# A Tight Linear Time (1/2)-Approximation for Unconstrained Submodular Maximization 

Niv Buchbinder ${ }^{\dagger}$<br>Statistics and Operations Research Dept. Tel Aviv University<br>Tel Aviv, Israel<br>niv.buchbinder@gmail.com

Moran Feldman<br>Computer Science Dept.<br>Technion<br>Haifa, Israel<br>moranfe@cs.technion.ac.il

Joseph (Seffi) Naor*<br>Computer Science Dept.<br>Technion<br>Haifa, Israel<br>naor@cs.technion.ac.il

Roy Schwartz<br>Theory Group<br>Microsoft Research<br>Redmond, WA<br>roysch@microsoft.com


#### Abstract

We consider the Unconstrained Submodular Maximization problem in which we are given a non-negative submodular function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}^{+}$, and the objective is to find a subset $S \subseteq \mathcal{N}$ maximizing $f(S)$. This is one of the most basic submodular optimization problems, having a wide range of applications. Some well known problems captured by Unconstrained Submodular Maximization include MaxCut, Max-DiCut, and variants of Max-SAT and maximum facility location. We present a simple randomized linear time algorithm achieving a tight approximation guarantee of $1 / 2$, thus matching the known hardness result of Feige et al. [11]. Our algorithm is based on an adaptation of the greedy approach which exploits certain symmetry properties of the problem. Our method might seem counterintuitive, since it is known that the greedy algorithm fails to achieve any bounded approximation factor for the problem.


Keywords-Submodular Functions, Approximation Algorithms

## I. Introduction

The study of combinatorial problems with submodular objective functions has recently attracted much attention, and is motivated by the principle of economy of scale, prevalent in real world applications. Submodular functions are also commonly used as utility functions in economics and algorithmic game theory. Submodular functions and submodular maximization play a major role in combinatorial optimization, where several well known examples of submodular functions include cuts in graphs and hypergraphs, rank functions of matroids, and covering functions. A function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}^{+}$is called submodular if for every $A \subseteq B \subseteq \mathcal{N}$ and $u \in \mathcal{N}, f(A \cup\{u\})-f(A) \geq f(B \cup\{u\})-f(B)^{1}$. Submodularity captures the principle of diminishing returns in economics.

Perhaps the most basic submodular maximization problem is Unconstrained Submodular Maximization (USM). Given a non-negative submodular fucntion $f$, the goal is to

[^0]find a subset $S \subseteq \mathcal{N}$ maximizing $f(S)$. Note that there is no restriction on the choice of $S$, as any subset of $\mathcal{N}$ is a feasible solution. USM captures many well studied problems such as Max-Cut, Max-DiCut [15], [20], [22], [25], [27], [36], variants of Max-SAT, and maximum facility location [1], [7], [8]. Moreover, USM has various applications in more practical settings such as marketing in social networks [21], revenue maximization with discrete choice [2], and algorithmic game theory [10], [34].

USM has been studied starting since the sixties in the operations research community [2], [6], [16], [17], [18], [26], [29], [33]. Not surprisingly, as USM captures NP-hard problems, all these works provide algorithms that either solve special cases of the problem, or provide exact algorithms whose time complexity cannot be efficiently bounded, or provide efficient algorithms whose output has no provable guarantee.

The first rigorous study of the USM problem was conducted by Feige et al. [11], who provided several constant approximation factor algorithms. They proved that a subset $S$ chosen uniformly at random is a ( $1 / 4$ )-approximation. Additionally, they also described two local search algorithms. The first uses $f$ as the objective function, and provides an approximation of $1 / 3$. The second uses a noisy version of $f$ as the objective function, and achieves an improved approximation guarantee of $2 / 5$. Gharan and Vondrák [14] showed that an extension of the last method, known as simulated annealing, can provide an improved approximation of roughly 0.41 . Their algorithm, like that of Feige et al. [11], uses local search with a noisy objective function. However, in [14] the noise decreases as the algorithm advances, as opposed to being constant as in [11]. Feldman et al. [12] observed that if the simulated annealing algorithm of [14] outputs a relatively poor solution, then it must generate at some point a set $S$ which is structurally similar to some optimal solution. Moreover, they showed that this structural similarity can be traded for value, providing an overall improved approximation of roughly 0.42 .

It is important to note that for many special cases of USM, better approximation factors are known. For example, the
seminal work of Goemans and Williamson [15] provides an 0.878 -approximation for Max-Cut based on a semidefinite programming formulation, and Ageev and Sviridenko [1] provide an approximation of 0.828 for the maximum facility location problem.

On the negative side, Feige et al. [11] studied the hardness of USM assuming the function $f$ is given via a value oracle ${ }^{2}$. They proved that for any constant $\varepsilon>0$, any algorithm achieving an approximation of $(1 / 2+\varepsilon)$ requires an exponential number of oracle queries. This hardness result holds even if $f$ is symmetric, in which case it is known to be tight [11]. Recently, Dobzinski and Vondrák [9] proved that even if $f$ has a compact representation (which is part of the input), the latter hardness still holds assuming $R P \neq N P$.

## A. Our Results

In this paper we resolve the approximability of the Unconstrained Submodular Maximization problem. We design a tight linear time randomized (1/2)-approximation for the problem. We begin by presenting a simple greedy-based deterministic algorithm that achieves a (1/3)-approximation for USM.

Theorem I.1. There exists a deterministic linear time (1/3)approximation algorithm for the Unconstrained Submodular Maximization problem.

We show that our analysis of the algorithm is tight by providing an instance for which the algorithm achieves an approximation of $1 / 3+\varepsilon$ for an arbitrary small $\varepsilon>0$. The improvement in the performance is achieved by incorporating randomness into the choices of the algorithm. We obtain a tight bound for USM without increasing the time complexity of the algorithm.
Theorem I.2. There exists a randomized linear time (1/2)approximation algorithm for the Unconstrained Submodular Maximization problem.

In both Theorems I. 1 and I. 2 we assume that a single query to the value oracle takes $O(1)$ time. Thus, a linear time algorithm is an algorithm which makes $O(n)$ oracle queries plus $O(n)$ other operations, where $n$ is the size of the ground set $\mathcal{N}$.

Building on the above two theorems, we provide two additional approximation algorithms for Submodular MaxSAT and Submodular Welfare with 2 players (for the exact definition of these problems please refer to Section IV). Both problems are already known to have a tight approximation [13]. However, our algorithms run in linear time, thus,

[^1]significantly improving the time complexity, while achieving the same performance guarantee.

Theorem I.3. There exists a randomized linear time (3/4)approximation algorithm for the Submodular Max-SAT problem.

Theorem I.4. There exists a randomized linear time (3/4)approximation algorithm for the Submodular Welfare problem with 2 players.

## B. Techniques

The algorithms we present are based on a greedy approach. It is known that a straightforward greedy approach fails for USM. In contrast, Feldman et al. [13] recently showed that a continuous greedy approach does provide a ( $1 / e$ )-approximation for maximizing any non-monotone submodular function over a matroid. Recall that better bounds than $1 / e$ are already known for USM.

To understand the main ideas of our algorithm, consider a non-negative submodular function $f$. Let us examine the complement of $f$, denoted by $\bar{f}$, defined as $\bar{f}(S) \triangleq$ $f(\mathcal{N} \backslash S)$, for any $S \subseteq \mathcal{N}$. Note that since $f$ is submodular, $\bar{f}$ is also submodular. Additionally, given an optimal solution $O P T \subseteq \mathcal{N}$ for USM with input $f, \mathcal{N} \backslash O P T$ is an optimal solution for $\bar{f}$, and both solutions have the exact same value. Consider, for example, the greedy algorithm. For $f$, it starts from an empty solution and iteratively adds elements to it in a greedy fashion. However, when applying the greedy algorithm to $\bar{f}$, one gets an algorithm for $f$ that effectively starts with the solution $\mathcal{N}$ and iteratively removes elements from it. Both algorithms are equally reasonable, but, unfortunately, both fail.

Despite the failure of the greedy algorithm when applied separately to either $f$ or $\bar{f}$, we show that a correlated execution on both $f$ and $\bar{f}$ provides a much better result. That is, we start with two solutions $\emptyset$ and $\mathcal{N}$. The algorithm considers the elements (in arbitrary order) one at a time. For each element it determines whether it should be added to the first solution or removed from the second solution. Thus, after a single pass over the ground set $\mathcal{N}$, both solutions completely coincide. This is the solution that the algorithm outputs. We show that a greedy choice in each step obtains an approximation guarantee of $1 / 3$, hence proving Theorem I.1. To get Theorem I.2, we use a natural extension of the greedy rule, together with randomization, thus improving the approximation guarantee to $1 / 2$.

