .*

Proof We recursively execute the following procedure: First, solve the LP to obtain a fractional solution (x, y) . Next, if this solution contains some fractional facility $i \in \mathcal{F}$ with $a(H_i) > \delta$, then we form two subproblems; in the first, we force $y_i = 0$ and in the latter, we force $y_i = 1$ and $x_{ij} = 1$ for $j \in H_i$.

Since there is an optimal integral solution, this branching process generates at least one feasible subproblem. Furthermore, each time we execute a branch with $y_i = 1$, the resulting subproblem has $a(\mathcal{C}_{\text{frac}})$ reduced by at least δ , where $\mathcal{C}_{\text{frac}}$ denotes the set of clients which are served by a fractional facility. So the search tree has depth at most $1/\delta$. At a leaf of this branching process, there holds $a(H_i) \leq \delta$ for every fractional facility i . Since each subproblem can be solved in $\text{poly}(n)$ time, the overall runtime is $n^{O(1/\delta)}$. \blacksquare

The next results show how δ -sparse fractional solutions are in certain senses “stable” under small modifications, and how this interacts with the KPR rounding process.

Proposition 39 *Suppose that (x, y) is a δ -sparse LP solution with respect to weighting function a , and vector y' is obtained by modifying t fractional entries of y . Then $\sum_{j \in \mathcal{C}} a_j Q(F_j, y') \leq t\delta + \sum_{j \in \mathcal{C}} a_j Q(F_j, y)$.*

Proof Let A denote the modified facilities and let $V = \bigcup_{i \in A} H_i$; these are the clients which are affected by the modified facilities. For $j \in \mathcal{C} - V$, we have $Q(F_j, y') = Q(F_j, y)$, as none of the clients $i \in F_j$ get modified. Thus, we can write:

$$\sum_{j \in \mathcal{C}} a_j Q(F_j, y') \leq \sum_{j \in \mathcal{C}} a_j Q(F_j, y) + \sum_{j \in V} a_j Q(F_j, y')$$

For the latter term, δ -sparsity implies $a(H_i) \leq \delta$ for each i , and so

$$\sum_{j \in V} a_j Q(F_j, y') \leq \sum_{j \in V} a_j \leq \sum_{i \in A} a(H_i) \leq t\delta$$

■

Proposition 40 *If (x, y) is a δ -sparse fractional LP solution with respect to weighting function a and $\tilde{Y} = \text{KPR}(\mathcal{G}, M, y, t)$ for some integer t with $12m \leq t < 2/\delta$, then we have*

$$\sum_j a_j \mathbf{E}[Q(F_j, \tilde{Y})] \leq 1/e + O(m^2 \delta \log \frac{1}{\delta t})$$

Proof It is convenient to take the following alternative, slowed-down view of the KPR rounding process, where $k \geq 0$ is a parameter to be determined.

- 1: set $y^0 = \text{KPR}(\mathcal{G}, M, y, t2^k)$
- 2: **for** $\ell = 1, \dots, k$ **do** set $y^\ell = \text{KPR}(\mathcal{G}, M, y^{\ell-1}, 2^{k-\ell}t)$
- 3: **Output** $\tilde{Y} = y^k$

Let us define $C_\ell = \sum_{j \in \mathcal{S}} a_j Q(F_j, y^\ell)$ for $\ell = 0, \dots, k$. Thus we need to estimate $C_k = \sum_j a_j \mathbf{E}[Q(F_j, \tilde{Y})]$. To do so, we will compute $\mathbf{E}[C_0]$ and $\mathbf{E}[C_\ell - C_{\ell-1}]$ for all $\ell = 1, \dots, k$.

For the C_0 term, we use Theorem 11(b) to get

$$\mathbf{E}[C_0] = \sum_{j \in \mathcal{C}} a_j \mathbf{E}[Q(F_j, y^0)] \leq \sum_{j \in \mathcal{C}} a_j (Q(F_j, y) + O(m^2/(t2^k))) \quad (15)$$

By Proposition 37, we have $Q(F_j, y) \leq 1/e$ and therefore $\mathbf{E}[C_0] \leq 1/e + O(m^2/(t2^k))$.

Next, suppose we condition on vector $y^{\ell-1}$ for some $\ell > 0$. By Theorem 11(b), every client $j \in \mathcal{C}$ has

$$\mathbf{E}[Q(F_j, y^\ell) \mid y^{\ell-1}] \leq Q(F_j, y^{\ell-1}) + O\left(\frac{m^2}{t2^{k-\ell}}\right).$$

Define the set of clients $A = \bigcup_{i: y^{\ell-1} \in (0,1)} H_i$. By Property (E5), the vector $y^{\ell-1}$ has at most $2t \cdot 2^{k-\ell}$ fractional entries i , and by δ -sparsity each such facility i has $a(H_i) \leq \delta$. Thus, $\sum_{j \in A} a_j \leq (2t \cdot 2^{k-\ell})\delta$. Now observe that for $j \notin A$ we have $Q(F_j, y^\ell) = Q(F_j, y^{\ell-1})$ with probability one (both are equal to zero or one). From this and Eq. (15), we get

$$\mathbf{E}[C_\ell - C_{\ell-1} \mid y^{\ell-1}] = \sum_{j \in A} a_j (\mathbf{E}[Q(F_j, y^\ell) \mid y^{\ell-1}] - Q(F_j, y^{\ell-1})) \leq \sum_{j \in A} a_j \cdot O\left(\frac{m^2}{t2^{k-\ell}}\right) = O(\delta m^2)$$

