

A New Framework for Distributed Submodular Maximization

Rafael da Ponte Barbosa* Alina Ene[†] Huy L. Nguyễn[‡] Justin Ward[§]

August 11, 2016

Abstract

A wide variety of problems in machine learning, including exemplar clustering, document summarization, and sensor placement, can be cast as constrained submodular maximization problems. A lot of recent effort has been devoted to developing distributed algorithms for these problems. However, these results suffer from high number of rounds, suboptimal approximation ratios, or both. We develop a framework for bringing existing algorithms in the sequential setting to the distributed setting, achieving near optimal approximation ratios for many settings in only a constant number of MapReduce rounds. Our techniques also give a fast sequential algorithm for non-monotone maximization subject to a matroid constraint.

1 Introduction

The general problem of maximizing a submodular function appears in a variety of contexts, both in theory and practice. From a theoretical perspective, the class of submodular functions is extremely rich, including examples as varied as cut functions of graphs and digraphs, the Shannon entropy function, weighted coverage functions, and log-determinants. Recently, there has been a great deal of interest in practical applications of submodular optimization, as well. Variants of facility location, sampling, sensor selection, clustering, influence maximization in social networks, and welfare maximization problems are all instances of submodular maximization. In practice, many of these applications involve processing enormous datasets requiring *efficient, distributed algorithms*.

In contrast, most successful approaches for submodular maximization have been based on *sequential* greedy algorithms, including the standard greedy algorithm [24, 16], the continuous greedy algorithm [9, 15], and the double greedy algorithm [7]. Indeed, such approaches attain the best-possible, tight approximation guarantees in a variety of settings [14, 26, 13], but unfortunately they all share a common limitation, inherited from the standard greedy algorithm: they are inherently sequential. This presents a seemingly fundamental barrier to obtaining efficient, highly parallel variants of these algorithms.

1.1 Our Contributions

As demonstrated by the extensive prior works on submodular maximization, the community has a good understanding of the problem under remarkably general types of constraints, which are handled by a small collection of general algorithms. In contrast, the existing works in the distributed setting

*Department of Computer Science and DIMAP, University of Warwick. rafael@dcs.warwick.ac.uk.

[†]Department of Computer Science, Boston University. a.ene@bu.edu.

[‡]College of Computer and Information Science, Northeastern University. hlnghuyen@cs.princeton.edu.

[§]School of Computer and Communication Sciences, EPFL, Lausanne, Switzerland. justin.ward@epfl.ch.

are either tailored to special cases, giving approximation factors far from optimal or requiring a large number of distributed rounds. One cannot help but wonder if, instead of retracing the individual advances made in the sequential setting over the last few decades, it may be possible to obtain a generic technique to carry over the algorithms in the sequential setting to the parallel world.

In this work, we present a significant step toward resolving the above question. Our main contribution is a *generic parallel algorithm* that allows us to parallelize a broad class of sequential algorithm with almost no loss in performance. The crux of our approach is a *common abstraction* that allows us to capture and parallelize both the standard and continuous greedy algorithms, and it provides a novel unifying perspective for these algorithmic paradigms. Our framework leads to the first distributed algorithms that nearly match the state of the art approximation guarantees for the sequential setting in only a constant number of rounds. In the following, we summarize our main contributions.

A parallel greedy algorithm. We obtain the following general result by parallelizing the standard greedy algorithm:

Theorem 5.2. *Let $f : 2^V \rightarrow \mathbb{R}_+$ be a submodular function, and $\mathcal{I} \subseteq 2^V$ be a hereditary set system¹. For any $\epsilon > 0$ there is a randomized distributed $O(1/\epsilon)$ -round algorithm that can be implemented in the MapReduce framework². The algorithm is an $(\alpha - O(\epsilon))$ -approximation with constant probability for the problem $\max_{S \in \mathcal{I}} f(S)$, where α is the approximation ratio of the standard, sequential greedy algorithm for the same problem.*

Our constant number of rounds is a significant improvement over the sample and prune technique of [20], which requires a number of rounds depending logarithmically on the *value* of the single best element. Remarkably, even for the especially simple case of a cardinality constraint, no previous work could get close to the approximation ratio of the simple sequential greedy algorithm in a constant number of rounds. Our framework nearly matches the approximation ratio of greedy in all situations in a constant number of rounds and immediately resolves this problem.

A parallel continuous greedy algorithm. We obtain new distributed approximation results for maximization over matroids, by using a heavily discretized variant of the measured continuous greedy algorithm, obtaining approximation guarantees nearly matching those attained by the continuous greedy in the sequential setting.

Theorem 6.3. *Let $f : 2^V \rightarrow \mathbb{R}_+$ be a submodular function, and $\mathcal{I} \subseteq 2^V$ be a matroid. For any $\epsilon > 0$ there is a randomized distributed $O(1/\epsilon)$ -round algorithm that can be implemented in the MapReduce framework. The algorithm is an $(\alpha - O(\epsilon))$ -approximation with constant probability for the problem $\max_{S \in \mathcal{I}} f(S)$, where α is $(1 - 1/e)$ for monotone f and $1/e$ for general f .*

Improved two-round algorithms and fast sequential algorithms. We also give improved two-round approximations for non-monotone submodular maximization under hereditary constraints. We make use of the same “strong greedy property” utilized in [12] but attain approximation guarantees strictly better than were given there. Our algorithm is based on a combination of the standard greedy algorithm Greedy and an additional, arbitrary algorithm Alg. Again, we suppose that f is a (not necessarily monotone) submodular function and \mathcal{I} is any hereditary constraint. In the following theorems and throughout the paper, $n := |V|$ is the size of the ground set, $k := \max_{S \in \mathcal{I}} |S|$

¹A set system is hereditary if for any $S \in \mathcal{I}$, all subsets of S are also in \mathcal{I} .

²We define the MapReduce model in Section 2.

Monotone functions				Non-monotone functions			
Constraint	Rounds	Approx.	Citation	Constraint	Rounds	Approx.	Citation
cardinality	$O(\frac{\log \Delta}{\epsilon})$	$1 - \frac{1}{e} - \epsilon$	[20]	cardinality	2	$\frac{1 - \frac{1}{m}}{2 + e}$	[21]
	2	0.545	[21]		2	$(1 - \frac{1}{m})\frac{1}{e}(1 - \frac{1}{e})$	Theorem D.1
	$O(\frac{1}{\epsilon})$	$1 - \frac{1}{e} - \epsilon$	Theorem 5.2				
matroid	$O(\frac{\log \Delta}{\epsilon})$	$\frac{1}{2} - \epsilon$	[20]	matroid	2	$\frac{1}{10}$	[12]
	2	$\frac{1}{4}$	[12]		2	$\frac{1 - \frac{1}{m}}{2 + e}$	Theorem 7.1
	$O(\frac{1}{\epsilon})$	$1 - \frac{1}{e} - \epsilon$	Theorem 6.3		$O(\frac{1}{\epsilon})$	$\frac{1}{e} - \epsilon$	Theorem 6.3
p -system	$O(\frac{\log \Delta}{\epsilon})$	$\frac{1}{p+1} - \epsilon$	[20]	p -system	2	$\frac{1}{2 + 4(p+1)}$	[12]
	2	$\frac{1}{2(p+1)}$	[12]		2	$\frac{3(1 - \frac{1}{m})}{5p + 7 + \frac{2}{p}}$	Theorem 7.1
	$O(\frac{1}{\epsilon})$	$\frac{1}{p+1} - \epsilon$	Theorem 5.2				

Table 1: New results for distributed submodular maximization. Here $\Delta = \max_{i \in V} f(\{i\})$ and m is the number of machines. In the results of [20], in the number of rounds, Δ can be replaced by the maximum size of a solution. All algorithms in previous works and ours are randomized and the approximation guarantees stated hold in expectation, and they can be strengthened to hold with high probability by repeating the algorithms in parallel.

is the maximum size of a solution, and m is the number of machines employed by the distributed algorithm.

Theorem 7.1. *Suppose that Greedy satisfies the strong greedy property with constant γ and that Alg is a β -approximation for the problem $\max_{S \in \mathcal{I}} f(S)$. Then there is a randomized, two-round distributed algorithm that achieves a $(1 - \frac{1}{m})\frac{\beta\gamma}{\beta + \gamma}$ approximation in expectation for $\max_{S \in \mathcal{I}} f(S)$.*

We show that by simulating the machines in this last distributed algorithm, we also obtain a fast, sequential algorithm for maximizing a non-monotone submodular function subject to a matroid constraint. Our algorithm shows that one can preprocess the instance in $O(\frac{n}{\epsilon} \log n)$ time and obtain a set X of size $O(k/\epsilon)$ so that *it suffices to solve the problem on X* . By using a variant of the continuous greedy algorithm on the resulting set X , we obtain the following result.

Theorem 7.2. *There is a sequential, randomized $(\frac{1}{2+e} - \epsilon)$ -approximation algorithm for the problem $\max_{S \in \mathcal{I}} f(S)$, where \mathcal{I} is any matroid constraint, running in time $O(\frac{n}{\epsilon} \log n) + \text{poly}(\frac{k}{\epsilon})$.*

As a final application of our techniques, we obtain a very simple two-round distributed algorithm for monotone maximization subject to a cardinality constraint.

Theorem 7.3. *There is a randomized, two-round, distributed algorithm achieving a $\frac{1}{2} - \epsilon$ approximation in expectation for $\max_{S: |S| \leq k} f(S)$, where f is a monotone function.*

1.2 Techniques

In contrast with the previous framework by [20] which is based on repeatedly eliminating bad elements, our framework is more in line with the greedy approach of identifying good elements. The algorithm maintains a pool of good elements that is grown over several rounds. In each round,

the elements are partitioned randomly into groups. Each group selects the best among its elements and the good pool using the sequential algorithm. Finally, the best elements from all groups are added to the good pool. The best solution among the ones found in the execution of the algorithm is returned at the end. The previous works based on 2 rounds of MapReduce such as [12] can be viewed as a single phase of our algorithm. The first phase can already identify a constant fraction of the weight of the solution, thus obtaining a constant factor approximation. However, it is not clear how to obtain the best approximation factor from such an approach. Our main insight is that, with a right measure of progress, we can grow the solution iteratively and obtain solutions that are arbitrarily close to those of sequential algorithms. We show that after only $O(\frac{1}{\epsilon})$ rounds, the pool of good elements already contains a good solution with constant probability.

1.3 Related Work

There has been a recent push toward obtaining fast, practical algorithms for submodular maximization problems arising in a variety of applied settings. Research in this direction has yielded a variety of techniques for speeding up the continuous greedy algorithm for monotone maximization [3, 22], as well as new approaches for non-monotone maximization based on insights from both the continuous greedy and double greedy algorithms [6, 8]. Of particular relevance to our results is the case of maximization under a matroid constraint. Here, for monotone functions the fastest current sequential algorithm gives a $1 - 1/e - \epsilon$ approximation using $O(\frac{\sqrt{kn}}{\epsilon^3} \ln^2(\frac{n}{\epsilon}) + \frac{k^2}{\epsilon})$ value queries. For non-monotone functions, Buchbinder *et al.* [8] give an $\frac{1+e^{-2}}{4} > 0.283$ -approximation in time $O(kn \log n + Mk)$, where M is the time required to compute a perfect matching on bipartite graph with k vertices per side. They also give a simple, combinatorial $1/4$ -approximation in time $O(kn \log n)$. In comparison, the sequential algorithm we present here is faster by a factor of $\Omega(k)$, at the cost of a slightly-weaker $\frac{1}{2+\epsilon} > 0.211$ -approximation.

Work on parallel and distributed algorithms for submodular maximization has been comparatively limited. Early results considered the special case of maximum k -coverage, and attained an $O(1 - 1/e - \epsilon)$ -approximation [11, 5]. Later, Kumar *et al.* [20] considered the more general problem of maximizing an arbitrary monotone submodular function subject to a matroid, knapsack, or p -system constraint. Their approach attains a $\frac{1}{2+\epsilon}$ approximation for matroids, and requires $O(\frac{1}{\epsilon} \log \Delta)$ MapReduce rounds, where Δ is the value of the best single element. More generally, they obtain a $\frac{1}{p+1+\epsilon}$ approximation for p -systems in $O(\frac{1}{\epsilon} \log \Delta)$ rounds. The factor of $\log \Delta$ in the number of rounds is inherent in their approach: they adapt the threshold greedy algorithm, which sequentially picks elements in $\log \Delta$ different thresholds. In another line of work, Mirzasoleiman *et al.* [23] introduced a simple, two-round distributed greedy algorithm for submodular maximization. While their algorithm is only an $O(\frac{1}{m})$ -approximation in the worst case, it performs very well in practice, and attains provable constant-factor guarantees for submodular functions exhibiting certain additional structure. Barbosa *et al.* [12] recently gave a more sophisticated analysis of this approach and showed that, if the initial distribution of elements is performed randomly, the algorithm indeed gives an expected, constant-factor guarantee for a variety of problems. Finally, Mirrokni and Zadimoghaddam [21] gave the currently-best 0.545-approximation for the cardinality constraint case using only 2 rounds of MapReduce.