## C. Related Work

Many works consider the problem of maximizing a nonnegative submodular function subject to various combinatorial constraints defining the feasible solutions [13], [19], [30], [31], [39], or minimizing a submodular function subject to such constraints [23], [24], [35]. Interestingly, the problem
of unconstrained submodular minimization can be solved in polynomial time [32].

Another line of work deals with maximizing normalized monotone submodular functions, again, subject to various combinatorial constraints. A continuous greedy algorithm was given by Calinescu et al. [3] for maximizing a normalized monotone submodular function subject to a matroid constraint. Later, Lee et al. [31] gave a local search algorithm achieving $1 / p-\varepsilon$ approximation for maximizing such functions subject to the intersection of $p$ matroids. Kulik et al. [28] showed a $1-1 / e-\varepsilon$ approximation algorithm for maximizing a normalized monotone submodular function subject to multiple knapsack constraints. Recently, Chekuri et al. [5] and Feldman et al. [13] gave non-monotone counterparts of the continuous greedy algorithm of [3], [38]. These results improve several non-monotone submodular optimization problems. Some of the above results were generalized by Chekuri et al. [4], who provide a dependent rounding technique for various polytopes, including matroid and matroid-intersection polytops. The advantage of this rounding technique is that it guarantees strong concentration bounds for submodular functions. Additionally, Chekuri et al. [5] define a contention resolution rounding scheme which allows one to obtain approximations for different combinations of constraints.

## II. A Deterministic (1/3)-Approximation Algorithm

In this section we present a deterministic linear time algorithm for USM. The algorithm proceeds in $n$ iterations that correspond to some arbitrary order $u_{1}, \ldots, u_{n}$ of the ground set $\mathcal{N}$. The algorithm maintains two solutions $X$ and $Y$. Initially, we set the solutions to $X_{0}=\emptyset$ and $Y_{0}=\mathcal{N}$. In the $i$ th iteration the algorithm either adds $u_{i}$ to $X_{i-1}$ or removes $u_{i}$ from $Y_{i-1}$. This decision is done greedily based on the marginal gain of each of the two options. Eventually, after $n$ iterations both solutions coincide, and we get $X_{n}=Y_{n}$; this is the output of the algorithm. A formal description of the algorithm appears as Algorithm 1.

```
Algorithm 1: DeterministicUSM \((f, \mathcal{N})\)
    \(X_{0} \leftarrow \emptyset, Y_{0} \leftarrow \mathcal{N}\).
    for \(i=1\) to \(n\) do
        \(a_{i} \leftarrow f\left(X_{i-1} \cup\left\{u_{i}\right\}\right)-f\left(X_{i-1}\right)\).
        \(b_{i} \leftarrow f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)-f\left(Y_{i-1}\right)\).
        if \(a_{i} \geq b_{i}\) then \(X_{i} \leftarrow X_{i-1} \cup\left\{u_{i}\right\}, Y_{i} \leftarrow Y_{i-1}\).
        else \(X_{i} \leftarrow X_{i-1}, Y_{i} \leftarrow Y_{i-1} \backslash\left\{u_{i}\right\}\).
    return \(X_{n}\) (or equivalently \(Y_{n}\) ).
```

The rest of this section is devoted for proving Theorem I.1, i.e., we prove that the approximation guarantee of Algorithm 1 is $1 / 3$. Denote by $a_{i}$ the change in value
of the first solution if element $u_{i}$ is added to it in the $i$ th iteration, i.e., $f\left(X_{i-1} \cup\left\{u_{i}\right\}\right)-f\left(X_{i-1}\right)$. Similarly, denote by $b_{i}$ the change in value of the second solution if element $u_{i}$ is removed from it in the $i$ th iteration, i.e., $f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)-f\left(Y_{i-1}\right)$. We start with the following useful lemma.

Lemma II.1. For every $1 \leq i \leq n, a_{i}+b_{i} \geq 0$.
Proof: Notice that $\left(X_{i-1} \cup\left\{u_{i}\right\}\right) \cup\left(Y_{i} \backslash\left\{u_{i}\right\}\right)=Y_{i-1}$ and $\left(X_{i-1} \cup\left\{u_{i}\right\}\right) \cap\left(Y_{i} \backslash\left\{u_{i}\right\}\right)=X_{i-1}$. By combining both observations with submodularity, one gets:

$$
\begin{aligned}
a_{i}+b_{i}= & {\left[f\left(X_{i-1} \cup\left\{u_{i}\right\}\right)-f\left(X_{i-1}\right)\right]+} \\
& {\left[f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)-f\left(Y_{i-1}\right)\right] } \\
= & {\left[f\left(X_{i-1} \cup\left\{u_{i}\right\}\right)+f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)\right]-} \\
& {\left[f\left(X_{i-1}\right)+f\left(Y_{i-1}\right)\right] \geq 0 . }
\end{aligned}
$$

Let $O P T$ denote an optimal solution. Define $O P T_{i} \triangleq$ $\left(O P T \cup X_{i}\right) \cap Y_{i}$. Thus, $O P T_{i}$ coincides with $X_{i}$ and $Y_{i}$ on elements $1, \ldots, i$, and it coincides with $O P T$ on elements $i+1, \ldots, n$. Note that $O P T_{0}=O P T$ and the output of the algorithm is $O P T_{n}=X_{n}=Y_{n}$. Examine the sequence $f\left(O P T_{0}\right), \ldots, f\left(O P T_{n}\right)$, which starts with $f(O P T)$ and ends with the value of the output of the algorithm. The main idea of the proof is to bound the total loss of value along this sequence. This goal is achieved by the following lemma which upper bounds the loss in value between every two consecutive steps in the sequence. Formally, the loss of value, i.e., $f\left(O P T_{i-1}\right)-f\left(O P T_{i}\right)$, is no more than the total increase in value of both solutions maintained by the algorithm, i.e., $\left[f\left(X_{i}\right)-f\left(X_{i-1}\right)\right]+\left[f\left(Y_{i}\right)-f\left(Y_{i-1}\right)\right]$.

Lemma II.2. For every $1 \leq i \leq n$,

$$
\begin{aligned}
f\left(O P T_{i-1}\right)- & f\left(O P T_{i}\right) \leq \\
& {\left[f\left(X_{i}\right)-f\left(X_{i-1}\right)\right]+\left[f\left(Y_{i}\right)-f\left(Y_{i-1}\right)\right] }
\end{aligned}
$$

Before proving Lemma II.2, let us show that Theorem I. 1 follows from it.

Proof of Theorem I.1: Summing up Lemma II. 2 for every $1 \leq i \leq n$ gives:

$$
\begin{aligned}
& \sum_{i=1}^{n}\left[f\left(O P T_{i-1}\right)-f\left(O P T_{i}\right)\right] \leq \\
& \sum_{i=1}^{n}\left[f\left(X_{i}\right)-f\left(X_{i-1}\right)\right]+\sum_{i=1}^{n}\left[f\left(Y_{i}\right)-f\left(Y_{i-1}\right)\right]
\end{aligned}
$$

The above sum is telescopic. Collapsing it, we get:

$$
\begin{aligned}
& f\left(O P T_{0}\right)-f\left(O P T_{n}\right) \leq \\
& {\left[f\left(X_{n}\right)-f\left(X_{0}\right)\right]+\left[f\left(Y_{n}\right)-f\left(Y_{0}\right)\right] \leq f\left(X_{n}\right)+f\left(Y_{n}\right)}
\end{aligned}
$$

Recalling the definitions of $O P T_{0}$ and $O P T_{n}$, we obtain that $f\left(X_{n}\right)=f\left(Y_{n}\right) \geq f(O P T) / 3$.

It is left to prove Lemma II.2.


Figure 1. Tight example for Algorithm 1. The weight of the dashed edges is $1-\varepsilon$. All other edges have weight of 1 .

Proof of Lemma II.2: Assume without loss of generality that $a_{i} \geq b_{i}$, i.e., $X_{i} \leftarrow X_{i-1} \cup\left\{u_{i}\right\}$ and $Y_{i} \leftarrow Y_{i-1}$ (the other case is analogous). Notice that in this case $O P T_{i}=\left(O P T \cup X_{i}\right) \cap Y_{i}=O P T_{i-1} \cup\left\{u_{i}\right\}$ and $Y_{i}=Y_{i-1}$. Thus the inequality that we need to prove becomes:
$f\left(O P T_{i-1}\right)-f\left(O P T_{i-1} \cup\left\{u_{i}\right\}\right) \leq f\left(X_{i}\right)-f\left(X_{i-1}\right)=a_{i}$.
We now consider two cases. If $u_{i} \in O P T$, then the left hand side of the last inequality is 0 , and all we need to show is that $a_{i}$ is non-negative. This is true since $a_{i}+b_{i} \geq 0$ by Lemma II.1, and $a_{i} \geq b_{i}$ by our assumption.