This implies that $\mathbf{E}[C_\ell - C_{\ell-1}] \leq O(\delta m^2)$, and so summing over ℓ gives

$$\mathbf{E}[C_k] = \mathbf{E}[C_0] + \sum_{\ell=1}^k \mathbf{E}[C_\ell - C_{\ell-1}] \leq 1/e + O\left(\frac{m^2}{t2^k}\right) + O(\delta k m^2)$$

Setting $k = \lceil \log_2 \frac{1}{\delta t} \rceil$ gives $\mathbf{E}[C_k] = 1/e + O(m^2 \delta \log \frac{1}{\delta t})$. (Note $k \geq 0$ by our assumption on t .) ■

7.3 Proof of Theorem 4

When $m = 1$ (the single-knapsack center problem), we use the following Algorithm 7 to satisfy the knapsack constraint with no violation while guaranteeing that all clients get an expected approximation ratio arbitrarily close to $1 + 2/e$.

Algorithm 7 SINGLEKNAPSACKCENTERROUND (M, a, δ)

- 1: Use Proposition 38 to obtain a δ -sparse fractional solution (x, y) with respect to a .
 - 2: Let $\tilde{Y} = \text{KPR}(\mathcal{G}, M, y, 12)$
 - 3: **for** each block G **do**
 - 4: open the facility $i \in G$ with $y_i > 0$ which has the *smallest* weight $M(i)$.
-

Theorem 4 For any $\gamma \in (0, 1)$, a γ -fair feasible solution can be obtained in $n^{\tilde{O}(1/\gamma)}$ runtime.

Proof Without loss of generality assume δ is sufficiently small. In light of Lemma 36, and by rescaling γ , it suffices to show that for any weighting function a , Algorithm 7 generates a solution set \mathcal{S} satisfying $\sum_j a_j d(j, \mathcal{S}) \leq 1 + 2/e + O(\gamma)$ and $\max_j d(j, \mathcal{S}) \leq 3$.

It is useful to view lines 3–4 of Algorithm 7 as a two-part process. First, we convert the fractional vector \tilde{Y} into an integral vector $z \in \{0, 1\}^n$, by moving all the mass in each block to the item with smallest weight. We then open all facilities i with $z_i = 1$. Note that at most 24 entries of \tilde{Y} are modified compared to z .

First, we claim that \mathcal{S} is feasible. For, the fractional solution (x, y) satisfies $My \leq 1$, and Proposition 16 ensures that the vector \tilde{Y} satisfies $M\tilde{Y} = My$. The modification process can only decrease Mz , so $M(\mathcal{S}) = Mz \leq M\tilde{Y}$. Finally, by Proposition 37, every client j has $d(j, \mathcal{S}) \leq 3$ with probability one.

We now turn to analyzing the connection cost. By Proposition 37 we have:

$$\sum_{j \in \mathcal{C}} a_j d(j, \mathcal{S}) \leq \sum_{j \in \mathcal{C}} a_j (1 + 2Q(F_j, z))$$

Since at most 24 entries of y are modified to get the integral vector z , Proposition 39 gives

$$\sum_{j \in \mathcal{C}} a_j d(j, \mathcal{S}) \leq O(\delta) + \sum_{j \in \mathcal{C}} a_j (1 + 2Q(F_j, y))$$

By Proposition 40 (noting that $m = 1$ and $t = 12$), we can take expectations of this quantity to get

$$\sum_{j \in \mathcal{C}} a_j \mathbf{E}[d(j, \mathcal{S})] \leq O(\delta) + O(\delta \log \frac{1}{\delta}) + \sum_{j \in \mathcal{C}} a_j (1 + 2/e)$$

Setting $\delta = \gamma / \log(1/\gamma)$, we get $\mathbf{E}[\sum_j a_j d(j, \mathcal{S})] \leq 1 + 2/e + O(\gamma)$. By running for an expected $O(1/\gamma)$ iterations of this process, we obtain a solution \mathcal{S} with $\sum_j a_j d(j, \mathcal{S}) \leq 1 + 2/e + O(\gamma)$. The overall runtime is $O(1/\gamma)n^{O(1/\delta)} = n^{O(\log(1/\gamma)/\gamma)}$. \blacksquare

7.4 Proof of Theorem 5

When there are multiple knapsack constraints, then the final rounding step cannot satisfy them all exactly. Instead, we use independent selection to obtain an additive pseudo-approximation, as shown in Algorithm 8.

Algorithm 8 MULTIKNAPSACKCENTERROUND (M, a, δ, t)

- 1: Use Proposition 38 to obtain a δ -sparse fractional solution x, y with respect to weighting function a .
 - 2: Let $Y = \text{FULLKPR}(G_\ell, M, y, t)$.
 - 3: **return** $\mathcal{S} = \{i \mid Y_i = 1\}$
-

Depending on the parameter δ , there is a trade-off between approximation ratio, budget violation, and running time. We summarize this as follows:

Theorem 5 Consider a multi-knapsack center instance with $m \geq 1$ constraints and let $\gamma, \epsilon \in (0, 1)$.

