Improved Approximations for k-Exchange Systems

(Extended Abstract)

Moran Feldman[1], Joseph (Seffi) Naor[1], Roy Schwartz[1], and Justin Ward[2]

[1] Computer Science Dept., Technion, Haifa, Israel
{moranfe,naor,schwartz}@cs.technion.ac.il
[2] Computer Science Dept., University of Toronto, Toronto, Canada
jward@cs.toronto.edu

Abstract. Submodular maximization and set systems play a major role in combinatorial optimization. It is long known that the greedy algorithm provides a $1/(k+1)$-approximation for maximizing a monotone submodular function over a k-system. For the special case of k-matroid intersection, a local search approach was recently shown to provide an improved approximation of $1/(k+\delta)$ for arbitrary $\delta > 0$. Unfortunately, many fundamental optimization problems are represented by a k-system which is not a k-intersection. An interesting question is whether the local search approach can be extended to include such problems.

We answer this question affirmatively. Motivated by the *b-matching* and *k-set packing* problems, as well as the more general *matroid k-parity* problem, we introduce a new class of set systems called k-exchange systems, that includes k-set packing, b-matching, matroid k-parity in strongly base orderable matroids, and additional combinatorial optimization problems such as: *independent set in* $(k+1)$*-claw free graphs*, *asymmetric TSP*, *job interval selection with identical lengths* and *frequency allocation on lines*. We give a natural local search algorithm which improves upon the current greedy approximation, for this new class of independence systems. Unlike known local search algorithms for similar problems, we use counting arguments to bound the performance of our algorithm.

Moreover, we consider additional objective functions and provide improved approximations for them as well. In the case of linear objective functions, we give a non-oblivious local search algorithm, that improves upon existing local search approaches for matroid k-parity.

1 Introduction

The study of combinatorial problems with submodular objective functions has attracted much attention recently, and is motivated by the principle of economy of scale, prevalent in real world applications. Additionally, submodular maximization plays a major role in combinatorial optimization since many optimization problems can be represented as constrained variants of submodular maximization. Often, the feasibility domain of such a variant is defined by a set system. A set system $(\mathcal{N}, \mathcal{I})$ is composed of a ground set $\mathcal{N}$ and a collection $\mathcal{I} \subseteq 2^{\mathcal{N}}$ of independent sets. For a constrained problem defined by $(\mathcal{N}, \mathcal{I})$, $\mathcal{I}$ is the collection of feasible solutions. Here is the known hierarchy of set systems:

$$k\text{-intersection} \subseteq k\text{-circuit bound} \subseteq k\text{-extendible} \subseteq k\text{-system}.$$

A set system belongs to the k-intersection class if it is the intersection of k matroids defined over a common ground set. The class of k-circuit bound contains all set systems in which adding a single element to an independent set creates at most k circuits (i.e., the resulting set contains at most k minimally dependent subsets). Recall that in a matroid, adding an element to an independent set creates at most a single circuit. Therefore, adding an element to an independent set in a k-intersection system closes at most k circuits, one per matroid. The class of k-extendible, intuitively, captures all set systems in which adding an element to an independent set requires throwing away at most k other elements from the set (in order to keep it independent). This generalizes k-circuit bound because in k-circuit bound we need to throw at most one element per circuit closed (i.e., up to k elements). The class of k-system contains all set systems in which for every set, not necessarily independent, the ratio of the sizes of the largest base of the set to the smallest base of the set is at most k (a base is a maximal independent subset).

Motivated by well studied problems such as *matroid k-parity* (which generalizes *matroid k-intersection*, *k-set packing*, *b-matching* and *k-dimensional matching*) as well as *independent set in $(k+1)$-claw free graphs*, we propose a new class of set systems which we call k-exchange systems. This class is general enough to capture various well studied combinatorial optimization problems, including matroid k-parity in strongly base orderable matroids, the other problems listed above, and additional problems such as: *job interval selection with identical lengths*, *asymmetric traveling salesperson* and *frequency allocation on lines*. Note that these last 3 problems, like independent set in $(k+1)$-claw free graphs, do *not* belong to the k-intersection class. On the other hand, we show that this class has a rich enough structure to enable us to present two local search algorithms with provable improved approximation guarantees for submodular and linear functions, respectively. Additionally, we relate the k-exchange class to the notion of strongly base orderable matroids, and show how it relates to the existing set systems hierarchy.

1.1 Our Results

Given a k-exchange system $(\mathcal{N}, \mathcal{I})$ and a function $f : 2^{\mathcal{N}} \to \mathbb{R}^+$, we provide approximation guarantees for the problem of finding an independent set $S \in \mathcal{I}$ maximizing $f(S)$ for several types of objective functions. The main application we consider is *strongly base orderable matroid k-parity*. This problem has two important special cases: *b-matching* and *k-set packing*.

Many interesting applications are k-exchange systems for *small* values of k, For example, the well studied *b-matching* problem is 2-exchange (but not 2-intersection) system. This lets us improve the best approximation of $1/3$ by

[15] to about $1/2$ for the case of a normalized monotone submodular f, and the best approximation of 0.0657 by [17] to about $1/4$ for the case of a general non-negative (not necessarily monotone) submodular f.

The types of objective functions considered in this work are: normalized monotone submodular, general non-negative (not necessarily monotone) submodular, linear and cardinality. Table 1 summarizes these results and the approximation ratio achieved for each application considered in this work.

We present 2 local search algorithms for maximizing submodular and linear objectives, respectively, subject to k-exchange systems. Our first algorithm is a very natural local search algorithm which is essentially identical to the algorithm of Lee et al. [26]. However, unlike the analysis of [26] which uses matroid intersection techniques, our analysis goes through a counting argument applied to an auxiliary graph. Our second algorithm yields improved results for the special case of linear objective functions, and is based on non-oblivious local search techniques employed by Berman [2] for the case of $(k+1)$-claw free graphs. This algorithm is guided by an auxiliary potential function, which considers the sum of the *squared* weights of the elements in the independent set.

It should be noted that as in previous local search algorithms, e.g., [26], the time complexity of our algorithms is exponential in k, thus, k is assumed to be a constant. As mentioned, k is indeed a small constant in our applications (refer to Table 1 for exact values of k).

