

Scalable Distributed Algorithms for Size-Constrained Submodular Maximization in the MapReduce and Adaptive Complexity Models

Yixin Chen

*Department of Computer Science & Engineering
Texas A&M University
College Station, TX*

CHEN777@TAMU.EDU

Tonmoy Dey

*Department of Computer Science
Florida State University
Tallahassee, FL, USA*

TDEY@FSU.EDU

Alan Kuhnle

*Department of Computer Science & Engineering
Texas A&M University
College Station, TX*

KUHNLE@TAMU.EDU

Abstract

Distributed maximization of a submodular function in the MapReduce (MR) model has received much attention, culminating in two frameworks that allow a centralized algorithm to be run in the MR setting without loss of approximation, as long as the centralized algorithm satisfies a certain consistency property – which had previously only been known to be satisfied by the standard greedy and continuous greedy algorithms. A separate line of work has studied parallelizability of submodular maximization in the adaptive complexity model, where each thread may have access to the entire ground set. For the size-constrained maximization of a monotone and submodular function, we show that several sublinearly adaptive (highly parallelizable) algorithms satisfy the consistency property required to work in the MR setting, which yields practical, parallelizable and distributed algorithms. Separately, we develop the first distributed algorithm with linear query complexity for this problem. Finally, we provide a method to increase the maximum cardinality constraint for MR algorithms at the cost of additional MR rounds.

1. Introduction

Submodular maximization has become an important problem in data mining and machine learning with real world applications ranging from video summarization (Mirzasoleiman et al., 2018) and mini-batch selection (Joseph et al., 2019) to more complex tasks such as active learning (Rangwani et al., 2021) and federated learning (Balakrishnan et al., 2022). In this work, we study the size-constrained maximization of a monotone, submodular function (SMCC), formally defined in Section 1.3. Because of the ubiquity of problems requiring the optimization of a submodular function, a vast literature on submodular optimization exists; we refer the reader to the surveys (Liu et al., 2020; Liu, 2020).

Reference	Ratio	MR Rounds	Adaptivity (Θ)	Queries (Ξ)	k_{max}
RG	$\frac{1}{2}(1 - \frac{1}{e})$	2	$O(k)$	$O(nk + k^2\ell)$	$O(\frac{n}{\ell^2})$
PALG	$1 - \frac{1}{e} - \varepsilon$	$O(\frac{1}{\varepsilon^2})$	$O(\frac{k}{\varepsilon^2})$	$O(\frac{nk}{\varepsilon^2} + \frac{k^2\ell^2}{\varepsilon^4})$	$O(\frac{n\varepsilon^2}{\ell^2})$
DDIST	$1 - \frac{1}{e} - \varepsilon$	$O(\frac{1}{\varepsilon})$	$O(\frac{k}{\varepsilon})$	$O(\frac{nk}{\varepsilon} + \frac{k^5\ell^2}{\varepsilon^2})$	$O(\frac{n\varepsilon}{\ell^2})$
BCG	$1 - \varepsilon$	r	$O(\frac{rk}{\varepsilon^{2/r}} \log^2(\frac{1}{\varepsilon^{1/r}}))$	$O(\frac{nk}{\varepsilon^{1/r}} + \frac{k^2\ell r m^2(\varepsilon^{1/r})}{\varepsilon^{3/r}})$	$O(\frac{n\varepsilon^{1/r}}{\ell^2})$
PG	$0.545 - \varepsilon$	2	$O(k^2)$	$O(nk + k^3)$	$O(\frac{n}{\ell^2})$
TG	$\frac{1}{2}$	2	$O(n \log(k))$	$O(n \log(k))$	$O(\frac{n}{\ell^2 \log(k)})$
R-DASH (Alg. 4)	$\frac{1}{2}(1 - \frac{1}{e} - \varepsilon)$	2	$O(\frac{\log(k) \log(n)}{\varepsilon^4})$	$O(\frac{n \log(k)}{\varepsilon^4})$	$O(\frac{n}{\ell^2})$
G-DASH (Alg. 5)	$1 - \frac{1}{e} - \varepsilon$	$O(\frac{1}{\varepsilon})$	$O(\frac{\log(k) \log(n)}{\varepsilon^5})$	$O(\frac{n \log(k)}{\varepsilon^5})$	$O(\frac{n\varepsilon}{\ell^2})$
T-DASH (Alg. 10)	$\frac{3}{8} - \varepsilon$	2	$O(\frac{\log(n)}{\varepsilon^3})$	$O(\frac{n \log(k)}{\varepsilon^4})$	$O(\frac{n\varepsilon}{\ell^2 \log(k)})$
L-DIST (Alg. 7)	$\frac{1}{8}$	2	$O(\frac{n}{\ell})$	$O(n)$	$O(\frac{n}{\ell^2 \log(n)})$

