

Submodular Function Maximization via the Multilinear Relaxation and Contention Resolution Schemes

Chandra Chekuri^{*}
Dept. of Computer Science
University of Illinois
Urbana, IL 61801
chekuri@cs.illinois.edu

Jan Vondrák
IBM Almaden Research Ctr.
650 Harry Road
San Jose, CA 95120
jvondrak@us.ibm.com

Rico Zenklusen[†]
Dept. of Mathematics
MIT
Cambridge, MA 02139
ricoz@math.mit.edu

ABSTRACT

We consider the problem of maximizing a non-negative submodular set function $f : 2^N \rightarrow \mathbb{R}_+$ over a ground set N subject to a variety of packing type constraints including (multiple) matroid constraints, knapsack constraints, and their intersections. In this paper we develop a general framework that allows us to derive a number of new results, in particular when f may be a *non-monotone* function. Our algorithms are based on (approximately) solving the multilinear extension F of f [5] over a polytope P that represents the constraints, and then effectively rounding the fractional solution. Although this approach has been used quite successfully in some settings [6, 22, 24, 13, 3], it has been limited in some important ways. We overcome these limitations as follows.

First, we give constant factor approximation algorithms to maximize F over an arbitrary down-closed polytope P that has an efficient separation oracle. Previously this was known only for monotone functions [36]. For non-monotone functions, a constant factor was known only when the polytope was either the intersection of a fixed number of knapsack constraints [24] or a matroid polytope [37, 30]. Second, we show that *contention resolution schemes* are an effective way to round a fractional solution, even when f is non-monotone. In particular, contention resolution schemes for different polytopes can be combined to handle the intersection of different constraints. Via LP duality we show that a contention resolution scheme for a constraint is related to the *correlation gap* [1] of weighted rank functions of the constraint. This leads to an optimal contention resolution scheme for the matroid polytope.

Our results provide a broadly applicable framework for maximizing linear and submodular functions subject to independence constraints. We give several illustrative examples. Contention resolution schemes may find other applications.

^{*}Partially supported by NSF grants CCF-0728782 and CCF-1016684.

[†]Supported by the Swiss National Science Foundation, grant number: PBEZP2-129524.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'11, June 6–8, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0691-1/11/06 ...\$10.00.

Categories and Subject Descriptors

F.2.2 [Non-numerical Algorithms and Problems]: Computations on Discrete Structures

General Terms

Algorithms, Theory

Keywords

Approximation algorithm, submodular function maximization, matroid, contention resolution scheme, independence constraint

1. INTRODUCTION

We consider the meta-problem of *maximizing* a non-negative submodular set function subject to independence constraints. Formally, let N be a finite ground set of cardinality n , and let $f : 2^N \rightarrow \mathbb{R}_+$ be a submodular set function over N .¹ Let $\mathcal{I} \subseteq 2^N$ be a downward-closed family² of subsets of N . Our problem is then $\max_{S \in \mathcal{I}} f(S)$. We are interested in independence families induced by natural and useful constraints such as matroid constraints, knapsack constraints, related special cases and their intersections. Throughout this paper we assume that f is given via a value oracle; that is, given a set $S \subseteq N$ the oracle returns $f(S)$. The function f could be monotone or non-monotone³; monotone functions typically allow better approximation results.

Submodular function maximization has recently attracted considerable attention in theoretical computer science. This is for a variety of reasons including applications, recognition of interesting algorithmic and structural properties, as well as the use of submodular functions as utility functions in algorithmic game theory. A number of well-known problems can be seen as special cases of submodular function maximization. For example, the APX-hard Max-Cut problem can be seen as (unconstrained) maximization of the cut function $f : 2^V \rightarrow \mathbb{R}$ of a graph $G = (V, E)$. (Note that f here is non-monotone.) Another well-known special case of our problem is the Max- k -Cover problem, which can be viewed as $\max\{f(S) : |S| \leq k\}$ where $f(S) = |\bigcup_{j \in S} A_j|$ is the coverage function for a collection of sets $\{A_i\}$. Max- k -Cover is hard to approximate to within a factor of $(1 - 1/e + \epsilon)$ for any fixed $\epsilon > 0$, unless $P = NP$ [15]. Hence we focus on approximation algorithms.⁴

¹A set function $f : 2^N \rightarrow \mathbb{R}$ is submodular iff $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$ for all $A, B \subseteq N$.

²A family of sets $\mathcal{I} \subseteq 2^N$ is downward-closed if for any $A \subset B \subseteq N$, $B \in \mathcal{I}$ implies that $A \in \mathcal{I}$.

³ f is *monotone* if $f(A) \leq f(B)$ whenever $A \subseteq B$.

⁴If f is not assumed to be non-negative, even the unconstrained

Classical work in submodular function maximization was based on combinatorial techniques such as the greedy algorithm and local search. We mention the work of Cornuejols, Fisher, Nemhauser and Wolsey [14, 29, 18, 28] from the late 70's which showed a variety of approximation bounds when f is monotone submodular and \mathcal{I} is induced by multiple matroid constraints. Recent algorithmic work has considerably extended and improved the classical results. Local-search methods have been identified as particularly useful, in particular, for non-monotone functions. This has led to the first constant factor approximation for the unconstrained submodular function maximization problem [16], and a variety of approximation results for knapsack and matroid constraints [24, 25]. The greedy algorithm has also been modified and made applicable to non-monotone functions [20].

Despite the above-mentioned results, combinatorial techniques have some limitations: (i) they have not been able to achieve optimal approximation results, except in the basic case of a single cardinality or knapsack constraint [29, 33]; (ii) they are not very flexible in terms of the ability to combine constraints and develop more general techniques (e.g., a $(1 - 1/e)$ -approximation was known for maximizing a monotone submodular function subject to 1 knapsack constraint [33], but little was known even for 2 knapsack constraints). A new approach which overcomes some of these obstacles and brings submodular function maximization closer to the world of polyhedral techniques is via the *multilinear relaxation*, introduced in this context in [5]: $F(\mathbf{x}) = \sum_{S \subseteq N} f(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j)$. The value $F(\mathbf{x})$ is equivalently the expected value of $f(R)$ where R is a random set obtained by picking each element i independently with probability x_i . We observe that if f is modular⁵ then F is simply a linear function.

Continuous extensions offer some advantages in the design of approximation algorithms. Suppose we have a polytope $P_{\mathcal{I}} \subseteq [0, 1]^N$ that is a relaxation for $\mathcal{I} \subseteq 2^N$ in the sense that $\{\mathbf{1}_I \mid I \in \mathcal{I}\} \subset P_{\mathcal{I}}$. Moreover suppose there is a polynomial-time separation oracle for $P_{\mathcal{I}}$ (we call such polytopes solvable). Then we can hope to (approximately) solve the continuous problem $\max_{\mathbf{x} \in P_{\mathcal{I}}} F(\mathbf{x})$ to find a fractional solution $\mathbf{x}^* \in P_{\mathcal{I}}$ and then round \mathbf{x}^* to an integral solution. This is a standard paradigm in approximation via linear and convex programming relaxations. Two natural questions arise in applying this paradigm to submodular functions, both due to the fact that the extension F is neither a convex nor concave function. First, can we (approximately) solve the problem $\max_{\mathbf{x} \in P_{\mathcal{I}}} F(\mathbf{x})$? Second, can we round a fractional solution effectively?

Recent work has addressed the above questions in several ways. First, Vondrák [36] gave a continuous greedy algorithm that gives an optimal $(1 - 1/e)$ -approximation for the problem $\max_{\mathbf{x} \in P} F(\mathbf{x})$ when f is monotone submodular and P is a solvable polytope. When f is non-monotone, the picture is less satisfactory. Lee et al. [24] gave a local-search based algorithm that gives a $(1/4 - \epsilon)$ -approximation to maximize F over the polytope induced by a fixed number of knapsack constraints. Vondrák [37] obtained a 0.309-approximation for maximizing F over a single matroid polytope, and this ratio has been recently improved to 0.325 [30]. However, no approximation algorithm was known to maximize F over a general solvable polytope P .

In terms of rounding a fractional solution \mathbf{x} , a natural strategy to preserve the value of $F(\mathbf{x})$ is to independently round each coordinate i to 1 with probability x_i . However, this rounding strategy

problem is inapproximable since deciding whether the optimum value is positive or zero requires an exponential number of queries.

⁵A function is modular if $f(A) + f(B) = f(A \cup B) + f(A \cap B)$ for all $A, B \subseteq N$. If f is modular then $f(A) = w_0 + \sum_{i \in A} w_i$ for some weight function $w : N \rightarrow \mathbb{R}$.

does not typically preserve the constraints imposed by \mathcal{I} . Various dependent rounding schemes have been proposed. It was shown in [5] that "pipage rounding" can be used to round solutions in the matroid polytope without losing in terms of the objective function $F(\mathbf{x})$ ([13] achieves the same via "swap-rounding"). In [22, 24, 3, 23] randomized rounding coupled with alteration was used for knapsack constraints. More recently, [13] showed concentration properties for rounding in a single matroid polytope when f is monotone, and [38] showed concentration for independent rounding even when f is non-monotone. These led to a few additional results. Despite this progress, the "integrality gap" of $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ has been so far unknown even when f is monotone and P the intersection of two matroid polytopes. (We remark that for pure intersections of matroids, combinatorial algorithms are known to yield good approximations [24, 25].) However, even for modular functions, combining constraints such as matroids and knapsack constraints has been difficult, and no general result was known that matched the best bounds one can get for them separately.

Our contribution at a high level: In this paper we overcome existing limitations by obtaining a general framework via the following results.

- We give the first constant factor approximation for the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ where P is any down-monotone solvable polytope and F the multilinear extension of any non-negative submodular function.
- We propose a general (dependent) randomized rounding framework for modular and submodular functions under independence constraints via what we call *contention resolution schemes* (CR schemes). A key advantage is the ability to easily combine schemes for different constraints to obtain a scheme for their intersection.
- We give an *optimal* $(1 - 1/e)$ -factor CR scheme for any matroid. Previously this was known only for the uniform matroid of rank 1 [17]. More generally, we give a tight connection between CR schemes and the *correlation gap* of the associated weighted rank functions.