1.2 Related Work

Extensive work has been conducted in recent years in the area of optimizing submodular functions under various constraints. We mention here the most relevant results. Historically, one of the very first problems examined was maximizing a monotone submodular function over a matroid. Several special cases of matroids and submodular functions were studied in [10, 18, 19, 22, 25] using the greedy approach. Recently, the general problem with an arbitrary matroid and an arbitrary submodular function was given a tight approximation of $(1 - 1/e)$ by Calinescu et al. [7]. A matching lower bound is due to [31, 32].

The problem of optimizing a normalized, *monotone* submodular function over the intersection of k matroids was considered by Fisher et al. [15] who gave a greedy algorithm with an approximation factor of $1/(k + 1)$, and state that their proof extends to the more general class of k-systems using the outline of Jenkyns [22] (the extended proof is explicitly given by Calinescu et al. [7]). For k-intersection systems, this result was improved by Lee et al. [26] to $1/(k+\delta)$, for any constant $\delta > 0$, while using a local search approach that exploits exchange properties of the underlying combinatorial structure. However, for optimizing a monotone submodular function over k-circuit bound and k-extendible set systems, the current best known approximation is still $1/(k + 1)$ [15].

Maximization of *non-monotone* submodular functions under various constraints has also attracted considerable attention in the last few years. The basic result in this area is an approximation factor of $2/5$, given by Feige et al. [12], for the unconstrained variant of the problem. This was recently improved twice,

Table 1. $k \geq 2$ is a constant, $\delta > 0$ is any given constant and $\beta = 1/(\alpha^{-1}+1)$, where α is the best known approximation for unconstrained maximization of a non-negative submodular function ($\alpha \geq 0.42$ see [13]).
f: NMS - normalized monotone submodular, NS - general non-negative submodular, L - linear, C - cardinality.

Maximization Problem	f	k	**This Paper**	**Previous Result**
k-exchange	NMS	k	$1/(k+\delta)$	$1/(k+1)$ [15]
	NS		$(k-1)/(k^2+\delta)$	$\beta/(k+2+1/k)$ [17]
	L[b]		$2/(k+1+\delta)$	$1/k$ [22]
	C[a]		$2/(k+\delta)$	$1/k$ [22]
Main Applications				
s.b.o.matroid k-parity	NMS	k	$1/(k+\delta)$	$1/k$ [15]
	NS		$(k-1)/(k^2+\delta)$	$\beta/(k+2+1/k)$ [17]
	L[b]		$2/(k+1+\delta)$	$1/k$ [22]
b-matching	NMS	2	$1/(2+\delta)$	$1/3$ [15]
	NS		$1/(4+\delta)$	$\beta/4.5$ [17]
k-set packing	NMS	k	$1/(k+\delta)$	$1/(k+1)$ [15]
	NS		$(k-1)/(k^2+\delta)$	$\beta/(k+2+1/k)$ [17]
Additional Applications				
independent set in $(k+1)$-claw free graphs	NMS	k	$1/(k+1+\delta)$	$1/k$ [15]
	NS		$(k-1)/(k^2+\delta)$	$\beta/(k+2+1/k)$ [17]
job interval selection identical lengths	NMS	2	$1/(2+\delta)$	$1/3$ [30]
	NS	3	$2/(9+\delta)$	$3\beta/16$ [17]
asymmetric traveling salesperson	NMS	3	$1/(3+\delta)$	$1/4$ [30]
	NS		$2/(9+\delta)$	$3\beta/16$ [17]
frequency allocation on lines	NMS	3	$1/(3+\delta)$	$1/4$ [15]
	NS		$2/(9+\delta)$	$3\beta/16$ [17]
	L		$1/(2+\delta)$	$1/3$ [22]
	C		$2/(3+\delta)$	$1-1/e$ [35]

[a] The result applies for $k \geq 3$.
[b] For $k = 2$, we also have a PTAS (see Corollary 1).

using a generalization of local search, called simulated annealing, by Gharan and Vondrák [16] and then by Feldman et al. [13]. Chekuri et al. [9] gave a 0.325 approximation for optimization over a matroid. This was very recently improved to roughly $e^{-1} \approx 0.368$ by Feldman et al. [14]. When optimizing over the intersection of k matroids, the current best result is $(k-1)/(k^2+\delta)$, and is due to Lee et al. [26]. A technique for using monotone submodular optimization for non-monotone submodular problems is given by [17]. Gupta et al. [17] use their technique to convert the greedy algorithm into an algorithm achieving an approximation ratio of $1/((\alpha^{-1}+1)(k+2+1/k))$ for k-system, where α is the best know approximation for unconstrained maximization of a non-negative submodular function ($\alpha \geq 0.42$, see [13]).

In the case of maximizing a *linear* objective function over k matroid constraints, an approximation of $1/k$ was given by Jenkyns [22] using a greedy algorithm. This was improved by [26] who gave an approximation of $1/(k-1+\delta)$, for any constant $\delta > 0$, using the same local search techniques as in the monotone submodular case. For the more general k-circuit bounded and k-extendible set systems, the current best known approximation is only $1/k$ [22] (as in the monotone submodular case this result is given for k-system). Hazan et al. [20] give a hardness result of $\Omega(\log k/k)$ that applies to this case.

Among the applications we consider, the most general is matroid k-parity in strongly base orderable matroids. The matroid k-parity problem, described in detail in Definition 3, is related to the matroid k-matching, which is a common generalization of matching and matroid intersection problems. In this problem, we are given a k-uniform hypergraph $H = (V, \mathcal{E})$ and a matroid $\mathcal{M}$ defined on the vertex set V of H. The goal is to find a matching $\mathcal{S}$ in H such that the set of elements covered by the edges in $\mathcal{S}$ are independent in $\mathcal{M}$. The matroid k-parity problem corresponds to the special case in which the edges of H are disjoint. It can be shown that matroid matching in a k-uniform hypergraph is reducible to matroid k-parity as well, and thus the two problems are equivalent [27].