2 The model

We adopt the most stringent MapReduce-style model among [19, 17, 4, 2], the Massively Parallel Communication (MPC) model from [4] as specified by [2]. Let N be the size of the input. In this model, there are M machines each with space S . The total memory of the system is $M \cdot S = O(N)$,

which is at most a constant factor more than the input size. Computation proceeds in synchronous rounds. In each round, each machine can perform local computation and at the end, it can send at most a total of $O(S)$ words to other machines. These $O(S)$ words could form a single message of size S , S messages of size 1, or any other combination whose sum is at most $O(S)$. Following [19], we restrict both $M, S < N^{1-\Omega(1)}$. The typical main complexity measure is the number of rounds.

Note that not all previous works on MapReduce-style algorithms for submodular maximization satisfy the strict requirements of the MPC model. For instance, as stated, the previous work by Kumar et al. [20] uses $\Theta(N \log N)$ total memory and thus it does not fit in this model (though it might be possible to modify their algorithms to satisfy this).

We assume that the size of the solution is at most N^{1-2c} for some constant $0 < c < 1/2$. Thus, an entire solution can be stored on a single machine in the model. This assumption is also used in previous work such as [21].

3 Preliminaries

A function $f : 2^V \rightarrow \mathbb{R}_+$ is submodular if and only if $f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B)$ for all $A \subseteq B$ and $e \notin B$. If $f(A \cup \{e\}) - f(A) \geq 0$ for all A and $e \notin A$ we say that f is *monotone*. Here we consider the general problem $\max\{f(S) : S \subseteq V, S \in \mathcal{I}\}$, where \mathcal{I} is any hereditary constraint (i.e., a downward-closed family of subsets of V).

Throughout the paper, $n := |V|$ is the size of the ground set, $k := \max_{S \in \mathcal{I}} |S|$ is the maximum size of a solution, and m is the number of machines employed by the distributed algorithm.

We shall consider both monotone and non-monotone submodular functions. However, the following simple observation shows that even non-monotone submodular functions are monotone when restricted to the optimal solution of a problem of the sort we consider.

Lemma 3.1. *Let f be a submodular function and $\text{OPT} = \arg \max_{S \in \mathcal{I}} f(S)$ for some hereditary constraint \mathcal{I} . Then, $f(A \cap \text{OPT}) \leq f(B \cap \text{OPT})$ for all $A \subseteq B$.*

Proof: Consider $X \subseteq \text{OPT}$ and $e \in \text{OPT} \setminus X$. By submodularity, $f(X \cup \{e\}) - f(X) \geq f(\text{OPT} \setminus \{e\}) - f(\text{OPT} \setminus \{e\})$. On the other hand, because \mathcal{I} is hereditary, $\text{OPT} \setminus \{e\}$ is feasible and thus $f(\text{OPT}) \geq f(\text{OPT} \setminus \{e\})$. Therefore $f(X \cup \{e\}) - f(X) \geq 0$ for all X and $e \in \text{OPT} \setminus X$. \square

Continuous extensions. In this paper, we work with two standard continuous extensions of submodular functions, the multilinear extension and the Lovász extension. The *multilinear extension* of f is the function $F : [0, 1]^V \rightarrow \mathbb{R}_+$ such that $F(\mathbf{x}) = \mathbf{E}[f(R(\mathbf{x}))]$, where $R(\mathbf{x})$ is a random subset of V in which each element e appears independently with probability x_e .

The *Lovász extension* of f is the function $f^- : [0, 1]^V \rightarrow \mathbb{R}_+$ such that $f^-(\mathbf{x}) = \mathbf{E}_{\theta \in \mathcal{U}(0,1)}[f(\{e : x_e \geq \theta\})]$, where $\mathcal{U}(0,1)$ is the uniform distribution on $[0,1]$. For any submodular function f , the Lovász extension f^- satisfies: $f^-(\mathbf{1}_S) = f(S)$ for all $S \subseteq V$; f^- is convex; and the restricted scale invariance property $f^-(c \cdot \mathbf{x}) \geq c \cdot f^-(\mathbf{x})$ for any $c \in [0, 1]$. We shall make use of the following lemmas.

Lemma 3.2 ([12], Lemma 1). *Let S be a random set with $\mathbf{E}[\mathbf{1}_S] = c \cdot \mathbf{p}$ (for $c \in [0, 1]$). Then, $\mathbf{E}[f(S)] \geq c \cdot f^-(\mathbf{p})$.*

Lemma 3.3. *Let $f : 2^V \rightarrow \mathbb{R}_+$ be a submodular function that is monotone when restricted to $X \subseteq V$. Further, let $T, S \subseteq X$, and let R be a random subset of T in which every element occurs with probability at least p . Then, $\mathbf{E}[f(R \cup S)] \geq p \cdot f(T \cup S) + (1 - p)f(S)$.*

Proof: Recall that f^- is the Lovász extension of f . Since f^- is convex,

$$\mathbf{E}[f(R \cup S)] = \mathbf{E}[f^-(\mathbf{1}_{R \cup S})] \geq f^-(\mathbf{E}[\mathbf{1}_{R \cup S}]) = f^-(\mathbf{E}[\mathbf{1}_{R \setminus S}] + \mathbf{1}_S).$$

Since every element of T occurs in R with probability at least p , we have $\mathbf{E}[\mathbf{1}_{R \setminus S}] \geq p \cdot \mathbf{1}_{T \setminus S}$. Then, since f is monotone with respect to $X \supseteq S \cup T$, we must have:

$$f^-(\mathbf{E}[\mathbf{1}_{R \setminus S}] + \mathbf{1}_S) \geq f^-(p \cdot \mathbf{1}_{T \setminus S} + \mathbf{1}_S).$$

Finally, from the definition of f^- , we have

$$f^-(p \cdot \mathbf{1}_{T \setminus S} + \mathbf{1}_S) = p \cdot f(T \cup S) + (1 - p)f(S).$$

□

4 Generic Parallel Algorithm for Submodular Maximization

In this section, we give a generic approach for parallelizing *any sequential algorithm* Alg for the problem $\max_{S \subseteq V: S \in \mathcal{I}} f(S)$, where $f : 2^V \rightarrow \mathbb{R}_+$ is a submodular function and $\mathcal{I} \subseteq 2^V$ is a hereditary constraint.

As a starting point, we need a common abstract description of existing sequential algorithms. Towards that end, we turn to the standard Greedy and Continuous Greedy algorithms for inspiration. The Greedy algorithm directly constructs a solution, whereas the Continuous Greedy algorithm first constructs a fractional solution \mathbf{x} which is then rounded to get an integral solution. In the common abstraction, we will need both the integral solution *and* the support of the fractional solution \mathbf{x} . To account for this, we will have the algorithm Alg return a pair of sets, $(\text{AlgSol}(V), \text{AlgRel}(V))$, where $\text{AlgSol}(V) \in \mathcal{I}$ is a feasible solution for the problem and $\text{AlgRel}(V)$ is a set providing additional information. When using the standard Greedy algorithm for Alg , $\text{AlgSol}(V)$ and $\text{AlgRel}(V)$ will both be equal to the Greedy solution. When using the Continuous Greedy algorithm for Alg , $\text{AlgSol}(V)$ will be the integral solution and $\text{AlgRel}(V)$ will be the support of the fractional solution constructed by the Continuous Greedy algorithm.

More importantly, we will need an abstraction that captures the greedy behavior of these algorithms. We encapsulate the crucial properties of greedy-like algorithms in the following definition. We believe that this framework is one of the most valuable and insightful contributions of this work, and it provides a general abstraction for a broader class of algorithms.

We assume that the algorithm Alg satisfies the following properties.

1. (α -Approximation) For every input $N \subseteq V$, $\text{AlgSol}(N)$ is an α -approximate solution to $\max_{S \subseteq N: S \in \mathcal{I}} f(S)$.
2. (Consistency) Let A and B be two disjoint subsets of V . Suppose that, for each element $e \in B$, we have $\text{AlgRel}(A \cup \{e\}) = \text{AlgRel}(A)$. Then $\text{AlgSol}(A \cup B) = \text{AlgSol}(A)$.

Armed with this definition, we can now describe our approach for parallelizing an abstract sequential algorithm Alg with almost no loss in the approximation guarantee.

Parallel algorithm ParallelAlg based on Alg. As before, let α be the approximation guarantee of the sequential algorithm Alg . Let $s := \max_{N \subseteq V} |\text{AlgSol}(N) \cup \text{AlgRel}(N)|$ be the maximum size of the sets returned by Alg . Let $\epsilon > 0$ be the desired accuracy, i.e., we will aim that ParallelAlg achieves an $(\alpha - \epsilon)$ approximation.

The algorithm uses $g := \Theta(1/(\alpha\epsilon))$ groups of machines with m machines in each group (and thus the total number of machines is gm). The number m of machines can be chosen arbitrarily and it will determine the amount of space needed on each machine, since the dataset is divided roughly equally among each of the m machines in each group. An optimal setting is $gm := O(\sqrt{n/s})$.

The algorithm performs $\Theta(1/\epsilon)$ runs. Throughout the process, we maintain two quantities: an *incumbent* solution S_{best} , which is the best solution produced on any single machine so far in the process, and a *pool* of elements $C \subseteq V$ (we assume that the incumbent solution is stored on one designated machine).

Each run of the algorithm proceeds as follows. Amongst each group of m machines, we partition V uniformly at random; each element e chooses an index $i \in [m]$ uniformly and independently at random and is assigned to the i th machine in the group. We do this separately for each group of machines, i.e., each element appears on exactly one machine in each group. For an individual machine $i \in [gm]$, let $X_{i,r}$ denote the set of elements that are assigned to i in run r by this procedure. Additionally, we place on each machine the same pool of elements C_{r-1} , constructed at the end of run $r - 1$.

Once the elements have been distributed as described above, on each machine i , we run the algorithm `Alg` on the input $X_{i,r} \cup C_{r-1}$ on the machine to obtain $(\text{AlgSol}(X_{i,r} \cup C_{r-1}), \text{AlgRel}(X_{i,r} \cup C_{r-1}))$. We update the incumbent solution S_{best} to be the better of the current solution S_{best} and the solutions $\text{AlgSol}(X_{i,r} \cup C_{r-1})$ constructed on each of the machines; this is achieved by having each machine send $\text{AlgSol}(X_{i,r} \cup C_{r-1})$ to some designated machine maintaining S_{best} , and this machine will update S_{best} in the next round. We update the pool by setting $C_r := C_{r-1} \cup_i \text{AlgRel}(X_{i,r} \cup C_{r-1})$; this is achieved by having each machine send $\text{AlgRel}(X_{i,r} \cup C_{r-1})$ to every other machine, and thus ensuring that the pool C_r is available on each machine during the next round.

At the end of the $\Theta(1/\epsilon)$ runs, the algorithm returns the incumbent solution S_{best} . This completes the description of our algorithm.

Avoiding duplicating the dataset. The algorithm above partitions the dataset over $\Theta(1/\epsilon)$ groups of machines and thus it duplicates the dataset $\Theta(1/\epsilon)$ times (this problem also applies to previous work [21]). This is done in order to achieve the best theoretical guarantee on the number of runs, but in practice it is undesirable to duplicate the data. Instead, we can use a single group of m machines and perform the computation of a single run sequentially over $\Theta(1/\epsilon)$ sub-run, where each sub-run performs the computation of one of the group of machines. This will lead to an algorithm that performs $\Theta(1/\epsilon^2)$ runs using m machines and it does not duplicate the dataset.

The analysis. We devote the rest of this section to the analysis of the algorithm `ParallelAlg`. We start by noting that, if we choose g and m so that $gm = O(\sqrt{n/s})$, the algorithm uses the following resources and thus it satisfies the requirements of the model in Section 2.

Lemma 4.1. *ParallelAlg can be implemented in the parallel model in Section 2 using the following resources.*

- The number of rounds is $O(1/\epsilon)$.
- The number of machines is $O(\sqrt{n/s})$.
- The amount of space used on each machine is $O(\sqrt{ns}/(\epsilon\alpha))$ with high probability.
- In each round, the total amount of communication from a machine to all other machines is $O(\sqrt{ns}/(\epsilon\alpha))$ with high probability. The total amount of communication over all machines in a given round is $O(n/(\epsilon\alpha))$.

Proof: We will choose $gm := \sqrt{n/s}$ as our number of machines. Using this choice, we can provide the guarantees stated in the lemma.