The above ingredients can be put together to give a variety of new results that we discuss in more detail in Section 2. We summarize some of our results in Table 1.

1.1 Maximizing the multilinear extension over a general polytope

We now give a more detailed description of our technical results and the general framework. First, we give a constant factor approximation for the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$, where F is the multilinear extension of a non-monotone submodular function f and P is a down-monotone solvable polytope; the monotone case admits a $(1 - 1/e)$ -approximation [36] as we mentioned already. The condition of down-monotonicity of the polytope is necessary for the non-monotone case; it follows from [37] that no constant factor approximation is possible for the matroid base polytope which is not down-monotone.

The main algorithmic technique for non-monotone functions is local search. Fractional local search with additional ideas has been the tool to solve the continuous problem in special cases of polytopes [24, 37, 30]. Previous fractional local search methods for a constant number of knapsack constraints ([24] and [37]) improved a current solution \mathbf{x} by considering moves along a small number of coordinates of \mathbf{x} . The analysis took advantage of the combinatorial structure of the underlying discrete structure (knapsacks or matroids) which was sufficiently simple that swaps along a few

Constraint type	Linear maximization	Monotone submod. max.	Non-negative submod. max.
$O(1)$ knapsacks	$[1 - \varepsilon]$	$[1 - 1/e - \varepsilon]$	$0.325 [0.25]$
k matroids & $\ell = O(1)$ knapsacks	$0.6/k$	$0.38/k [\Omega(1/(k + \ell))]$	$0.19/k [\Omega(1/(k + \ell))]$
k -matroid & ℓ -sparse PIP	$\Omega(1/(k + \ell))$	$\Omega(1/(k + \ell)) [\Omega(1/k\ell)]$	$\Omega(1/(k + \ell)) [\Omega(1/k\ell)]$
Unsplittable flow in paths and trees	$[\Omega(1)]$	$\Omega(1)$	$\Omega(1)$

Table 1: Approximation factors for different types of constraints and objective functions. Results in brackets were previously known.

coordinates sufficed. How do we obtain an algorithm that works for any polytope P ?

A new insight: Our key high-level idea is simple yet insightful. Any point $\mathbf{x} \in P$ can be written as a convex combination of the vertices of P . We view the problem of $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ as optimizing a submodular function over the ground set consisting of the (exponentially many) vertices of P (duplicated many times in the limit). This effectively reduces the polytope constraint to a cardinality constraint and one can apply known algorithmic ideas for the *discrete* problem. From this viewpoint we obtain a new fractional local search procedure: given a current point \mathbf{x} , a local swap corresponds to removing a vertex in the convex combination of \mathbf{x} and adding a new vertex of P (with appropriate scalar multipliers). To implement this efficiently we can use linear optimization over P . (We remark that the continuous greedy algorithm for the monotone case [36] can also be interpreted with this insight.)

Our algorithms are derived using the above high-level idea. We note that when specialized to the matroid polytope or knapsack polytope which have combinatorial structure, our algorithms become simpler and in fact resemble previous algorithms. This is perhaps not a coincidence; it could be argued that our interpretation via the vertices of the underlying polytope P is perhaps the “right” view. Our algorithms and proofs of approximation guarantees are in fact simpler than the previously given proofs for particular polytopes [24, 37, 30]. We present three algorithms of varying complexity. The first algorithm is close in spirit to the local-search algorithm of Lee et al. for knapsack constraints [24] and gives a 0.25-approximation. The second algorithm uses some ideas of [37] for the case of a matroid polytope and gives a 0.309-approximation with respect to the best *integer* solution in P . The most involved algorithm is a generalization of a recent algorithm inspired by simulated annealing [30] which gives a 0.325-approximation, also with respect to the best integer solution in P . We remark that a known limit on the approximability of $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ is a hardness of 0.478-approximation in the value oracle model, even in the special case of a matroid polytope, also due to [30]. We summarize our results in the following theorem.

THEOREM 1.1. *For any nonnegative submodular function f and a solvable down-monotone polytope P , there is a 0.25-approximation algorithm for the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ where F is the multilinear extension of f . There is also an algorithm for this problem which returns a solution $\mathbf{y} \in P$ of value $F(\mathbf{y}) \geq 0.325 \cdot \max\{F(\mathbf{x}) : \mathbf{x} \in P \cap \{0, 1\}^N\}$.*

1.2 Contention resolution schemes

We show that a certain natural class of rounding schemes that we call *contention resolution schemes* (CR schemes) provides a useful and general framework for rounding submodular functions. For a ground set N , let $P_{\mathcal{I}}$ be a convex relaxation of the constraints imposed by $\mathcal{I} \subseteq 2^N$, and let $\mathbf{x} \in P_{\mathcal{I}}$. From the definition of F , a natural strategy to round a point \mathbf{x} is to independently round the coordinates; however, this is unlikely to preserve the constraints imposed by \mathcal{I} . Let $R(\mathbf{x}) \subseteq N$ be a random set obtained by including each element $i \in N$ independently with probability x_i .

The set $R(\mathbf{x})$ is not necessarily feasible. We would like to remove (randomly) some elements from $R(\mathbf{x})$, so that we obtain a feasible set $I \subseteq R(\mathbf{x})$. The property we would like to achieve is that every element i appears in I with probability at least cx_i for some parameter $c > 0$. We call such a scheme “ c -balanced contention resolution” for $P_{\mathcal{I}}$. We stress that the scheme needs to work for all $\mathbf{x} \in P_{\mathcal{I}}$. In several settings we need to first scale down the fractional solution which calls for a more general definition below.

DEFINITION 1.2. *A (b, c) -balanced CR scheme for $P_{\mathcal{I}}$ is a scheme such that for any $\mathbf{x} \in P_{\mathcal{I}}$, the scheme selects an independent subset $I \subseteq R(b\mathbf{x})$ with the following property: $\Pr[i \in I \mid i \in R(b\mathbf{x})] \geq c$ for every element i . The scheme is said to be monotone if $\Pr[i \in I \mid R(b\mathbf{x}) = R_1] \geq \Pr[i \in I \mid R(b\mathbf{x}) = R_2]$ whenever $i \in R_1 \subseteq R_2$. A scheme is said to be strict if $\Pr[i \in I \mid i \in R(b\mathbf{x})] = c$ for every i .*

We emphasize that a CR scheme is defined with respect to a specific polyhedral relaxation $P_{\mathcal{I}}$ of \mathcal{I} . We observe that several previous rounding procedures for packing (and also covering) problems rely on the well-known technique of *alteration* of a set obtained via independent rounding (see [32, 4, 7, 12, 3] for some examples), and are examples of CR schemes. However, these schemes are oblivious in that they do not depend on \mathbf{x} itself (other than in picking the random set R), and the alteration is also deterministic. Our definition is inspired by the “fair contention resolution scheme” in [17] which considered the special case of contention for a single item. The dependence on \mathbf{x} is essential for matroids.

We observe that monotonicity of the CR scheme is a necessary property for submodular functions while it is not required for modular functions. We prove the following theorem via the FKG inequality. We note that a similar theorem was shown earlier for monotone submodular functions in [3].

THEOREM 1.3. *Let $f : 2^N \rightarrow \mathbb{R}_+$ be a non-negative submodular function and \mathbf{x} be a point in $P_{\mathcal{I}}$, a convex relaxation for $\mathcal{I} \subseteq 2^N$. Let $I(\mathbf{x}) \in \mathcal{I}$ be the random output of a monotone (b, c) -balanced CR scheme on $\mathbf{x} \in P_{\mathcal{I}}$. If f is non-monotone, let us assume in addition that the CR scheme is strict. Then $\mathbf{E}[f(I)] \geq c\mathbf{E}[F(b\mathbf{x})]$.*

REMARK 1.4. *One can show that the strictness assumption for non-monotone CR schemes can be dropped in the above theorem, when applying the following pruning operation to the set I obtained by a monotone (b, c) -balanced CR scheme: Go once through all elements of I in any order, and keep an element i in I only if its marginal value with respect to the previously selected elements is positive. A possible drawback of this pruning step is that the obtained procedure to construct I from \mathbf{x} is not oblivious to f .*

For a given $P_{\mathcal{I}}$ it is natural to look for the largest c such that there is a c -balanced scheme. However, when one is interested in combining several schemes we need more flexibility. In particular, we will be interested in (b, c) -balanced schemes that have the property that as $b \rightarrow 0$ we have $c \rightarrow 1$. Note that this is a natural property since scaling down the fractional solution and then independently rounding should make it more likely that the resulting set belongs to \mathcal{I} (since \mathcal{I} is downward closed). We show the following lemma.

LEMMA 1.5. Let $\mathcal{I} = \cap_{i=1}^k \mathcal{I}_i$ and $P_{\mathcal{I}} = \cap_i P_{\mathcal{I}_i}$. Suppose each $P_{\mathcal{I}_i}$ has a monotone (b, c_i) -balanced CR scheme. Then $P_{\mathcal{I}}$ has a monotone $(b, \prod_i c_i)$ -balanced CR scheme. In the special case that each element of N participates in at most k constraints and $c_i = c$ for all i then $P_{\mathcal{I}}$ has a monotone (b, c^k) CR scheme. Moreover, if the scheme for each $P_{\mathcal{I}_i}$ is implementable in poly-time time then the combined scheme for $P_{\mathcal{I}}$ can be implemented in poly-time.

REMARK 1.6. The combined CR scheme works with respect to the natural combination of constraint relaxations — an intersection of the respective polytopes. This ensures that the relaxed problem is still tractable and we can apply our optimization framework.

Contention resolution via correlation gap and an optimal scheme for matroids: A natural question is how one proves the existence of a contention resolution scheme. As we mentioned, several existing rounding schemes are based on deterministic and oblivious alteration to a set obtained via independent rounding. Most of these schemes have been applied to constraint systems induced by linear inequalities of the form $A\mathbf{x} \leq \mathbf{b}$ where A is a non-negative matrix. Until recently there was no contention resolution scheme for the matroid polytope; an optimal $(b, \frac{1-e^{-b}}{b})$ -scheme was previously known for the very special case of the uniform matroid of rank one [17]. We note that the recent work of Chawla et al. [9, 10] implicitly contains a $(b, 1-b)$ -balanced scheme for matroids, although their motivation was different. In this paper we develop an optimal scheme for an arbitrary matroid.