If the matroid $\mathcal{M}$ is given by an independence oracle, there are instances of matroid matching problem (and hence also matroid parity) for which obtaining an optimal solution requires an exponential number of oracle calls, even when $k = 2$ and all weights are 1 [28, 23]. These instances can be modified to show that matroid parity is NP-complete (via a reduction from MaxClique) [34]. In the unweighted case, Lovász [28] obtained a polynomial time algorithm for matroid 2-matching in *linear* matroids. More recently, Lee et al. [27] gave a PTAS for matroid 2-parity in arbitrary matroids, and a $k/2 + \epsilon$ approximation for matroid k-parity in arbitrary matroids.

In the weighted case, it can be shown (see [7]) that the greedy algorithm provides a k-approximation. Although this remains the best known result for general matroids, some improvement has been made in the case of $k = 2$ for restricted classes of matroids. Tong et. al give an exact polynomial time algorithm for weighted matroid 2-parity in gammoids [37]. This result has recently been extended by Soto [36] to a PTAS for the class of all *strongly base orderable* matroids, which strictly includes gammoids. Additionally, Soto shows that matroid 2-matching remains NP-hard even in this restricted case.

An important special case of matroid 2-parity in strongly base orderable matroids is the b-matching problem, in which we are given a maximum degree for each vertex in a graph and seek a collection of edges satisfying all vertices' degree constraints. Many exact algorithms were given for maximum weight linear b-matching problem with an improving dependence of the time complexity on the maximal value of b (see, e.g., [33, 29, 1]). Mestre [30] gave a linear time approximation algorithm for this problem, and Kalyanasundaram and Pruhs [24] considered an online version of it.

Another important special case of both strongly base orderable matroid k-parity is the k-set packing problem. For linear objective functions, this problem has been considered extensively. For cardinality objective, $2/(k + \delta)$ approximation was already given by Hurkens and Schrijver [21], and this approximation ratio was extended relatively recently to general linear functions by Berman [2], building on an earlier work by Chandra and Halldórsson [8], in the more general context of $(k+1)$-claw free graphs.

Organization. Section 2 contains formal definitions and a description of the relationship of k-exchange systems to existing set systems. In Section 3 we present our main applications. Section 4 contains our local search algorithms and analyses for the case of submodular and linear objective function functions. Due to space limitations, the analysis for the cardinality and non-monotone submodular objective functions is omitted from this extended abstract.

2 Preliminaries

Let $\mathcal{N}$ be a given ground set and $\mathcal{I} \subseteq 2^{\mathcal{N}}$ a non-empty, downward closed collection of subsets of $\mathcal{N}$. We call such a system $(\mathcal{N}, \mathcal{I})$ an independence system. We use the standard terminology for discussing independence systems. Given an independence system $(\mathcal{N}, \mathcal{I})$ we say that a set $S \subseteq \mathcal{N}$ is *independent* if $S \in \mathcal{I}$, and call the inclusion-wise maximal independent sets in $\mathcal{I}$ *bases*.

Algorithmically, an independence system $(\mathcal{N}, \mathcal{I})$ might not be given explicitly since the size of $\mathcal{I}$ might be exponential in the size of the ground set. Therefore, we assume access to an independence oracle that given $S \subseteq \mathcal{N}$ determines whether $S \in \mathcal{I}$. Given an independence system $(\mathcal{N}, \mathcal{I})$, and a function $f : 2^{\mathcal{N}} \to \mathbb{R}^+$, we are concerned with the problem of finding a set $S \in \mathcal{I}$ that maximizes f. In particular, we consider a variety of restricted classes of functions f. In the most restricted setting, $f(S) = |S|$ is simply the *cardinality* of S. A natural generalization is the weighted, or *linear* case in which each element $e \in \mathcal{N}$ is assigned a weight $w(e)$, and $f(S) = \sum_{e \in S} w(e)$. More generally we consider *submodular* function f. A function $f : 2^{\mathcal{N}} \to \mathbb{R}^+$ is *submodular* if for every $A, B \subseteq \mathcal{N}$: $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$. Equivalently, f is submodular if for every $A \subseteq B \subseteq \mathcal{N}$ and $e \in \mathcal{N}$: $f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B)$. Additionally, a submodular function f is *monotone* if $f(A) \leq f(B)$ for every $A \subseteq B \subseteq \mathcal{N}$, and *normalized* if $f(\emptyset) = 0$.

The description of a submodular function f might be exponential in the size of the ground set. In this paper we assume the *value oracle* model for accessing f, in which an algorithm is given access to an oracle that returns $f(S)$ for a given set $S \subseteq \mathcal{N}$. This is the model commonly used throughout the literature.

The greedy algorithm provides a k-approximation for maximum weighted independent set in all k-systems. Unfortunately, even the more restricted variants of k-systems, such as k-extendible and k-circuit bounded systems, appear too weak to obtain similar (or potentially stronger) results for local search algorithms. With this goal in mind, we propose the following new class of independence systems, which we call *k-exchange systems*:

Definition 1 (k-exchange system). *An independence system $(\mathcal{N}, \mathcal{I})$ is a k-exchange system if, for all S and T in $\mathcal{I}$, there exists a multiset $Y = \{Y_e \subseteq S \setminus T \mid e \in T \setminus S\}$ such that:*

(K1) $|Y_e| \leq k$ *for each* $e \in T \setminus S$.

(K2) *Every* $e' \in T \setminus S$ *appears in at most* k *sets of* Y.
(K3) *For all* $T' \subseteq T \setminus S$, $(S \setminus (\bigcup_{e \in T'} Y_e)) \cup T' \in \mathcal{I}$

The following theorem follows immediately from the definitions of k-extendible and k-exchange systems.

Theorem 1. *Every k-exchange system is a k-extendible system.*

Mestre shows that the 1-extendible systems are exactly matroids [30]. We can provide a similar motivation for k-exchange systems in terms of *strongly base orderable matroids*, which derive from work by Brualdi and Scrimger on exchange systems [6, 4, 5].

Definition 2 (strongly base orderable matroid [5]). *A matroid $\mathcal{M}$ is strongly base orderable if for all bases S and T of $\mathcal{M}$ there exists a bijection $\pi : S \to T$ such that for all $S' \subseteq S$, $(T \setminus \pi(S')) \cup S'$ is a base.*

If we restrict S' to be a singleton set in this definition—thus considering only single replacements—we obtain a well-known result of Brualdi [3] that holds for all matroids. In a strongly base orderable matroid, we require additionally that any set of these individual replacements can be performed simultaneously to obtain an independent set. This simultaneous replacement property is exactly what we want for local search, as it allows us to extend the local analysis for single replacements to larger replacements needed by our algorithms.