Note that we can combine the update step of the incumbent solution and the pool of a given run with the next run's distribution of elements into a single round of communication. Specifically, each machine computes a new random assignment for each element of its sample $X_{i,r}$, assigns all of its new pool elements to *all* machines, and sends its solution to the designated machine. Thus each run corresponds to a round of communication. In each round, a machine communicates its sample $X_{i,r}$, which has size $O(n/m) = O(\sqrt{ns}/(\epsilon\alpha))$ with high probability, and the sets $\text{AlgSol}(X_{i,r} \cup C_{r-1})$ and $\text{AlgRel}(X_{i,r} \cup C_{r-1})$ that have size $O(s)$ to all other machines. Thus the total amount that a machine communicates is $O(\sqrt{ns}/(\epsilon\alpha) + s \cdot gm) = O(\sqrt{ns}/(\epsilon\alpha))$ with high probability, and the total amount that all machines communicate is $O(n + n/m \cdot gm) = O(n/(\epsilon\alpha))$.

In every round, the space used on a given machine is the size of its sample $X_{i,r}$, which is $O(n/m) = O(\sqrt{ns}/(\epsilon\alpha))$ with high probability; the size of the incumbent solution, which is $O(s)$; and the size of the pool, which is $O(gm \cdot s/\epsilon) = O(\sqrt{ns}/\epsilon)$. Therefore the total amount of space used on each machine is $O(\sqrt{ns}/(\epsilon\alpha))$ with high probability. \square

Thus it remains to analyze the quality of the solution constructed by the algorithm. In the remainder of this section, we show that, if Alg satisfies the α -approximation and consistency properties defined above, the parallel algorithm ParallelAlg achieves an $(\alpha - O(\epsilon))$ approximation. For simplicity, in this section we assume that Alg is deterministic; in Section B, we extend our approach to the setting in which Alg is randomized. We start by introducing some notation. Let $\mathcal{V}(1/m)$ denote the distribution over random subsets of V where each element is included independently with probability $1/m$. Let OPT be an optimal solution. Recall that $X_{i,r} \sim \mathcal{V}(1/m)$ is the random sample placed on machine i at the beginning of run r and C_{r-1} is the pool of elements at the beginning of run r . The following theorem is the crux of our analysis.

Theorem 4.2. *Consider a run $r \geq 1$ of the algorithm. Let $\widehat{C}_{r-1} \subseteq V$. Then one of the following must hold:*

- (1) $\mathbf{E}_{X_{1,r}}[f(\text{AlgSol}(C_{r-1} \cup X_{1,r})) \mid C_{r-1} = \widehat{C}_{r-1}] \geq (1 - \epsilon)^2 \alpha \cdot f(\text{OPT})$, or
- (2) $\mathbf{E}[f(C_r \cap \text{OPT}) \mid C_{r-1} = \widehat{C}_{r-1}] - f(\widehat{C}_{r-1} \cap \text{OPT}) \geq \frac{\epsilon}{2} \cdot f(\text{OPT})$.

Intuitively, Theorem 4.2 shows that, in expectation, if we have not found a good solution on some machine after $O(1/\epsilon)$ runs, then the current pool C , available to every machine, must satisfy $f(C \cap \text{OPT}) = f(\text{OPT})$, and so each machine in the next run will in fact return a solution of quality at least $\alpha f(\text{OPT})$. The following theorem, whose proof we give in Section A, makes this formal.

Theorem 4.3. *ParallelAlg achieves an $(1 - \epsilon)^3 \alpha$ approximation with constant probability.*

We devote the rest of this section to the proof of Theorem 4.2. Consider a run r of the algorithm. Let $\widehat{C}_{r-1} \subseteq V$. In the following, we condition on the event that $C_{r-1} = \widehat{C}_{r-1}$.

For each element $e \in V$, let $p_r(e) = \Pr_{X \sim \mathcal{V}(1/m)}[e \in \text{AlgRel}(\widehat{C}_{r-1} \cup X \cup \{e\})]$ if $e \in \text{OPT} \setminus \widehat{C}_{r-1}$, and 0 otherwise. As shown in the following lemma, the probability $p_r(e)$ gives us a handle on the probability that e is in the union of the relevant sets.

Lemma 4.4. *For each element $e \in \text{OPT} \setminus \widehat{C}_{r-1}$,*

$$\Pr[e \in \cup_{1 \leq i \leq gm} \text{AlgRel}(\widehat{C}_{r-1} \cup X_{i,r})] = 1 - (1 - p_r(e))^g,$$

where g is the number of groups into which the machines are partitioned.

Proof: For each group G_j , we can show that e is *not* in the union of the relevant sets for that group with probability $1 - p_r(e)$. Since different groups have independent partitions, e is not in the union of the relevant sets for all machines with probability $(1 - p_r(e))^g$, and the lemma follows. More precisely, for each group G_j , let \bar{Y}_j be the event that $e \notin \bigcup_{i \in G_j} \text{AlgRel}(\widehat{C}_{r-1} \cup X_{i,r})$. Let $G_{j,\ell}$ denote the ℓ th machine in G_j . We have

$$\begin{aligned} \Pr[\bar{Y}_j] &= \frac{1}{m} \sum_{\ell=1}^m \Pr[\bar{Y}_j \mid e \text{ is on } G_{j,\ell}] = \frac{1}{m} \sum_{\ell=1}^m \Pr[e \notin \text{AlgRel}(\widehat{C}_{r-1} \cup X_{\ell,r}) \mid e \in X_{\ell,r}] \\ &= \frac{1}{m} \sum_{\ell=1}^m \Pr_{X \sim \mathcal{V}(1/m)}[e \notin \text{AlgRel}(\widehat{C}_{r-1} \cup X \cup \{e\})] = 1 - p_r(e), \end{aligned}$$

where the first equality follows from the fact that e assigned to a machine in G_j chosen independently and uniformly at random, and the third from the fact that the distribution of $X_{\ell,r} \sim \mathcal{V}(1/m)$ conditioned on $e \in X_{\ell,r}$ is identical to the distribution of $X \cup \{e\}$ with $X \sim \mathcal{V}(1/m)$. Since the events $\{Y_j : 1 \leq j \leq g\}$ are mutually independent, $\Pr[\bigwedge_{1 \leq j \leq g} \bar{Y}_j] = \prod_{j=1}^g \Pr[\bar{Y}_j] = (1 - p_r(e))^g$. \square

Returning to the proof of Theorem 4.2, we define a partition (P_r, Q_r) of $\text{OPT} \setminus \widehat{C}_{r-1}$ as follows:

$$P_r = \{e \in \text{OPT} \setminus \widehat{C}_{r-1} : p_r(e) < \epsilon\} \quad Q_r = \{e \in \text{OPT} \setminus \widehat{C}_{r-1} : p_r(e) \geq \epsilon\}$$

The following subsets of P_r and Q_r are key to our analysis (recall that $X_{i,r}$ is the random sample placed on machine i at the beginning of the run r):

$$P'_r = \{e \in P_r : e \notin \text{AlgRel}(\widehat{C}_{r-1} \cup X_{1,r} \cup \{e\})\} \quad Q'_r = Q_r \cap \left(\bigcup_{i=1}^{gm} \text{AlgRel}(\widehat{C}_{r-1} \cup X_{i,r}) \right).$$

Note that each element $e \in P_r$ is in P'_r with probability $1 - p_r(e) \geq 1 - \epsilon$. Further, by Lemma 4.4, each element $e \in Q_r$ is in Q'_r with probability $1 - (1 - p_r(e))^g \geq 1 - \frac{1}{e} \geq \frac{1}{2}$.

It follows from the definition of P'_r and the consistency property of Alg that

$$\text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r}) = \text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r} \cup P'_r).$$

Let $\text{OPT}_{r-1} = \widehat{C}_{r-1} \cap \text{OPT}$ be the part of OPT in this iteration's pool. Then, since Alg is an α approximation and $P'_r \cup \text{OPT}_{r-1} \subseteq \text{OPT}$ is a feasible solution, we have

$$f(\text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r})) \geq \alpha \cdot f(P'_r \cup \text{OPT}_{r-1}).$$

Taking expectations on both sides, we have:

$$\mathbf{E}[f(\text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r}))] \geq \alpha \cdot \mathbf{E}[f(P'_r \cup \text{OPT}_{r-1})] \geq (1 - \epsilon)\alpha \cdot f(P_r \cup \text{OPT}_{r-1}), \quad (1)$$

where the final inequality follows from Lemma 3.3, since f is monotone when restricted to $\text{OPT} \supseteq P_r \cup \text{OPT}_{r-1}$, and P'_r contains every element of P_r with probability at least $(1 - \epsilon)$.

Note that $Q'_r \subseteq (\text{OPT} \cap C_r) \setminus \text{OPT}_{r-1}$. As before, f is monotone when restricted to OPT . Additionally, Q'_r contains every element of Q_r with probability at least $1/2$. Thus,

$$\mathbf{E}[f(C_r \cap \text{OPT}) \mid C_{r-1} = \widehat{C}_{r-1}] \geq \mathbf{E}[f(Q'_r \cup \text{OPT}_{r-1})] \geq \frac{1}{2} \cdot f(Q_r \cup \text{OPT}_{r-1}) + \frac{1}{2} \cdot f(\text{OPT}_{r-1}),$$

where the final inequality follows from Lemma 3.3. Rearranging this inequality using the condition $C_{r-1} = \widehat{C}_{r-1}$ and the definition $\text{OPT}_{r-1} = \widehat{C}_{r-1} \cap \text{OPT}$ we obtain:

$$\begin{aligned} \mathbf{E}[f(C_r \cap \text{OPT}) - f(C_{r-1} \cap \text{OPT}) \mid C_{r-1} = \widehat{C}_{r-1}] &\geq \frac{1}{2} (f(Q_r \cup \text{OPT}_{r-1}) - f(\text{OPT}_{r-1})) \\ &\geq \frac{1}{2} (f(P_r \cup Q_r \cup \text{OPT}_{r-1}) - f(P_r \cup \text{OPT}_{r-1})) = \frac{1}{2} (f(\text{OPT}) - f(P_r \cup \text{OPT}_{r-1})), \end{aligned} \quad (2)$$

where the second inequality follows from submodularity.

Now, if $f(P_r \cup (\widehat{C}_{r-1} \cap \text{OPT})) \geq (1-\epsilon) \cdot f(\text{OPT})$ then this fact together with (1) imply that the first property in the statement of Theorem 4.2 must hold. Otherwise, $f(\text{OPT}) - f(P_r \cup (\widehat{C}_{r-1} \cap \text{OPT})) \geq \epsilon \cdot f(\text{OPT})$; this fact together with (2) implies that the second property must hold.

This completes the description of our generic approach. In the following sections, we instantiate the algorithm Alg with the standard Greedy algorithm and a heavily discretized Continuous Greedy algorithm, and obtain our main results stated in the introduction.

5 A Parallel Greedy Algorithm

In this section, we combine the generic approach from Section 4 with the standard greedy algorithm, and give our results for monotone maximization stated in Theorem 5.2.

We let Alg be the standard Greedy algorithm. We let $\text{AlgRel}(N) = \text{AlgSol}(N) = \text{Greedy}(N)$. It was shown in previous work that the Greedy algorithm satisfies the consistency property.

Lemma 5.1 ([12], Lemma 2). *Let $A \subseteq V$ and $B \subseteq V$ be two disjoint subsets of V . Suppose that, for each element $e \in B$, we have $\text{Greedy}(A \cup \{e\}) = \text{Greedy}(A)$. Then $\text{Greedy}(A \cup B) = \text{Greedy}(A)$.*

Informally, this simply means that if Greedy rejects some element e when presented with input $A \cup \{e\}$, then adding other similarly rejected elements to $A \cup \{e\}$ cannot cause e to be accepted. This allows us to immediately apply the result from Section 4 and obtain the following result.

Theorem 5.2. *Let $f : 2^V \rightarrow \mathbb{R}_+$ be a submodular function, and $\mathcal{I} \subseteq 2^V$ be a hereditary set system. For any $\epsilon > 0$ there is a randomized distributed $O(1/\epsilon)$ -round algorithm that can be implemented in the model described in Section 2. The algorithm is an $(\alpha - O(\epsilon))$ -approximation with constant probability for the problem $\max_{S \in \mathcal{I}} f(S)$, where α is the approximation ratio of the standard, sequential greedy algorithm for the same problem.*

6 A Parallel Continuous Greedy Algorithm

For monotone maximization subject to a matroid constraint, Theorem 5.2 guarantees only a $(1/2 - \epsilon)$ approximation, due to the limitations of the standard greedy algorithm. We obtain a nearly optimal $(1 - 1/e - \epsilon)$ approximation by instantiating the framework in Section 4 with the DCGreedy algorithm shown in Algorithm 1.