THEOREM 1.7. There is an optimal $(b, \frac{1-e^{-b}}{b})$ -balanced contention resolution scheme for any matroid polytope. Moreover the scheme is monotone and efficiently implementable.

We use randomized schemes and view them abstractly as a convex combination of deterministic schemes. This allows us, via LP duality, to show that the best contention resolution scheme for a constraint system is related to the notion of correlation gap for weighted rank functions of the underlying constraint. We reiterate that the scheme depends on the fractional solution \mathbf{x} that we wish to round; the alteration of the random set $R(\mathbf{x})$ is itself a randomized procedure that is tailored to \mathbf{x} , and is found by solving a linear program. We are inspired to make the general connection to correlation gap due to the recent work of Yan [39]; he applied a similar idea in the context of greedy posted-price ordering schemes for Bayesian mechanism design, improving the bounds of [9, 10].

1.3 A framework for rounding via contention resolution schemes

We now describe our framework for the problem $\max_{S \in \mathcal{I}} f(S)$. The framework assumes the following: (i) there is a polynomial-time value oracle for f , and (ii) that there is a polytope $P_{\mathcal{I}}$ that contains the set $\{\mathbf{1}_S \mid S \in \mathcal{I}\}$ and moreover that there is a polynomial-time separation oracle for $P_{\mathcal{I}}$, and (iii) there is a strict and monotone (b, c) -balanced contention resolution scheme for $P_{\mathcal{I}}$. We then have the following simple algorithm:

1. Using an approximation algorithm, obtain in polynomial time a point $\mathbf{x}^* \in P_{\mathcal{I}}$ such that $F(\mathbf{x}^*) \geq \alpha \cdot \max_{S \in \mathcal{I}} \{F(\mathbf{x}) \mid \mathbf{x} \in P_{\mathcal{I}} \cap \{0, 1\}^N\} \geq \alpha \cdot \max_{S \in \mathcal{I}} f(S)$.
2. Use a strict and monotone (b, c) -balanced contention resolution scheme for $P_{\mathcal{I}}$ on \mathbf{x}^* to output a random set $I(\mathbf{x}^*) \in \mathcal{I}$.

THEOREM 1.8. The above algorithm is a randomized $(0.325bc)$ -approximation algorithm for $\max_{S \in \mathcal{I}} f(S)$. If f is monotone then the approximation ratio is $(1 - 1/e)bc$. If f is modular then the ratio is bc and the contention resolution scheme is not restricted to be monotone.

PROOF. We have $F(\mathbf{x}^*) \geq \alpha \text{OPT}$ with $\text{OPT} = \max_{S \in \mathcal{I}} f(S)$. Theorem 1.3 shows that if we apply a strict and monotone (b, c) -balanced contention resolution scheme to \mathbf{x}^* then the random set I output by it has the property that $\mathbf{E}[f(I)] \geq bcF(\mathbf{x}^*)$, hence we have that $\mathbf{E}[f(I)] \geq \alpha \cdot (bc)\text{OPT}$.

For non-monotone submodular functions, Theorem 1.1 gives $\alpha = 0.325$. For monotone submodular functions, [36] gives $\alpha = 1 - 1/e$. For modular f , $F(\mathbf{x})$ is a linear function, and hence $\alpha = 1$ can be obtained by linear programming. Moreover, if $F(\mathbf{x})$ is a linear function then by linearity of expectation, $\mathbf{E}[f(I)] \geq bcF(\mathbf{x}^*)$ without any monotonicity assumption on the scheme. \square

Combining schemes for different constraints: We are particularly interested in the case when $\mathcal{I} = \cap_{i=1}^k \mathcal{I}_i$ is the intersection of several different independence systems on N ; each system corresponds to a different set of constraints that we would like to impose. Assuming that we can apply the above framework to each \mathcal{I}_i separately, it is straightforward to obtain an algorithm for \mathcal{I} as follows.

Let $P_{\mathcal{I}_i}$ be a polytope that is the relaxation of \mathcal{I}_i . In other words $\{\mathbf{1}_S \mid S \in \mathcal{I}_i\}$ is contained in $P_{\mathcal{I}_i}$. Let $P_{\mathcal{I}} = \cap_i P_{\mathcal{I}_i}$. It follows that $\{\mathbf{1}_S \mid S \in \mathcal{I}\}$ is contained in $P_{\mathcal{I}}$ and also that there is a polynomial-time separation oracle for $P_{\mathcal{I}}$ if there is one for each $P_{\mathcal{I}_i}$. Now suppose there is a monotone (b, c_i) contention resolution scheme for $P_{\mathcal{I}_i}$ for some common choice of b . It follows from Lemma 1.5 that $P_{\mathcal{I}}$ has a monotone $(b, \prod_i c_i)$ contention resolution scheme. We can then apply Theorem 1.8 to obtain a $(\alpha b \prod_i c_i)$ -randomized approximation for $\max_{S \in \mathcal{I}} f(S)$ where α depends on whether f is modular, monotone or non-monotone.

Organization: The rest of the paper is divided into three parts. Some illustrative applications of our framework are discussed in Section 2. Constant factor approximation algorithms for maximizing F over a solvable polytope are described in Section 3. The connection between contention resolution schemes and correlation gap and its use in deriving optimal schemes for matroids are discussed in Section 4, as well as contention resolution schemes for knapsack constraints, sparse packing systems, and UFP in paths and trees. Several proofs are omitted due to space constraints.

2. APPLICATIONS

Contention resolution. In this section we briefly outline some concrete results that can be obtained via our framework. Our approximation results for various constraints are direct consequences of the respective CR schemes, so we focus on these first. We note that the schemes are with respect to the natural polyhedral relaxations.

Matroids and matchoids: For every matroid constraint $\mathcal{M} = (N, \mathcal{I})$, we develop an optimal $(1 - 1/e)$ -balanced CR scheme. More generally, for any $b \in (0, 1]$ we design a $(b, \frac{1-e^{-b}}{b})$ -balanced CR scheme, which lends itself well to combinations with other constraints. For the intersection of k matroids, or more generally for a k -uniform matchoid (a common generalization of k -set packing and intersection of k matroids [26]), we obtain a $(b, (\frac{1-e^{-b}}{b})^k)$ -balanced CR scheme for any $b \in (0, 1]$. The choice of $b = \frac{2}{k+1}$ gives a $\frac{2}{e(k+1)}$ -balanced CR scheme for every k -uniform matchoid.

Knapsack / linear packing constraints: Given a non-negative $m \times n$ matrix A and non-negative vector \mathbf{b} , let $\mathcal{I} = \{S \mid A\mathbf{1}_S \leq \mathbf{b}\}$ where $\mathbf{1}_S$ is the indicator vector of set $S \subseteq N$. The width of the system of inequalities is defined as $W = \lfloor \min_{i,j} b_j / A_{i,j} \rfloor$. Some special cases of interest are (i) A is a $\{0, 1\}$ -matrix, (ii) A is column-restricted, that is, all non-zero entries in each column are the same and (iii) A is k -column sparse, that is at most k non-zero entries

in each column. Many combinatorial problems can be captured by these constraints. Previous results implicitly contain the following:

- For a constant number of knapsack constraints ($m = O(1)$), by guessing and enumeration tricks, one can “effectively” get a $(1 - \varepsilon, 1 - \varepsilon)$ -balanced CR scheme for any fixed $\varepsilon > 0$.
- When A is k -sparse, we derive a $(b, 1 - 2kb)$ -balanced CR scheme. If A has in addition width $W \geq 2$, we get a $(b, 1 - k(2eb)^{W-1})$ CR scheme for any $b \in (0, 1)$. These results follow from [3].
- When A is a $\{0, 1\}$ -matrix induced by the problem of routing unit-demand paths in a capacitated path or tree, there is a $(b, 1 - O(b))$ CR scheme implicit in [4, 7, 12]. This can be extended to the unsplittable flow problem (UFP) in capacitated paths and trees via grouping and scaling techniques [21, 12, 11].

We discuss additional details of the preceding CR schemes in Section 4. We mention that several rounding schemes in the literature for packing problems, typically developed for modular functions, can be reinterpreted as CR schemes. Our framework then can be used to obtain algorithms for non-negative submodular set functions. See [8] for an illuminating example.

Approximation algorithms. Our CR schemes with suitable choices of parameters imply the following results for problems of the form $\max_{S \in \mathcal{I}} f(S)$, where f is non-negative submodular:

- If \mathcal{I} is the intersection of a constant number of knapsack constraints, we achieve a 0.325-approximation, improving the $(0.2 - \varepsilon)$ -approximation from [24] and a recent $(0.25 - \varepsilon)$ -approximation [23].
- If \mathcal{I} is the intersection of a k -uniform matchoid and ℓ knapsack constraints with ℓ a fixed constant, we obtain an $\Omega(\frac{1}{k})$ approximation (constant independent of ℓ), which improves the bound of $\Omega(\frac{1}{k+\ell})$ from [19]. We remark that this is a new result even for linear objective functions.
- If \mathcal{I} is the intersection of a k -uniform matchoid and an ℓ -sparse knapsack constraint system of width W , we give an $\Omega(\frac{1}{k+\ell^{1/W}})$ -approximation, improving $\Omega(\frac{1}{k\ell})$ from [19].
- We obtain a constant factor approximation for maximizing a non-negative submodular function of routed requests in a capacitated path or tree. An $O(1)$ approximation is known for modular functions [4, 7, 12, 11] but no prior approach that we are aware of could obtain a constant factor for non-monotone submodular functions.