Table 1: Theoretical comparison of our algorithms to MapReduce model algorithms: RG of Barbosa et al. (2015), PALG of Barbosa et al. (2016), DDIST of Kazemi et al. (2021), BCG of Epasto et al. (2017), PG of Mirrokni and Zadimoghaddam (2015), and TG of Liu and Vondrak (2019). The column k_{max} provides the maximum cardinality constraint that the algorithm can compute.

A foundational result of Nemhauser, Wolsey, and Fisher (1978) shows that a simple greedy algorithm (GREEDY, pseudocode in Appendix H) achieves the optimal approximation ratio for SMCC of $1 - 1/e \approx 0.63$ in the value query model, in which the submodular function f is available to the algorithm as an oracle that returns $f(S)$ when queried with set S . However, because of the modern revolution in data (Mao et al., 2021; Ettinger et al., 2021), there is a need for more scalable algorithms. Specifically, the standard greedy algorithm is a centralized algorithm that needs access to the whole dataset, makes many adaptive rounds (defined below) of queries to the submodular function, and has quadratic total query complexity in the worst case.

Distributed Algorithms for SMCC. Massive data sets are often too large to fit on a single machine and are distributed over a cluster of many machines. In this context, there has been a line of work developing algorithms for SMCC in the MapReduce model, defined formally in Section 1.3 (see Table 1 and the Related Work section below). Of these, PALG (Barbosa et al., 2015) is a general framework that allows a centralized algorithm ALG to be converted to the MapReduce distributed setting with nearly the same theoretical approximation ratio, as long as ALG satisfies a technical, randomized consistency property (RCP), defined in Section 1.3.

In addition, DISTRIBUTEDDISTORTED (Kazemi et al., 2019), is also a general framework that can adapt any ALG satisfying RCP, as we show¹ in Section 4.2. However, to the best of our knowledge, the only centralized algorithms to satisfy this property are the GREEDY (Nemhauser et al., 1978) and the continuous greedy algorithm of Calinescu, Chekuri, Pal, and Vondrak (2007).

Parallelizable Algorithms for SMCC. A separate line of works, initiated by Balkanski and Singer (2018) has taken an orthogonal approach to scaling up the standard greedy

¹In Kazemi et al. (2019), an instantiated version of DISTRIBUTEDDISTORTED with a distorted greedy algorithm is analyzed.

algorithm: parallelizable algorithms for SMCC as measured by the *adaptive complexity* of the algorithm, formally defined in Section 1.3. In this model, each thread may have access to the entire ground set and hence these algorithms do not apply in a distributed setting. Altogether, these algorithms have exponentially improved the adaptivity of standard greedy from $O(n)$ to $O(\log(n))$ while obtaining the nearly the same $1 - 1/e$ approximation ratio (Fahrbach, Mirrokni, & Zadimoghaddam, 2019b; Balkanski, Rubinfeld, & Singer, 2019; Ene & Nguyen, 2019; Chekuri & Quanrud, 2019). Recently, highly practical sublinearly adaptive algorithms have been developed without compromising theoretical guarantees (Breuer, Balkanski, & Singer, 2020; Chen, Dey, & Kuhnle, 2021). However, none of these algorithms have been shown to satisfy the RCP.

Combining Parallelizability and Distributed Settings. In this work, we study parallelizable, distributed algorithms for SMCC; each node of the distributed cluster likely has many processors, so we seek to take advantage of this situation by parallelizing the algorithm within each machine in the cluster. In the existing MapReduce algorithms, the number of processors ℓ could be set to the number of processors available to ensure the usage of all available resources – this would consider each processor as a separate machine in the cluster. However, there are a number of disadvantages to this approach: 1) A large number of machines severely restricts the size of the cardinality constraint k – since, in all of the MapReduce algorithms, a set of size $k\ell$ must be stored on a single machine; therefore, we must have $k \leq O(\Psi/\ell) = O(n/\ell^2)$. For example, with $n = 10^6$, $\ell = 256$ implies that $k \leq 15 \cdot C$, for some C . However, if each machine has 8 processor cores, the number of machines ℓ can be set to 32 (and then parallelized on each machine), which enables the use of 256 processors with $k \leq 976 \cdot C$. As the number of cores and processors per physical machine continues to grow, large practical benefits can be obtained by a parallelized and distributed algorithm. 2) Modern CPU architectures have many cores that all share a common memory, and it is inefficient to view communication between these cores to be as expensive as communication between separate machines in a cluster.