The following theorem is easily obtained by equating Y_e in Definition 1 with the singleton set $\{\pi(e)\}$, where π is as in Definition 2.

Theorem 2. *An independence system $(\mathcal{N}, \mathcal{I})$ is a strongly base orderable matroid if and only if it is a 1-exchange system.*

The cycle matroid on K_4 is not strongly base orderable, and so provides an example of a matroid that is not a 1-exchange system. Conversely, it is possible to find examples of 2-exchange systems that are not 2-circuit bounded.

3 Applications

In this section, we discuss some applications of k-exchange systems. Due to space constraints, we describe only the application for strongly base orderable matroid k-parity. This problem generalizes our other two main applications: k-set packing and b-matching.

Definition 3. *In the matroid k-parity problem, we are given a collection $\mathcal{E}$ of disjoint k-element subsets from a ground set $\mathcal{G}$ and a matroid $(\mathcal{G}, \mathcal{M})$ defined on the ground set. The goal is to find a collection $\mathcal{S}$ of subsets in $\mathcal{E}$ maximizing a function $f : \mathcal{E} \to \mathbb{R}^+$, subject to the constraint that $\bigcup \mathcal{S} \in \mathcal{M}$.*

We consider matroid k-matching in the special case in which the given matroid is strongly base orderable. For clarity, we use calligraphic letters to denote sets of sets from the partition $\mathcal{E}$ and capital letters to denote sets of elements from V, (including, in particular, each of the sets in $\mathcal{E}$). Then, matroid k-parity can be expressed as the independence system $(\mathcal{E}, \mathcal{I})$ where $\mathcal{I} = \{\mathcal{S} \subseteq \mathcal{E} : \bigcup \mathcal{S} \in \mathcal{M}\}$.

Algorithm 1: LS-k-EXCHANGE$((\mathcal{N}, \mathcal{I}), f, \varepsilon, p)$

$e \leftarrow \text{argmax}\{f(\{e\}) \mid e \in \mathcal{N}\}$.
$S \leftarrow \{e\}$.
Let $\mathcal{G} = (\mathcal{I}, \mathcal{E})$ be the p-exchange graph of $(\mathcal{N}, \mathcal{I})$ and $\varepsilon_n = \varepsilon/|\mathcal{N}|$.
while $\exists(S \to T) \in \mathcal{E}$ *such that* $f(T) \geq (1 + \varepsilon_n) f(S)$ **do** $S \leftarrow T$.
Output S.

Theorem 3. *Strongly base orderable matroid k-parity is a k-exchange system.*

Proof. Consider two solutions $\mathcal{S}, \mathcal{T} \in \mathcal{I}$. We must have $\bigcup \mathcal{S}$ and $\bigcup \mathcal{T}$ in $\mathcal{M}$. Let $\pi : \bigcup \mathcal{S} \to \bigcup \mathcal{T}$ be the bijection guaranteed by Definition 2, and for any set $E \in \mathcal{S}$ define $Y_E = \{A \in \mathcal{T} : A \cap \pi(E) \neq \emptyset\}$. The sets in $\mathcal{E}$ are disjoint and contain at most k elements. Since π is a bijection, we must therefore have $|Y_E| \leq k$ for all $E \in \mathcal{S}$ and each $A \in \mathcal{T}$ appears in at most k sets Y_E. Thus, Y satisfies Properties (K1) and (K2). Consider a set $\mathcal{C} \subseteq \mathcal{S}$, and let $\mathcal{S}' = (\mathcal{S} \setminus \bigcup\{Y_E : E \in \mathcal{C}\}) \cup \mathcal{C}$. From the definition of π we have $(\bigcup \mathcal{S} \setminus \pi(\bigcup \mathcal{C})) \cup \bigcup \mathcal{C} \in \mathcal{M}$, and $\bigcup \mathcal{S}'$ is a subset of this set, so $\bigcup \mathcal{S}' \in \mathcal{M}$ and hence $\mathcal{S}' \in \mathcal{I}$, showing that (K3) is satisfied.

4 Combinatorial Local Search Approximation Algorithms

Let us define a directed graph representing improvements considered by our local search algorithms:

Definition 4. *Given a k-exchange system $(\mathcal{N}, \mathcal{I})$, $S, T \in \mathcal{I}$, and $p \in \mathbb{N}$, T is* p-reachable *from S if the following conditions are satisfied:*

1. $|T \setminus S| \leq p$.
2. $|S \setminus T| \leq (k-1)p + 1$.

Definition 5. *Given a k-exchange system $(\mathcal{N}, \mathcal{I})$, and $p \in \mathbb{N}$, the* p-exchange graph *of $(\mathcal{N}, \mathcal{I})$ is a directed graph $\mathcal{G} = (\mathcal{I}, \mathcal{E})$ where $(S \to T) \in \mathcal{E}$ if and only if T is p-reachable from S.*

Our algorithms are local search algorithms starting from a vertex in $\mathcal{G}$ and touring the graph arbitrarily until they find a sink vertex S. The algorithms than output S.
The start point of the algorithms is the singleton of maximum value. It is important to note that this start point is an independent set. Otherwise, the element of the singleton does not belong to any independent set (recall that $(\mathcal{N}, \mathcal{I})$ is monotone), and therefore, can be removed from $(\mathcal{N}, \mathcal{I})$.