If $u_{i} \notin O P T$, then also $u_{i} \notin O P T_{i-1}$, and thus:

$$
\begin{aligned}
& f\left(O P T_{i-1}\right)-f\left(O P T_{i-1} \cup\left\{u_{i}\right\}\right) \leq \\
& \\
& f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)-f\left(Y_{i-1}\right)=b_{i} \leq a_{i}
\end{aligned}
$$

The first inequality follows by submodularity: $O P T_{i-1}=$ $\left(\left(O P T \cup X_{i-1}\right) \cap Y_{i-1}\right) \subseteq Y_{i-1} \backslash\left\{u_{i}\right\}$ (recall that $u_{i} \in Y_{i-1}$ and $\left.u_{i} \notin O P T_{i-1}\right)$. The second inequality follows from our assumption that $a_{i} \geq b_{i}$.

## A. Tight Example

We now show that the analysis of Algorithm 1 is tight.
Theorem II.3. For an arbitrarily small constant $\varepsilon>0$, there exists a submodular function for which Algorithm 1 provides only $(1 / 3+\varepsilon)$-approximation.

Proof: The proof follows by analyzing Algorithm 1 on the cut function of the weighted digraph given in Figure 1. The maximum weight cut in this digraph is $\left\{u_{1}, u_{4}, u_{5}\right\}$. This cut has weight of $6-2 \varepsilon$. We claim that Algorithm 1 outputs the cut $\left\{u_{2}, u_{3}, u_{4}, u_{5}\right\}$, whose value is only 2 (assuming the nodes of the graph are considered at the given order). Hence, the approximation guarantee of Algorithm 1 on the above instance is:

$$
\frac{2}{6-2 \varepsilon}=\frac{1}{3-\varepsilon} \leq \frac{1}{3}+\varepsilon
$$

Let us track the execution of the algorithm. Let $X_{i}, Y_{i}$ be the solutions maintained by the algorithm. Initially $X_{0}=$ $\emptyset, Y_{0}=\left\{u_{1}, u_{2}, u_{3}, u_{4}, u_{5}\right\}$. Note that in case of a tie $\left(a_{i}=\right.$ $b_{i}$ ), Algorithm 1 takes node $u_{i}$.

1) In the first iteration the algorithm considers $u_{1}$. Adding this node to $X_{0}$ increases the value of this solution by $2-2 \varepsilon$. On the other hand, removing this node from $Y_{0}$ increases the value of $Y_{0}$ by 2 . Hence, $X_{1} \leftarrow$ $X_{0}, Y_{1} \leftarrow Y_{0} \backslash\left\{u_{1}\right\}$.
2) Let us inspect the next two iterations in which the algorithm considers $u_{2}, u_{3}$. One can easily verify that these two iterations are independent, hence, we consider only $u_{2}$. Adding $u_{2}$ to $X_{1}$ increases its value by 1 . On the other hand, removing $u_{2}$ from $Y_{1}=\left\{u_{2}, u_{3}, u_{4}, u_{5}\right\}$ also increases the value of $Y_{1}$ by 1 . Thus, the algorithm adds $u_{2}$ to $X_{1}$. Since $u_{2}$ and $u_{3}$ are symmetric, the algorithm also adds $u_{3}$ to $X_{1}$. Thus, at the end of these two iterations $X_{3}=$ $X_{1} \cup\left\{u_{2}, u_{3}\right\}=\left\{u_{2}, u_{3}\right\}, Y_{3}=\left\{u_{2}, u_{3}, u_{4}, u_{5}\right\}$.
3) Finally, the algorithm considers $u_{4}$ and $u_{5}$. These two iterations are also independent so we consider only $u_{4}$. Adding $u_{4}$ to $X_{3}$ does not increase the value of $X_{3}$. Also removing $u_{4}$ from $Y_{3}$ does not increase the value of $Y_{3}$. Thus, the algorithm adds $u_{4}$ to $X_{3}$. Since $u_{4}$ and $u_{5}$ are symmetric, the algorithm also adds $u_{5}$ to $X_{3}$. Thus, we get $X_{5}=Y_{5}=\left\{u_{2}, u_{3}, u_{4}, u_{5}\right\}$.

## III. A Randomized (1/2)-Approximation

## Algorithm

In this section we present a randomized algorithm achieving a tight ( $1 / 2$ )-approximation for USM. Algorithm 1 compared the marginal profits $a_{i}$ and $b_{i}$, and based on this comparison it made a greedy deterministic decision whether to include or exclude $u_{i}$ from its output. The random algorithm we next present makes a "smoother" decision, based on the values $a_{i}$ and $b_{i}$. In each step, it randomly chooses whether to include or exclude $u_{i}$ from the output with probability derived from the values $a_{i}$ and $b_{i}$. A formal description of the algorithm appears as Algorithm 2.

The rest of this section is devoted to proving that Algorithm 2 provides an approximation guarantee of $1 / 2$ to USM. In Appendix A, we describe another algorithm which can be thought of as the continuous counterpart of Algorithm 2. This algorithm achieves the same approximation ratio (up to low order terms), but keeps a fractional inner state.

Let us first introduce some notation. Notice that for every $1 \leq i \leq n, X_{i}$ and $Y_{i}$ are random variables denoting the sets of elements in the two solutions generated by the algorithm at the end of the $i$ th iteration. As in Section II, let us define the following random variables: $O P T_{i} \triangleq\left(O P T \cup X_{i}\right) \cap Y_{i}$. Note that as before, $X_{0}=\emptyset$, $Y_{0}=\mathcal{N}$ and $O P T_{0}=O P T$. Moreover, the following

```
Algorithm 2: RandomizedUSM \((f, \mathcal{N})\)
    \(X_{0} \leftarrow \emptyset, Y_{0} \leftarrow \mathcal{N}\).
    for \(i=1\) to \(n\) do
        \(a_{i} \leftarrow f\left(X_{i-1} \cup\left\{u_{i}\right\}\right)-f\left(X_{i-1}\right)\).
        \(b_{i} \leftarrow f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)-f\left(Y_{i-1}\right)\).
        \(a_{i}^{\prime} \leftarrow \max \left\{a_{i}, 0\right\}, b_{i}^{\prime} \leftarrow \max \left\{b_{i}, 0\right\}\).
        with probability \(a_{i}^{\prime} /\left(a_{i}^{\prime}+b_{i}^{\prime}\right)^{*}\) do:
        \(X_{i} \leftarrow X_{i-1} \cup\left\{u_{i}\right\}, Y_{i} \leftarrow Y_{i-1}\).
        else (with the compliment probability \(\left.b_{i}^{\prime} /\left(a_{i}^{\prime}+b_{i}^{\prime}\right)\right)\)
        do: \(X_{i} \leftarrow X_{i-1}, Y_{i} \leftarrow Y_{i-1} \backslash\left\{u_{i}\right\}\).
    return \(X_{n}\) (or equivalently \(Y_{n}\) ).
    \({ }^{*}\) If \(a_{i}^{\prime}=b_{i}^{\prime}=0\), we assume \(a_{i}^{\prime} /\left(a_{i}^{\prime}+b_{i}^{\prime}\right)=1\).
```

always holds: $O P T_{n}=X_{n}=Y_{n}$. The proof idea is similar to that of the deterministic algorithm in Section II when considering the sequence $\mathbb{E}\left[f\left(O P T_{0}\right)\right], \ldots, \mathbb{E}\left[f\left(O P T_{n}\right)\right]$. This sequence starts with $f(O P T)$ and ends with the expected value of the algorithm's output. The following lemma upper bounds the loss of two consecutive elements in the sequence. Formally, $\mathbb{E}\left[f\left(O P T_{i-1}\right)-f\left(O P T_{i}\right)\right]$ is upper bounded by the average expected change in the value of the two solutions maintained by the algorithm, i.e., $1 / 2 \cdot \mathbb{E}\left[\left(f\left(X_{i}\right)-f\left(X_{i-1}\right)+f\left(Y_{i}\right)-f\left(Y_{i-1}\right)\right]\right.$.