The DCGreedy algorithm is a heavily discretized version of the measured continuous greedy approach of [15], and it first constructs an approximate fractional solution to the problem $\max_{\mathbf{x} \in P(\mathcal{I})} F(\mathbf{x})$ of maximizing the multilinear extension F of f subject to the constraint that \mathbf{x} is in the matroid polytope $P(\mathcal{I})$, and then rounds the fractional solution without loss using pipage rounding or swap rounding [1, 10].

In this section, we combine the generic approach from Section 4 with the DCGreedy algorithm. We use DCGreedy as Alg; the relevant set $\text{AlgRel}(N)$ is the set of elements in the support of the fractional solution $\mathbf{x}(1/\epsilon)$, and $\text{AlgSol}(N)$ is the integral solution obtained by rounding $\mathbf{x}(1/\epsilon)$.

Algorithm 1: Discretized Continuous Greedy (DCGreedy).

Input: $N \subseteq V$

- 1 $\mathbf{x}(0) \leftarrow \mathbf{0}$
- 2 **for** $t \leftarrow 1$ to $1/\epsilon$ **do**
- 3 $\mathbf{y}(t) \leftarrow \text{GreedyStep}(N, \mathbf{x}(t))$
- 4 $\mathbf{x}(t) \leftarrow \mathbf{x}(t-1) + \mathbf{y}(t)$
- 5 $S \leftarrow \text{SwapRounding}(\mathbf{x}(1/\epsilon), \mathcal{I})$
- 6 Let T be the support of $\mathbf{x}(1/\epsilon)$
- 7 **return** (S, T)

Algorithm 2: Greedy Update Step (GreedyStep).

Input: $N \subseteq V, \mathbf{x} \in [0, 1]^N$

- 1 $W \leftarrow \emptyset, \mathbf{y} \leftarrow \mathbf{0}$
- 2 **repeat**
- 3 $D \leftarrow \{e \in N \setminus W : W \cup \{e\} \in \mathcal{I}\}$
- 4 **foreach** $e \in D$ **do**
- 5 $w_e \leftarrow \mathbf{E}[f(R(\mathbf{x} + \mathbf{y}) \cup \{e\}) - f(R(\mathbf{x} + \mathbf{y}))]$
- 6 Let $e^* = \arg \max_{e \in D} w_e$
- 7 **if** $D = \emptyset$ or $w_{e^*} < 0$ **then return** \mathbf{y}
- 8 **else**
- 9 $y_{e^*} \leftarrow y_{e^*} + \epsilon(1 - x_{e^*})$
- 10 $W \leftarrow W \cup \{e^*\}$

Figure 1: The discretized continuous greedy algorithm. On line 5 of Algorithm 2, for a vector $\mathbf{z} \in [0, 1]^N$, we use $R(\mathbf{z})$ to denote a random subset of N that contains each element e independently with probability z_e . The weights on line 5 cannot be computed exactly in polynomial time, but they can be efficiently approximated using random samples.

Note that it is necessary to ensure that the fractional solution has small support so that the size of $\text{AlgRel}(N)$ is small. We achieve this by heavily discretizing the continuous greedy algorithm, thereby limiting the number of support updates performed in lines 3 and 4 of DCGreedy. Unfortunately, performing this discretization naively introduces an error in the approximation that is too large. Thus, we make use of a key idea from [3], which can be applied in the case of a matroid constraint. This allows us to show the following lemma whose proof is deferred to Section C.

Lemma 6.1. *The DCGreedy algorithm achieves an $(1 - 1/e - O(\epsilon))$ approximation for monotone functions and an $(1/e - O(\epsilon))$ approximation for non-monotone functions.*

The lemma above provides us with the desired approximation guarantees for DCGreedy, and thus it remains to show the consistency property. Before doing so, we must address how the weights are computed on line 5 of the GreedyStep algorithm (see Algorithm 2). Computing the weights exactly requires exponential time, but they can be approximated in polynomial time using random samples. In order to illustrate the main ideas behind the proof of consistency, we assume that the weights are computed exactly, since this will keep the algorithm deterministic. In the Appendix, we remove this assumption and we analyze the resulting randomized algorithm using an extension of our framework.

Lemma 6.2. *Let A and B be two disjoint subsets of V . Suppose that, for each element $e \in B$, we have $\text{DCGreedyRel}(A \cup \{e\}) = \text{DCGreedyRel}(A)$. Then $\text{DCGreedySol}(A \cup B) = \text{DCGreedySol}(A)$.*

Proof: We will show that the GreedyStep algorithm picks the same set W on input (A, \mathbf{x}) and $(A \cup B, \mathbf{x})$, which implies the lemma. Suppose for contradiction that the algorithm makes different choices on input (A, \mathbf{x}) and $(A \cup B, \mathbf{x})$. Consider the first iteration where the two runs differ, and let e be the element added to W in that iteration on input $(A \cup B, \mathbf{x})$. Note that $e \notin A$ and thus we

have $e \in B$. But then e will be added to W on input $(A \cup \{e\}, \mathbf{x})$. Thus $e \in \text{DCGreedyRel}(A \cup \{e\})$, which contradicts the fact that $e \in B$. \square

Thus we can apply the result from Section 4 and obtain the following result.

Theorem 6.3. *Let $f : 2^V \rightarrow \mathbb{R}_+$ be a submodular function, and $\mathcal{I} \subseteq 2^V$ be a matroid. For any $\epsilon > 0$ there is a randomized distributed $O(1/\epsilon)$ -round algorithm that can be implemented in the model described in Section 2. The algorithm is an $(\alpha - O(\epsilon))$ -approximation with constant probability for the problem $\max_{S \in \mathcal{I}} f(S)$, where α is $(1 - 1/e)$ for monotone f and $1/e$ for general f .*

7 Faster Algorithms

In this section, we build on the techniques from the previous sections to give a distributed algorithm for non-monotone maximization that requires only two rounds, rather than $O(1/\epsilon)$ rounds, and achieves an improved approximation guarantee over the two-round algorithm proposed in [12]. In the case of non-monotone maximization over a matroid, we show that our techniques can be used to obtain a new, fast *sequential* algorithm as well.

7.1 Two-Round Algorithms For Non-Monotone Maximization

We first give an improved two-round algorithm for non-monotone maximization subject to a any hereditary constraint. The algorithm is similar to that of [12] for *monotone* maximization; perhaps surprisingly, we show that this approach achieves a good approximation even for non-monotone functions. We randomly partition the elements onto the m machines, and run Greedy on the elements V_i on machine i to pick a set S_i . We place the sets S_i on a single machine and we run any algorithm Alg on $B := \bigcup_i S_i$ to find a solution T . We return the best solution amongst S_1, \dots, S_m, T .

We analyze the algorithm for any hereditary constraint \mathcal{I} for which the Greedy algorithm satisfies the following property (for some γ), which we refer to as the strong greedy property:

$$\forall S \in \mathcal{I} : f(\text{Greedy}(V)) \geq \gamma \cdot f(\text{Greedy}(V) \cup S) \quad (\text{GP})$$

By the standard Greedy analysis, we have $\gamma = 1/2$ for a matroid constraint and $\gamma = 1/(p+1)$ for a p -system constraint.

Theorem 7.1. *Suppose that Greedy satisfies the strong greedy property with constant γ and let Alg be any β -approximation for the problem $\max_{S \in \mathcal{I}} f(S)$. Then there is a randomized, two-round distributed algorithm that achieves a $(1 - \frac{1}{m}) \frac{\beta\gamma}{\beta+\gamma}$ approximation in expectation for $\max_{S \in \mathcal{I}} f(S)$.*

Proof: For each element e , we let probability $p_e = \Pr_{X \sim \mathcal{V}(1/m)}[e \in \text{Greedy}(X \cup \{e\})]$, if $e \in \text{OPT}$, and 0 otherwise. Then, let $\mathbf{p} \in [0, 1]^V$ denote the vector whose entries are given by the probabilities p_e .

We first analyze the expected value of the Greedy solution S_1 . Let

$$O = \{e \in \text{OPT} : e \notin \text{Greedy}(V_1 \cup \{e\})\}.$$

By Lemma 5.1, $\text{Greedy}(V_1 \cup O) = \text{Greedy}(V_1) = S_1$, and by (GP), $f(S_1) \geq \gamma \cdot f(S_1 \cup O)$. Therefore

$$\begin{aligned} \mathbf{E}[f(S_1)] &\geq \gamma \cdot \mathbf{E}[f(S_1 \cup O)] \\ &= \gamma \cdot \mathbf{E}[f^-(\mathbf{1}_{S_1 \cup O})] \\ &\geq \gamma \cdot f^-(\mathbf{E}[\mathbf{1}_{S_1 \cup O}]) \\ &= \gamma \cdot f^-(\mathbf{E}[\mathbf{1}_{S_1}] + (\mathbf{1}_{\text{OPT}} - \mathbf{p})). \end{aligned} \quad (3)$$

On line three, we have used the fact that f^- is convex and on line four we have used the fact that $\mathbf{E}[\mathbf{1}_{S_1 \cup O}] = \mathbf{E}[\mathbf{1}_{S_1}] + (\mathbf{1}_{\text{OPT}} - \mathbf{p})$.

Now consider the solution T . Since Alg is a β -approximation, we have

$$\begin{aligned} \mathbf{E}[f(T)] &\geq \beta \cdot \mathbf{E}[f(B \cap \text{OPT})] \\ &= \beta \cdot \mathbf{E}[f^-(\mathbf{1}_{B \cap \text{OPT}})] \\ &\geq \beta \cdot f^-(\mathbf{E}[\mathbf{1}_{B \cap \text{OPT}}]) \\ &= \beta \cdot f^-(\mathbf{p}). \end{aligned} \tag{4}$$

Similarly to above, we have used the convexity of f^- and the fact that $\mathbf{E}[\mathbf{1}_{B \cap \text{OPT}}] = \mathbf{p}$.

By combining (3) and (4), and using convexity of f^- , we obtain

$$\frac{1}{\gamma} \mathbf{E}[f(S_1)] + \frac{1}{\beta} \mathbf{E}[f(T)] \geq f^-(\mathbf{E}[\mathbf{1}_{S_1}] + (\mathbf{1}_{\text{OPT}} - \mathbf{p})) + f^-(\mathbf{p}) \geq 2 \cdot f^-\left(\frac{\mathbf{E}[\mathbf{1}_{S_1}] + \mathbf{1}_{\text{OPT}}}{2}\right).$$

Since $S_1 \subseteq V_1$ and V_1 is a $1/m$ sample of V , we have $\mathbf{E}[\mathbf{1}_{S_1}] \leq \frac{1}{m} \cdot \mathbf{1}_V$. Therefore, using the definition of f^- and the non-negativity of f , we obtain

$$2 \cdot f^-\left(\frac{\mathbf{E}[\mathbf{1}_{S_1}] + \mathbf{1}_{\text{OPT}}}{2}\right) \geq \left(1 - \frac{1}{m}\right) f(\text{OPT}).$$

Thus

$$\max\{\mathbf{E}[f(S_1)], \mathbf{E}[f(T)]\} \geq \left(1 - \frac{1}{m}\right) \frac{\beta\gamma}{\beta + \gamma} \cdot f(\text{OPT}).$$

□

Examples of results. We conclude this section with some examples of approximation guarantees that we can obtain using Theorem 7.1. For a matroid constraint, we have $\gamma = 1/2$ and, if we use the measured Continuous Greedy algorithm for Alg, we have $\beta = 1/e$; thus we obtain a $(1 - \frac{1}{m}) \frac{1}{2+e}$ approximation. We remark that, for a cardinality constraint, one can strengthen the proof of Theorem 7.1 slightly and obtain a $(1 - \frac{1}{m}) \frac{1}{e} (1 - \frac{1}{e})$ approximation; we give the details in Section D.

For a p -system constraint, we have $\gamma = 1/(p+1)$. We can use the algorithm of Gupta *et al.* [18] for Alg that achieves an approximation $\beta = 3 / \left(2p + 4 + \frac{2}{p}\right)$ when combined with the algorithm of [7] for unconstrained non-monotone maximization. Thus we obtain a $3 \left(1 - \frac{1}{m}\right) / \left(5p + 7 + \frac{2}{p}\right)$ approximation.

7.2 A Fast Sequential Algorithm for Matroid Constraints

We now show how our approach can be used to obtain a fast *sequential* algorithm for non-monotone maximization subject to a matroid constraint. The analysis given in Theorem 7.1 only relies on the following two properties of the Greedy algorithm: it satisfies (GP) and Lemma 5.1. Thus we can replace the Greedy algorithm by any algorithm satisfying these two properties. In particular, the Descending Thresholds Greedy (shown in Figure 2 as DThreshGreedy) of [20, 3] satisfies these conditions with $\gamma = 1/2 - \epsilon$.