3. SOLVING THE MULTILINEAR RELAXATION FOR NON-NEGATIVE SUBMODULAR FUNCTIONS

In this section, we address the question of solving the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ where F is the multilinear extension of a submodular function. As we already mentioned, due to [36, 6], there is a $(1 - 1/e)$ -approximation for the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ whenever F is the multilinear extension of a monotone submodular function and P is any solvable polytope. Here, we consider the maximization of a possibly *non-monotone submodular function* over a down-monotone solvable polytope.⁶ We assume

⁶As we noted in the introduction, [37] implies that there is no constant factor approximation for maximizing non-monotone submodular functions over general polytopes. The approximation that can be achieved for matroid base polytopes is proportional to $1 - 1/\nu$ where ν is the fractional packing number of bases, and in fact this trade-off generalizes to arbitrary solvable polytopes. We defer, to a full version, the details of the theorem one can prove in this direction.

in the following that $P \subseteq [0, 1]^N$ is a down-monotone solvable polytope and $F : [0, 1]^N \rightarrow \mathbb{R}_+$ is the multilinear extension of a submodular function. We present three algorithms for this problem.

3.1 Continuous local-search 0.25-approximation

First we consider the following natural local-search algorithm.

ALGORITHM 3.1. *Initialize $\mathbf{x} := 0$. As long as there is $\mathbf{y} \in P$ such that $(\mathbf{y} - \mathbf{x}) \cdot \nabla F(\mathbf{x}) > 0$ (which can be found by linear programming), move \mathbf{x} continuously in the direction $\mathbf{y} - \mathbf{x}$. If there is no such $\mathbf{y} \in P$, return \mathbf{x} .*

Naturally, a polynomial-time implementation of this algorithm would move in discrete steps and continue only as long as the improvements are sufficiently large. One possible implementation would maintain a poly-size convex combination $\mathbf{x} = \frac{1}{q} \sum_{i=1}^q \mathbf{v}_i$ where \mathbf{v}_i are certain vertices of P (with possible repetition). Each discrete step corresponds to replacing a vertex in the convex combination by another. This is in line with our intuitive description in Section 1.

We ignore the implementation details here. We assume that when the algorithm terminates, we have $(\mathbf{y} - \mathbf{x}) \cdot \nabla F(\mathbf{x}) \leq 0$ for every $\mathbf{y} \in P$. The basic lemma in the analysis of this algorithm as well as the improved algorithms is the following.

LEMMA 3.2. *For any two points $\mathbf{x}, \mathbf{y} \in [0, 1]^N$: $(\mathbf{y} - \mathbf{x}) \cdot \nabla F(\mathbf{x}) \geq F(\mathbf{x} \vee \mathbf{y}) + F(\mathbf{x} \wedge \mathbf{y}) - 2F(\mathbf{x})$.*

PROOF. By submodularity, F is concave along any line with a nonnegative direction vector, such as $(\mathbf{x} \vee \mathbf{y}) - \mathbf{x} \geq 0$. Therefore,

$$F(\mathbf{x} \vee \mathbf{y}) - F(\mathbf{x}) \leq ((\mathbf{x} \vee \mathbf{y}) - \mathbf{x}) \cdot \nabla F(\mathbf{x}), \text{ and similarly}$$

$$F(\mathbf{x} \wedge \mathbf{y}) - F(\mathbf{x}) \leq ((\mathbf{x} \wedge \mathbf{y}) - \mathbf{x}) \cdot \nabla F(\mathbf{x}),$$

because of the concavity of F along $(\mathbf{x} \wedge \mathbf{y}) - \mathbf{x} \leq 0$. Adding up these two inequalities, we get $F(\mathbf{x} \vee \mathbf{y}) + F(\mathbf{x} \wedge \mathbf{y}) - 2F(\mathbf{x}) \leq ((\mathbf{x} \vee \mathbf{y}) + (\mathbf{x} \wedge \mathbf{y}) - 2\mathbf{x}) \cdot \nabla F(\mathbf{x})$. It remains to observe that $(\mathbf{x} \vee \mathbf{y}) + (\mathbf{x} \wedge \mathbf{y}) = \mathbf{x} + \mathbf{y}$, which proves the lemma. \square

COROLLARY 3.3. *If \mathbf{x} is a local optimum such that $(\mathbf{y} - \mathbf{x}) \cdot \nabla F(\mathbf{x}) \leq 0$, then $2F(\mathbf{x}) \geq F(\mathbf{x} \vee \mathbf{y}) + F(\mathbf{x} \wedge \mathbf{y})$.*

Next, we show that if we combine this local optimum with a suitable “complementary solution”, we get a $1/4$ -approximation to the global optimum. The following is our algorithm.

ALGORITHM 3.4. *Using Algorithm 3.1, find a local optimum \mathbf{x} in P . Then define $Q = \{\mathbf{y} \in P : \mathbf{y} \leq 1 - \mathbf{x}\}$ and again using Algorithm 3.1, find a local optimum \mathbf{y} in Q . If $F(\mathbf{x}) \geq F(\mathbf{y})$ return \mathbf{x} , otherwise \mathbf{y} .*

We use the following property of the multilinear extension of a submodular function. Let us replace each coordinate by a $[0, 1]$ interval and let us represent a certain value x_i of the i 'th coordinate by a subset of $[0, 1]$ of the corresponding measure.

DEFINITION 3.5. *Let $\mathcal{X} \in \mathcal{L}^N$, where \mathcal{L} denotes the set of all measurable subsets of $[0, 1]$. We say that \mathcal{X} represents a vector $\mathbf{x} \in [0, 1]^N$, if \mathcal{X}_i has measure x_i for each $i \in N$.*

From a “discrete point of view”, we can imagine that each coordinate is replaced by some large number of elements M and a value of x_i is represented by any subset of size Mx_i . This can be carried out if all the vectors we work with are rational. In the following, we consider functions on subsets of this new ground set. We show a natural property, namely that a function derived from the multilinear extension of a submodular function is again submodular. (An analogous property in the discrete case was proved in [27, 24].)

LEMMA 3.6. Let $F : [0, 1]^N \rightarrow \mathbb{R}$ be a multilinear extension of a submodular function f . Define a function F^* on \mathcal{L}^N , by $F^*(\mathcal{X}) = F(\mathbf{x})$, where $\mathbf{x} \in [0, 1]^N$ is the vector represented by \mathcal{X} . Then F^* is submodular:

$$F^*(\mathcal{X} \cup \mathcal{Y}) + F^*(\mathcal{X} \cap \mathcal{Y}) \leq F^*(\mathcal{X}) + F^*(\mathcal{Y}),$$

where the union and intersection is interpreted component-wise.

PROOF. We have $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$ where $\hat{x}_i = 1$ independently with probability x_i . An equivalent way to generate $\hat{\mathbf{x}}$ is to choose any set $\mathcal{X} \in \mathcal{L}^N$ representing \mathbf{x} , generate uniformly and independently a number $r_i \in [0, 1]$ for each $i \in N$, and set $\hat{x}_i = 1$ iff $r_i \in \mathcal{X}_i$. Since the measure of \mathcal{X}_i is x_i , $\hat{x}_i = 1$ with probability exactly x_i . Therefore,

$$F^*(\mathcal{X}) = F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})] = \mathbf{E}[f(\{i : r_i \in \mathcal{X}_i\})].$$

Similarly,

$$F^*(\mathcal{Y}) = \mathbf{E}[f(\{i : r_i \in \mathcal{Y}_i\})].$$

This also holds for $\mathcal{X} \cup \mathcal{Y}$ and $\mathcal{X} \cap \mathcal{Y}$: since $(\mathcal{X} \cup \mathcal{Y})_i = \mathcal{X}_i \cup \mathcal{Y}_i$ and $(\mathcal{X} \cap \mathcal{Y})_i = \mathcal{X}_i \cap \mathcal{Y}_i$, we get

$$F^*(\mathcal{X} \cup \mathcal{Y}) = \mathbf{E}[f(\{i : r_i \in \mathcal{X}_i\} \cup \{i : r_i \in \mathcal{Y}_i\})]$$

and

$$F^*(\mathcal{X} \cap \mathcal{Y}) = \mathbf{E}[f(\{i : r_i \in \mathcal{X}_i\} \cap \{i : r_i \in \mathcal{Y}_i\})].$$

Hence, by the submodularity of f ,

$$\begin{aligned} & F^*(\mathcal{X} \cup \mathcal{Y}) + F^*(\mathcal{X} \cap \mathcal{Y}) \\ &= \mathbf{E}[f(\{i : r_i \in \mathcal{X}_i\} \cup \{i : r_i \in \mathcal{Y}_i\}) \\ &\quad + f(\{i : r_i \in \mathcal{X}_i\} \cap \{i : r_i \in \mathcal{Y}_i\})] \\ &\leq \mathbf{E}[f(\{i : r_i \in \mathcal{X}_i\}) + f(\{i : r_i \in \mathcal{Y}_i\})] \\ &= F^*(\mathcal{X}) + F^*(\mathcal{Y}). \end{aligned}$$

□

From here, we obtain our main lemma - the average of the two fractional local optima is at least $\frac{1}{4}OPT$.

LEMMA 3.7. Let $OPT = \max\{F(x) : x \in P\}$. Let \mathbf{x} be a local optimum in P , and \mathbf{y} a local optimum in $Q = \{\mathbf{y} \in P : \mathbf{y} \leq \mathbf{1} - \mathbf{x}\}$. Then $2F(\mathbf{x}) + 2F(\mathbf{y}) \geq OPT$.

PROOF. Let $OPT = F(\mathbf{z})$ where $\mathbf{z} \in P$. By Corollary 3.3, the local optimum $\mathbf{x} \in P$ satisfies

$$2F(\mathbf{x}) \geq F(\mathbf{x} \vee \mathbf{z}) + F(\mathbf{x} \wedge \mathbf{z}). \quad (1)$$

In the restricted polytope $Q = \{\mathbf{y} \in P : \mathbf{y} \leq \mathbf{1} - \mathbf{x}\}$, consider the point $\mathbf{z}' = (\mathbf{z} - \mathbf{x}) \vee 0 \in Q$. Again by Corollary 3.3, the local optimum $\mathbf{y} \in Q$ satisfies

$$2F(\mathbf{y}) \geq F(\mathbf{y} \vee \mathbf{z}') + F(\mathbf{y} \wedge \mathbf{z}'). \quad (2)$$

Now we use a representation of vectors by subsets as described in Def. 3.5. We choose $\mathcal{X}, \mathcal{Y}, \mathcal{Z}, \mathcal{Z}' \in \mathcal{L}^N$ to represent $\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{z}'$ as follows: for each $i \in N$, $\mathcal{X}_i = [0, x_i]$, $\mathcal{Y}_i = [x_i, x_i + y_i]$ (note that $x_i + y_i \leq 1$), $\mathcal{Z}_i = [0, z_i]$ and $\mathcal{Z}'_i = [0, z'_i] = [0, \max\{z_i - x_i, 0\}]$. Note that $(\mathcal{X} \cap \mathcal{Y})_i = \emptyset$ for $i \in N$.