Algorithm 1 attempts to maximizing the objective function itself while touring $\mathcal{G}$. In the case of a linear f, and $k > 2$, we can improve the approximation ratio of algorithm by using the following technique. Since f is linear, we can assign weights $w(e) = f(\{e\})$ for each element $e \in \mathcal{N}$ and express $f(S)$ as $\sum_{e \in S} w(e)$. We construct a new linear objective function $w^2(S) = \sum_{e \in S} w(e)^2$,

Algorithm 2: NON-OBLIVIOUS-LS-k-EXCHANGE$((\mathcal{N}, \mathcal{I}), f, \varepsilon)$

$e \leftarrow \text{argmax}\,\{w(e) \mid e \in \mathcal{N}\}$.
$S \leftarrow \{e\}$.
Round the weights w down to integer multiples of $f(S)\epsilon/n$.
Let $\mathcal{G} = (\mathcal{I}, \mathcal{E})$ be the k-exchange graph of $(\mathcal{N}, \mathcal{I})$. **while** $\exists(S \to T) \in \mathcal{E}$ *such that* $w^2(T) > w^2(S)$ **do** $S \leftarrow T$.
Output S.

and use this function to guide our search. Additionally, we ensure convergence by rounding the weights w using the initial solution, as in [2]. Note that in this algorithm, we always search in the k-exchange graph, rather than the p-exchange graph for some given p.

The graph $\mathcal{G}$, like the set $\mathcal{I}$, might be exponential. However, standard techniques can be used to show that the algorithms terminate in polynomial time.

Theorem 4. *For any constants k, $0 < \varepsilon < k$ and $p \in \mathbb{N}$, Algorithms 1 and 2 terminate in polynomial time.*

4.1 Analysis of Algorithm 1

Our analysis of Algorithm 1 proceeds by considering a particular subset of improvements considered by the algorithm in line 4. Let $(\mathcal{N}, \mathcal{I})$ be a k-exchange system, and let $S \in \mathcal{I}$ be the independent set produced by Algorithm 1 and $T \in \mathcal{I}$ be any other independent set. We construct the following bipartite graph $G_{S,T} = (S \setminus T, T \setminus S, E)$, where $E = \{(e, e') \mid e \in T \setminus S, e' \in Y_e\}$. Note that Properties (K1) and (K2) imply that the maximum degree in $G_{S,T}$ is at most k.

The following theorem is a key ingredient in the analysis of Algorithm 1, allowing us to decompose $G_{S,T}$ into a collection of paths, from which we will obtain the improvements considered in our analysis. Like $G_{S,T}$, its use is restricted only to the analysis itself, as no actual construction of $\mathcal{P}(G, k, h)$ is needed.

Theorem 5. *Let G be an undirected graph whose maximum degree is at most $k \geq 2$. Then, for every $h \in \mathbb{N}$ there exists a multiset $\mathcal{P}(G, k, h)$ of simple paths in G and a labeling $\ell : V \times \mathcal{P}(G, k, h) \to \{\emptyset, 1, 2, \ldots, h\}$ such that:*

1. *For every $P \in \mathcal{P}(G, k, h)$, the labeling ℓ of the nodes of P is consecutive and increasing with labels from $\{1, 2, \ldots, h\}$. Vertices not in P receive label $\emptyset$.*
2. *For every $P \in \mathcal{P}(G, k, h)$ and v in P, if $deg_G(v) = k$ and $\ell(v, P) \notin \{1, h\}$, then at least two of the neighbors of v are in P.*
3. *For every $v \in V$ and label $i \in \{1, 2, \ldots, h\}$, there are $n(k, h) = k \cdot (k-1)^{h-2}$ paths $P \in \mathcal{P}(G, k, h)$ for which $\ell(v, P) = i$.*

Note for condition 2, v might be an end vertex of a path P, but still have a label different from 1 and h. This might happen since paths might contain less than h vertices and start with a label different from 1.

Due to space constraints we omit the full proof of Theorem 5, but let us provide some intuition as to why the construction of $\mathcal{P}(G,k,h)$ is possible. Assume that the degree of every vertex in G is exactly k and that G's girth is at least h. Construct the multiset $\mathcal{P}(G,k,h)$ in the following way. ¿From every vertex $u \in V$, choose all possible paths starting at v and containing exactly h vertices. Number the vertices of these paths consecutively, starting from 1 up to h. First, note that all these paths are simple since the girth of G is at least h and all paths contain exactly h vertices. Second, the number of paths starting from u is: $n(k,h) = k \cdot (k-1)^{h-2}$, since all vertices have degree of exactly k. Third, the number of times each label is given in the graph is exactly $n \cdot n(k,h)$. Since the number of vertices at distance i, $1 \le i \le h$, from each vertex u is identical, the labels are distributed equally among the vertices. Thus, the number of paths in which a given vertex u appears with a given label, is exactly $n(k,h)$. This concludes the proof of the theorem in case G has the above properties.

Applying Theorem 5 to $G_{S,T}$ with $h = 2p$, gives a multiset $\mathcal{P}(G_{S,T},k,2p)$ of simple paths in $G_{S,T}$ and a labeling $\ell : (S \triangle T) \times \mathcal{P}(G_{S,T},k,2p) \to \{\emptyset, 1, 2, \ldots, 2p\}$ with all the properties guaranteed by Theorem 5. We use $\mathcal{P}(G_{S,T},k,2p)$ to construct a new multiset $\mathcal{P}'$ of subsets of vertices of $G_{S,T}$. For each path in $P \in \mathcal{P}$ that contains at least one vertex from T [3], we add $(P \cup N(P))$ to $\mathcal{P}'$, where $N(P)$ is the set of all vertices in $G_{S,T}$ neighboring some vertex in P. Intuitively, $\mathcal{P}'$ is the collection of all paths in $\mathcal{P}(G_{S,T},k,2p)$ with an extra "padding" of vertices from S that surround P, excluding "paths" composed of a single vertex from S.

Lemma 1. *Every vertex $e \in T \setminus S$ appears in $2p \cdot n(k,2p)$ sets of $\mathcal{P}'$, and every vertex $e' \in S \setminus T$ appears in at most $2\left((k-1)p+1\right) \cdot n(k,2p)$ sets of $\mathcal{P}'$.*

Proof. By property 3 of Theorem 5, every $e \in T \setminus S$ appears in $n(k,2p)$ paths of $\mathcal{P}(G_{S,T},k,2p)$ for every possible label. Since there are $2p$ possible labels, the number of appearances is exactly $2p \cdot n(k,2p)$. In the creation of $\mathcal{P}'$ from $\mathcal{P}(G_{S,T},k,2p)$, no $e \in T \setminus S$ is added or removed from any path $P \in \mathcal{P}(G_{S,T},k,2p)$, thus, this is also the number of appearances of every $e \in T \setminus S$ in $\mathcal{P}'$.