Our algorithm proceeds as follows. We randomly partition the elements into $m := 1/\epsilon$ samples V_1, V_2, \dots, V_m . On each sample, we run the Descending Thresholds Greedy algorithm on V_i to obtain a solution S_i . Let $A := \operatorname{argmax}_{i \in [m]} f(S_i)$ and $B := \bigcup_i S_i$. Then, $|B| \leq k/\epsilon$, where k is the rank of the matroid. We run any β -approximation algorithm Alg on B to find a solution B' , and we return the better of A and B' . We obtain the following result.

Algorithm 3: Descending Thresholds Greedy algorithm (DThreshGreedy).

Input: $N \subseteq V$

- 1 $S \leftarrow \emptyset, d \leftarrow \max_{e \in N} f(\{e\})$
- 2 **for** $w = d; w \geq \frac{\epsilon}{n}d; w \leftarrow w(1 - \epsilon)$ **do**
- 3 **foreach** $e \in N$ **do**
- 4 **if** $S \cup \{e\} \in \mathcal{I}$
- 5 **and** $f(S \cup \{e\}) - f(S) \geq w$ **then**
- 6 $S \leftarrow S \cup \{e\}$
- 7 **return** S

Figure 2: The Descending Thresholds Greedy algorithm of [20, 3].

Theorem 7.2. *There is a sequential, randomized $(\frac{1}{2+\epsilon} - \epsilon)$ -approximation algorithm for the problem $\max_{S \in \mathcal{I}} f(S)$, where \mathcal{I} is any matroid constraint, running in time $O(\frac{n}{\epsilon} \log n) + \text{poly}(\frac{k}{\epsilon})$.*

Proof: The running time of the Descending Thresholds Greedy algorithm on a ground set of size s is $O(\frac{s}{\epsilon} \log(\frac{s}{\epsilon}))$. Each random sample has size $O(\epsilon n)$ with high probability, and thus the total time needed to construct B is $O(\frac{n}{\epsilon} \log n)$ with high probability. It follows from the analysis in Theorem 7.1 that the best of the two solutions A and a β -approximation to $\max_{S \subseteq B: S \in \mathcal{I}} f(S)$ is a $\frac{1}{2+\frac{1}{\beta}} - \epsilon$ approximation. We can then use any $1/e$ -approximation algorithm as Alg. \square

7.3 Monotone Maximization with a Cardinality Constraint

Here, we show how the previous techniques give a simple two-round algorithm that achieves a $1/2 - \epsilon$ approximation for monotone maximization subject to a cardinality constraint.

Theorem 7.3. *There is a randomized, two-round, distributed algorithm achieving a $\frac{1}{2} - \epsilon$ approximation in expectation for $\max_{S: |S| \leq k} f(S)$, where f is a monotone function.*

The algorithm. Let $\epsilon > 0$ be a parameter. The algorithm uses $\Theta(\log(1/\epsilon)/\epsilon)$ groups of machines with m machines in each group (and thus the total number of machines is $O(m \log(1/\epsilon)/\epsilon)$).

We randomly distribute the ground set V to the machines as follows. Amongst each group of m machines, we partition V uniformly at random; each element e chooses an index $i \in [m]$ uniformly and independently at random and is assigned to the i th machine in the group. We do this separately for each group of machines, i.e., each element appears on exactly one machine in each group.

We run Greedy on each of the machines to select a set of k elements. Let S be the union of all of the Greedy solutions. We place S on a single machine together with a random sample $X \sim \mathcal{V}(1/m)$. On this machine, we pick the final solution as follows. For each value $a \in \{0, 1, \dots, k\}$, we select a solution T_a as follows. Let $\text{Greedy}(f, X, a)$ denote the first a elements chosen from the random sample X using the greedy algorithm on objective function f .

Then, let $T_a^1 = \text{Greedy}(f, X, a)$ and define $g(A) = f(T_a^1 \cup A) - f(T_a^1)$ for each $A \subseteq V$. Note that g is a non-negative, monotone submodular function. Let $T_a^2 = \text{Greedy}(g, S, k - a)$; that is, we pick $k - a$ elements from S using the Greedy algorithm with the function g as input. We set $T_a = T_a^1 \cup T_a^2$. The final solution T is the better of the $k + 1$ solutions T_a , where $a \in \{0, 1, \dots, k\}$.

The analysis. In the following, we show that the algorithm above is a $1/2 - \epsilon$ approximation. For each element e , we define a probability $p_e = \Pr_{X \sim \mathcal{V}(1/m)}[e \in \text{Greedy}(X \cup \{e\})]$, if $e \in \text{OPT}$ and 0 otherwise. We define a partition (O_1, O_2) of OPT as follows:

$$O_1 = \{e \in \text{OPT} \mid p_e < \epsilon\}, \quad O_2 = \{e \in \text{OPT} \mid p_e \geq \epsilon\}.$$

Let $a = |O_1|$ and let

$$O'_1 = \{e \in O_1 \mid e \notin \text{Greedy}(f, X \cup \{e\}, a)\}.$$

By the consistency property of the greedy algorithm (Lemma 5.1),

$$T_a^1 = \text{Greedy}(f, X, a) = \text{Greedy}(f, X \cup O'_1, a).$$

Additionally, for a cardinality constraint, **Greedy** satisfies (GP) with $\gamma = 1/2$ (see Subsection 7.1). Therefore

$$f(T_a^1) \geq \frac{1}{2}f(T_a^1 \cup O'_1), \quad (5)$$

$$g(T_a^2) \geq \frac{1}{2}g(T_a^2 \cup (O_2 \cap S)). \quad (6)$$

The inequality (6) can be rewritten as

$$f(T_a^1 \cup T_a^2) - f(T_a^1) \geq \frac{1}{2}(f(T_a^1 \cup T_a^2 \cup (O_2 \cap S)) - f(T_a^1)). \quad (7)$$

Adding (5) and (7), we obtain

$$\begin{aligned} f(T_a) &\geq \frac{1}{2}(f(T_a^1 \cup O'_1) + f(T_a^1 \cup T_a^2 \cup (O_2 \cap S)) - f(T_a^1)) \\ &\geq \frac{1}{2}f(T_a \cup O'_1 \cup (O_2 \cap S)) \\ &\geq \frac{1}{2}f(O'_1 \cup (O_2 \cap S)), \end{aligned}$$

where the last two inequalities follow from submodularity and monotonicity. Note that each element $e \in O_1$ is in O'_1 with probability $1 - p_e \geq 1 - \epsilon$. Each element $e \in O_2$ is in the union of the **Greedy** solutions from a given group of machines with probability $p_e \geq \epsilon$; since there are $\Theta(\log(1/\epsilon)/\epsilon)$ groups of machines and the groups have independent partitions, e is in S with probability at least $1 - \epsilon$. Therefore

$$\mathbf{E}[\mathbf{1}_{O'_1 \cup (O_2 \cap S)}] \geq (1 - \epsilon)\mathbf{1}_{\text{OPT}}.$$

Thus

$$\mathbf{E}[f(T_a)] \geq \frac{1}{2}f^-(\mathbf{E}[\mathbf{1}_{O'_1 \cup (O_2 \cap S)}]) \geq (1 - \epsilon)\frac{1}{2}f(\text{OPT}).$$

In the last inequality, we have used that if $x \geq y$ component-wise and f is monotone, $f^-(x) \geq f^-(y)$.

Acknowledgments. This work was done in part while A.E. was with the Computer Science department at the University of Warwick and a visitor to the Toyota Technological Institute at Chicago, H.N. was with the Toyota Technological Institute at Chicago, and J.W. was with the Computer Science department at the University of Warwick. J.W. was supported by EPSRC grant EP/J021814/1 and ERC Starting Grant 335288-OptApprox.

References

- [1] Alexander Ageev and Maxim Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328, 2004.
- [2] Alexandr Andoni, Aleksandar Nikolov, Krzysztof Onak, and Grigory Yaroslavtsev. Parallel algorithms for geometric graph problems. In *ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS)*, pages 574–583, 2014.
- [3] Ashwinkumar Badanidiyuru and Jan Vondrák. Fast algorithms for maximizing submodular functions. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1497–1514, 2014.
- [4] Paul Beame, Paraschos Koutris, and Dan Suciu. Communication steps for parallel query processing. In *ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS)*, pages 273–284, 2013.
- [5] Guy E. Blelloch, Richard Peng, and Kanat Tangwongsan. Linear-work greedy parallel approximate set cover and variants. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 23–32, 2011.
- [6] Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. Submodular maximization with cardinality constraints. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1433–1452, 2014.
- [7] Niv Buchbinder, Moran Feldman, Joseph (Seffi) Naor, and Roy Schwartz. A tight linear time $(1/2)$ -approximation for unconstrained submodular maximization. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1384–1402, 2012.
- [8] Niv Buchbinder, Moran Feldman, and Roy Schwartz. Comparing apples and oranges: Query tradeoff in submodular maximization. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1149–1168, 2015.
- [9] Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a submodular set function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.
- [10] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 575–584, 2010.
- [11] Flavio Chierichetti, Ravi Kumar, and Andrew Tomkins. Max-cover in map-reduce. In *International World Wide Web Conference (WWW)*, pages 231–240, 2010.
- [12] Rafael da Ponte Barbosa, Alina Ene, Huy Le Nguyen, and Justin Ward. The power of randomization: Distributed submodular maximization on massive datasets. In *International Conference on Machine Learning (ICML)*, pages 1236–1244, 2015.
- [13] Shahar Dobzinski and Jan Vondrák. From query complexity to computational complexity. In *ACM Symposium on Theory of Computing (STOC)*, pages 1107–1116, 2012.

- [14] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998.
- [15] Moran Feldman, Joseph (Seffi) Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 570–579, 2011.
- [16] M L Fisher, G L Nemhauser, and L A Wolsey. An analysis of approximations for maximizing submodular set functions—ii. *Mathematical Programming Studies*, 8:73–87, 1978.
- [17] Michael T Goodrich, Nodari Sitchinava, and Qin Zhang. Sorting, searching, and simulation in the mapreduce framework. In *International Symposium on Algorithms and Computation*, pages 374–383, 2011.
- [18] Anupam Gupta, Aaron Roth, Grant Schoenebeck, and Kunal Talwar. Constrained non-monotone submodular maximization: Offline and secretary algorithms. In *Conference on Web and Internet Economics (WINE)*, pages 246–257, 2010.
- [19] Howard Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for mapreduce. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 938–948, 2010.
- [20] Ravi Kumar, Benjamin Moseley, Sergei Vassilvitskii, and Andrea Vattani. Fast greedy algorithms in mapreduce and streaming. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 1–10, 2013.
- [21] Vahab S. Mirrokni and Morteza Zadimoghaddam. Randomized composable core-sets for distributed submodular maximization. In *ACM Symposium on Theory of Computing (STOC)*, pages 153–162, 2015.
- [22] Baharan Mirzasoleiman, Ashwinkumar Badanidiyuru, Amin Karbasi, Jan Vondrák, and Andreas Krause. Lazier than lazy greedy. In *AAAI Conference on Artificial Intelligence*, pages 1812–1818, 2015.
- [23] Baharan Mirzasoleiman, Amin Karbasi, Rik Sarkar, and Andreas Krause. Distributed submodular maximization: Identifying representative elements in massive data. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2049–2057, 2013.
- [24] G L Nemhauser, L A Wolsey, and M L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14(1):265–294, 1978.
- [25] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.
- [26] Jan Vondrák. Symmetry and approximability of submodular maximization problems. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 651–670, 2009.