In addition to the MapReduce and adaptive complexity models, the total query complexity of the algorithm to the submodular oracle is also relevant to the scalability of the algorithm. Frequently, evaluation of the submodular function is expensive, so the time spent answering oracle queries dominates the other parts of the computation. Therefore, the following questions are posed: *Q1: Is it possible to design constant-factor approximation distributed algorithms with 1) constant number of MapReduce rounds; 2) sublinear adaptive complexity (highly-parallelizable); and 3) nearly linear total query complexity? Q2: Can we design practical distributed algorithms that also meet these three criteria?*

1.1 Contributions

In overview, the contributions of the paper are: 1) an exponential improvement in adaptivity (from $\Omega(n)$ to $O(\log n)$) for an algorithm with a constant number of MR rounds; 2) the first linear-time, constant-factor algorithm with constant MR rounds; 3) a way to increase the size constraint limitation at the cost of additional MR rounds; and 4) an empirical evaluation on a cluster of 64 machines that demonstrates an order of magnitude improvement in runtime is obtained by combining the MR and adaptive complexity models.

To develop MR algorithms with sublinear adaptivity, we first modify and analyze the low-adaptive algorithm THRESHSEQMOD from Chen et al. (2021) and show that our modification satisfies the RCP (Section 2). Next, THRESHSEQMOD is used to create a low-adaptive greedy procedure LAG; LAG then satisfies the RCP, so can be used within PALG (Barbosa et al., 2016) to yield a $1 - 1/e - \varepsilon$ approximation in $O(1/\varepsilon^2)$ MR rounds, with adaptivity $O(\log n)$. We show in Section 4.2 that DISTRIBUTEDDISTORTED (Kazemi et al., 2021) also works with any ALG satisfying RCP, which yields an improvement in the number of MR rounds to achieve the same ratio. The resulting algorithm we term G-DASH (Section 4), which achieves ratio $1 - 1/e - \varepsilon$ with $O(\log^2 n)$ adaptivity and $O(1/\varepsilon)$ MR rounds.

Although G-DASH achieves nearly the optimal ratio in constant MR rounds and sublinear adaptivity, the number of MR rounds is high. To obtain truly practical algorithms, we develop constant-factor algorithms with 2 MR rounds: R-DASH and T-DASH with $O(\log(k) \log(n))$ and $O(\log n)$ adaptivity, respectively, and with approximation ratios of $\frac{1}{2}(1 - 1/e - \varepsilon)$ ($\simeq 0.316$) and $3/8 - \varepsilon$. R-DASH is our most practical algorithm and may be regarded as a parallelized version of RANDGREEDI (Barbosa et al., 2015), the first MapReduce algorithm for SMCC and the state-of-the-art MR algorithm in terms of empirical performance. T-DASH is a novel algorithm that improves the theoretical properties of R-DASH at the cost of empirical performance (see Table 1 and Section 8). Notably, the adaptivity of T-DASH is $O(\log(n))$ which is close to the best known adaptivity for a constant factor algorithm of Chen et al. (2021), which has adaptivity of $O(\log(n/k))$.

Although our MR algorithms are nearly linear time (within a polylogarithmic factor of linear), all of the existing MR algorithms are superlinear time, which raises the question of if an algorithm exists that has a linear query complexity and has constant MR rounds and constant approximation factor. We answer this question affirmatively by adapting a linear-time algorithm of Kuhnle (2021), Chen et al. (2021), and showing that our adaptation satisfies RCP (Section 3). Subsequently, we develop the first MR algorithm (L-DIST) with an overall linear query complexity.

Our next contribution is MED, a general plug-in framework for distributed algorithms, which increases the size of the maximum cardinality constraint at the cost of more MR rounds. As discussed above, the maximum k value of any prior MR algorithm for SMCC is $O(n/\ell^2)$, where ℓ is the number of machines; MED increases this to $O(n/\ell)$. We also show that under certain assumptions on the objective function (which are satisfied by all of the empirical applications evaluated in Section 8), MED can be run with $k_{max} = n$, which removes any cardinality restrictions. When used in conjunction with a γ -approximation MR algorithm, MED provides an $(1 - e^{-\gamma})$ -approximate solution.