Defining F^* as in Lemma 3.6, we have $F^*(\mathcal{X}) = F(\mathbf{x})$, $F^*(\mathcal{Y}) = F(\mathbf{y})$, $F^*(\mathcal{Z}) = F(\mathbf{z}) = OPT$ and $F^*(\mathcal{Z}') = F(\mathbf{z}')$. Using relations like $[0, x_i] \cup [0, z_i] = [0, \max\{x_i, z_i\}]$, we also get $F^*(\mathcal{X} \cup \mathcal{Z}) = F(\mathbf{x} \vee \mathbf{z})$ and $F^*(\mathcal{X} \cap \mathcal{Z}) = F(\mathbf{x} \wedge \mathbf{z})$. Furthermore, we have $(\mathcal{Z}_i \setminus \mathcal{X}_i) \cup \mathcal{Y}_i = [x_i, \max\{x_i + y_i, z_i\}] = [x_i, x_i + \max\{y_i, z'_i\}]$. This is an interval of length $\max\{y_i, z'_i\} = (\mathbf{y} \vee \mathbf{z}')_i$ and hence $F^*((\mathcal{Z} \setminus \mathcal{X}) \cup \mathcal{Y}) = F(\mathbf{y} \vee \mathbf{z}')$.

The property of the first local optimum (1) can be thus written as $2F(\mathbf{x}) \geq F^*(\mathcal{X} \cup \mathcal{Z}) + F^*(\mathcal{X} \cap \mathcal{Z})$. The property of the complementary local optimum (2) can be written as $2F(\mathbf{y}) \geq F^*((\mathcal{Z} \setminus \mathcal{X}) \cup \mathcal{Y})$ (we discarded the nonnegative term $F(\mathbf{y} \wedge \mathbf{z}')$ which is not useful in the following). By Lemma 3.6, F^* is submodular. Hence we get

$$\begin{aligned} F^*(\mathcal{X} \cap \mathcal{Z}) + F^*((\mathcal{Z} \setminus \mathcal{X}) \cup \mathcal{Y}) &\geq F^*((\mathcal{X} \cap \mathcal{Z}) \cup ((\mathcal{Z} \setminus \mathcal{X}) \cup \mathcal{Y})) \\ &= F^*(\mathcal{Z} \cup \mathcal{Y}) \end{aligned}$$

(we discarded the intersection term). Finally, using the fact that $\mathcal{X} \cap \mathcal{Y} = \emptyset$ and again the submodularity of F^* , we get

$$F^*(\mathcal{Z} \cup \mathcal{X}) + F^*(\mathcal{Z} \cup \mathcal{Y}) \geq F^*((\mathcal{Z} \cup \mathcal{X}) \cap (\mathcal{Z} \cup \mathcal{Y})) = F^*(\mathcal{Z})$$

(we discarded the union term). To summarize,

$$\begin{aligned} 2F(\mathbf{x}) + 2F(\mathbf{y}) &\geq F^*(\mathcal{X} \cap \mathcal{Z}) + F^*(\mathcal{X} \cup \mathcal{Z}) + F^*((\mathcal{Z} \setminus \mathcal{X}) \cup \mathcal{Y}) \\ &\geq F^*(\mathcal{Z}) = OPT. \end{aligned}$$

□

COROLLARY 3.8. For any down-monotone polytope $P \subseteq [0, 1]^N$ and multilinear extension of a submodular function $F : [0, 1]^N \rightarrow \mathbb{R}_+$, Algorithm 3.4 is a $\frac{1}{4}$ -approximation to the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$.

3.2 Restricted local-search 0.309-approximation

Next, we present a modified local-search algorithm which is a generalization of the algorithm for matroid polytopes from [37]. We remark that this algorithm is in fact simpler than the $\frac{1}{4}$ -approximation from the previous section, in that it does not require a second-stage complementary local search.

ALGORITHM 3.9. Initialize $\mathbf{x} := 0$ and fix a parameter $t \in [0, 1]$. As long as there is $\mathbf{y} \in P \cap [0, t]^N$ such that $(\mathbf{y} - \mathbf{x}) \cdot \nabla F(\mathbf{x}) > 0$ (which can be found by linear programming), move \mathbf{x} continuously in the direction $\mathbf{y} - \mathbf{x}$. Return \mathbf{x} .

This algorithm also works for any down-monotone polytope P . With the choice of $t = \frac{1}{2}(3 - \sqrt{5})$, it achieves a $\frac{1}{4}(-1 + \sqrt{5}) \simeq 0.309$ -approximation (with respect to the optimal 0-1 solution; we are not sure currently whether the analysis extends to optimal fractional solutions). We omit the analysis in this extended abstract.

3.3 Simulated annealing 0.325-approximation

Finally, we present the algorithm with the best ratio, based on the ideas of simulated annealing and the recent work of [30]. This algorithm can be seen as an extension of the 0.309-approximation, where local search is applied to a restricted polytope $P \cap [0, t]^N$. Here, we vary the "temperature parameter" t continuously from 0 to 1, while performing local search in the restricted polytope.

ALGORITHM 3.10. Initialize $\mathbf{x} := 0$ and $t := 0$. As long as $t \leq 1$, repeat the following:

1. Run a local search inside $P \cap [0, t]^N$, until $(\mathbf{y} - \mathbf{x}) \cdot \nabla F(\mathbf{x}) \leq 0$ for all $\mathbf{y} \in P \cap [0, t]^N$.
2. Generate λ uniformly at random in $[0, t]$, let $T_{>\lambda}(\mathbf{x}) = \{i : x_i > \lambda\}$ and $Q = \{\mathbf{z} \in P : \forall i \in T_{>\lambda}(\mathbf{x}); z_i = 0\}$. Initialize $\mathbf{z} := 0$ and run a local search inside Q , to find an auxiliary local optimum $\mathbf{z} \in Q$. Remember the best auxiliary local optimum, maximizing $F(\mathbf{z})$.
3. Find $\mathbf{y} \in P$ maximizing $\mathbf{y} \cdot \nabla F(\mathbf{x})$, modify $\mathbf{x} := \mathbf{x} + \frac{\delta}{1-t}(\mathbf{y} - \mathbf{x})$, $t := t + \delta$, and go to step 1.

Eventually, return the better of \mathbf{x} and the best auxiliary local optimum \mathbf{z} .

Note that the point \mathbf{x} evolves throughout the process, while the search for \mathbf{z} starts separately in each iteration. In Step 3, we look for a point \mathbf{y} in the full polytope P rather than the restricted polytope $P \cap [0, t]^N$. Since $\mathbf{x} \in [0, t]^N$ and $\mathbf{y} \in [0, 1]^N$, the modified point $\mathbf{x} + \frac{\delta}{1-t}(\mathbf{y} - \mathbf{x})$ has coordinates $x_i + \frac{\delta}{1-t}(y_i - x_i) \leq t + \delta$, and so it is contained in $P \cap [0, t + \delta]^N$. Coming back to Step 1, we continue local search from this point. We defer the analysis to the full version of this paper.

4. CONTENTION RESOLUTION AND THE CORRELATION GAP

In this section we highlight a close connection between CR schemes and a concept known as *correlation gap* [1], and discuss how to obtain an asymptotically optimal (b, c) -balanced CR scheme for matroids.

DEFINITION 4.1. For a set function $f : 2^N \rightarrow \mathbb{R}_+$, the correlation gap is defined as

$$\kappa(f) = \inf_{\mathbf{x} \in [0, 1]^N} \frac{\mathbf{E}[f(\hat{\mathbf{x}})]}{f^+(\mathbf{x})},$$

where $f^+(\mathbf{x}) = \max\{\sum_S \alpha_S f(S) : \sum_S \alpha_S \mathbf{1}_S = \mathbf{x}, \sum_S \alpha_S = 1, \alpha_S \geq 0\}$ is the maximum possible expectation of f over distributions of expectation \mathbf{x} , and $\hat{\mathbf{x}}$ is the product distribution with expectation \mathbf{x} . Furthermore, for a class of functions \mathcal{C} , the correlation gap $\kappa(\mathcal{C})$ is the infimum of correlation gaps over all functions in \mathcal{C} .

In other words, the correlation gap is the worst-case ratio between the multilinear extension $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$ and the concave closure $f^+(\mathbf{x})$. We define the correlation gap as a number $\kappa \in [0, 1]$, to be in line with the parameter c in our notion of a c -balanced CR scheme (the higher the better). The definition in [1] uses the inverse ratio.