A Proof of Theorem 4.3

Theorem 4.3. *ParallelAlg achieves an $(1 - \epsilon)^3 \alpha$ approximation with constant probability.*

Proof: Let $R = c/\epsilon$ be the total number of runs, and $\mathcal{C} = (C_0, C_1, \dots, C_R)$. Let $I_r(C_{r-1}) \in \{0, 1\}$ be equal to 1 if and only if

$$\mathbf{E}_{X_{1,r}}[f(\text{AlgSol}(C_{r-1} \cup X_{1,r}))] \geq (1 - \epsilon)^2 \alpha \cdot f(\text{OPT}).$$

Let

$$\begin{aligned} \Phi_r(\mathcal{C}) &= I_r(C_{r-1}) + \frac{2(f(C_r \cap \text{OPT}) - f(C_{r-1} \cap \text{OPT}))}{\epsilon f(\text{OPT})}, \\ \Phi(\mathcal{C}) &= \sum_{r=1}^R \Phi_r(C_{r-1}) \leq \sum_{r=1}^R I_r(C_{r-1}) + \frac{2f(C_R \cap \text{OPT})}{\epsilon f(\text{OPT})} \leq \sum_{r=1}^R I_r(C_{r-1}) + \frac{2}{\epsilon}. \end{aligned}$$

Taking expectation over the random choices of \mathcal{C} , we have

$$\mathbf{E}_{\mathcal{C}}[\Phi(\mathcal{C})] \leq \sum_{r=1}^R \mathbf{E}[I_r(C_{r-1})] + \frac{2}{\epsilon}$$

On the other hand, by Theorem 4.2, $\mathbf{E}[\Phi_r(C_{r-1})] \geq 1$ and therefore $\mathbf{E}[\Phi(\mathcal{C})] \geq R$. Thus

$$\frac{2}{\epsilon} + \sum_{r=1}^R \mathbf{E}[I_r(C_{r-1})] \geq \Phi(\mathcal{C}) \geq R.$$

Since $R > 6/\epsilon$, we have

$$\sum_{r=1}^R \mathbf{E}[I_r(\hat{C}_{r-1})] \geq \frac{2R}{3}.$$

Therefore, with probability at least $2/3$, there exists a run r such that $I_r(C_{r-1}) = 1$. Fix the randomness up to the first such run, i.e., condition on a fixed $C_{r-1} = \hat{C}_{r-1}$ such that $I_r(\hat{C}_{r-1}) = 1$ and C_r, \dots, C_R remain random. Assume for contradiction that with probability at least $1 - \epsilon\alpha(1 - \epsilon)^2$ over the choices of $X_{1,r}$,

$$f(\text{AlgSol}(C_{r-1} \cup X_{1,r})) < (1 - \epsilon)^3 \alpha \cdot f(\text{OPT}).$$

Then we have

$$\begin{aligned} \mathbf{E}[f(\text{AlgSol}(C_{r-1} \cup X_{1,r}))] &< (\epsilon\alpha(1 - \epsilon)^2 + (1 - \epsilon\alpha(1 - \epsilon)^2)(1 - \epsilon)^3 \alpha) f(\text{OPT}) \\ &= (\epsilon + (1 - \epsilon\alpha(1 - \epsilon)^2)(1 - \epsilon)) (1 - \epsilon)^2 \alpha f(\text{OPT}) \\ &< (1 - \epsilon)^2 \alpha f(\text{OPT}), \end{aligned}$$

contradicting our assumption on C_{r-1} . Thus, with probability at least $\epsilon\alpha(1 - \epsilon)^2$, we have

$$f(\text{AlgSol}(C_{r-1} \cup X_{1,r})) \geq (1 - \epsilon)^3 \alpha \cdot f(\text{OPT}).$$

Notice that the above argument applies not only to machine 1 in run r but also the first machine in each of the g groups in the same run r and their random samples $X_{i,r}$ are independent. Thus, since $g \geq c/(\epsilon\alpha)$ for a sufficiently large constant c , with probability at least $5/6$, we have $\max_i f(\text{AlgSol}(C_{r-1} \cup X_{i,r})) \geq (1 - \epsilon)^3 \alpha \cdot f(\text{OPT})$. Overall, the algorithm succeeds with probability at least $2/3 \cdot 5/6 = 5/9$. \square

B A Framework for Parallelizing Randomized Algorithms

In this section, we extend the framework from Section 4 to the setting in which the sequential algorithm Alg is randomized.

We represent the randomness of Alg as a vector $\mathbf{b} \sim \mathcal{D}$ drawn from some distribution \mathcal{D} . It is convenient to have the randomness \mathbf{b} given as input to the algorithm. More precisely, we assume that Alg takes as input a subset $N \subseteq V$ and a random vector $\mathbf{b} \sim \mathcal{D}$ and returns a pair of sets, $\text{AlgSol}(N, \mathbf{b})$ and $\text{AlgRel}(N, \mathbf{b})$. We assume that the size of \mathbf{b} depends only on the size of V , and hence is independent of the size of N . Finally, we assume that Alg has the following properties.

1. ((α, ϵ, δ)-Approximation) Let $\text{OPT} = \operatorname{argmax}_{S \in \mathcal{I}, S \subseteq V} f(S)$ be an optimal solution over the entire ground set V . Let $A \subseteq V$ and $\mathbf{b} \sim \mathcal{D}$. Let $B \subseteq \text{OPT}$ be a subset such that, for each $e \in B$, $e \notin \text{AlgRel}(A \cup \{e\}, \mathbf{b})$. We have

$$\Pr_{\mathbf{b} \sim \mathcal{D}} \left[f(\text{AlgSol}(A \cup B, \mathbf{b})) \geq \alpha \cdot f((A \cup B) \cap \text{OPT}) - \epsilon f(\text{OPT}) \right] \geq 1 - \delta.$$

2. (Consistency) Let \mathbf{b} be any fixed vector. Let A and B be two disjoint subsets of V . Suppose that, for each element $e \in B$, we have $\text{AlgRel}(A \cup \{e\}, \mathbf{b}) = \text{AlgRel}(A, \mathbf{b})$. Then $\text{AlgSol}(A \cup B, \mathbf{b}) = \text{AlgSol}(A, \mathbf{b})$.

Note that our assumption that the length of \mathbf{b} is independent of the size of the input subset allows expressions such as $\text{Alg}(A \cup \{e\}, \mathbf{b})$ and $\text{Alg}(A, \mathbf{b})$ to both make sense despite the fact that $|A \cup \{e\}| \neq |A|$.

Our algorithm works exactly as that described in Section 4, with the exception that each machine i in round r now additionally samples a random vector $\mathbf{b}_{i,r} \sim \mathcal{D}$. Then, on each machine, we run Alg on the set $V_{i,r} := X_{i,r} \cup C$ of elements on the machine and obtain $\text{AlgSol}(V_{i,r}, \mathbf{b}_{i,r})$ and $\text{AlgRel}(V_{i,r}, \mathbf{b}_{i,r})$. As in Section 4, the union $\bigcup_i \text{AlgRel}(V_{i,r}, \mathbf{b}_{i,r})$ of relevant elements is added to C , and the solution S_{best} is replaced by the best solution among $\{\text{AlgSol}(V_{i,r}, \mathbf{b}_{i,r}) : 1 \leq i \leq M\}$ and S_{best} .

In the final round we place C on a single machine, sample a random vector $\mathbf{b} \sim \mathcal{D}$, and run Alg on C , and \mathbf{b} to obtain the solution $\text{AlgSol}(C, \mathbf{b})$. The final solution is the best among $\text{AlgSol}(C, \mathbf{b})$ and S_{best} .

Analysis. The number of rounds, number of machines, space per machine, and amount of communication are the same as in Section 4. Thus, we focus on the approximation guarantee of the parallel algorithm. Using Theorem B.1 instead of Theorem 4.2, we can then finish the analysis in almost the same way as the deterministic case. The only difference is that instead of arguing that the algorithm works well with most of the random choices for $X_{i,r}$ as before, the proof now argues that the algorithm works well with most of the random choices for both $X_{i,r}$ and $\mathbf{b}_{i,r}$. Nonetheless, the same proof except for this substitution works.

Theorem B.1. *Consider a run $r > 1$ of the algorithm. Let $\widehat{C}_{r-1} \subseteq V$. Then one of the following must hold:*

- (1) $\mathbb{E}_{X_{1,r}, \mathbf{b}_{1,r}} [f(\text{AlgSol}(C_{r-1} \cup X_{1,r}, \mathbf{b}_{1,r})) \mid C_{r-1} = \widehat{C}_{r-1}] \geq (\alpha - O(\epsilon)) \cdot f(\text{OPT})$, or
- (2) $\mathbb{E}[f(C_r \cap \text{OPT}) \mid C_{r-1} = \widehat{C}_{r-1}] - f(\widehat{C}_{r-1} \cap \text{OPT}) \geq \frac{\epsilon}{2} \cdot f(\text{OPT})$.

Proof: Consider a run r of the algorithm. Let $\widehat{C}_{r-1} \subseteq V$. In the following, we condition on the event that $C_{r-1} = \widehat{C}_{r-1}$.

For each element $e \in V$, let $p_r(e) = \Pr_{X \sim \mathcal{V}(1/m), \mathbf{b} \sim \mathcal{D}}[e \in \text{AlgRel}(\widehat{C}_{r-1} \cup X \cup \{e\}, \mathbf{b})]$, if $e \in \text{OPT} \setminus \widehat{C}_{r-1}$, and 0 otherwise. The proof of the following lemma is exactly the same as Lemma 4.4 and thus is omitted.

Lemma B.2. *For each element $e \in \text{OPT} \setminus \widehat{C}_{r-1}$,*

$$\Pr[e \in \cup_{1 \leq i \leq gm} \text{AlgRel}(\widehat{C}_{r-1} \cup X_{i,r}, \mathbf{b}_{i,r})] = 1 - (1 - p_r(e))^g,$$

where g is the number of groups into which the machines are partitioned.

We define a partition (P_r, Q_r) of $\text{OPT} \setminus \widehat{C}_{r-1}$ as follows:

$$P_r = \{e \in \text{OPT} \setminus \widehat{C}_{r-1} : p_r(e) < \epsilon\}, \quad Q_r = \{e \in \text{OPT} \setminus \widehat{C}_{r-1} : p_r(e) \geq \epsilon\}.$$

The following subsets of P_r and Q_r are key to our analysis (recall that $X_{i,r}$ is the random sample placed on machine i at the beginning of the run and $\mathbf{b}_{i,r}$ is the random vector sampled by machine i in round r):

$$P'_r = \{e \in P_r : e \notin \text{AlgRel}(\widehat{C}_{r-1} \cup X_{1,r} \cup \{e\}, \mathbf{b}_{1,r})\}, \quad Q'_r = Q_r \cap (\cup_{i=1}^{gm} \text{AlgRel}(\widehat{C}_{r-1} \cup X_{i,r}, \mathbf{b}_{i,r})).$$

Note that each element $e \in P_r$ is in P'_r with probability $1 - p_r(e) \geq 1 - \epsilon$. Further, by Lemma B.2, each element $e \in Q_r$ is in Q'_r with probability $1 - (1 - p_r(e))^g \geq 1 - \frac{1}{e} \geq \frac{1}{2}$.

It follows from the definition of P'_r and the consistency property of Alg that

$$\text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r}, \mathbf{b}_{1,r}) = \text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r} \cup P'_r, \mathbf{b}_{1,r}). \quad (8)$$

Let $\text{OPT}_{r-1} = \widehat{C}_{r-1} \cap \text{OPT}$ be the part of OPT in this iteration's pool. We apply the $(\alpha, \epsilon, \delta)$ -approximation property with $A = \widehat{C}_{r-1} \cup X_{1,r}$, $\mathbf{b} = \mathbf{b}_{1,r}$, and $B = P'_r$ to obtain

$$\Pr_{\mathbf{b}_{1,r}} \left[\text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r} \cup P'_r, \mathbf{b}_{1,r}) \geq \alpha \cdot f((\widehat{C}_{r-1} \cup X_{1,r} \cup P'_r) \cap \text{OPT}) - \epsilon f(\text{OPT}) \right] \geq 1 - \delta.$$

Since f is monotone when restricted to OPT , and $P'_r \cup \text{OPT}_{r-1} \subseteq (\widehat{C}_{r-1} \cup X_{1,r} \cup P'_r) \cap \text{OPT}$, this inequality implies that

$$\Pr_{\mathbf{b}_{1,r}} \left[\text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r} \cup P'_r, \mathbf{b}_{1,r}) \geq \alpha \cdot f(P'_r \cup \text{OPT}_{r-1}) - \epsilon f(\text{OPT}) \right] \geq 1 - \delta.$$

Therefore, equation (8) gives

$$\Pr_{\mathbf{b}_{1,r}} \left[\text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r}, \mathbf{b}_{1,r}) \geq \alpha \cdot f(P'_r \cup \text{OPT}_{r-1}) - \epsilon f(\text{OPT}) \right] \geq 1 - \delta.$$

Taking expectation on both sides gives

$$\begin{aligned} \mathbb{E}_{X_{1,r}, \mathbf{b}_{1,r}} [f(\text{AlgSol}(\widehat{C}_{r-1} \cup X_{1,r}, \mathbf{b}_{1,r}))] &\geq (1 - \delta)\alpha \cdot \mathbb{E}_{X_{1,r}} [f(P'_r \cup \text{OPT}_{r-1})] - \epsilon f(\text{OPT}) \\ &\geq \alpha \cdot \mathbb{E}_{X_{1,r}} [f(P'_r \cup \text{OPT}_{r-1})] - (\epsilon + \alpha\delta)f(\text{OPT}) \\ &\geq (1 - \epsilon)\alpha \cdot f(P_r \cup \text{OPT}_{r-1}) - (\epsilon + \alpha\delta)f(\text{OPT}) \\ &\geq (1 - \epsilon)\alpha \cdot f(P_r \cup \text{OPT}_{r-1}) - 2\epsilon f(\text{OPT}). \end{aligned} \quad (9)$$

Here, the second inequality follows from the fact that f is monotone restricted to $\text{OPT} \supseteq (P_r \cup \text{OPT}_{r-1})$, the third from Lemma 3.3 and the fact that every element of P_r appears in P'_r with probability at least $(1 - \epsilon)$, and the last from our assumption that $\alpha\delta \leq \epsilon$.