- (a) A γ -fair $O\left(\frac{m\sqrt{\log(m/\gamma)}}{\sqrt{\gamma}}\right)$ -additive pseudo-solution can be obtained in $\text{poly}(n/\gamma)$ runtime.
- (b) A γ -fair $O(\sqrt{m \log m})$ -additive pseudo-solution can be obtained in $n^{\tilde{O}(m^2/\gamma)}$ runtime.
- (c) A γ -fair ϵ -multiplicative pseudo-solution can be obtained in $n^{\tilde{O}(m^{3/2}/\epsilon+m^2/\gamma)}$ runtime.

Proof In light of Lemma 36, and by rescaling γ , it suffices to show that for any weighting function a , we can get a solution \mathcal{S} of the above form satisfying $\sum_j a_j d(j, \mathcal{S}) \leq 1 + 2/e + O(\gamma)$ and $\max_j d(j, \mathcal{S}) \leq 3$.

For result (a), we use Algorithm 8 with $t = m^2/\gamma$ and $\delta = 1$. By Proposition 37 we have $\sum_j a_j d(j, \mathcal{S}) \leq \sum_j a_j (1 + 2Q(W, Y))$. By Theorem 11(b) we have $\mathbf{E}[Q(F_j, Y)] \leq Q(F_j, y) + O(m^2/t)$. By Proposition 37, this in turn is at most $1/e + O(m^2/t)$. Since $m^2/t = \gamma$, we have

$$\mathbf{E}\left[\sum_j a_j d(j, \mathcal{S})\right] \leq \sum_j a_j (1 + 2/e + O(\gamma)) = 1 + 2/e + O(\gamma)$$

Let \mathcal{E} be the desired event that \mathcal{S} is a q -additive pseudo-solution for $q = O(\sqrt{t \log(m/\gamma)})$. By Theorem 19, we have $\Pr(\mathcal{E}) \geq 1 - \gamma$ and so $\mathbf{E}[\sum_j a_j d(j, \mathcal{S}) \mid \mathcal{E}] \leq \frac{1+2/e+O(\gamma)}{1-\gamma} \leq 1 + 2/e + O(\gamma)$. Thus, after an expected $O(1/\gamma)$ repetitions, we get a solution \mathcal{S} which is a q -additive pseudo-solution and which has $\sum_j a_j d(j, \mathcal{S}) \leq 1 + 2/e + O(\gamma)$. With this choice of δ , Algorithm 8 runs in $\text{poly}(n/\gamma)$ time.

For result (b), we use Algorithm 8 with parameters $t = 12m^2$ and $\delta = \frac{\gamma}{m^2 \log(1/\gamma)}$. We can break the process of generating Y into two steps: we first generate $\tilde{Y} = \text{KPR}(\mathcal{G}, M, y, 12m^2)$ and then generate $Y = \text{INDSELECT}(\mathcal{G}, \tilde{Y})$. We may assume γ is smaller than any needed constant, so $t < 2/\delta$. Therefore, Proposition 40 gives

$$\mathbf{E}\left[\sum_j a_j (1 + 2Q(F_j, \tilde{Y}))\right] \leq 1 + 2/e + O(m^2 \delta \log \frac{1}{\delta t}).$$

Our choice of δ and t ensures this is at most $1 + 2/e + O(\gamma)$. So after an expected $\Omega(1/\gamma)$ repetitions the solution \tilde{Y} satisfies $\sum_j a_j(1 + 2Q(F_j, \tilde{Y})) \leq 1 + 2/e + O(\gamma)$. Now suppose that this event has occurred, and let us condition on the fixed vector \tilde{Y} . Proposition 37 shows that $\sum_j a_j d(j, \mathcal{S}) \leq \sum_j a_j(1 + 2Q(F_j, Y))$. The vector Y is derived by modifying at most $2t$ fractional entries of vector \tilde{Y} . Therefore, by Proposition 39, we have

$$\sum_j a_j(1 + 2Q(F_j, Y)) \leq 2\delta t + \sum_j a_j(1 + 2Q(F_j, \tilde{Y})) \leq 1 + 2/e + O(\gamma)$$

Furthermore, by Theorem 19, the resulting solution Y is an $O(\sqrt{m \log m})$ -additive pseudo-solution with probability $\Omega(1)$. (We are using the fact that this result holds even after conditioning on \tilde{Y}). Integrating over \tilde{Y} , we see that the vector Y has the desired properties with probability $\Omega(1/\gamma)$. Since Algorithm 8 takes $n^{O(1/\delta)}$ time, the overall expected runtime is $O(1/\gamma)n^{O(1/\delta)} = n^{O(m^2 \log(1/\gamma)/\gamma)}$.

For result (c), say facility i is *big* if $M_k(i) \geq \rho = \Theta(\frac{\epsilon}{\sqrt{m \log m}})$ for any constraint k . We can guess the big facilities in an optimal solution in $n^{O(m/\rho)}$ time. We remove all other big facilities, and apply result (b) to the residual instance. This gives a solution \mathcal{S} with $M(\mathcal{S}) \leq (1 + O(\sqrt{m \log m})) \cdot \rho \vec{1} \leq (1 + \epsilon) \vec{1}$. The runtime is $n^{O(m/\rho + m^2 \log(1/\gamma)/\gamma)} = n^{O(m^{3/2} \log m/\epsilon + m^2 \log(1/\gamma)/\gamma)}$. ■

By contrast, independent rounding would require $n^{O(m \log(m/\gamma)/\epsilon^2)}$ time for an ϵ -multiplicative pseudo-solution, which is a significantly worse dependence upon ϵ .