Relation between CR schemes and the correlation gap: The relationship between CR schemes and correlation gap arises as follows. Let $\mathcal{I} \subseteq 2^N$ denote the set of feasible solutions. Consider a product distribution on 2^N with expectation $\mathbf{p} \in P_{\mathcal{I}}$, in other words a random set R which contains elements independently with probabilities p_i . Let Π be the family of all deterministic CR schemes π , i.e. ways to choose a subset $\pi(R) \subseteq R$ such that $\pi(R) \in \mathcal{I}$. (Although the scheme is deterministic, there is randomness here due to R .) Any randomized CR scheme can be written as a convex combination of such deterministic schemes; let us denote the coefficients by λ_π . Define $q_{i,\pi} = \Pr_R[i \in \pi(R)]$, the probability that element i is chosen in the scheme π . Hence, when executing a randomized CR scheme with coefficients λ , first with probability λ_π a deterministic CR scheme π is chosen, and then $\pi(R)$ is returned. The goal of our randomized scheme is to achieve the property that every element i appears in $\pi(R)$ with overall probability at least cp_i . Let us write down a linear program describing the optimal randomized CR scheme, and its dual.

$$\begin{aligned} \text{(LP1)} \quad & \max \quad c \\ & \text{s.t.} \quad \sum_{\pi \in \Pi} q_{i,\pi} \lambda_\pi \geq p_i c \quad \forall i \in N \\ & \quad \quad \sum_{\pi \in \Pi} \lambda_\pi = 1 \\ & \quad \quad \lambda_\pi \geq 0 \quad \forall \pi \in \Pi \\ \text{(DP1)} \quad & \min \quad \mu \\ & \text{s.t.} \quad \sum_{i \in N} q_{i,\pi} y_i \leq \mu \quad \forall \pi \in \Pi \\ & \quad \quad \sum_{i \in N} p_i y_i = 1 \\ & \quad \quad y_i \geq 0 \quad \forall i \in N \end{aligned}$$

We can interpret the dual as follows. Given an assignment to the variables y_i , the value of the dual is $\max_{\pi \in \Pi} \sum_i q_{i,\pi} y_i = \max_{\pi \in \Pi} \sum_i y_i \Pr_R[i \in \pi(R)] = \max_{\pi \in \Pi} \mathbf{E}_R[\sum_{i \in \pi(R)} y_i]$. Since π can choose an arbitrary feasible subset for each R , the optimal π is given by choosing for each R the maximum-weight subset $\pi(R)$ under the weights y_i , and the dual value is $\mathbf{E}_R[\max_{S \subseteq R, S \in \mathcal{I}} \sum_{i \in S} y_i]$. In words, this is the expected value one can extract from a random set R with marginals p_i , when the weights are normalized by $\sum_i p_i y_i = 1$. Minimizing over the choices of weights y_i , we obtain what we call the correlation gap of the solution set \mathcal{I} ,

DEFINITION 4.2. For $\mathcal{I} \subseteq 2^N$, we define the correlation gap as $\kappa(\mathcal{I}) = \inf_{\mathbf{p} \in P_{\mathcal{I}}, \mathbf{y} \geq 0} \frac{1}{\sum_i p_i y_i} \mathbf{E}[\max_{S \subseteq R, S \in \mathcal{I}} \sum_{i \in S} y_i]$, where R contains element i independently with probability p_i .

THEOREM 4.3. The correlation gap of \mathcal{I} is equal to the maximum c such that \mathcal{I} admits a c -balanced CR scheme.

PROOF. As discussed above, the correlation gap of \mathcal{I} is equal to the optimum value of the dual LP. By LP duality, this is equal to the optimum of the primal LP, which is the best value of c for which there is a c -balanced CR scheme. \square

The following lemma shows a close connection between the correlation gap of a solution set \mathcal{I} and the correlation gap of the respective rank function. More precisely, the correlation gap of \mathcal{I} corresponds to the worst (i.e. smallest) correlation gap of the respective rank function over all weight vectors.

LEMMA 4.4. For $\mathcal{I} \subseteq 2^N$ and weight vector $\mathbf{y} \geq 0$, let $r_{\mathbf{y}}(R) = \max_{S \subseteq R, S \in \mathcal{I}} \sum_{i \in S} y_i$ denote the associated weighted rank function. Then $\kappa(\mathcal{I}) = \inf_{\mathbf{y} \geq 0} \kappa(r_{\mathbf{y}})$.

PROOF. Using the notation $r_{\mathbf{y}}(R)$ for the weighted rank functions with weights \mathbf{y} , the correlation gap of \mathcal{I} can be rewritten as $\kappa(\mathcal{I}) = \inf_{\mathbf{p} \in P_{\mathcal{I}}, \mathbf{y} \geq 0} \frac{\mathbf{E}[r_{\mathbf{y}}(R)]}{\sum_i p_i y_i}$, where R contains elements independently with probabilities p_i . We first observe that for any $\mathbf{p} \in P_{\mathcal{I}}$, we have $r_{\mathbf{y}}^+(\mathbf{p}) = \sum_i p_i y_i$. Consider a convex combination $\mathbf{p} = \sum_{S \in \mathcal{I}} \alpha_S \mathbf{1}_S$, $\sum \alpha_S = 1$, $\alpha_S \geq 0$ with $r_{\mathbf{y}}^+(\mathbf{p}) = \sum_{S \in \mathcal{I}} \alpha_S y(S)$. Since the weighted rank function of a feasible set $S \in \mathcal{I}$ is simply its weight we obtain

$$r_{\mathbf{y}}^+(\mathbf{p}) = \sum_{S \in \mathcal{I}} \alpha_S y(S) = \mathbf{y} \cdot \sum_{S \in \mathcal{I}} \alpha_S \mathbf{1}_S = \mathbf{y} \cdot \mathbf{p} = \sum_i p_i y_i,$$

and hence $\kappa(\mathcal{I}) = \inf_{\mathbf{p} \in P_{\mathcal{I}}, \mathbf{y} \geq 0} \frac{\mathbf{E}[r_{\mathbf{y}}(R)]}{\sum_i p_i y_i} = \inf_{\mathbf{p} \in P_{\mathcal{I}}, \mathbf{y} \geq 0} \frac{\mathbf{E}[r_{\mathbf{y}}(R)]}{r_{\mathbf{y}}^+(\mathbf{p})}$.

To prove the claim it remains to show that

$$\inf_{\mathbf{p} \in P_{\mathcal{I}}, \mathbf{y} \geq 0} \frac{\mathbf{E}[r_{\mathbf{y}}(R)]}{r_{\mathbf{y}}^+(\mathbf{p})} = \inf_{\mathbf{p} \in [0, 1]^N, \mathbf{y} \geq 0} \frac{\mathbf{E}[r_{\mathbf{y}}(R)]}{r_{\mathbf{y}}^+(\mathbf{p})}. \quad (3)$$

Let $\mathbf{y} \geq 0$. We will prove (3) by showing that for any point $\mathbf{p} \in [0, 1]^N$ there is a point $\mathbf{p}' \in P_{\mathcal{I}}$ with $\mathbf{p}' \leq \mathbf{p}$ (coordinate-wise), and satisfying $r_{\mathbf{y}}^+(\mathbf{p}') \geq r_{\mathbf{y}}^+(\mathbf{p})$. Since $r_{\mathbf{y}}$ is monotone, we then obtain $\mathbf{E}[r_{\mathbf{y}}(R)]/r_{\mathbf{y}}^+(\mathbf{p}) \geq \mathbf{E}[r_{\mathbf{y}}(R)]/r_{\mathbf{y}}^+(\mathbf{p}')$, showing that the infimum over \mathbf{p} on the right-hand side of (3) can indeed be restricted to the polytope $P_{\mathcal{I}}$. Let $\mathbf{p} = \sum_{S \subseteq N} \alpha_S \mathbf{1}_S$, $\sum_{S \subseteq N} \alpha_S = 1$, $\alpha_S \geq 0$ be a convex combination of \mathbf{p} such that $r_{\mathbf{y}}^+(\mathbf{p}) = \sum_{S \subseteq N} \alpha_S r_{\mathbf{y}}(S)$. For every $S \subseteq N$, let $I(S) \subseteq S$ be a maximum weight independent set, hence $r_{\mathbf{y}}(S) = y(I(S))$. The point $\mathbf{p}' = \sum_{S \subseteq N} \alpha_S \mathbf{1}_{I(S)}$ clearly satisfies $\mathbf{p}' \leq \mathbf{p}$, and furthermore

$$\begin{aligned} r_{\mathbf{y}}^+(\mathbf{p}') & \geq \sum_{S \in \mathcal{I}} \left(\sum_{W \subseteq N, I(W)=S} \alpha_S \right) r_{\mathbf{y}}(S) = \sum_{S \subseteq N} \alpha_S r_{\mathbf{y}}(S) \\ & = r_{\mathbf{y}}^+(\mathbf{p}). \end{aligned}$$

\square

Monotonicity, efficiency and strictness: In the discussion above, we have ignored two issues: the monotonicity of our CR scheme, and the question whether we can find it efficiently. These issues can be also related to the concept of correlation gap, using LP duality.

If we want to obtain a monotone CR scheme, we can simply define Π to be the family of all deterministic monotone CR schemes. (It is not true that all monotone randomized CR schemes can be obtained as convex combinations of deterministic ones, but certainly this construction yields monotone randomized CR schemes.) LP duality implies that if there is a family of *monotone CR schemes* $\Pi' \subset \Pi$ that certifies that (DP1) is lower-bounded by c , i.e. for any weight $\mathbf{y} \geq 0$ with $\sum_{i \in N} p_i y_i = 1$ we have $\max_{\pi \in \Pi'} \sum_{i \in N} q_{i, \pi} y_i \geq c$, then there exists a monotone c -balanced CR scheme which is a convex combination of schemes in Π' . Rephrased in the context of correlation gaps, Π' is a family certifying that the correlation gap of \mathcal{I} is lower-bounded by c , since $\kappa(\mathcal{I}) \geq \inf_{\mathbf{p} \in P_{\mathcal{I}}, \mathbf{y} \geq 0} \frac{1}{\sum_i p_i y_i} \max_{\pi \in \Pi'} \mathbf{E}[y(\pi(R))] \geq c$ where the first inequality follows from $r_{\mathbf{y}}(R) \geq \mathbf{E}[y(\pi(R))]$ for all $\pi \in \Pi$.

Similarly, the question of efficiency translates into the dual as follows. If for each weight vector \mathbf{y} we have an efficient procedure to compute an efficient CR scheme $\pi : 2^N \rightarrow \mathcal{I}$ with $\sum_{i \in N} q_{i, \pi} y_i \geq c$, then we can use this procedure to approximately separate over the dual. This allows us to find efficiently a polynomial-sized collection of constraints that certify that the dual optimum is at least c . Hence, by solving (LP1) only over the variables corresponding to those constraints, a c -balanced CR scheme can be obtained efficiently. Notice that the thus obtained CR scheme is efficient, since it is a mixture of a polynomial number of efficient schemes. Without further details, we formulate these extensions in the following theorem.