Let $e' \in S \setminus T$. By the construction of $\mathcal{P}'$, a set in $\mathcal{P}'$ that contains e' must contain a vertex $e \in T \setminus S$ where $(e,e') \in E$ (e is a neighbor of e' in $G_{S,T}$). Every such neighboring vertex e, by the first part of the lemma, appears in exactly $2p \cdot n(k,2p)$ sets in $\mathcal{P}'$. Therefore, the number of appearances of e' in sets of $\mathcal{P}'$ is at most: $deg_{G_{S,T}}(e') \cdot 2p \cdot n(k,2p) \le 2pk \cdot n(k,2p)$ (recall that the maximum degree of $G_{S,T}$ is at most k). Furthermore, $\ell(e',P) \ne \{1,2p\}$, by property 2 of Theorem 5, e' has at least two neighbors in $T \setminus S$ which belong to P itself. Hence, $P \cup N(P) \in \mathcal{P}'$ should be counted only once while in the above counting it was counted at least twice. The number of such $P \in \mathcal{P}(G_{S,T},k,2p)$ is exactly $2(p-1) \cdot n(k,2p)$ (by property 3 of Theorem 5). Removing the double counting from the bound, we can conclude that for every $e' \in S \setminus T$, the number of sets in $\mathcal{P}'$ it appears in is at most $2pk \cdot n(k,2p) - 2(p-1) \cdot n(k,2p) = 2\left((k-1)p+1\right) \cdot n(k,2p)$.

[3] Note that this implies that this vertex is from $T \setminus S$ since $G_{S,T}$ does not contain vertices from $S \cap T$.

Note: The proof of Theorem 6 assumes each vertex $e' \in S \setminus T$ appears in *exactly* $2\left((k-1)p+1\right) \cdot n(k, 2p)$ sets in $\mathcal{P}'$. This can be achieved by adding "dummy" sets to $\mathcal{P}'$ containing e' alone.

We are now ready to state our main theorem. In the proof, we use of the following two technical lemmata from [26]:

Lemma 2 (Lemma 1.1 in [26]). *Let f be a non-negative submodular function of $\mathcal{N}$. Let $S' \subseteq S \subseteq \mathcal{N}$ and let $\{T_\ell\}_{\ell=1}^t$ be a collection of subsets of $S \setminus S'$ such that every elements of $S \setminus S'$ appears in exactly k of these subsets. Then, $\sum_{\ell=1}^t \left[f(S) - f(S \setminus T_\ell)\right] \leq k\left(f(S) - f(S')\right)$.*

Lemma 3 (Lemma 1.2 in [26]). *Let f be a non-negative submodular function of $\mathcal{N}$. Let $S \subseteq \mathcal{N}$, $C \subseteq \mathcal{N}$ and let $\{T_\ell\}_{\ell=1}^t$ be a collection of subsets of $C \setminus S$ such that every elements of $C \setminus S$ appears in exactly k of these subsets. Then, $\sum_{\ell=1}^t \left[f(S \cup T_\ell) - f(S)\right] \geq k\left(f(S \cup C) - f(S)\right)$.*

Theorem 6. *For every $T \in \mathcal{I}$ and every submodular f:*

$$f(S \cup T) + \left(k - 1 + \frac{1}{p}\right) \cdot f(S \cap T) \leq \left(k + \frac{1}{p} + k\varepsilon\right) \cdot f(S) \ .$$

Proof. Note that by construction, the symmetric difference of $S \triangle P'$ is an independent set, for any $P' \in \mathcal{P}'$ and furthermore $f(S \triangle P')$ p-reachable from S. Since Algorithm 1 terminated with S, it must be the case that S is approximately "locally optimal", and therefore,

$$f(S \triangle P') < (1 + \varepsilon_n)\, f(S) \ . \tag{1}$$

for all $P' \in \mathcal{P}'$. By submodularity of f, the fact that $S \setminus P' \subseteq S \triangle P'$ and the fact that all vertices in $S \cap P'$ do not belong to either $S \setminus P'$ or $S \triangle P'$, we get:

$$f(S \cup P') - f(S \triangle P') \leq f(S) - f(S \setminus P') \ . \tag{2}$$

Adding Inequalities 1 and 2 gives:

$$f(S \cup P') - (1 + \varepsilon_n)\, f(S) \leq f(S) - f(S \setminus P') \ . \tag{3}$$

Inequality 3 holds for every $P' \in \mathcal{P}'$. Summing over all such sets yields:

$$\sum_{P' \in \mathcal{P}'} \left[f(S \cup P') - f(S)\right] - \varepsilon_n |\mathcal{P}'| f(S) \leq \sum_{P' \in \mathcal{P}'} \left[f(S) - f(S \setminus P')\right] \ . \tag{4}$$

Now, we note that any given P' contains only vertices from $S \triangle T$. Thus, Inequality 4 is equivalent to:

$$\sum_{P' \in \mathcal{P}'} \left[f(S \cup (P' \cap (T \setminus S))) - f(S)\right] - \varepsilon_n |\mathcal{P}'| f(S) \leq \sum_{P' \in \mathcal{P}'} \left[f(S) - f(S \setminus (P' \cap (S \setminus T)))\right] \ . \tag{5}$$

By Lemma 1 (and the note after it), each vertex in $S \setminus T$ appears in exactly $2\left((k-1)p+1\right) \cdot n(k,2p)$ sets in $\mathcal{P}'$, while every vertex of $T \setminus S_{ALG}$ appears in exactly $2p \cdot n(k,2p)$ sets in $\mathcal{P}'$. Thus, applying Lemma 2 to the right of Inequality 5 and Lemma 3 to the left gives:

$$2p \cdot n(k,2p)(f(S \cup T) - f(S)) - \varepsilon_n |\mathcal{P}'| f(S) \leq \\ 2\left((k-1)p+1\right) n(k,2p) \left(f(S) - f(S \cap T)\right) \ .$$

Rearranging terms and using the definition of ε_n we obtain:

$$f(S \cup T) + \left(k - 1 + \frac{1}{p}\right) f(S \cap T) \leq \left(k + \frac{1}{p}\right) f(S) + \frac{\varepsilon |\mathcal{P}'|}{2p \cdot n(k,2p)|\mathcal{N}|} f(S) \ . \tag{6}$$

Finally, we note every set in $P' \in \mathcal{P}'$ contains at least 1 vertex from $G_{S,T}$, and, by Lemma 1, every vertex in $G_{S_{ALG},T}$ appears in exactly $2p \cdot n(k,2p)$ or $2\left((k-1)p+1\right) \cdot n(k,2p)$ sets of $\mathcal{P}'$ (depending on whether the vertex is in $T \setminus S_{ALG}$ or $S_{ALG} \setminus T$). Therefore, in the worst case $|\mathcal{P}'| \leq |S_{ALG} \triangle T| \cdot 2n(k,2p) \max\{p, (k-1)p+1\} \leq 2|\mathcal{N}| n(k,2p) kp$.