8. Further correlation bounds for KPR

In this section, we show some additional near-negative-correlation properties for KPR. Although these are not directly used by our clustering algorithms, they may be useful elsewhere. We also remark that the dependent rounding algorithm of Bansal (2019) has been specifically designed to give concentration bounds, which would be similar to (and incomparable in strength with) the ones we develop here.

8.1 Analyzing property (E1) for small $Q(W, y)$

The additive gap in Theorem 11(b) can make it unsuitable when $Q(W, y)$ is small. Theorem 11(c) has a multiplicative gap, but has an undesirable dependence on the size of W . Although we cannot achieve a multiplicative gap independent of the size of W , we can get something which is somewhat in between an additive and multiplicative gap. We show the following main result:

Theorem 41 *Let $W \subseteq U$ and $t \geq 12m$, and define $\theta = m^2/t$. Let $\tilde{Y} = \text{KPR}(\mathcal{G}, M, y, t)$. Then, for any $b > 2$ there holds $\mathbf{E}[Q(W, \tilde{Y})] \leq e^{O(\theta^2 \log^3 b)}(Q(W, y) + \theta/b)$.*

Let us fix W, t, y for the remainder of this section and define $\theta = m^2/t$. Our overall strategy will be to solve a recurrence relation on $Q(W, y)$. One significant complication, which requires much technical delicacy, is that the intermediate values of $Q(W, y)$ are random variables.

Consider the main loop of KPR, and define y^i to be the state vector after the i^{th} iteration of applying KPR-ITERATION. Let us select an integer parameter $k > 20$, whose role will be clarified

later. For each $u = 1, \dots, k$ we define I_u to be the first iteration i such that $T(y^i) \leq t$ or $Q(W, y^i) > \alpha_u$, where we define $\alpha_u = 2^{u-k}$. We also define $\alpha_0 = 0$ and $I_0 = 0$. Thus, y^0 is the vector after INTRABLOCKREDUCE. Also y^k is the final output vector \tilde{Y} since $\alpha_k = 1$.

Throughout we write $S_i = Q(W, y^i)$ and $T_i = T(y^i)$. For each value $u = 0, \dots, k$ we define $H_u = S_{I_u}$, and so $H_0 = Q(W, y^0)$ while H_k is the final value $Q(W, \tilde{Y})$.

There are three stages to analyze the evolution of $\mathbf{E}[S_i]$. First, we analyze the change in a single round (going from S_i to S_{i+1}). Second, we analyze the change over each value of u (going from H_u to H_{u+1}). Finally, we analyze the total change from H_0 to H_k .

Proposition 42 *Let $u \geq 0$, and suppose that we condition on all state up to round i with $i < I_u$. If $t \geq 100mk$, then we have $\mathbf{E}[S_{i+1}] \leq S_i + O(\alpha_u m^2 k^2 / T_i^2)$.*

Proof By Lemma 18, we have:

$$\begin{aligned} \mathbf{E}[S_{i+1}] &\leq S_i \cosh\left(\frac{6m}{T_i} \sum_G y^i(G \cap W)\right) \\ &\leq S_i \cosh\left(\frac{6m}{T_i} \sum_G \ln(1 - y^i(G \cap W))\right) = S_i \cosh\left(\frac{6m \ln S_i}{T_i}\right) \end{aligned}$$

where the second inequality uses the fact that $x \leq -\ln(1 - x)$ for $x \in [0, 1]$. Also, by definition of I_u , we have $S_i \leq \alpha_u$ and $T_i \geq t$. To finish, we claim that

$$S_i \cosh\left(\frac{6m \ln S_i}{T_i}\right) - S_i \leq O\left(\frac{\alpha_u m^2 k^2}{T_i^2}\right). \quad (16)$$

To show this, let $\beta = 6m/T_i$, and observe that since $T_i \geq t \geq 100mk$ we have $\beta \leq 1/k$. Now consider the function $f(s) = s \cosh(\beta \ln s) - s$. Simple analysis shows that $f(s)$ is an increasing function for $s \leq w = (\frac{1-\beta}{1+\beta})^{1/\beta}$. Furthermore, the restriction that $\beta \leq 1/k \leq 1/20$ ensures that $w = \Theta(1)$. So if $\alpha_u \leq w$, then $f(S_i) \leq f(\alpha_u) = \alpha_u \cosh(\beta \ln \alpha_u) - \alpha_u$, and a second-order Taylor series for \cosh then shows that $f(\alpha_u) \leq O(\alpha_u \beta^2 \log^2 \alpha_u) \leq O(\alpha_u \beta^2 k^2)$, implying Eq. (16). If $\alpha_u \geq w$, then Proposition 48 shows that $f(S_i) \leq O(\beta^2) = O(\alpha_u \beta^2)$ and again Eq. (16) holds. ■

Proposition 43 *If $t \geq 100mk$, then $\mathbf{E}[H_{u+1} \mid H_u] \leq H_u + [[H_u \geq \alpha_u]] \cdot O(\alpha_{u+1} \theta k^2)$ for each u .*

Proof Suppose we condition on I_u as well as all state up to iteration $i = I_u$, including the random variable H_u . By definition of I_u , we must have either $S_i \geq \alpha_u$ or $T_i < t$. In the latter case, we immediately have $I_{u+1} = i$ as well and so $H_{u+1} = H_u$. So, let us assume that $S_i \geq \alpha_u$ and we want to show that $\mathbf{E}[H_{u+1}] \leq H_u + O(\alpha_{u+1} \theta k^2)$.