Next, note that $Q'_r \subseteq (C_r \cap \text{OPT}) \setminus \text{OPT}_{r-1}$. This together with monotonicity of f restricted to r imply:

$$\begin{aligned} \mathbf{E}[f(C_r \cap \text{OPT}) \mid C_{r-1} = \widehat{C}_{r-1}] &\geq \mathbf{E}[f(Q'_r \cup \text{OPT}_{r-1})] \\ &\geq \frac{1}{2} \cdot f(Q_r \cup \text{OPT}_{r-1}) + \frac{1}{2} \cdot f(\text{OPT}_{r-1}), \end{aligned}$$

where the last inequality follows from Lemma 3.3 and Lemma B.2. Rearranging this inequality using the condition $C_{r-1} = \widehat{C}_{r-1}$ and the definition $\text{OPT}_{r-1} = \widehat{C}_{r-1} \cap \text{OPT}$ we obtain:

$$\begin{aligned} \mathbf{E}[f(C_r \cap \text{OPT}) - f(C_{r-1} \cap \text{OPT}) \mid C_{r-1} = \widehat{C}_{r-1}] &\geq \frac{1}{2} (f(Q_r \cup \text{OPT}_{r-1}) - f(\text{OPT}_{r-1})) \\ &\geq \frac{1}{2} (f(P_r \cup Q_r \cup \text{OPT}_{r-1}) - f(P_r \cup \text{OPT}_{r-1})) \\ &= \frac{1}{2} \left(f(\text{OPT}) - f(P_r \cup (\widehat{C}_{r-1} \cap \text{OPT})) \right), \quad (10) \end{aligned}$$

where the second inequality follows from submodularity.

Now, if $f(P_r \cup (\widehat{C}_{r-1} \cap \text{OPT})) \geq (1 - \epsilon) \cdot f(\text{OPT})$ then this fact together with (9) imply the first property in the statement of Theorem B.1 must hold. Otherwise, $f(\text{OPT}) - f(P_r \cup (\widehat{C}_{r-1} \cap \text{OPT})) \geq \epsilon \cdot f(\text{OPT})$; this fact together with (10) imply that the second property must hold. \square

C Analysis of DCGreedy for the Application of the Randomized Framework

In this section, we show that we can instantiate the randomized framework from Section B with a modified DCGreedy algorithm, and obtain the results stated in Section 6. Specifically, we extend the DCGreedy algorithm to the setting in which the weights w_e on line 5 of GreedyStep are evaluated approximately via samples. The resulting GreedyStep is shown in Algorithm 4.

We devote the rest of this section to proving the following result.

Theorem C.1. *The modified DCGreedy algorithm with approximate evaluation of w_e 's satisfies the consistency property and the $(\alpha, \epsilon, \delta)$ -approximation property with $\delta = 1/n$ and $\alpha = 1/e - O(\epsilon)$ for non-monotone functions and $\alpha = 1 - 1/e - O(\epsilon)$ for monotone functions.*

We begin by verifying that the consistency property holds. Consider a vector \mathbf{b} and two subsets $A, B \subseteq V$ such that, for each element $e \in B$, we have $\text{DCGreedyRel}(A \cup \{e\}, \mathbf{b}) = \text{DCGreedyRel}(A, \mathbf{b})$. We shall show that the approximate GreedyStep algorithm always picks the same set W on input $(A, \mathbf{x}, \mathbf{b})$ and $(A \cup B, \mathbf{x}, \mathbf{b})$. (Note that, since the two runs have the same randomness \mathbf{b} , they will use the same approximate weights.) Suppose for contradiction that the algorithm makes different choices on input $(A, \mathbf{x}, \mathbf{b})$ and $(A \cup B, \mathbf{x}, \mathbf{b})$. Consider the first iteration where the two runs differ, and let e be the element added to W in that iteration on input $(A \cup B, \mathbf{x}, \mathbf{b})$. Note that $e \notin A$ and thus we have $e \in B$. But then e would be added to W on input $(A \cup \{e\}, \mathbf{x}, \mathbf{b})$, as well. Thus $e \in \text{DCGreedyRel}(A \cup \{e\}, \mathbf{b})$, which contradicts the fact that $e \in B$. Thus the consistency property holds.

Now we verify that the $(\alpha, \epsilon, \delta)$ -approximation property holds. The analysis of the modified DCGreedy algorithm is similar to the analyses in [15, 3].

Algorithm 4: Greedy Update Step (GreedyStep) with randomized approximation of w_e 's.

Input: $N \subseteq V, \mathbf{x} \in [0, 1]^N$

- 1 $W \leftarrow \emptyset$
- 2 $\mathbf{y} \leftarrow \mathbf{0}$
- 3 **repeat**
- 4 Let $D \leftarrow \{e \in N \setminus W : W \cup \{e\} \in \mathcal{I}\}$
- 5 Pick $\ell = \Theta\left(\frac{s \log n}{\epsilon^2}\right)$ independent random samples for $R(\mathbf{x} + \mathbf{y})$
- 6 **foreach** $e \in D$ **do**
- 7 $w_e \leftarrow$ approximation of $\mathbf{E}[f(R(\mathbf{x} + \mathbf{y}) \cup \{e\}) - f(R(\mathbf{x} + \mathbf{y}))]$
 via above samples
- 8 Let $e^* = \arg \max_{e \in D} \{w_e\}$
- 9 **if** $D = \emptyset$ **or** $w_{e^*} < 0$ **then**
- 10 **return** \mathbf{y}
- 11 **else**
- 12 $y_{e^*} \leftarrow y_{e^*} + \epsilon(1 - x_{e^*})$
- 13 $W \leftarrow W \cup \{e\}$

Figure 3: Discretized continuous greedy (DCGreedy) with approximate evaluation of the w_e 's on line 7. The weight w_e is estimated as follows. Given ℓ independent random sets R_1, \dots, R_ℓ (the samples for $R(\mathbf{x} + \mathbf{y})$), w_e is set to $\frac{1}{\ell} \sum_{i=1}^{\ell} (f(R_i \cup \{e\}) - f(R_i))$.

Lemma C.2. *Let \mathcal{I} be matroid on V and $\text{OPT} = \arg \max_{S \in \mathcal{I}, S \subseteq V} f(S)$. Let $A \subseteq V$ and $\mathbf{b} \sim \mathcal{D}$. Let $B \subseteq \text{OPT}$ be a subset such that, for each $e \in B$, $e \notin \text{DCGreedy}(A \cup \{e\}, \mathbf{b})$. Then, we have*

$$\Pr_{\mathbf{b} \sim \mathcal{D}} [F(\text{DCGreedy}(A \cup B, \mathbf{b})) \geq \alpha \cdot f((A \cup B) \cap \text{OPT}) - \epsilon \cdot f(\text{OPT})] \geq 1 - 1/n,$$

where $\alpha = (1 - 1/e - O(\epsilon))$ for monotone f and $(1/e - O(\epsilon))$ for general f .

In the remainder of this section, we prove Lemma C.2. If $\mathbf{x}, \mathbf{y} \in [0, 1]^N$, we denote by $\mathbf{x} \vee \mathbf{y}$ the vector such that $(\mathbf{x} \vee \mathbf{y})_i = \max\{\mathbf{x}_i, \mathbf{y}_i\}$. Similarly, $\mathbf{x} \wedge \mathbf{y}$ is the vector such that $(\mathbf{x} \wedge \mathbf{y})_i = \min\{\mathbf{x}_i, \mathbf{y}_i\}$. Let $\text{OPT} = \arg \max_{S \subseteq V, S \in \mathcal{I}} f(S)$ be an optimal solution over the entire ground set V , and consider the execution of $\text{DCGreedy}(A \cup \text{OPT}, \mathbf{b})$. Let Z be the set of vectors that $\text{DCGreedy}(A \cup \text{OPT}, \mathbf{b})$ considers when computing the weights of elements, i.e., the set of all vectors $\mathbf{z} := \mathbf{x} + \mathbf{y}$, where $\mathbf{x} = \mathbf{x}(t)$ for some iteration t of $\text{DCGreedy}(A \cup \text{OPT}, \mathbf{b})$ and \mathbf{y} is the vector constructed by previous iterations of $\text{GreedyStep}(A \cup \text{OPT}, \mathbf{x}, \mathbf{b})$. Formally, we associate the vector $\mathbf{z}_j \in Z$ with the j th execution of GreedyStep 's main loop (counted across all the iterations of DCGreedy). Note that $|Z| \leq s/\epsilon$, since GreedyStep 's loop is executed at most s times for each of the $1/\epsilon$ iterations of DCGreedy . For each sample, the random string \mathbf{b} can simply store $|V|$ random thresholds in $[0, 1]$. For a given vector \mathbf{z} , these thresholds can be used to round \mathbf{z} to an integral indicator vector (a sample of $R(\mathbf{z})$) in order to estimate $\mathbf{E}[f(R(\mathbf{z}) \cup \{e\}) - f(R(\mathbf{z}))]$.

Consider the j th time GreedyStep executes line 7, and suppose that for each element $e \in A \cup \text{OPT}$ we compute a weight $w_e(\mathbf{z}_j, \mathbf{b})$, by using ℓ random samples encoded by \mathbf{b} to estimate $R(\mathbf{z}_j)$, as in

GreedyStep. We say that $w_e(\mathbf{z}_j, \mathbf{b})$ is a good estimate if

$$|w_e(\mathbf{z}_j, \mathbf{b}) - \mathbf{E}[f(R(\mathbf{z}_j) \cup \{e\}) - f(R(\mathbf{z}_j))]| \leq \frac{\epsilon}{2s} f(\text{OPT}) + \frac{\epsilon}{2} \mathbf{E}[f(R(\mathbf{z}_j) \cup \{e\}) - f(R(\mathbf{z}_j))].$$

We say that \mathbf{b} is good if all of the weights $\{w_e(\mathbf{z}_j, \mathbf{b}) : \mathbf{z}_j \in Z, e \in A \cup \text{OPT}\}$ are good estimates.

Lemma C.3. *The randomness \mathbf{b} is good with probability at least $1 - 1/n$.*

Proof: Let $d = \max_{e \in V} f(e) \leq f(\text{OPT})$. Consider a weight w_e and let R_1, \dots, R_ℓ denote the independent random sets used to compute w_e in line 7 of **GreedyStep**. For each $i \in [\ell]$, let $w_{e,i} = f(R_i \cup \{e\}) - f(R_i)$. Note that, by submodularity, $w_{e,i} \leq d \leq f(\text{OPT})$. We use the following version of the Chernoff bound.

Lemma C.4 (Lemma 2.3 in [3]). *Let X_1, \dots, X_m be independent random variables such that for each i , $X_i \in [0, 1]$. Let $X = \frac{1}{m} \sum_{i=1}^m X_i$ and $\mu = \mathbf{E}[X]$. Then*

$$\Pr[X > (1 + \alpha)\mu + \beta] \leq \exp\left(-\frac{m\alpha\beta}{3}\right),$$

$$\Pr[X < (1 - \alpha)\mu - \beta] \leq \exp\left(-\frac{m\alpha\beta}{2}\right).$$

If we choose an appropriately large constant in the definition of ℓ , then setting $m = \ell$, $X_i = w_{e,i}/d$, $\alpha = \epsilon/2$, and $\beta = \epsilon/2s$ in Lemma C.4, we obtain that w_e is a good estimate with probability at least $1 - 1/n^4 \geq 1 - \epsilon/(sn^2)$. The size of Z is at most s/ϵ and for each element of Z , there are at most n weights to be estimated, so the lemma follows by the union bound. \square

Now, note that if some random string \mathbf{b} is good, then all weights calculated by **DCGreedy**($A \cup B, \mathbf{b}$) in are good also, since $A \cup B \subseteq A \cup \text{OPT}$, and, as we have noted, the consistency property implies that **GreedyStep** picks the same set on inputs $(A \cup B, \mathbf{x}, \mathbf{b})$ and $(A, \mathbf{x}, \mathbf{b})$ in each iteration. We now fix some good \mathbf{b} , and show that for *any* $B \subseteq \text{OPT} \setminus A$, we must have:

$$F(\text{DCGreedy}(A \cup B, \mathbf{b})) \geq \alpha \cdot f((A \cup B) \cap \text{OPT}). \quad (11)$$

Where $\alpha = 1/e - O(\epsilon)$ for non-monotone functions and $\alpha = 1 - 1/e - O(\epsilon)$ for monotone functions. When f is monotone, this follows from previous work [3]. Thus we focus on the non-monotone case. This will finish the proof Lemma C.2.