Lemma III.1. For every $1 \leq i \leq n$,

$$
\begin{align*}
& \mathbb{E}\left[f\left(O P T_{i-1}\right)-f\left(O P T_{i}\right)\right] \leq \\
& \frac{1}{2} \cdot \mathbb{E}\left[\left(f\left(X_{i}\right)-f\left(X_{i-1}\right)+f\left(Y_{i}\right)-f\left(Y_{i-1}\right)\right]\right. \tag{1}
\end{align*}
$$

where expectations are taken over the random choices of the algorithm.

Before proving Lemma III.1, let us show that Theorem I. 2 follows from it.

Proof of Theorem I.2: Summing up Lemma III. 1 for every $1 \leq i \leq n$ yields:

$$
\begin{aligned}
& \sum_{i=1}^{n} \mathbb{E}\left[f\left(O P T_{i-1}\right)-f\left(O P T_{i}\right)\right] \leq \\
& \frac{1}{2} \cdot \sum_{i=1}^{n} \mathbb{E}\left[f\left(X_{i}\right)-F\left(X_{i-1}\right)+f\left(Y_{i}\right)-F\left(Y_{i-1}\right)\right]
\end{aligned}
$$

The above sum is telescopic. Collapsing it, we get:

$$
\begin{aligned}
& \mathbb{E}\left[f\left(O P T_{0}\right)-f\left(O P T_{n}\right)\right] \leq \\
& \frac{1}{2} \cdot \mathbb{E}\left[f\left(X_{n}\right)-f\left(X_{0}\right)+f\left(Y_{n}\right)-f\left(Y_{0}\right)\right] \leq \\
& \frac{\mathbb{E}\left[f\left(X_{n}\right)+f\left(Y_{n}\right)\right]}{2}
\end{aligned}
$$

Recalling the definitions of $O P T_{0}$ and $O P T_{n}$, we obtain that $\mathbb{E}\left[f\left(X_{n}\right)\right]=\mathbb{E}\left[f\left(Y_{n}\right)\right] \geq f(O P T) / 2$.

It is left to prove Lemma III.1.

Proof of Lemma III.1: Notice that it suffices to prove Inequality (1) conditioned on any event of the form $X_{i-1}=$ $S_{i-1}$, when $S_{i-1} \subseteq\left\{u_{1}, \ldots, u_{i-1}\right\}$ and the probability that $X_{i-1}=S_{i-1}$ is non-zero. Hence, fix such an event for a given $S_{i-1}$. The rest of the proof implicitly assumes everything is conditioned on this event. Notice that due to the conditioning, the following quantities become constants:

- $Y_{i-1}=S_{i-1} \cup\left\{u_{i}, \ldots, u_{n}\right\}$.
- $O P T_{i-1} \triangleq\left(O P T \cup X_{i-1}\right) \cap Y_{i-1}=$

$$
S_{i-1} \cup\left(O P T \cap\left\{u_{i}, \ldots, u_{n}\right\}\right)
$$

- $a_{i}$ and $b_{i}$.

Moreover, by Lemma II.1, $a_{i}+b_{i} \geq 0$. Thus, it cannot be that both $a_{i}, b_{i}$ are strictly less than zero. Hence, we only need to consider the following 3 cases:

Case $1\left(\mathbf{a}_{\mathbf{i}} \geq \mathbf{0}\right.$ and $\left.\mathbf{b}_{\mathbf{i}} \leq \mathbf{0}\right)$ : In this case $a_{i}^{\prime} /\left(a_{i}^{\prime}+\right.$ $\left.b_{i}^{\prime}\right)=1$, and so the following always happen: $Y_{i}=Y_{i-1}=$ $S_{i-1} \cup\left\{u_{i}, \ldots, u_{n}\right\}$ and $X_{i} \leftarrow S_{i-1} \cup\left\{u_{i}\right\}$. Hence, $f\left(Y_{i}\right)-$ $f\left(Y_{i-1}\right)=0$. Also, by our definition $O P T_{i}=\left(O P T \cup X_{i}\right) \cap$ $Y_{i}=O P T_{i-1} \cup\left\{u_{i}\right\}$. Thus, we are left to prove that:

$$
\begin{aligned}
& f\left(O P T_{i-1}\right)-f\left(O P T_{i-1} \cup\left\{u_{i}\right\}\right) \leq \\
& \frac{1}{2} \cdot\left[f\left(X_{i}\right)-f\left(X_{i-1}\right)\right]=\frac{a_{i}}{2}
\end{aligned}
$$

If $u_{i} \in O P T$, then the left hand side of the last expression is 0 , which is clearly no larger than the non-negative $a_{i} / 2$. If $u_{i} \notin O P T$, then:

$$
\begin{aligned}
& f\left(O P T_{i-1}\right)-f\left(O P T_{i-1} \cup\left\{u_{i}\right\}\right) \leq \\
& f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)-f\left(Y_{i-1}\right)=b_{i} \leq 0 \leq a_{i} / 2
\end{aligned}
$$

The first inequality follows from submodularity since $O P T_{i-1} \triangleq\left(O P T \cup X_{i-1}\right) \cap Y_{i-1} \subseteq Y_{i-1} \backslash\left\{u_{i}\right\}$ (note that $u_{i} \in Y_{i-1}$ and $\left.u_{i} \notin O P T_{i-1}\right)$.

Case $2\left(\mathbf{a}_{\mathbf{i}}<\mathbf{0}\right.$ and $\left.\mathbf{b}_{\mathbf{i}} \geq \mathbf{0}\right)$ : This case is analogous to the previous one, and therefore, we omit its proof.

Case $3\left(\mathbf{a}_{\mathbf{i}} \geq \mathbf{0}\right.$ and $\left.\mathbf{b}_{\mathbf{i}}>\mathbf{0}\right)$ : In this case $a_{i}^{\prime}=a_{i}, b_{i}^{\prime}=$ $b_{i}$. Therefore, with probability $a_{i} /\left(a_{i}+b_{i}\right)$ the following events happen: $X_{i} \leftarrow X_{i-1} \cup\left\{u_{i}\right\}$ and $Y_{i} \leftarrow Y_{i-1}$, and with probability $b_{i} /\left(a_{i}+b_{i}\right)$ the following events happen: $X_{i} \leftarrow X_{i-1}$ and $Y_{i} \leftarrow Y_{i-1} \backslash\left\{u_{i}\right\}$. Thus,

$$
\begin{align*}
& \mathbb{E}\left[f\left(X_{i}\right)-f\left(X_{i-1}\right)+f\left(Y_{i}\right)-f\left(Y_{i-1}\right)\right]= \\
& \frac{a_{i}}{a_{i}+b_{i}} \cdot\left[f\left(X_{i-1} \cup\left\{u_{i}\right\}\right)-f\left(X_{i-1}\right)\right]+ \\
& \\
& \frac{b_{i}}{a_{i}+b_{i}} \cdot\left[f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)-f\left(Y_{i-1}\right)\right]  \tag{2}\\
& = \\
& \frac{a_{i}^{2}+b_{i}^{2}}{a_{i}+b_{i}} .
\end{align*}
$$

Next, we upper bound $\mathbb{E}\left[f\left(O P T_{i-1}\right)-f\left(O P T_{i}\right)\right]$. As $O P T_{i}=\left(O P T \cup X_{i}\right) \cap Y_{i}$, we get,

$$
\begin{align*}
& \mathbb{E}\left[f\left(O P T_{i-1}\right)-f\left(O P T_{i}\right)\right]= \\
& \quad \frac{a_{i}}{a_{i}+b_{i}} \cdot\left[f\left(O P T_{i-1}\right)-f\left(O P T_{i-1} \cup\left\{u_{i}\right\}\right)\right]+ \\
& \quad \frac{b_{i}}{a_{i}+b_{i}} \cdot\left[f\left(O P T_{i-1}\right)-f\left(O P T_{i-1} \backslash\left\{u_{i}\right\}\right)\right] \\
& \leq \frac{a_{i} b_{i}}{a_{i}+b_{i}} \tag{3}
\end{align*}
$$

The final inequality follows by considering two cases. Note first that $u_{i} \in Y_{i-1}$ and $u_{i} \notin X_{i-1}$. If $u_{i} \notin O P T_{i-1}$ then the second term of the left hand side of the last inequality equals zero. Moreover, $O P T_{i-1} \triangleq\left(O P T \cup X_{i-1}\right) \cap Y_{i-1} \subseteq$ $Y_{i-1} \backslash\left\{u_{i}\right\}$, and therefore, by submodularity,

$$
\begin{aligned}
f\left(O P T_{i-1}\right)- & f\left(O P T_{i-1} \cup\left\{u_{i}\right\}\right) \leq \\
& f\left(Y_{i-1} \backslash\left\{u_{i}\right\}\right)-f\left(Y_{i-1}\right)=b_{i}
\end{aligned}
$$