For each $j \geq i$, define $\tilde{S}_j = S_{\min(j, I_{u+1})}$. Note that $\tilde{S}_i = S_i = H_u$ and $\lim_{j \rightarrow \infty} \tilde{S}_j = H_{u+1}$. We claim that, for all $v \geq i$, there holds

$$\mathbf{E}[\tilde{S}_{v+1} \mid \tilde{S}_v] \leq \tilde{S}_v + [[T_v > t]] \cdot O(\alpha_{u+1} m^2 k^2 / T_v^2) \quad (17)$$

For, if $v \geq I_{u+1}$ or $T_v \leq t$, we have $\tilde{S}_{v+1} = \tilde{S}_v$ with probability one; otherwise, we have $S_v = \tilde{S}_v$ and so this follows from Proposition 42. Since $\tilde{S}_i = H_u$, we can sum Eq. (17) over

$v = i, \dots, j - 1$ and use iterated expectations to obtain:

$$\mathbf{E}[\tilde{S}_j] \leq \tilde{S}_i + O(m^2 \alpha_u k^2) \cdot \sum_{v=i}^{j-1} \frac{\Pr(T_v > t)}{T_v^2} = H_u + O(m^2 \alpha_{u+1} k^2) \sum_{\ell=t+1}^{\infty} \sum_{v=i}^j \frac{\Pr(T_v = \ell)}{\ell^2}$$

By Proposition 14, the vector y gains an integral entry with probability at least 0.24 in each round. This implies that the number of iterations v with $T_v = \ell$, is stochastically dominated by a Geometric(0.24) random variable. Thus $\sum_{v=i}^{\infty} \Pr(T_v = \ell) \leq 1/0.24 \leq 5$ and so

$$\mathbf{E}[\tilde{S}_j] \leq H_u + O(m^2 \alpha_{u+1} k^2) \cdot \sum_{\ell=t+1}^{\infty} 5/\ell^2 \leq H_u + O(m^2 \alpha_{u+1} k^2/t)$$

Since $H_{u+1} = \lim_{j \rightarrow \infty} \tilde{S}_j$ and $m^2/t = \theta$, this implies that $\mathbf{E}[H_{u+1}] \leq H_u + O(\alpha_{u+1} \theta k^2)$. ■

Proposition 44 We have $\mathbf{E}[H_k] \leq (\mathbf{E}[H_0] + O(\theta k^2 2^{-k})) e^{O(\theta k^3)}$.

Proof If $t < 100mk$, then $(\mathbf{E}[H_0] + \theta k^2 2^{-k}) e^{\theta k^3} \geq e^{-\Omega(k)} \cdot e^{0.01mk^2} \geq \Omega(1)$. Since $H_k \leq 1$, the claimed bound will then hold vacuously. So assume that $t \geq 100mk$. By Proposition 43 and iterated expectations, we have for each $u \geq 0$:

$$\mathbf{E}[H_{u+1}] \leq \mathbf{E}[H_u] + \Pr(H_u \geq \alpha_u) \cdot O(\alpha_{u+1} \theta k^2) \quad (18)$$

For $u = 0$, recall that $\alpha_0 = 0, \alpha_1 = 2^{1-k}$; thus Eq. (18) implies $\mathbf{E}[H_1] \leq \mathbf{E}[H_0] + O(\frac{\theta k^2}{2^k})$. For each $u \geq 1$, Markov's inequality applied to Eq. (18) gives

$$\mathbf{E}[H_{u+1}] \leq \mathbf{E}[H_u] + O\left(\frac{\mathbf{E}[H_u]}{\alpha_u} \cdot \alpha_{u+1} \theta k^2\right) = \mathbf{E}[H_u](1 + O(\theta k^2))$$

Combining these bounds for $u = 0, \dots, k - 1$ gives

$$\mathbf{E}[H_k] \leq (\mathbf{E}[H_0] + O(\frac{\theta k^2}{2^k})) (1 + O(\theta k^2))^k \leq (\mathbf{E}[H_0] + O(\frac{\theta k^2}{2^k})) e^{O(\theta k^3)}$$

■

Proof [Proof of Theorem 41] Set $k = \lceil c \log_2 b \rceil$, for some constant $c > 20$. By Proposition 44, we have

$$\mathbf{E}[H_k] \leq (\mathbf{E}[H_0] + \frac{K \theta c^2 \log^2 b}{b^c}) e^{K c^3 \theta \log^3 b}$$

for an absolute constant $K \geq 1$. Choosing c to be a sufficiently large constant gives $K c^2 \log^2 b / (b^c) \leq 1/b$ for all $b > 2$, and so $\mathbf{E}[H_k] \leq (\mathbf{E}[H_0] + \theta/b) e^{K c^3 \theta \log^3 b}$. Finally, recall that $Q(W, \tilde{Y}) = H_k$ and by Proposition 12 we have $\mathbf{E}[H_0] = Q(W, y)$. ■

8.2 Concentration bounds for FULLKPR

The analysis in Section 8.1 can be used show a lower-tail concentration bound for FULLKPR.