Finally, an extensive empirical evaluation of our algorithms and the current state-of-the-art on a 64-machine cluster with 32 cores each and data instances ranging up to 5 million nodes show that R-DASH is orders of magnitude faster than state-of-the-art MR algorithms and demonstrates an exponential improvement in scaling with larger k . Moreover, we show that MR algorithms slow down with increasing number of machines past a certain point, even if enough memory is available, which further motivates distributing a parallelizable algorithm over smaller number of machines. This observation also serves as a motivation for the development of the MED+ALG framework. In our evaluation, we found that

the MED+ALG framework delivers solutions significantly faster than the ALG for large constraints, showcasing its superior performance and efficiency.

A previous version of this work was published in a conference (Dey, Chen, & Kuhnle, 2023). In that version, RCP was claimed to be shown for the THRESHOLDSEQ algorithm of Chen et al. (2021). Unfortunately, RCP does not hold for this algorithm. In this version, we provide a modified version of the THRESHOLDSEQ algorithm of Chen et al. (2021), for which we show RCP. Additionally, in this version, we introduce the first linear time MR algorithm, LINEAR-TIME-DISTRIBUTED (L-DIST). Finally, in our empirical evaluation, we expand upon the conference version by conducting two additional experiments using a larger cluster comprising 64 machines. These additional experiments aim to evaluate the performance of L-DIST and investigate the effectiveness of MED in handling increasing cardinality constraints, respectively.

1.2 Organization

The rest of this paper is organized as follows. In Section 2, we present the low-adaptive procedures and show they satisfy the RCP. In Section 3, we analyze an algorithm with linear query complexity and show it satisfies the RCP. In Section 4, we show how to use algorithms satisfying the RCP to obtain MR algorithms. In Section 5, we improve the theoretical properties of our 2-round MR algorithm. In Section 6, we detail the linear-time MR algorithm. In Section 7, we show how to increase the supported constraint size by adding additional MR rounds. In Section 8, we conduct our empirical evaluation.

1.3 Preliminaries

A submodular set function captures the diminishing gain property of adding an element to a set that decreases with increase in size of the set. Formally, let N be a finite set of size n . A non-negative set function $f : 2^N \rightarrow \mathbb{R}^+$ is submodular iff for all sets $A \subseteq B \subseteq N$ and $x \in N \setminus B$, $f(A \cup x) - f(A) \geq f(B \cup x) - f(B)$. The submodular function is monotone if $f(A) \leq f(B)$ for all $A \subseteq B$. We use $\Delta(x | S)$ to denote the marginal game of element x to set S : $\Delta(x | S) = f(S \cup \{x\}) - f(S)$. In this paper, we study the following optimization problem for submodular optimization under cardinality constraint (SMCC):

$$O \leftarrow \arg \max_{S \subseteq N} f(S); \text{ subject to } |S| \leq k \tag{1}$$

MapReduce Model. The MapReduce (MR) model is a formalization of distributed computation into MapReduce rounds of communication between machines. A dataset of size n is distributed on ℓ machines, each with memory to hold at most Ψ elements of the ground set. The total memory of the machines is constrained to be $\Psi \cdot \ell = O(n)$. After each round of computation, a machine may send $O(\Psi)$ amount of data to other machines. We assume $\ell \leq n^{1-c}$ for some constant $c \geq 1/2$.

Next, we formally define the RCP needed to use the two frameworks to convert a centralized algorithm to the MR setting, as discussed previously.

Property 1 (Randomized Consistency Property of Barbosa et al. (2016)). *Let \mathbf{q} be a fixed sequence of random bits; and ALG be a randomized algorithm with randomness determined*

by \mathbf{q} . Suppose $\text{ALG}(\mathcal{N}, \mathbf{q})$ returns a pair of sets, $(\text{ALGSOL}(\mathcal{N}, \mathbf{q}), \text{ALGREL}(\mathcal{N}, \mathbf{q}))$, where $\text{ALGSOL}(\mathcal{N}, \mathbf{q})$ is the feasible solution and $\text{ALGREL}(\mathcal{N}, \mathbf{q})$ is a set providing additional information. Let A and B be two disjoint subsets of \mathcal{N} , and that for each $b \in B$, $\text{ALGREL}(A \cup \{b\}, \mathbf{q}) = \text{ALGREL}(A, \mathbf{q})$. Also, suppose that $\text{ALG}(A, \mathbf{q})$ terminates successfully. Then $\text{ALG}(A \cup B, \mathbf{q})$ terminates successfully and $\text{ALGSOL}(A \cup B, \mathbf{q}) = \text{ALGSOL}(A, \mathbf{q})$.