If $u_{i} \in O P T_{i-1}$ then the first term of the left hand side of the inequality equals zero, and we also have $X_{i-1} \subseteq$ $\left(\left(O P T \cup X_{i-1}\right) \cap Y_{i-1}\right) \backslash\left\{u_{i}\right\}=O P T_{i-1} \backslash\left\{u_{i}\right\}$. Thus, by submodularity,

$$
\begin{aligned}
& f\left(O P T_{i-1}\right)-f\left(O P T_{i-1} \backslash\left\{u_{i}\right\}\right) \leq \\
& \\
& \quad f\left(X_{i-1} \cup\left\{u_{i}\right\}\right)-f\left(X_{i-1}\right)=a_{i}
\end{aligned}
$$

Plugging (2) and (3) into Inequality (1), we get the following:

$$
\frac{a_{i} b_{i}}{a_{i}+b_{i}} \leq \frac{1}{2} \cdot\left(\frac{a_{i}^{2}+b_{i}^{2}}{a_{i}+b_{i}}\right)
$$

which can be easily verified.

## IV. Approximation Algorithms for Submodular Max-SAT and Submodular Welfare with 2 Players

In this section we build upon ideas from the previous sections to obtain linear time tight (3/4)-approximation algorithms for both the Submodular Max-SAT (SSAT) and Submodular Welfare (SW) with 2 players problems. In contrast to USM, tight approximations are already known for the above two problems [13]. However, the algorithms we present in this work, in addition to providing tight approximations, also run in linear time. We require the definition of a monotone submodular function: a submodular function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}^{+}$is monotone if for every $A \subseteq B \subseteq \mathcal{N}$, $f(A) \leq f(B)$. Moreover, a monotone submodular function $f$ is normalized if $f(\emptyset)=0$.

## A. A (3/4)-Approximation for Submodular Max-SAT

In the Submodular Max-SAT problem we are given a CNF formula $\Psi$ with a set $\mathcal{C}$ of clauses over a set $\mathcal{N}$ of variables, and a normalized monotone submodular function $f: 2^{\mathcal{C}} \rightarrow \mathbb{R}^{+}$over the set of clauses. Given an assignment $\phi: \mathcal{N} \rightarrow\{0,1\}$, let $C(\phi) \subseteq \mathcal{C}$ be the set of clauses satisfied
by $\phi$. The goal is to find an assignment $\phi$ that maximizes $f(C(\phi))$.

Usually, an assignment $\phi$ can give each variable exactly a single truth value. However, for the sake of the algorithm we extend the notion of assignments, and think of an extended assignment $\phi^{\prime}$ which is a relation $\phi^{\prime} \subseteq \mathcal{N} \times\{0,1\}$. That is, the assignment $\phi^{\prime}$ can assign up to 2 truth values to each variable. A clause $C$ is satisfied by an (extended) assignment $\phi^{\prime}$, if there exists a positive literal in the clause which is assigned to truth value 1 , or there exists a negative literal in the clause which is assigned a truth value 0 . Note again that it might happen that some variable is assigned both 0 and 1 . Note also, that an assignment is a feasible solution to the original problem if and only if it assigns exactly one truth value to every variable of $\mathcal{N}$. Let $C\left(\phi^{\prime}\right)$ be the set of clauses satisfied by $\phi^{\prime}$. We define $g: \mathcal{N} \times\{0,1\} \rightarrow \mathbb{R}^{+}$ using $g\left(\phi^{\prime}\right) \triangleq f\left(C\left(\phi^{\prime}\right)\right)$. Using this notation, we can restate SSAT as the problem of maximizing the function $g$ over the set of feasible assignments.

## Observation IV.1. The function $g$ is a normalized monotone submodular function.

The algorithm we design for SSAT conducts $n$ iterations. It maintains at each iteration $1 \leq i \leq n$ two assignments $X_{i}$ and $Y_{i}$ which always satisfy that: $X_{i} \subseteq Y_{i}$. Initially, $X_{0}=\emptyset$ assigns no truth values to the variables, and $Y_{0}=\mathcal{N} \times\{0,1\}$ assigns both truth values to all variables. The algorithm considers the variables in an arbitrary order $u_{1}, u_{2}, \ldots, u_{n}$. For every variable $u_{i}$, the algorithm evaluates the marginal profit from assigning it only the truth value 0 in both assignments, and assigning it only the truth value 1 in both assignments. Based on these marginal values, the algorithm makes a random decision on the truth value assigned to $u_{i}$. After the algorithm considers a variable $u_{i}$, both assignments $X_{i}$ and $Y_{i}$ agree on the single truth value assigned to $u_{i}$. Thus, when the algorithm terminates both assignments are identical and feasible. A formal statement of the algorithm appears as Algorithm 3.

The proof of the Theorem I. 3 uses similar ideas to previous proofs in this paper, and is therefore, deferred to Appendix B. Note, however, that unlike for the previous algorithms, here, the fact that the algorithm runs in linear time requires a proof. The source of the difficulty is that we have an oracle access to $f$, but use the function $g$ in the algorithm. Therefore, we need to prove that we may implement all oracle queries of $g$ that are conducted during the execution of the algorithm in linear time.

A Note on Max-SAT: The well known Max-SAT problem is in fact a special case of SSAT where $f$ is a linear function. We note that Algorithm 3 can be applied to Max-SAT in order to achieve a $(3 / 4)$-approximation in linear time, however, this is not immediate. This result is summarized in the following theorem. The proof of this theorem is deferred to a full version of the paper.

```
Algorithm 3: RandomizedSSAT \((f, \Psi)\)
    \(X_{0} \leftarrow \emptyset, Y_{0} \leftarrow \mathcal{N} \times\{0,1\}\).
    for \(i=1\) to \(n\) do
        \(a_{i, 0} \leftarrow g\left(X_{i-1} \cup\left\{u_{i}, 0\right\}\right)-g\left(X_{i-1}\right)\).
        \(a_{i, 1} \leftarrow g\left(X_{i-1} \cup\left\{u_{i}, 1\right\}\right)-g\left(X_{i-1}\right)\).
        \(b_{i, 0} \leftarrow g\left(Y_{i-1} \backslash\left\{u_{i}, 0\right\}\right)-g\left(Y_{i-1}\right)\).
        \(b_{i, 1} \leftarrow g\left(Y_{i-1} \backslash\left\{u_{i}, 1\right\}\right)-g\left(Y_{i-1}\right)\).
        \(s_{i, 0} \leftarrow \max \left\{a_{i, 0}+b_{i, 1}, 0\right\}\).
        \(s_{i, 1} \leftarrow \max \left\{a_{i, 1}+b_{i, 0}, 0\right\}\).
        with probability \(s_{i, 0} /\left(s_{i, 0}+s_{i, 1}\right)^{*}\) do:
        \(X_{i} \leftarrow X_{i-1} \cup\left\{u_{i}, 0\right\}, Y_{i} \leftarrow Y_{i-1} \backslash\left\{u_{i}, 1\right\}\).
        else (with the compliment probability
        \(\left.s_{i, 1} /\left(s_{i, 0}+s_{i, 1}\right)\right)\) do:
        \(X_{i} \leftarrow X_{i-1} \cup\left\{u_{i}, 1\right\}, Y_{i} \leftarrow Y_{i-1} \backslash\left\{u_{i}, 0\right\}\).
    return \(X_{n}\) (or equivalently \(Y_{n}\) ).
    \({ }^{*}\) If \(s_{i, 0}=s_{i, 1}=0\), we assume \(s_{i, 0} /\left(s_{i, 0}+s_{i, 1}\right)=1\).
```

Theorem IV.2. Algorithm 3 has a linear time implementation for instances of Max-SAT.

## B. A (3/4)-Approximation for Submodular Welfare with 2 Players

The input for the Submodular Welfare problem consists of a ground set $\mathcal{N}$ of $n$ elements and $k$ players, each equipped with a normalized monotone submodular utility function $f_{i}: 2^{\mathcal{N}} \rightarrow \mathbb{R}^{+}$. The goal is to divide the elements among the players while maximizing the social welfare. Formally, the objective is to partition $\mathcal{N}$ into $\mathcal{N}_{1}, \mathcal{N}_{2}, \ldots, \mathcal{N}_{k}$ while maximizing $\sum_{i=1}^{k} f_{i}\left(\mathcal{N}_{i}\right)$.