Theorem 45 *Let $Y = \text{FULLKPR}(\mathcal{G}, M, y, t)$ for $t > 12m$ and let $\theta = m^2/t$. If $y \bullet w \geq \mu$ for some parameters $\mu \geq 1$ and $w \in [0, 1]^U$, then*

$$\Pr(Y \bullet w \leq \mu(1 - \delta)) \leq e^{O(\theta(1+\delta\mu)^3)} \cdot \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^\mu$$

Proof Let \mathcal{E} be the event $Y \bullet w \leq \mu(1 - \delta)$. Consider forming a random set W , wherein each $j \in U$ goes into W independently with probability qw_j for some parameter $q \in [0, 1]$ to be specified. We will compute $\mathbf{E}[Q(W, Y)]$ in two different ways. First, suppose we condition on the event \mathcal{E} , as well as all the random variables Y . Then by Proposition 47,

$$\mathbf{E}[Q(W, Y) \mid \mathcal{E}, Y] = \mathbf{E}\left[\prod_{G \in \mathcal{G}} (1 - q \sum_{j \in G} w_j Y_j) \mid \mathcal{E}, Y\right] \geq \mathbf{E}[(1 - q)^{Y \bullet w} \mid \mathcal{E}, Y] \geq (1 - q)^{\mu(1-\delta)}$$

Therefore, $\mathbf{E}[Q(W, Y)] \geq \Pr(\mathcal{E})(1 - q)^{\mu(1-\delta)}$. On the other hand, if we condition on the random variable W then Theorem 41 gives $\mathbf{E}[Q(W, Y) \mid W] \leq e^{O(\theta \log^3 b)}(Q(W, y) + \theta/b)$ for any parameter $b > 2$. By the way we form W , we can compute

$$\mathbf{E}[Q(W, y)] = \mathbf{E}\left[\prod_{G \in \mathcal{G}} (1 - \sum_{j \in G} y_j [j \in W])\right] = \prod_{G \in \mathcal{G}} (1 - q \sum_{j \in G} w_j y_j) \leq e^{-\sum_j q w_j y_j} = e^{-q\mu}$$

and hence we have

$$\mathbf{E}[Q(W, Y)] \leq e^{O(\theta \log^3 b)}(e^{-q\mu} + \theta/b)$$

Putting these two bounds together, we see $\Pr(\mathcal{E}) \leq \frac{e^{O(\theta \log^3 b)}(e^{-q\mu} + \theta/b)}{(1 - q)^{\mu(1-\delta)}}$. Let us now set $q = \delta$ and $b = e^{1+\delta\mu} > 2$. With these parameters, we have

$$\Pr(\mathcal{E}) \leq \frac{e^{O(\theta(1+\delta\mu)^3)}(e^{-\delta\mu} + \theta e^{-1-\delta\mu})}{(1 - \delta)^{\mu(1-\delta)}} \leq e^{O(\theta(1+\delta\mu)^3)} \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^\mu$$

■

To show upper tail bounds, we use an approach of Schmidt et al. (1995) based on symmetric polynomials. We also use a number of extremal bounds from that paper for such polynomials.

Theorem 46 *Let $Y = \text{FULLKPR}(\mathcal{G}, M, y, t)$ and let $\theta = m^2/t$. If $y \bullet w \leq \mu$ and $t > 10000m(1 + \delta\mu)$ for some parameters $\mu \geq 1$ and $w \in [0, 1]^U$, then*

$$\Pr(Y \bullet w \geq \mu(1 + \delta)) \leq e^{O(\theta(1+\delta\mu)^2)} \cdot \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu$$

Proof Let us consider the random variable defined by

$$H = \sum_{\substack{L \subseteq U \\ |L|=k}} \prod_{j \in L} w_j Y_j$$

where $k = \lceil \mu\delta \rceil$. We use Proposition 20 to compute the expectation of H :

$$\mathbf{E}[H] = \sum_{\substack{L \subseteq U \\ |L|=k}} \left(\prod_{j \in L} w_j \right) \mathbf{E} \left[\prod_{j \in L} Y_j \right] \leq e^{O(\theta k^2 k^2)} \sum_{\substack{L \subseteq U \\ |L|=k}} \prod_{j \in L} w_j y_j .$$

This is precisely the k^{th} symmetric polynomial applied to the quantities $w_j y_j$. As shown in Schmidt et al. (1995), there holds:

$$\sum_{\substack{L \subseteq U \\ |L|=k}} \prod_{j \in L} w_j y_j \leq \frac{(\sum_j w_j y_j)^k}{k!} = \frac{\mu^k}{k!}$$

As shown in Schmidt et al. (1995), whenever $k \leq \lfloor a \rfloor$ and $Y \bullet w \geq a$ for a real number $a \geq 0$, there holds $H \geq \binom{a}{k}$. We use this fact with $a = \mu(1 + \delta)$; note $k \leq \lfloor \mu(1 + \delta) \rfloor$ since $\mu \geq 1$. Applying Markov's inequality to H gives

$$\Pr(Y \bullet w \geq \mu(1 + \delta)) \leq \frac{\mathbf{E}[H]}{\binom{\mu(1+\delta)}{k}} \leq e^{O(\theta k^2)} \frac{\mu^k / k!}{\binom{\mu(1+\delta)}{k}}$$