By setting $1/p + k\epsilon \leq \delta$ and using basic properties of monotone submodular and linear functions we obtain the following.

Corollary 1. *Given a set function $f : 2^{\mathcal{N}} \to \mathbb{R}^+$ and any $\delta > 0$, Algorithm 1 is a $1/(k+\delta)$ approximation algorithm if f is a normalized monotone submodular function and a $1/(k-1+\delta)$ approximation algorithm if f is a linear function.*

4.2 Analysis of Algorithm 2

The analysis of Algorithm 2 closely follows Berman's analysis for $(k+1)$-claw free graphs [2]. Like the analysis of Algorithm 1, this analysis also considers a subset of the possible improvements considered by the algorithm. Here, however, we consider improvements of the following form. Consider a k-exchange system $(\mathcal{I}, \mathcal{N})$ and the rounded weight function $w : \mathcal{N} \to \mathbb{R}^+$ produced in line 3 of the algorithm (we shall use the rounded weights for the remainder of our analysis). Let S be the solution produced by Algorithm 2, and T be any independent set in $\mathcal{I}$. For each element $x \in S$, let P_x be the set of all elements $e \in T$ such that $x = \arg\max_{y \in Y_e} w(y)$. For elements $e \in S \cap T$, we define $Y_e = \{e\}$, so $P_e = \{e\}$. Then, $\mathcal{P} = \{P_x\}_{x \in S}$ is a partition of T. We consider improvements $P_x \cup N(P_x)$ where $N(P_x) = \bigcup_{e \in P_x} Y_e$. Note that for all $y \in N(P_x)$, we have $w(y) \leq w(x)$. The following theorem from Berman's analysis allows us to relate the value of w^2 to that of w.

Lemma 4. *For all $x \in S$, $e \in P_x$: $w^2(e) - w^2(Y_e \setminus \{x\}) \geq w(x) \cdot (2w(e) - w(Y_e))$.*

We now prove our main theorem regarding Algorithm 2.

Theorem 7. *$\frac{k+1}{2} w(S) \geq w(T)$.*

Proof. For each element $x \in S$ we consider the improvement $P_x \cup N(P_x)$. By construction, the symmetric difference of $S \triangle (P_x \cup N(P_x))$ is an independent set, for any $x \in S$, and moreover, $S \triangle (P_x \cup N(P_x))$ is k-reachable from S. Since Algorithm 2 terminated, producing solution S, it must be the case that S is locally optimal, and so $w^2(S \triangle (P_x \cup N(P_x))) \leq w^2(S)$ for each $x \in S$. Combining this with the linearity of w^2 and the fact that $x \in Y_e$ for all $e \in P_x$, we have:

$$w^2(P_x) \leq w^2(N(P_x)) \leq w(x)^2 + \sum_{e \in P_x} w^2(Y_e \setminus \{x\}) \ . \tag{7}$$

Rearranging Inequality 7 using $w^2(P_x) = \sum_{e \in P_x} w(e)^2$ we obtain:

$$\sum_{e \in P_x} w(e)^2 - w^2(Y_e \setminus \{x\}) \leq w(x)^2 \ . \tag{8}$$

Applying Lemma 4 to each term on the left of Inequality (8) gives $\sum_{e \in P_x} w(x) \cdot (2w(e) - f(Y_e)) \leq w(x)^2$ Dividing both sides by $w(x)$ we obtain:

$$\sum_{e \in P_x} (2w(e) - w(Y_e)) \leq w(x) \ . \tag{9}$$

Inequality (9) holds for all $x \in S$. Thus, summing over all $x \in S$, we have

$$\sum_{x \in S} w(x) \geq \sum_{x \in S} \sum_{e \in P_x} [2w(e) - w(Y_e)] \ . \tag{10}$$

We note that $\sum_{x \in S} w(x) = w(S)$, and P_x is a partition of T, so Inequality 10 is equivalent to:

$$\begin{aligned} w(S) \geq \sum_{e \in T} [2w(e) - w(Y_e)] &= 2w(T) - \sum_{e \in T} w(Y_e) \\ &\geq 2w(T) - k \sum_{x \in S} w(x) = 2w(T) - kw(S) \ , \end{aligned} \tag{11}$$

where the second inequality follows from (K2).

Corollary 2. *Algorithm 2 is a $2/(k+1+\delta)$ approximation algorithm for maximizing a linear function $f : 2^{\mathcal{N}} \to \mathbb{R}^+$ for any $\delta > 0$.*

Proof. It can be shown (as in [2]) that for any ϵ, our rounding operation introduces at most an extra multiplicative error of ϵ with respect to the original weight function, while ensuring that the algorithm makes only poly(n, ϵ^{-1}) improvements.

5 Open Questions

The k-coverable class intersects the k-intersection class, and for both classes the same results are achieved by Algorithm 1, though different insights are used to

analysis each case. Finding a common generalization of both classes admitting a uniform analysis is an intriguing question. A more concrete question is whether there is an exact algorithm for maximizing a *linear* function over the 2-coverable class (analogously to Edmonds exact algorithm for 2-intersection [11]). Finally, it would be interesting to see whether Algorithm 2 can be applied to monotone submodular functions. In the submodular case, we no longer have weights to square in the non-oblivious potential function, but one possible approach is to consider the sum of the squared marginals of the submodular function.
Acknowledgments: The last author thanks Julián Mestre and Allan Borodin for providing comments on a preliminary version of the paper.