We give below two different short proofs of Theorem I. 4 via reductions to SSAT and USM, respectively. The second proof is due to Vondrák [37].

Proof of Theorem I.4: We provide here two proofs.
Proof (1): Given an instance of SW with 2 players, construct an instance of SSAT as follows:

1) The set of variables is $\mathcal{N}$.
2) The CNF formula $\Psi$ consists of $2|\mathcal{N}|$ singleton clauses; one for every possible literal.
3) The objective function $f: 2^{\mathcal{C}} \rightarrow \mathbb{R}^{+}$is defined as following. Let $P \subseteq \mathcal{C}$ be the set of clauses of $\Psi$ consisting of positive literals. Then, $f(C)=f_{1}(C \cap$ $P)+f_{2}(C \cap(\mathcal{C} \backslash P))$.
Every assignment $\phi$ to this instance of SSAT corresponds to a solution of SW using the following rule: $\mathcal{N}_{1}=\{u \in$ $\mathcal{N} \mid \phi(u)=0\}$ and $\mathcal{N}_{2}=\{u \in \mathcal{N} \mid \phi(u)=1\}$. One can easily observe that this correspondence is reversible, and that each assignment has the same value as the solution it corresponds to. Hence, the above reduction preserves approximation ratios.

Moreover, queries of $f$ can be answered in constant time using the following technique. We track for every subset
$C \subseteq \mathcal{C}$ in the algorithm the subsets $C \cap P$ and $C \cap(\mathcal{C} \backslash P)$. For Algorithm 3 this can be done without effecting its running time. Then, whenever the value of $f(C)$ is queried, answering it simply requires making two oracle queries: $f_{1}(C \cap P)$ and $f_{2}(C \cap(\mathcal{C} \backslash P))$.

Proof (2): In any feasible solution to SW with two players, the set $\mathcal{N}_{1}$ uniquely determines the set $\mathcal{N}_{2}=\mathcal{N}-$ $\mathcal{N}_{1}$. Hence, the value of the solution as a function of $\mathcal{N}_{1}$ is given by $g\left(\mathcal{N}_{1}\right)=f_{1}\left(\mathcal{N}_{1}\right)+f_{2}\left(\mathcal{N}-\mathcal{N}_{1}\right)$. Thus, SW with two players can be restated as the problem of maximizing the function $g$ over the subsets of $\mathcal{N}$.

Observe that the function $g$ is a submodular function, but unlike $f_{1}$ and $f_{2}$, it is possibly non-monotone. Moreover, we can answer queries to the function $g$ using only two oracle queries to $f_{1}$ and $f_{2}{ }^{3}$. Thus, we obtain an instance of USM. We apply Algorithm 2 to this instance. Using the analysis of Algorithm 2 as is, provides only a (1/2)approximation for our problem. However, by noticing that $g(\varnothing)+g(\mathcal{N}) \geq f_{1}(\mathcal{N})+f_{2}(\mathcal{N}) \geq g(O P T)$, the claimed (3/4)-approximation is obtained.

## References

[1] A. A. Ageev and M. I. Sviridenko. An 0.828 approximation algorithm for the uncapacitated facility location problem. Discrete Appl. Math., 93:149-156, July 1999.
[2] Shabbir Ahmed and Alper Atamtürk. Maximizing a class of submodular utility functions. Mathematical Programming, 128:149-169, 2011.
[3] Gruia Calinescu, Chandra Chekuri, Martin Pal, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. To appear in SIAM Journal on Computing.
[4] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In FOCS, pages 575-584, 2010.
[5] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In STOC, pages 783-792, 2011.
[6] V. P. Cherenin. Solving some combinaotiral problems of optimal planning by the method of successive calculations. Proceedings of the Conference on Experiences and Perspectives on the Applications of Mathematical Methods and Electronic Computers in Planning, Mimeograph, Novosibirsk, 1962 (in Russian).
[7] G. Cornuejols, M. L. Fisher, and G. L. Nemhauser. Location of bank accounts to optimize float: an analytic study of exact and approximate algorithms. Management Sciences, 23:789-810, 1977.

[^2][8] G. Cornuejols, M. L. Fisher, and G. L. Nemhauser. On the uncapacitated location problem. Annals of Discrete Mathematics, 1:163-177, 1977.
[9] Shahar Dobzinski and Jan Vondrák. From query complexity to computational complexity. In STOC, 2012.
[10] Shaddin Dughmi, Tim Roughgarden, and Mukund Sundararajan. Revenue submodularity. In Proceedings of the 10th ACM conference on Electronic commerce, EC '09, pages 243-252, 2009.
[11] Uriel Feige, Vahab S. Mirrokni, and Jan Vondrák. Maximizing non-monotone submodular functions. In FOCS, pages 461471, 2007.
[12] Moran Feldman, Joseph (Seffi) Naor, and Roy Schwartz. Nonmonotone submodular maximization via a structural continuous greedy algorithm. In ICALP, pages 342-353, 2011.
[13] Moran Feldman, Joseph (Seffi) Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In FOCS, 2011.
[14] Shayan Oveis Gharan and Jan Vondrák. Submodular maximization by simulated annealing. In SODA, pages 1098-1117, 2011.
[15] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. Journal of the ACM, 42(6):1115-1145, 1995.
[16] Boris Goldengorin and Diptesh Ghosh. A multilevel search algorithm for the maximization of submodular functions applied to the quadratic cost partition problem. J. of Global Optimization, 32(1):65-82, May 2005.
[17] Boris Goldengorin, Gerard Sierksma, Gert A. Tijssen, and Michael Tso. The data-correcting algorithm for the minimization of supermodular functions. Manage. Sci., 45(11):15391551, November 1999.
[18] Boris Goldengorin, Gert A. Tijssen, and Michael Tso. The maximization of submodular functions : old and new proofs for the correctness of the dichotomy algorithm. Research Report 99A17, University of Groningen, Research Institute SOM (Systems, Organisations and Management), 1999.
[19] Anupam Gupta, Aaron Roth, Grant Schoenebeck, and Kunal Talwar. Constrained non-monotone submodular maximization: offline and secretary algorithms. In Proceedings of the 6th international conference on Internet and network economics, WINE'10, pages 246-257. Springer-Verlag, 2010.
[20] Eran Halperin and Uri Zwick. Combinatorial approximation algorithms for the maximum directed cut problem. In SODA, pages 1-7, 2001.
[21] Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In Proceedings of the 17th international conference on World Wide Web, WWW '08, pages 189-198, 2008.
[22] Johan Hȧstad. Some optimal inapproximability results. J. ACM, 48:798-859, July 2001.
[23] Satoru Iwata and Kiyohito Nagano. Submodular function minimization under covering constraints. In FOCS, pages 671680, 2009.
[24] Stefanie Jegelka and Jeff Bilmes. Cooperative cuts: Graph cuts with submodular edge weights. Technical Report 189-03-2010, Max Planck Institute for Biological Cybernetics, Tuebingen, 2010.
[25] Richard M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, Complexity of Computer Computations, pages 85-103. Plenum Press, 1972.
[26] V. R. Khachaturov. Some problems of the consecutive calculation method and its applications to location problems. Ph.D. thesis, Central Economics \& Mathematics Institute, Russian Academy of Sciences, Moscow, 1968 (in Russian).
[27] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2 -variable csps? SIAM J. Comput., 37:319-357, April 2007.
[28] Ariel Kulik, Hadas Shachnai, and Tami Tamir. Maximizing submodular set functions subject to multiple linear constraints. In SODA, pages 545-554, 2009.
[29] Heesang Lee, George L. Nemhauser, and Yinhua Wang. Maximizing a submodular function by integer programming: Polyhedral results for the quadratic case. European Journal of Operational Research, 94(1):154-166, 1996.
[30] Jon Lee, Vahab S. Mirrokni, Viswanath Nagarajan, and Maxim Sviridenko. Maximizing non-monotone submodular functions under matroid or knapsack constraints. SIAM Journal on Discrete Mathematics, (4):2053-2078, 2010.
[31] Jon Lee, Maxim Sviridenko, and Jan Vondrák. Submodular maximization over multiple matroids via generalized exchange properties. In APPROX, pages 244-257, 2009.
[32] L. Lovász M. Grötschel and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. Combinatoria, 1(2):169-197, 1981.
[33] Michel Minoux. Accelerated greedy algorithms for maximizing submodular set functions. In J. Stoer, editor, Optimization Techniques, volume 7 of Lecture Notes in Control and Information Sciences, pages 234-243. Springer Berlin / Heidelberg, 1978.
[34] A. S. Schulz and N. A. Uhan. Approximating the least core and least core value of cooperative games with supermodular costs. Working paper, 2010.
[35] Zoya Svitkina and Lisa Fleischer. Submodular approximation: Sampling-based algorithms and lower bounds. In FOCS, pages 697-706, 2008.
[36] Luca Trevisan, Gregory B. Sorkin, Madhu Sudan, and David P. Williamson. Gadgets, approximation, and linear programming. SIAM J. Comput., 29:2074-2097, April 2000.
[37] Jan Vondrák. personal communication.
[38] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In STOC, pages 67-74, 2008.
[39] Jan Vondrák. Symmetry and approximability of submodular maximization problems. In $F O C S$, pages 651-670, 2009.