Finally, as shown in Schmidt et al. (1995), the value $k = \lceil \mu\delta \rceil$ ensures that $\frac{\mu^k / k!}{\binom{\mu(1+\delta)}{k}} \leq \left(\frac{e^\delta}{(1+\delta)^{1+\delta}} \right)^\mu$. So for $t > 10000m(1 + \delta\mu)$, we have shown that

$$\Pr(Y \bullet w \geq \mu(1 + \delta)) \leq e^{O(\theta(1+\delta\mu)^2)} \cdot \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu$$

■

9. Acknowledgments

We thank Nikhil Bansal, Chandra Chekuri, Shi Li, and the referees of the conference and journal versions of this paper for their helpful suggestions.

Appendix A. Some technical lemmas

Proposition 47 For $Y \in \{0, 1\}^n$, $w \in [0, 1]^n$, $\lambda \in [0, 1]$, we have $(1 - \lambda)^{Y \bullet w} \leq \prod_{G \in \mathcal{G}} (1 - \lambda \sum_{j \in G} w_j Y_j)$.

Proof Since Y is an integral vector and $Y(G) \leq 1$, we have $1 - \lambda \sum_{j \in G} w_j Y_j = \prod_{j \in G} (1 - \lambda w_j Y_j)$ for any block G . Therefore, we get

$$\begin{aligned} \prod_{G \in \mathcal{G}} (1 - \lambda \sum_{j \in G} w_j Y_j) &= \prod_{G \in \mathcal{G}} (1 - \lambda \sum_{j \in G} w_j Y_j) = \prod_{G \in \mathcal{G}} \prod_{j \in G} (1 - \lambda w_j Y_j) = \prod_{i=1}^n (1 - \lambda w_i Y_i) \\ &\geq \prod_{i=1}^n (1 - \lambda)^{w_i Y_i} \quad \text{as } (1 + ab) \geq (1 + a)^b \text{ for } a \geq -1 \text{ and } b \in [0, 1] \\ &= (1 - \lambda)^{Y \bullet w} \end{aligned}$$

■

Proposition 48 For $u \in \mathbb{Z}_{\geq 0}$ and $a \in [0, \frac{1}{u+1}]$ and $s \in [0, 1]$, we have $s \cosh(a(u - \ln s)) - s \leq a^2(u + 1)^2$.

Proof Let $f(s) = s \cosh(a(u - \ln s)) - s$. The critical points of function $f(s)$ occur at $s_0 = e^u$ and $s_1 = e^{u(\frac{1-a}{1+a})^{1/a}}$. Since s_0 is outside the allowed parameter range, this means that maximum value of $f(s)$ in the interval $[0, 1]$ must occur at either $s = 0$, $s = s_1$, or $s = 1$.

At $s = 0$, we have $f(s) = 0$. At $s = 1$, we have $f(1) = \cosh(au) - 1$, which is at most $(au)^2$ since $au \leq 1$. Let us now bound s_1 . One can check that $(\frac{1-a}{1+a})^{1/a}$ is a decreasing function of a . Thus, as $a \leq 1/(u + 1)$, we have $s_1 \geq e^{u(\frac{1-1/(u+1)}{1+1/(u+1)})^{u+1}} = e^{u(u/(u+2))^{u+1}}$. This is larger than 1 for $u \geq 3$, so it is out of the range of interest.

So we only need to check $f(s_1) \leq a^2(u + 1)^2$ for $u = 0, 1, 2$; these are all routine calculations. ■

Proposition 49 For real numbers T, t, a with $T \geq t \geq 2\sqrt{a} > 0$, we have

$$e^{a(1/t-1/(T-1))} \leq e^{a(1/t-1/T)} - \frac{ae^{a(1/t-1/T)}}{2T^2}$$

Proof Dividing both sides by $e^{a(1/t-1/T)}$, we need to show that $e^{-a/(T-1)+a/T} \leq (1 - a/(2T^2))$. Since $T \geq 2\sqrt{a}$, we have $a/(2T^2) \leq 1/8$. Therefore, $1 - a/(2T^2) \geq e^{-3/4(a/T^2)}$. So it suffices to show that $-a/(T - 1) + a/T \leq -(3/4)a/T^2$. It is routine to verify this holds for $a > 0$ and $T \geq 2$. ■

References

- A. Ageev and M. Sviridenko. Pipeage rounding: a new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328, 2004.
- Noga Alon and Joel Spencer. *The Probabilistic Method*. Wiley, 4th edition, 2016.
- Noga Alon, Hao Huang, and Benny Sudakov. Nonnegative k -sums, fractional covers, and probability of small deviations. *Journal of Combinatorial Theory, Series B*, 102(3):784–796, 2012.
- Nikhil Bansal. On a generalization of iterated and randomized rounding. In *Proc. 51st ACM Symposium on Theory of Computing (STOC)*, pages 1125–1135, 2019.
- Nikhil Bansal and Viswanath Nagarajan. Approximation-friendly discrepancy rounding. In *A Journey Through Discrete Mathematics*, pages 89–114. Springer, 2017.
- J. Beck and T. Fiala. “Integer-making” theorems. *Discrete Applied Mathematics*, 3:1–8, 1981.