References

1. R. P. Anstee. A polynomial algorithm for b-matchings: An alternative approach. *Information Processing Letters*, 24(3):153 – 157, 1987.
2. P. Berman. A d/2 approximation for maximum weight independent set in d-claw free graphs. *Nordic J. of Computing*, 7:178–184, September 2000.
3. R. A. Brualdi. Comments on bases in dependence structures. *Bull. of the Australian Math. Soc.*, 1(02):161–167, 1969.
4. R. A. Brualdi. Common transversals and strong exchange systems. *J. of Combinatorial Theory*, 8(3):307–329, Apr. 1970.
5. R. A. Brualdi. Induced matroids. *Proc. of the American Math. Soc.*, 29:213–221, 1971.
6. R. A. Brualdi and E. B. Scrimger. Exchange systems, matchings, and transversals. *J. of Combinatorial Theory*, 5(3):244–257, Nov. 1968.
7. G. Calinescu, C. Chekuri, M. Pal, and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. *To appear in SIAM Journal on Computing, Special Issue for STOC 2008*, 2008.
8. B. Chandra and M. Halldórsson. Greedy local improvement and weighted set packing approximation. In *SODA*, pages 169–176, 1999.
9. C. Chekuri, J. Vondrák, and R. Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *STOC*, pages 783–792, 2011.
10. M. Conforti and G. Cornuèjols. Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Disc. Appl. Math.*, 7(3):251–274, 1984.
11. J. Edmonds. Matroid intersection. In P. Hammer, E. Johnson, and B. Korte, editors, *Discrete Optimization I, Proceedings of the Advanced Research Institute on Discrete Optimization and Systems Applications of the Systems Science Panel of NATO and of the Discrete Optimization Symposium*, volume 4 of *Annals of Discrete Mathematics*, pages 39–49. Elsevier, 1979.
12. U. Feige, V. S. Mirrokni, and J. Vondrák. Maximizing non-monotone submodular functions. In *FOCS*, pages 461–471, 2007.
13. M. Feldman, J. S. Naor, and R. Schwartz. Nonmonotone submodular maximization via a structural continuous greedy algorithm. In *ICALP*, pages 342–353, 2011.
14. M. Feldman, J. S. Naor, and R. Schwartz. A unified continuous greedy algorithm for submodular maximization. *To appear in FOCS 2011*, 2011.
15. M. Fisher, G. Nemhauser, and L. Wolsey. An analysis of approximations for maximizing submodular set functions - ii. *Math. Prog. Study*, 8:73–87, 1978.

16. S. O. Gharan and J. Vondrák. Submodular maximization by simulated annealing. In *SODA*, pages 1098–1116, 2011.
17. A. Gupta, A. Roth, G. Schoenebeck, and K. Talwar. Constrained non-monotone submodular maximization: Offline and secretary algorithms. In *WINE*, 2010.
18. D. Hausmann and B. Korte. K-greedy algorithms for independence systems. *Oper. Res. Ser. A-B*, 22(1):219–228, 1978.
19. D. Hausmann, B. Korte, and T. Jenkyns. Worst case analysis of greedy type algorithms for independence systems. *Math. Prog. Study*, 12:120–131, 1980.
20. E. Hazan, S. Safra, and O. Schwartz. On the complexity of approximating k-set packing. *Comput. Complex.*, 15:20–39, May 2006.
21. C. A. J. Hurkens and A. Schrijver. On the size of systems of sets every t of which have an sdr, with an application to the worst case ratio of heuristics for packing problems. *SIAM J. Disc. Math.*, 2(1):68–72, 1989.
22. T. Jenkyns. The efficacy of the greedy algorithm. *Cong. Num.*, 17:341–350, 1976.
23. P. M. Jensen and B. Korte. Complexity of matroid property algorithms. *SIAM J. Computing*, 11(1):184, 1982.
24. B. Kalyanasundaram and K. Pruhs. An optimal deterministic algorithm for online b-matching. In V. Chandru and V. Vinay, editors, *Foundations of Software Technology and Theoretical Computer Science*, volume 1180 of *Lecture Notes in Computer Science*, pages 193–199. Springer Berlin / Heidelberg, 1996.
25. B. Korte and D. Hausmann. An analysis of the greedy heuristic for independence systems. *Annals of Discrete Math.*, 2:65–74, 1978.
26. J. Lee, M. Sviridenko, and J. Vondrák. Submodular maximization over multiple matroids via generalized exchange properties, 2009. To appear in Mathematics of Operations Research.
27. J. Lee, M. Sviridenko, and J. Vondrak. Matroid matching: the power of local search. In *STOC*, pages 369–378, 2010.
28. L. Lovász. The matroid matching problem. In L. Lovász and V. T. Sós, editors, *Algebraic Methods in Graph Theory*, Amsterdam, 1981.
29. A. B. Marsh, III. *Matching algorithms*. PhD thesis, The Johns Hopkins University, 1979.
30. J. Mestre. Greedy in approximation algorithms. In *ESA*, pages 528–539, 2006.
31. G. Nemhauser and L. Wolsey. Best algorithms for approximating the maximum of a submodular set function. *Math. Oper. Res.*, 3(3):177–188, 1978.
32. G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions - i. *Math. Prog.*, 14(1):265–294, 1978.
33. W. Pulleyblank. *Faces of matching polyhedra*. PhD thesis, Deptartment of Combinatorics and Optimization, University of Waterloo, 1973.
34. A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.
35. H. Simon. Approximation algorithms for channel assignment in cellular radio networks. In J. Csirik, J. Demetrovics, and F. Gcseg, editors, *Fundamentals of Computation Theory*, volume 380 of *Lecture Notes in Computer Science*, pages 405–415. Springer Berlin / Heidelberg, 1989.
36. J. A. Soto. A simple PTAS for weighted matroid matching on strongly base orderable matroids. *To appear in LAGOS*, 2011.
37. P. Tong, E. L. Lawler, and V. V. Vazirani. Solving the weighted parity problem for gammoids by reduction to graphic matching. Technical Report UCB/CSD-82-103, EECS Department, University of California, Berkeley, Apr 1982.