## Appendix

## A. A (1/2)-Approximation for USM using Fractional Values

In this section we present Algorithm 4, which is the continuous counterpart of Algorithm 2, presented in Section III.

In order to describe Algorithm 4, we need some notation. Given a submodular function $f: \mathcal{N} \rightarrow \mathbb{R}^{+}$, its multilinear extension is a function $F:[0,1]^{\mathcal{N}} \rightarrow \mathbb{R}^{+}$, whose value at point $x \in[0,1]^{\mathcal{N}}$ is the expected value of $f$ over a random subset $R(x) \subseteq \mathcal{N}$. The subset $R(x)$ contains each element $u \in \mathcal{N}$ independently with probability $x_{u}$. Formally, for every $x \in[0,1]^{\mathcal{N}}, F(x) \triangleq$ $\mathbb{E}[R(x)]=\sum_{S \subseteq \mathcal{N}} f(S) \prod_{u \in S} x_{u} \prod_{u \notin S}\left(1-x_{u}\right)$. For two vectors $x, y \in[\overline{0}, 1]^{\mathcal{N}}$, we use $x \vee y$ and $x \wedge y$ to denote the coordinate-wise maximum and minimum, respectively, of $x$ and $y$ (formally, $(x \vee y)_{u}=\max \left\{x_{u}, y_{u}\right\}$ and $\left.(x \wedge y)_{u}=\min \left\{x_{u}, y_{u}\right\}\right)$.

We abuse notations both in the description of the algorithm and in its analysis, and unify a set with its characteristic vector. We also assume that we have an oracle access to $F$. If this is not the case, then one can approximate the value of $F$ arbitrarily well using sampling. This makes the approximation ratio deteriorate by a low order term only (see, e.g., [3] for details).

```
Algorithm 4: MultilinearUSM \((f, \mathcal{N})\)
\(x_{0} \leftarrow \emptyset, y_{0} \leftarrow \mathcal{N}\).
    for \(i=1\) to \(n\) do
        \(a_{i}^{\prime} \leftarrow F\left(x_{i-1}+\left\{u_{i}\right\}\right)-F\left(x_{i-1}\right)\).
        \(b_{i}^{\prime} \leftarrow F\left(y_{i-1}-\left\{u_{i}\right\}\right)-F\left(y_{i-1}\right)\).
        \(a_{i}^{\prime} \leftarrow \max \left\{a_{i}, 0\right\}, b_{i}^{\prime} \leftarrow \max \left\{b_{i}, 0\right\}\).
        \(x_{i} \leftarrow x_{i-1}+\frac{a_{i}^{\prime}}{a_{i}^{\prime}+b_{i}^{\prime}} \cdot\left\{u_{i}\right\}^{*}\).
        \(y_{i} \leftarrow y_{i-1}-\frac{b_{i}^{\prime}}{a_{i}^{\prime}+b_{i}^{\prime}} \cdot\left\{u_{i}\right\}^{*}\).
    return a random set \(R\left(x_{n}\right)\) (or equivalently \(R\left(y_{n}\right)\) ).
    * If \(a_{i}^{\prime}=b_{i}^{\prime}=0\), we assume \(a_{i}^{\prime} /\left(a_{i}^{\prime}+b_{i}^{\prime}\right)=1\) and
    \(b_{i}^{\prime} /\left(a_{i}^{\prime}+b_{i}^{\prime}\right)=0\).
```

The main difference between Algorithms 2 and 4 is that Algorithm 2 chooses each element with some probability, whereas Algorithm 4 assigns a fractional value to the element. This requires the following modifications to the algorithm:

- The sets $X_{i}, Y_{i} \subseteq 2^{\mathcal{N}}$ are replaced by the vectors $x_{i}, y_{i} \in[0,1]^{\mathcal{N}}$.
- Algorithm 4 uses the multilinear extension $F$ instead of the original submodular function $f$.

Theorem A.1. If one has an oracle access to $F$, Algorithm 4 is a (1/2)-approximation algorithm for NSM. Otherwise, it can be implemented using sampling, and achieves an approximation ratio of $(1 / 2)-o(1)$.

The rest of this section is devoted to proving Theorem A.1. Similarly to Section III, define $O P T_{i} \triangleq(O P T \vee$ $\left.x_{i}\right) \wedge y_{i}$. Examine the sequence $F\left(O P T_{0}\right), \ldots, F\left(O P T_{n}\right)$. Notice that $O P T_{0}=O P T$, i.e., the sequence starts with the value of an optimal solution, and that $O P T_{n}=x_{n}=y_{n}$, i.e., the sequence ends at a fractional point whose value is the expected value of the algorithm's output. The following lemma upper bounds the loss of two consecutive elements in the sequence. Formally, $F\left(O P T_{i-1}\right)-F\left(O P T_{i}\right)$ is upper bounded by the average change in the value of the two solutions maintained by the algorithm, i.e.,

$$
1 / 2 \cdot\left[F\left(x_{i}\right)-F\left(x_{i-1}\right)+F\left(y_{i}\right)-F\left(y_{i-1}\right)\right] .
$$

Lemma A.2. For every $1 \leq i \leq n$,

$$
\begin{aligned}
F\left(O P T_{i-1}\right)- & F\left(O P T_{i}\right) \leq \\
& \frac{1}{2} \cdot\left[F\left(x_{i}\right)-F\left(x_{i-1}\right)+F\left(y_{i}\right)-F\left(y_{i-1}\right)\right]
\end{aligned}
$$

Before proving Lemma A. 2, let us show that Theorem A. 1 follows from it.

Proof of Theorem A.1: Summing up Lemma A. 2 for every $1 \leq i \leq n$ gives:

$$
\begin{aligned}
& \sum_{i=1}^{n}\left[F\left(O P T_{i-1}\right)-F\left(O P T_{i}\right)\right] \leq \\
& \frac{1}{2} \cdot \sum_{i=1}^{n}\left[F\left(x_{i}\right)-F\left(x_{i-1}\right)\right]+\frac{1}{2} \cdot \sum_{i=1}^{n}\left[F\left(y_{i}\right)-F\left(y_{i-1}\right)\right]
\end{aligned}
$$

The above sum is telescopic. Collapsing it, we get:

$$
\begin{aligned}
& F\left(O P T_{0}\right)-F\left(O P T_{n}\right) \leq \\
& \frac{1}{2} \cdot\left[F\left(x_{n}\right)-F\left(x_{0}\right)\right]+\frac{1}{2} \cdot\left[F\left(y_{n}\right)-F\left(y_{0}\right)\right] \leq \\
& \frac{F\left(x_{n}\right)+F\left(y_{n}\right)}{2}
\end{aligned}
$$

Recalling the definitions of $O P T_{0}$ and $O P T_{n}$, we obtain that $F\left(x_{n}\right)=F\left(y_{n}\right) \geq f(O P T) / 2$.

It is left to prove Lemma A.2.
Proof of Lemma A.2: By Lemma II.1, $a_{i}+b_{i} \geq 0$, therefore, it cannot be that both $a_{i}, b_{i}$ are strictly less than zero. Thus, we have 3 cases to consider.