THEOREM 4.5. *There is an efficient c -balanced CR scheme for \mathcal{I} iff there is an efficient algorithm which for any weight vector $\mathbf{y} \geq 0$, and a set R containing elements independently with probabilities $\mathbf{p} \in P(\mathcal{I})$, returns a feasible subset $\pi(R) \subseteq R$ such that $\mathbf{E}_R \left[\sum_{i \in \pi(R)} y_i \right] \geq c \sum_i y_i p_i$. In addition, if the algorithm can be chosen so that $\pi(R)$ is a monotone function of R (i.e., if $i \in R_1 \subseteq R_2$, and $i \in \pi(R_2)$, then $i \in \pi(R_1)$), then there is an efficient monotone c -balanced CR scheme for \mathcal{I} .*

Finally, we comment on the issue of strictness (i.e. obtaining an exact conditional probability $\Pr[i \in I \mid i \in R] = c$), which is needed in the case of non-monotone submodular functions. If we have a contention resolution scheme guaranteeing a lower bound $\Pr[i \in I \mid i \in R] \geq c$, then we can simulate this scheme for a given distribution of R and estimate the actual probability for each element, $c'_i = \Pr[i \in I \mid i \in R] \geq c$, within a polynomially small error (assuming that c is a constant). Then we can modify the contention resolution scheme by removing element i with probability $1 - c/c'_i$. The resulting scheme is arbitrarily close to being strict, and the approximation factor will not be affected significantly. We omit the details.

The above framework easily extends to (b, c) -balanced CR schemes by restricting \mathbf{p} to be in the scaled-down polytope $b \cdot P_{\mathcal{I}}$. In the following, we discuss how for any fixed $b > 0$ an asymptotically optimal $(b, \frac{1-e^{-b}}{b})$ -balanced and monotone CR schemes for matroids can be obtained using the above approach.

4.1 Contention resolution for matroids

Let $\mathcal{M} = (N, \mathcal{I})$ be a matroid, $b \in (0, 1]$, and let $\mathbf{x} \in P_{\mathcal{I}}$ be the given point for which we want to find a (b, c) -balanced CR scheme for c as large as possible. We denote by $\mathbf{p} = b \cdot \mathbf{x}$ the scaled-down point, which puts us notation-wise in the same setting as discussed

above. Let $R(b)$ be the random set including each element $i \in N$ independently with probability $p_i = bx_i$. Consider the separation problem for (DP1), which asks for a given weight vector $\mathbf{y} \geq 0$ with $\sum_i p_i y_i = 1$ and some μ , whether $\max_{\pi \in \Pi} \sum_i q_{i, \pi} y_i \leq \mu$. As discussed above, the maximum is achieved for any CR scheme π that returns for any set $R(b)$ a maximum weight subset with respect to \mathbf{y} . However, in the case of matroids, such a CR scheme π corresponds exactly to the greedy algorithm $\pi_{\mathbf{y}}$ for finding a maximum weight independent set with respect to the weights \mathbf{y} . Hence, to separate over the dual, it suffices to compute $q_{i, \pi_{\mathbf{y}}}$ for $i \in N$ and check whether $\sum_{i \in N} q_{i, \pi_{\mathbf{y}}} y_i \leq \mu$. Using sample average approximations we can, for any $\varepsilon > 0$, check with high probability whether $\sum_{i \in N} q_{i, \pi_{\mathbf{y}}} y_i \leq \mu + \varepsilon$ in time polynomial in the input and $1/\varepsilon^7$. Using this approximate separation oracle for the dual we get the following result due to the ellipsoid method, where we get rid of the “with high probability” statement by absorbing the small probability of an unsuccessful estimate in the ε of the claimed $(b, c - \varepsilon)$ -balanced CR scheme.

THEOREM 4.6. *For any $\varepsilon > 0$ and any matroid \mathcal{M} that admits a (b, c) -balanced CR scheme, we can obtain a $(b, c - \varepsilon)$ -balanced and monotone CR scheme for \mathcal{M} running in time polynomial in the input and $1/\varepsilon$.*

A consequence of the fact that it is sufficient to consider greedy algorithms for dual separation, is that all constraints in the dual, that do not correspond to greedy algorithms, are redundant. Hence, for the case of matroids, convex combinations of greedy algorithms lead to the strongest CR schemes. Since all greedy CR schemes $\pi_{\mathbf{y}}$ are monotone, this implies that only considering monotone CR schemes is not restrictive in the case of matroids.

To convert Theorem 4.6 into a concrete statement about the value c , it suffices to prove the existence of a good (b, c) -balanced CR scheme. The existence of a $(1 - 1/e)$ -balanced CR scheme for matroids follows by the fact that the correlation gap of monotone submodular functions is $1 - 1/e$ [5]; by Lemma 4.4 this implies that the correlation gap of the independent sets of any matroid is bounded by $1 - 1/e$, and the result follows by applying Theorem 4.3. The result about the correlation gap of monotone submodular functions can be refined to obtain the following statement about the existence of (b, c) -balanced CR schemes for matroids.

THEOREM 4.7. *For any matroid \mathcal{M} on n elements, $b \in (0, 1]$, and $\mathbf{x} \in P_{\mathcal{I}}$, there exists a $(b, \frac{1-e^{-b} + \Omega(\frac{1}{\text{poly}(n)})}{b})$ -balanced CR scheme.*

By combining Theorem 4.6 and 4.7, and choosing $\varepsilon = O(\frac{1}{b \cdot \text{poly}(n)})$, we obtain our main result for CR schemes in the context of matroids.

COROLLARY 4.8. *For any matroid \mathcal{M} , $b \in (0, 1]$, and $\mathbf{x} \in P_{\mathcal{I}}$, we can efficiently construct an efficient $(b, \frac{1-e^{-b}}{b})$ -balanced and monotone CR scheme.*

As shown by the following theorem, the CR schemes that can be obtained according to Corollary 4.8 are, up to an additive ε , asymptotically optimal.

THEOREM 4.9. *For any $b \in (0, 1]$ and $\varepsilon > 0$, there is no $(b, \frac{1-e^{-b}}{b} + \varepsilon)$ -balanced CR scheme for uniform matroids of rank one.*

⁷Exact computation of the $q_{i, \pi_{\mathbf{y}}}$ can be shown to be $\#P$ -hard even for graphic matroids by a reduction from the s - t reliability problem.

4.2 Contention resolution for knapsacks

Here we sketch a contention resolution scheme for knapsack constraints. This essentially follows from known techniques; we remark that Kulik, Shachnai and Tamir [23] showed how to round a fractional solution to the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ for any constant number of knapsack constraints and any non-negative submodular function, while losing a $(1 - \varepsilon)$ factor for an arbitrarily small $\varepsilon > 0$. Our goal is to show that these techniques can be implemented in a black-box fashion and integrated in our framework. We prove the following lemma.

LEMMA 4.10. *For any $\delta, \varepsilon > 0$ and a knapsack constraint $\mathcal{F} = \{S : \sum_{i \in S} a_i \leq 1\}$ such that $a_i \leq \delta$ for all i , there is a monotone $(1 - \varepsilon, 1 - e^{-\Omega(\varepsilon^2/\delta)})$ -balanced contention resolution scheme.*

This contention resolution scheme is directly applicable only if the item sizes are relatively small compared to the knapsack capacity. However, standard enumeration tricks allow us to apply this scheme to general instances as well. This can be done for any constant number of knapsack constraints. We formulate this as follows.

COROLLARY 4.11. *For any constant $k \geq 1$ and $\varepsilon > 0$, there is a constant n_0 (that depends on k, ε) such that for any submodular maximization instance involving k knapsack constraints (and possibly other constraints), there is a set F of at most n_0 elements and a residual instance on the remaining elements such that*

- Any α -approximate solution to the residual instance together with F is an $\alpha(1 - k\varepsilon)$ -approximate solution to the original instance.
- In the residual instance, each knapsack constraint admits a $(1 - \varepsilon, 1 - \varepsilon)$ -balanced CR scheme.

An advantage of this black box approach is that knapsack constraints can be combined arbitrarily with other types of constraints. They do not affect the approximation ratio significantly. However, the enumeration stage affects the running time by an $O(n^{n_0})$ factor.

4.3 Sparse packing systems

We now consider packing constraints of the type $A\mathbf{x} \leq \mathbf{b}$, where $\mathbf{x} \in \{0, 1\}^N$ is the indicator vector of a solution. We can assume without loss of generality that the right-hand side is $\mathbf{b} = \mathbf{1}$. We say that the system is k -sparse, if each column of A has at most k nonzero entries (i.e., each element participates in at most k linear constraints). The approximation algorithms in [3] can be seen to give a contention resolution scheme for k -sparse packing systems.

CR scheme for k -sparse packing systems:

- We say that element j participates in constraint i , if $a_{ij} > 0$. We call an element j *big* for this constraint, if $a_{ij} > 1/2$. Otherwise we call element j *small* for this constraint.
- Sample R with probabilities bx_i .
- For each constraint i : if there is exactly one big element in R that participates in i , mark all the small elements in R for this constraint for deletion; otherwise check whether $\sum_{j \in R} a_{ij} > 1$ and if so, mark all elements participating in i for deletion.
- Define I to be R minus the elements marked for deletion.

Based on the analysis in [3], we obtain the following.

LEMMA 4.12. *For any $b \in (0, \frac{1}{2k})$, the above is a monotone $(b, 1 - 2kb)$ -balanced CR scheme for k -sparse packing systems.*

Recall the notion of width for a packing system: $W = \lfloor \frac{1}{\max_{i,j} a_{ij}} \rfloor$, where a_{ij} are the entries of the packing matrix (recall that we normalize the right-hand side to be $\mathbf{b} = \mathbf{1}$). Assuming that $W \geq 2$, one can use a simpler CR scheme and improve the parameters.

CR scheme for k -sparse packing systems of width W :

- Sample R with probabilities bx_i .
- For each constraint i for which $\sum_{j \in R} a_{ij} > 1$, mark all elements participating in i for deletion.
- Define I to be R minus the elements marked for deletion.