- Jaroslav Byrka, Aravind Srinivasan, and Chaitanya Swamy. Fault-tolerant facility location: a randomized dependent LP-rounding algorithm. In *Proc. 14th international conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 244–257, 2010.
- Jarosław Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k -median and positive correlation in budgeted optimization. *ACM Transactions on Algorithms (TALG)*, 13(2):1–31, 2017.
- Gruia Călinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011. doi: 10.1137/080733991. URL <http://dx.doi.org/10.1137/080733991>.
- Moses Charikar and Shi Li. A dependent LP-rounding approach for the k -median problem. *Automata, Languages, and Programming (ICALP)*, pages 194–205, 2012.
- Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *Proc. 51st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 575–584, 2010. doi: 10.1109/FOCS.2010.60. URL <http://dx.doi.org/10.1109/FOCS.2010.60>.
- Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Multi-budgeted matchings and matroid intersection via dependent rounding. In *Proc. 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1080–1097, 2011. doi: 10.1137/1.9781611973082.82. URL <http://dx.doi.org/10.1137/1.9781611973082.82>.
- Danny Z. Chen, Jian Li, Hongyu Liang, and Haitao Wang. Matroid and knapsack center problems. *Algorithmica*, 75(1):27–52, 2016. doi: 10.1007/s00453-015-0010-1.
- Devdatt P. Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *Random Structures & Algorithms*, 13(2):99–124, 1998.
- Talya Eden, Dana Ron, and C Seshadhri. Sublinear time estimation of degree distribution moments: The arboricity connection. *SIAM Journal on Discrete Mathematics*, 33(4):2267–2285, 2019.
- Uriel Feige. On sums of independent random variables with unbounded variance and estimating the average degree in a graph. *SIAM Journal on Computing*, 35(4):964–984, 2006.
- R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM*, 53(3):324–360, 2006.
- Brian Garnett. Small deviations of sums of independent random variables. *Journal of Combinatorial Theory, Series A*, 169, Article #105119, 2020.
- Oded Goldreich and Dana Ron. Approximating average parameters of graphs. *Random Structures & Algorithms*, 32(4):473–493, 2008.
- David G Harris and Aravind Srinivasan. The Moser–Tardos framework with partial resampling. *Journal of the ACM*, 66(5):36, 2019.
- David G. Harris, Shi Li, Thomas Pensyl, Aravind Srinivasan, and Khoa Trinh. Approximation algorithms for stochastic clustering. *Journal of Machine Learning Research*, 20(153):1–33, 2019.

- Simai He, Jiawei Zhang, and Shuzhong Zhang. Bounding probability of small deviation: a fourth moment approach. *Mathematics of Operations Research*, 35(1):208–232, 2010.
- Dorit S. Hochbaum and David B. Shmoys. A unified approach to approximation algorithms for bottleneck problems. *Journal of the ACM*, 33(3):533–550, 1986. doi: 10.1145/5925.5933. URL <http://doi.acm.org/10.1145/5925.5933>.
- Kamal Jain, Mohammad Mahdian, and Amin Saberi. A new greedy approach for facility location problems. In *Proc. 34th ACM Symposium on Theory of Computing (STOC)*, pages 731–740, 2002.
- Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.
- R. M. Karp, F. T. Leighton, R. L. Rivest, C. D. Thompson, U. V. Vazirani, and V. V. Vazirani. Global wire routing in two-dimensional arrays. *Algorithmica*, 2:113–129, 1987.
- Samir Khuller, Robert Pless, and Yoram J. Sussmann. Fault tolerant k -center problems. *Theoretical Computer Science*, 242(1-2):237–245, 2000. doi: 10.1016/S0304-3975(98)00222-9. URL [http://dx.doi.org/10.1016/S0304-3975\(98\)00222-9](http://dx.doi.org/10.1016/S0304-3975(98)00222-9).
- Ravishankar Krishnaswamy, Amit Kumar, Viswanath Nagarajan, Yogish Sabharwal, and Barna Saha. The matroid median problem. In *Proc. 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1117–1130, 2011.
- Ravishankar Krishnaswamy, Shi Li, and Sai Sandeep. Constant approximation for k -median and k -means with outliers via iterative rounding. In *Proc. 50th ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 646–659, 2018.
- Frank Thomson Leighton, Chi-Jen Lu, Satish Rao, and Aravind Srinivasan. New algorithmic aspects of the Local Lemma with applications to routing and partitioning. *SIAM Journal on Computing*, 31(2):626–641, 2001.
- Shi Li and Ola Svensson. Approximating k -median via pseudo-approximation. *Journal of the ACM*, 45(2):530–547, 2016.
- Stephen Mitchell Samuels. On a Chebyshev-type inequality for sums of independent random variables. *The Annals of Mathematical Statistics*, 37(1):248–259, 1966.
- Jeanette P Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff–Hoeffding bounds for applications with limited independence. *SIAM Journal on Discrete Mathematics*, 8(2):223–250, 1995.
- A. Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 588–597, 2001.
- Aravind Srinivasan. An extension of the Lovász Local Lemma, and its applications to integer programming. *SIAM Journal on Computing*, 36(3):609–634, 2006.
- Chaitanya Swamy. *Approximation algorithms for clustering problems*. Ph.D. dissertation, Cornell University, 2004.

David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011. ISBN 978-0-521-19527-0.