Case $1\left(\mathbf{a}_{\mathbf{i}} \geq \mathbf{0}\right.$ and $\left.\mathbf{b}_{\mathbf{i}} \leq \mathbf{0}\right)$ : In this case $a_{i}^{\prime} /\left(a_{i}^{\prime}+\right.$ $\left.b_{i}^{\prime}\right)=1$, and so $y_{i}=y_{i-1}, x_{i} \leftarrow x_{i-1} \vee\left\{u_{i}\right\}$. Hence, $F\left(y_{i}\right)-F\left(y_{i-1}\right)=0$. Also, by our definition $O P T_{i}=$ $\left(O P T \vee x_{i}\right) \wedge y_{i}=O P T_{i-1} \vee\left\{u_{i}\right\}$. Thus, we are left to prove that:

$$
\begin{aligned}
F\left(O P T_{i-1}\right)- & F\left(O P T_{i-1} \vee\left\{u_{i}\right\}\right) \leq \\
& \frac{1}{2} \cdot\left[F\left(x_{i}\right)-F\left(x_{i-1}\right)\right]=a_{i} / 2
\end{aligned}
$$

If $u_{i} \in O P T$, then the left hand side of the above equation is 0 , which is clearly no larger than the non-negative $a_{i} / 2$. If $u_{i} \notin O P T$, then:

$$
\begin{aligned}
& F\left(O P T_{i-1}\right)-F\left(O P T_{i-1} \vee\left\{u_{i}\right\}\right) \leq \\
& \quad F\left(y_{i-1}-\left\{u_{i}\right\}\right)-F\left(y_{i-1}\right)=b_{i} \leq 0 \leq a_{i} / 2
\end{aligned}
$$

The first inequality follows from submodularity since $O P T_{i-1}=\left(\left(O P T \vee x_{i-1}\right) \wedge y_{i-1}\right) \leq y_{i-1}-\left\{u_{i}\right\}$ (note that in this case $\left(y_{i-1}\right)_{u_{i}}=1$ and $\left.\left(O P T_{i-1}\right)_{u_{i}}=0\right)$.

Case $2\left(\mathbf{a}_{\mathbf{i}}<\mathbf{0}\right.$ and $\left.\mathbf{b}_{\mathbf{i}} \geq \mathbf{0}\right)$ : This case is analogous to the previous one, and therefore, we omit its proof.

Case $3\left(\mathbf{a}_{\mathbf{i}} \geq \mathbf{0}\right.$ and $\left.\mathbf{b}_{\mathbf{i}}>\mathbf{0}\right)$ : In this case $a_{i}^{\prime}=a_{i}, b_{i}^{\prime}=$ $b_{i}$ and so, $x_{i} \leftarrow \bar{x}_{i-1}+\frac{a_{i}}{a_{i}+b_{i}} \cdot\left\{u_{i}\right\}$ and $y_{i} \leftarrow y_{i-1}-\frac{b_{i}}{a_{i}+b_{i}}$. $\left\{u_{i}\right\}$. Therefore, we have,

$$
\begin{align*}
& F\left(x_{i}\right)-F\left(x_{i-1}\right)= \\
& \frac{a_{i}}{a_{i}+b_{i}} \cdot\left[F\left(x_{i-1} \vee\left\{u_{i}\right\}\right)-F\left(x_{i-1}\right)\right]=\frac{a_{i}^{2}}{a_{i}+b_{i}} . \tag{4}
\end{align*}
$$

A similar argument shows that:

$$
\begin{equation*}
F\left(y_{i}\right)-F\left(y_{i-1}\right)=\frac{b_{i}^{2}}{a_{i}+b_{i}} \tag{5}
\end{equation*}
$$

Next, we upper bound $F\left(O P T_{i-1}\right)-F\left(O P T_{i}\right)$. For simplicity, let us assume $u_{i} \notin O P T$ (the proof for the other case is similar). Recall, that $O P T_{i}=\left(O P T \vee x_{i}\right) \wedge y_{i}$.

$$
\begin{align*}
& F\left(O P T_{i-1}\right)-F\left(O P T_{i}\right)= \\
& \frac{a_{i}}{a_{i}+b_{i}} \cdot\left[F\left(O P T_{i-1}\right)-F\left(O P T_{i-1} \vee\left\{u_{i}\right\}\right)\right] \leq \\
& \frac{a_{i}}{a_{i}+b_{i}} \cdot\left[F\left(y_{i-1}-\left\{u_{i}\right\}\right)-F\left(y_{i-1}\right)\right]=\frac{a_{i} b_{i}}{a_{i}+b_{i}} \tag{6}
\end{align*}
$$

The inequality follows from the submodularity of $f$ since,

$$
O P T_{i-1}=\left(\left(O P T \vee x_{i-1}\right) \wedge y_{i-1}\right) \leq y_{i-1}-\left\{u_{i}\right\}
$$

(note again that in this case $\left(y_{i-1}\right)_{u_{i}}=1$ and $\left.\left(O P T_{i-1}\right)_{u_{i}}=0\right)$. Plugging (4), (5) and (6) into the inequality that we need to prove, we get the following:

$$
\frac{a_{i} b_{i}}{a_{i}+b_{i}} \leq \frac{1}{2} \cdot \frac{a_{i}^{2}+b_{i}^{2}}{a_{i}+b_{i}}
$$

which can be easily verified.

## B. Proof of Theorem I. 3

In this section we prove that Algorithm 3 is a randomized linear time (3/4)-approximation algorithm for SSAT. As a first step we state the following useful lemma.

Lemma A.3. For every $1 \leq i \leq n$,

$$
\begin{align*}
& \mathbb{E}\left[g\left(O P T_{i-1}\right)-g\left(O P T_{i}\right)\right] \leq \\
& \frac{1}{2} \cdot \mathbb{E}\left[\left(g\left(X_{i}\right)-g\left(X_{i-1}\right)+g\left(Y_{i}\right)-g\left(Y_{i-1}\right)\right]\right. \tag{7}
\end{align*}
$$

where expectations are taken over the random choices of the algorithm.

The proof of this lemma is deferred to a full version of the paper. Let us just note that the proof follows the lines of Lemma III.1. Let us show that the approximation guarantee of Theorem I. 3 follows from the above lemma. The proof that Algorithm 3 can be implemented to run in linear time is deferred to a full version of the paper.

Proof of approximation guarantee of Algorithm 3: Summing up Lemma A. 3 for every $1 \leq i \leq n$ we get,

$$
\begin{aligned}
& \sum_{i=1}^{n} \mathbb{E}\left[g\left(O P T_{i-1}\right)-g\left(O P T_{i}\right)\right] \leq \\
& \quad \frac{1}{2} \cdot \sum_{i=1}^{n} \mathbb{E}\left[g\left(X_{i}\right)-g\left(X_{i-1}\right)+g\left(Y_{i}\right)-g\left(Y_{i-1}\right)\right]
\end{aligned}
$$

The above sum is telescopic. Collapsing it, we get:

$$
\begin{aligned}
& \mathbb{E}\left[g\left(O P T_{0}\right)-g\left(O P T_{n}\right)\right] \\
& \leq \frac{1}{2} \cdot \mathbb{E}\left[g\left(X_{n}\right)-g\left(X_{0}\right)+g\left(Y_{n}\right)-g\left(Y_{0}\right)\right] \\
& \leq \frac{\mathbb{E}\left[g\left(X_{n}\right)+g\left(Y_{n}\right)-g\left(Y_{0}\right)\right]}{2}
\end{aligned}
$$

Recalling the definitions of $O P T_{0}$ and $O P T_{n}$, we obtain that:

$$
\mathbb{E}\left[g\left(X_{n}\right)\right]=\mathbb{E}\left[g\left(Y_{n}\right)\right] \geq g(O P T) / 2+g\left(Y_{0}\right) / 4
$$

The approximation ratio now follows from the observation that $Y_{0}$ satisfies all clauses of $\Psi$, and therefore, $g\left(Y_{0}\right) \geq$ $g(O P T)$.


[^0]:    ${ }^{\dagger}$ Research supported in part by ISF grant $954 / 11$, and by BSF grant 2010426.
    *Research supported in part by the Google Inter-university center for Electronic Markets, by ISF grant 954/11, and by BSF grant 2010426.
    ${ }^{1}$ An equivalent definition is that for every $A, B \subseteq \mathcal{N}: f(A)+f(B) \geq$ $f(A \cup B)+f(A \cap B)$.

[^1]:    ${ }^{2}$ The explicit representation of a submodular function might be exponential in the size of its ground set. Thus, it is standard to assume that the function is accessed via a value oracle. For a submodular function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}^{+}$, given a set $S \subseteq \mathcal{N}$, the value oracle returns the value of $f(S)$.

[^2]:    ${ }^{3}$ For every algorithm, assuming a representation of sets allowing addition and removal of only a single element at a time, one can maintain the complement sets of all sets maintained by the algorithm without changing the running time. Hence, we need not worry about the calculation of $\mathcal{N}$ $\mathcal{N}_{1}$.