Let $N = A \cup B$ for some $B \subseteq \text{OPT} \setminus A$ and consider the restricted maximization problem $\max_{S \subseteq N, S \in \mathcal{I}} f(S)$. We will need the following two lemmas from previous work. The first lemma is well-known and it follows from the exchange property of a matroid (see for example [25]).

Lemma C.5. *Let $\mathcal{M} = (N, \mathcal{I})$ be a matroid and let $B_1, B_2 \in \mathcal{I}$ be two bases in the matroid. There is a bijection $\pi : B_1 \rightarrow B_2$ such that for every element $e \in B_1$ we have $B_1 \setminus \{e\} \cup \{\pi(e)\} \in \mathcal{I}$.*

Lemma C.6 ([15]). *Consider a vector $\mathbf{x} \in [0, 1]^N$. Assuming $x_e \leq a$ for every $e \in V$, then for every set $S \subseteq N$, $F(\mathbf{x} \vee \mathbf{1}_S) \geq (1 - a)f(S)$.*

Now, we begin by showing that **DCGreedy** improves the current solution by a large amount in each step.

Lemma C.7. *Suppose that the randomness \mathbf{b} is good. In each iteration t of **DCGreedy**, $F(\mathbf{x}(t)) - F(\mathbf{x}(t-1)) \geq \epsilon(1 - \epsilon)((1 - \epsilon)^t f(N \cap \text{OPT}) - F(\mathbf{x}(t))) - \epsilon^2 f(\text{OPT})$.*

Proof: Fix an iteration t , and for brevity denote $\mathbf{x} = \mathbf{x}(t-1)$, $\mathbf{x}' = \mathbf{x}(t)$. Let W be the set of elements selected by the GreedyStep for this update, and let \mathbf{y} be the associated update vector. We suppose without loss of generality that $|W| = s$, where s is the rank of the matroid \mathcal{I} , since if $|W| < s$ we can simply add $s - |W|$ dummy elements to W . Let e_i be the i th element added to W by GreedyStep and let $\mathbf{y}(i)$ be the value of \mathbf{y} after i elements have been added to W .

By Lemma C.5, there is a bijective mapping $\pi : N \cap \text{OPT} \rightarrow W'$ between $N \cap \text{OPT}$ and a subset $W' \subseteq W$ of size $|N \cap \text{OPT}|$ such that, for each element $o \in N \cap \text{OPT}$, $W \setminus \{\pi(o)\} \cup \{o\} \in \mathcal{I}$. For each $i \in [s]$, let $o_i := \pi^{-1}(e_i)$ if $e_i \in W'$ and $o_i := e_i$ otherwise.

For each i , we have $w_{e_i} \geq w_{o_i}$, since o_i is a candidate element during the iteration of GreedyStep that picked e_i . Thus, since all the weights are good estimates, we have

$$\begin{aligned} & \mathbf{E}[f(R(\mathbf{x} + \mathbf{y}(i-1)) \cup \{e_i\}) - f(R(\mathbf{x} + \mathbf{y}(i-1)))] \\ & \geq (1 - \epsilon) \mathbf{E}[f(R(\mathbf{x} + \mathbf{y}(i-1)) \cup \{o_i\}) - f(R(\mathbf{x} + \mathbf{y}(i-1)))] - \frac{\epsilon}{s} f(\text{OPT}). \end{aligned} \quad (12)$$

for all \mathbf{y} and i . Then, we have:

$$\begin{aligned} F(\mathbf{x}') - F(\mathbf{x}) &= F(\mathbf{x} + \mathbf{y}) - F(\mathbf{x}) \\ &= \sum_{i=1}^s (F(\mathbf{x} + \mathbf{y}(i)) - F(\mathbf{x} + \mathbf{y}(i-1))) \\ &= \sum_{i=1}^s \epsilon(1 - x_{e_i}) \frac{\partial F}{\partial x_{e_i}} \Big|_{\mathbf{x} + \mathbf{y}(i-1)} \\ &= \sum_{i=1}^s \epsilon \mathbf{E}[f(R(\mathbf{x} + \mathbf{y}(i-1)) \cup \{e_i\}) - f(R(\mathbf{x} + \mathbf{y}(i-1)))] \\ &\geq \sum_{i=1}^s \epsilon \left((1 - \epsilon) \mathbf{E}[f(R(\mathbf{x} + \mathbf{y}(i-1)) \cup \{o_i\}) - f(R(\mathbf{x} + \mathbf{y}(i-1)))] - \frac{\epsilon}{s} f(\text{OPT}) \right) \\ &\geq \sum_{i=1}^s \epsilon \left((1 - \epsilon) \mathbf{E}[f(R(\mathbf{x}') \cup \{o_i\}) - f(R(\mathbf{x}'))] - \frac{\epsilon}{s} f(\text{OPT}) \right) \\ &\geq \epsilon(1 - \epsilon)(F(\mathbf{x}' \vee \mathbf{1}_{N \cap \text{OPT}}) - F(\mathbf{x}')) - \epsilon^2 f(\text{OPT}), \end{aligned} \quad (13)$$

where the first inequality follows from (12) and the last two from submodularity.

We relate the value $F(\mathbf{x}' \vee \mathbf{1}_{N \cap \text{OPT}})$ to $f(\text{OPT})$ using Lemma C.6. At each step, we increase each coordinate e of \mathbf{x} by at most $\epsilon(1 - x_e(t))$. Thus, for any step $0 \leq j \leq 1/\epsilon$, we have

$$x_e(j+1) - x_e(j) \leq (1 - x_e(j))\epsilon,$$

or, equivalently,

$$x_e(j+1) - (1 - \epsilon)x_e(j) \leq \epsilon.$$

Thus, for each time step $t \leq 1/\epsilon$, we have

$$x_e(t) \leq \sum_{j=1}^t \epsilon(1 - \epsilon)^{t-1-j} = 1 - (1 - \epsilon)^t$$

By combining the inequality above with Lemma C.6, we obtain

$$F(\mathbf{x}(t) \vee \mathbf{1}_{N \cap \text{OPT}}) \geq (1 - \epsilon)^t f(N \cap \text{OPT}).$$

Plugging this bound into (13) then completes the proof. \square

Lemma C.8. *Suppose that the randomness \mathbf{b} is good. The final solution $\mathbf{x}(1/\epsilon)$ constructed by $\text{DCGreedy}(N, \mathbf{b})$ satisfies $F(\mathbf{x}(1/\epsilon)) \geq (1/e - \epsilon)f(N \cap \text{OPT}) - \epsilon f(\text{OPT})$. Therefore the integral solution S satisfies $f(S) \geq (1/e - \epsilon)f(N \cap \text{OPT}) - \epsilon f(\text{OPT})$.*

Proof: By rearranging the inequality from Lemma C.7, we obtain

$$\begin{aligned} F(\mathbf{x}(t)) &\geq \frac{\epsilon(1 - \epsilon)^{t+1}f(N \cap \text{OPT}) + F(\mathbf{x}(t-1)) - \epsilon^2 f(\text{OPT})}{1 + \epsilon} \\ &\geq \epsilon(1 - \epsilon)^{t+2}f(N \cap \text{OPT}) + (1 - \epsilon)F(\mathbf{x}(t-1)) - \epsilon^2 f(\text{OPT}). \end{aligned}$$

It follows by induction that

$$F(\mathbf{x}(t)) \geq t\epsilon(1 - \epsilon)^{t+2}f(N \cap \text{OPT}) - t\epsilon^2 f(\text{OPT}).$$

Thus

$$F(\mathbf{x}(1/\epsilon)) \geq (1 - \epsilon)^{\frac{1}{\epsilon}+2}f(N \cap \text{OPT}) - \epsilon f(\text{OPT}) \geq \left(\frac{1}{e} - \epsilon\right)f(N \cap \text{OPT}) - \epsilon f(\text{OPT}).$$

\square

Combining Lemmas C.3 and C.8 then complete the proof of Lemma C.2.

D Improved Analysis of the Two-Round Algorithm for Non-monotone Maximization with a Cardinality Constraint

In this section, we show that for a cardinality constraint, we can improve the analysis slightly of the algorithm given in Subsection 7.1.

Theorem D.1. *If \mathcal{I} is a cardinality constraint, the two-round algorithm from Subsection 7.1 achieves a $\left(1 - \frac{1}{m}\right) \frac{1 - \frac{1}{e}}{1 + \frac{1}{\beta}(1 - \frac{1}{e})}$ approximation in expectation for non-monotone maximization, where β is the approximation guarantee of Alg.*

Proof: The analysis is similar to the one in the proof of Theorem 7.1, and we describe the main changes in this section. We define O as before, and modify the analysis of the solution S_1 as follows. Let S_1^j be the subset of S_1 consisting of the first j elements picked by Greedy, with $S_1^0 = \emptyset$. By the standard analysis of the Greedy algorithm for a cardinality constraint, for each $j \in [k]$, we have

$$f(S_1^j) - f(S_1^{j-1}) \geq \frac{f(S_1^{j-1} \cup O) - f(S_1^{j-1})}{k},$$

and therefore

$$f(S_1) \geq \sum_{j=0}^{k-1} \frac{1}{k} \left(1 - \frac{1}{k}\right)^{k-1-j} f(S_1^j \cup O).$$

Now, using $\mathbf{E}[\mathbf{1}_{S_1^j \cup O}] = \mathbf{E}[\mathbf{1}_{S_1^j}] + \mathbf{E}[\mathbf{1}_O] = \mathbf{E}[\mathbf{1}_{S_1^j}] + \mathbf{1}_{\text{OPT}} - \mathbf{p}$, and Lemma 3.2, we obtain:

$$\mathbf{E}[f(S_1^j \cup O)] \geq f^-(\mathbf{E}[\mathbf{1}_{S_1^j}] + \mathbf{1}_{\text{OPT}} - \mathbf{p}).$$

Therefore

$$\mathbf{E}[f(S_1)] \geq \sum_{j=0}^{k-1} \frac{1}{k} \left(1 - \frac{1}{k}\right)^{k-1-j} f^-(\mathbf{E}[\mathbf{1}_{S_1^j}] + \mathbf{1}_{\text{OPT}} - \mathbf{p}). \quad (14)$$

We analyze the expected value of the solution T as before, obtaining (4) from page 13. By combining (14) and (4), we get

$$\begin{aligned} \mathbf{E}[f(S_1)] + \frac{1}{\beta} \left(1 - \left(1 - \frac{1}{k}\right)^k\right) \mathbf{E}[f(T)] &\geq \sum_{j=0}^{k-1} \frac{1}{k} \left(1 - \frac{1}{k}\right)^{k-1-j} \left(f^-(\mathbf{E}[\mathbf{1}_{S_1^j}] + \mathbf{1}_{\text{OPT}} - \mathbf{p}) + f^-(\mathbf{p})\right) \\ &\geq \sum_{j=0}^{k-1} \frac{1}{k} \left(1 - \frac{1}{k}\right)^{k-1-j} 2 \cdot f^-\left(\frac{\mathbf{E}[\mathbf{1}_{S_1^j}] + \mathbf{1}_{\text{OPT}}}{2}\right), \end{aligned}$$

where the last inequality follows from the convexity of f^- . Since $S_1^j \subseteq V_1$ and V_1 is a $1/m$ sample of V , we have $\mathbf{E}[\mathbf{1}_{S_1^j}] \leq \frac{1}{m} \cdot \mathbf{1}_V$. Therefore, using the definition of f^- and the non-negativity of f , we obtain

$$2 \cdot f^-\left(\frac{\mathbf{E}[\mathbf{1}_{S_1^j}] + \mathbf{1}_{\text{OPT}}}{2}\right) \geq \left(1 - \frac{1}{m}\right) f(\text{OPT}).$$

and thus

$$\mathbf{E}[f(S_1)] + \frac{1}{\beta} \left(1 - \left(1 - \frac{1}{k}\right)^k\right) \mathbf{E}[f(T)] \geq \left(1 - \left(1 - \frac{1}{k}\right)^k\right) \left(1 - \frac{1}{m}\right) f(\text{OPT}).$$

It follows that

$$\begin{aligned} \max\{\mathbf{E}[f(S_1)], \mathbf{E}[f(T)]\} &\geq \left(1 - \frac{1}{m}\right) \frac{\left(1 - \left(1 - \frac{1}{k}\right)^k\right)}{1 + \frac{1}{\beta} \left(1 - \left(1 - \frac{1}{k}\right)^k\right)} f(\text{OPT}) \\ &\geq \left(1 - \frac{1}{m}\right) \frac{1 - \frac{1}{e}}{1 + \frac{1}{\beta} \left(1 - \frac{1}{e}\right)} f(\text{OPT}). \end{aligned}$$

□