LEMMA 4.13. *For any $b \in (0, \frac{1}{2e})$, the above is a monotone $(b, 1 - k(2eb)^{W-1})$ -balanced CR scheme for any k -sparse system of packing constraints of width $W \geq 2$.*

4.4 UFP in paths and trees

We consider the following routing/packing problem. Let $T = (V, E)$ be a capacitated tree with u_e denoting the capacity of edge $e \in E$. We are given k node pairs $s_1 t_1, \dots, s_k t_k$ with pair i having a non-negative demand d_i ; we assume $d_{\max} = \max_i d_i \leq u_{\min} = \min_e u_e$ (the no-bottleneck assumption). Let $N = \{1, \dots, k\}$. We say that $S \subseteq N$ is routable if for each $i \in S$ a demand d_i is routed along the unique path from s_i to t_i , and the total flow on any edge e is at most u_e . Previously an $O(1)$ -approximation has been given for the problem of finding a maximum weight subset of routable demands [12]; the problem is APX-hard even for unit-demands and unit-weights. Let $\mathcal{I} = \{S \subseteq N \mid S \text{ is routable}\}$. Here we consider $\max_{S \in \mathcal{I}} f(S)$ for a non-negative submodular function f . A natural (packing) LP relaxation for $P_{\mathcal{I}}$ has a variable $x_i \in [0, 1]$ for each pair i and a constraint $\sum_{i: e \in Q_i} d_i x_i \leq u_e$ for each edge e where Q_i is the unique s_i - t_i path in T .

CR scheme for unit-demands:

- Root T arbitrarily. Let depth of pair $s_i t_i$ be the depth of the least common ancestor of s_i and t_i in T .
- Sample R with probabilities bx_i . Let $I = \emptyset$.
- Consider pairs in R in increasing order of depth.
- Add i to I if $I \cup \{i\}$ is routable, otherwise reject i .
- Output I .

The techniques in [7, 12] give the following lemma.

LEMMA 4.14. *There is an absolute constant ρ such that for any $b \in (0, 1)$ the above is a $(b, 1 - \rho b)$ -balanced CR scheme.*

CR scheme for general demands: A CR scheme for general demands can be obtained as follows. The linear program $P_{\mathcal{I}}$ is a packing LP of the form $A\mathbf{x} \leq \mathbf{b}$, $\mathbf{x} \in [0, 1]$ where A is column-restricted (all the non-zero values in a column have the same value). For such column-restricted packing integer programs (CPIPs), when demands satisfy the no-bottleneck assumption, one can use grouping and scaling techniques first suggested by Kolliopoulos and Stein [21] (see also [12]) to show that the integrality gap for a CPIP A is at most a fixed constant factor worse than that of the underlying 0-1 matrix A' (obtained from A by placing a 1 in each non-zero entry). Note that in the context of the UFP problem, the matrix A corresponds to the problem with arbitrary demands while the matrix A' corresponds to the one with unit-demands. One can use the same grouping and scaling techniques to show that a $(b, 1 - b')$ CR scheme for A' can be used to obtain a $(b, 1 - \rho b')$ CR scheme for A where ρ is an absolute constant. Using this general conversion theorem and Lemma 4.14, one can obtain a $(b, 1 - O(b))$ -balanced CR scheme for UFP in trees for any $b > 0$.

Without the no-bottleneck assumption the linear program has an $\Omega(n)$ integrality gap even for UFP on paths [7]. One can still apply the grouping and scaling techniques without the no-bottleneck assumption under a mild restriction; we refer the reader to [11].

Acknowledgments: We thank Mohit Singh for helpful discussions on contention resolution schemes for matroids.

5. REFERENCES

- [1] S. Agrawal, Y. Ding, A. Saberi and Y. Ye. Correlation robust stochastic optimization. *Proc. of 21st ACM-SIAM SODA*, 1087–1096, 2010.
- [2] N. Alon and J. Spencer. *The Probabilistic Method*. 2nd Edition, John Wiley & Sons, 2000.
- [3] N. Bansal, N. Korula, V. Nagarajan and A. Srinivasan. On k -column sparse packing programs. *Proc. of 14th IPCO*, 369–382, 2010.
- [4] G. Calinescu, A. Chakrabarti, H. Karloff and Y. Rabani. Improved approximation algorithms for resource allocation. In *Proc. of 8th IPCO*, 439–456, 2001.
- [5] G. Calinescu, C. Chekuri, M. Pál and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. *Proc. of 12th IPCO*, 182–196, 2007.
- [6] G. Calinescu, C. Chekuri, M. Pál and J. Vondrák. Maximizing a monotone submodular set function subject to a matroid constraint. To appear in *SIAM J. on Computing* special issue for *ACM STOC 2008*.
- [7] A. Chakrabarti, C. Chekuri, A. Gupta and A. Kumar. Approximation algorithms for the unsplittable flow problem. *Algorithmica*, 47(1):53–78, 2007.
- [8] T. Chan and S. Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. Full version of a paper that appeared in *Proc. of ACM SoCG*, 2009. Available at <http://arxiv.org/abs/1103.1431>.
- [9] S. Chawla, J. Hartline, D. Malec and B. Sivan. Sequential posted pricing and multi-parameter mechanism design. Manuscript, 2009, available at <http://arxiv.org/abs/0907.2435v1>.
- [10] S. Chawla, J. Hartline, D. Malec and B. Sivan. Sequential posted pricing and multi-parameter mechanism design. *Proc. of 42nd ACM STOC*, 311–320, 2010.
- [11] C. Chekuri, A. Ene and N. Korula. UFP in paths and trees and column-restricted packing integer programs. *Proc. of 12th APPROX*, 42–55, 2009.
- [12] C. Chekuri, M. Mydlarz and F. B. Shepherd. Multicommodity demand flow in a tree and packing integer programs. *ACM Trans. on Algorithms*, to appear. Preliminary version in *Proc. of ICALP*, 410–425, 2003.
- [13] C. Chekuri, J. Vondrák and R. Zenklus. Dependent randomized rounding via exchange properties of combinatorial structures. *Proc. of 51st IEEE FOCS*, 575–584, 2010.
- [14] G. Cornuejols, M. L. Fisher and G. Nemhauser. Location of bank accounts to optimize float: An analytic study of exact and approximate algorithms. *Management Science*, 23:789–810, 1977.
- [15] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [16] U. Feige, V. Mirrokni and J. Vondrák. Maximizing a non-monotone submodular function. *Proc. of 48th IEEE FOCS*, 461–471, 2007.
- [17] U. Feige and J. Vondrák. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. *Proc. of 47th IEEE FOCS*, 667–676, 2006.
- [18] M. L. Fisher, G. L. Nemhauser and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions - II. *Math. Prog. Study*, 8:73–87, 1978.
- [19] A. Gupta, V. Nagarajan and R. Ravi. Thresholded covering algorithms for robust and max-min optimization. *Proc. of 37th ICALP*, 262–274, 2010. A more detailed version available at <http://arxiv.org/abs/1003.0722>.
- [20] A. Gupta, A. Roth, G. Schoenebeck and K. Talwar. Constrained non-monotone submodular maximization: Offline and secretary algorithms. *CoRR*, abs/1003.1517, 2010.
- [21] S. G. Kolliopoulos and C. Stein. Approximating disjoint-path problems using greedy algorithms and Packing Integer Programs. *Math. Programming A*, (99):63–87, 2004. Preliminary version in *Proc. of IPCO*, 1998.
- [22] A. Kulik, H. Shachnai and T. Tamir. Maximizing submodular set functions subject to multiple linear constraints. *Proc. of 20th ACM-SIAM SODA*, 545–554, 2009.
- [23] A. Kulik, H. Shachnai and T. Tamir. Approximations for monotone and non-monotone submodular maximization with knapsack constraints. Manuscript, 2010.
- [24] J. Lee, V. Mirrokni, V. Nagarajan and M. Sviridenko. Maximizing nonmonotone submodular functions under matroid and knapsack constraints. *SIAM J. on Disc. Math.*, 23(4): 2053–2078, 2010. Preliminary version in *Proc. of 41st ACM STOC*, 323–332, 2009.
- [25] J. Lee, M. Sviridenko and J. Vondrák. Submodular maximization over multiple matroids via generalized exchange properties. *Proc. of 12th APPROX*, 244–257, 2009.
- [26] J. Lee, M. Sviridenko and J. Vondrák. Matroid matching: the power of local search. *Proc. of 42th ACM STOC*, 369–378, 2010.
- [27] V. Mirrokni, M. Schapira and J. Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. *Proc. of 9th ACM EC*, 70–77, 2008.
- [28] G. L. Nemhauser and L. A. Wolsey. Best algorithms for approximating the maximum of a submodular set function. *Math. Oper. Research*, 3(3):177–188, 1978.
- [29] G. L. Nemhauser, L. A. Wolsey and M. L. Fisher. An analysis of approximations for maximizing submodular set functions - I. *Math. Prog.*, 14:265–294, 1978.
- [30] S. Oveis Gharan and J. Vondrák. Submodular maximization by simulated annealing. *Proc. of 22nd ACM-SIAM SODA*, 1098–1116, 2011.
- [31] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
- [32] A. Srinivasan. New approaches to covering and packing problems. *Proc. of 12th ACM-SIAM SODA*, 567–576, 2001.
- [33] M. Sviridenko. A note on maximizing a submodular set function subject to knapsack constraint. *Operations Research Letters* 32, 41–43, 2004.
- [34] L.G. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing* 8(3):410–421, 1979.
- [35] J. Vondrák. Submodularity in combinatorial optimization. *PhD Thesis*, Charles University, 2007.
- [36] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. *Proc. of 40th ACM STOC*, 67–74, 2008.
- [37] J. Vondrák. Symmetry and approximability of submodular maximization problems. *Proc. of 50th IEEE FOCS*, 251–270, 2009.
- [38] J. Vondrák. A note on concentration of submodular functions. Manuscript, 2010, available at <http://arxiv.org/abs/1005.2791>.
- [39] Q. Yan. Mechanism design via correlation gap. *Proc. of 22th ACM-SIAM SODA*, 710–719, 2011.