Adaptive Complexity. The adaptive complexity of an algorithm is the minimum number of sequential *adaptive rounds* of at most polynomially many queries, in which the queries to the submodular function can be arranged, such that the queries in each round only depend on the results of queries in previous rounds.

1.4 Related Work

MapReduce Algorithms. Addressing the coverage maximization problem (a special case of SMCC), Chierichetti et al. (2010) proposed an approximation algorithm achieving a ratio of $1 - 1/e$, with polylogarithmic MapReduce (MR) rounds. This was later enhanced by Blesloch, Simhadri, and Tangwongsan (2012), who reduced the number of rounds to $\log^2 n$. Kumar et al. (2013) contributed further advancements with a $1 - 1/e$ algorithm, significantly reducing the number of MR rounds to logarithmic levels. Additionally, they introduced a $1/2 - \epsilon$ approximation algorithm that operates within $O(1/\delta)$ MR rounds, albeit with a logarithmic increase in communication complexity.

For SMCC, Mirzasoleiman et al. (2013) introduced the two-round distributed greedy algorithm (GREEDI), demonstrating its efficacy through empirical studies across various machine learning applications. However, its worst-case approximation guarantee is $1/\Theta(\min\{\sqrt{k}, \ell\})$. Subsequent advancements led to the development of constant-factor algorithms by Barbosa et al. (2015). Barbosa et al. (2015) notably introduced randomization into the distributed greedy algorithm GREEDI, resulting in the RANDGREEDI algorithm. This algorithm achieves a $\frac{1}{2}(1 - 1/e)$ approximation ratio with two MR rounds and $O(nk)$ queries. Mirrokni and Zadimoghaddam (2015) also introduced an algorithm with two MR rounds, where the first round found randomized composable core-sets and the second round applied a tie-breaking rule. This algorithm improved the approximation ratio from $\frac{1}{2}(1 - 1/e)$ to $0.545 - \epsilon$. Later on, Barbosa et al. (2016) proposed a $O(1/\epsilon^2)$ round algorithm with the optimal $(1 - 1/e - \epsilon)$ approximation ratio without data duplication. Then, Kazemi et al. (2021) improved its space and communication complexity by a factor of $O(1/\epsilon)$ with the same approximation ratio. Another two-round algorithm, proposed by Liu and Vondrak (2019), achieved an $1/2 - \epsilon$ approximation ratio, but requires data duplication with four times more elements distributed to each machine in the first round. As a result, distributed setups have a more rigid memory restriction when running this algorithm.

Parallelizable Algorithms. A separate line of works consider the parallelizability of the algorithm, as measured by its adaptive complexity. Balkanski and Singer (2018) introduced the first $O(\log(n))$ -adaptive algorithm, achieving a $(1/3 - \epsilon)$ approximation ratio with $O(nk^2 \log^3(n))$ query complexity for SMC. Then, Balkanski et al. (2019) enhanced the approximation ratio to $1 - 1/e - \epsilon$ with $O(nk^2 \log^2(n))$ query complexity while maintaining the same sublinear adaptivity. In the meantime, Ene and Nguyen (2019), Chekuri and Quanrud (2019), and Fahrback et al. (2019b) achieved the same approximation ratio and

adaptivity but with improved query complexities: $O(n \text{poly}(\log(n)))$ for Ene and Nguyen (2019), $O(n \log(n))$ for Chekuri and Quanrud (2019), and $O(n)$ for Fahrbach et al. (2019b). Recently, highly practical sublinearly adaptive algorithms have been developed without compromising on the approximation ratio. FAST, introduced by Breuer et al. (2020), operates with $O(\log(n) \log^2(\log(k)))$ adaptive rounds and $O(n \log(\log(k)))$ queries, while LS+PGB, proposed by Chen et al. (2021), runs with $O(\log(n))$ adaptive rounds and $O(n)$ queries. However, none of these algorithms have been shown to satisfy the RCP.

Fast (Low Query Complexity) Algorithms. Since a query to the oracle for the submodular function is typically an expensive operation, the total query complexity of an algorithm is highly relevant to its scalability. The work by Badanidiyuru and Vondrak (2014) reduced the query complexity of the standard greedy algorithm from $O(kn)$ to $O(n \log(n))$, while nearly maintaining the ratio of $1 - 1/e$. Subsequently, Mirzasoleiman et al. (2015) utilized random sampling on the ground set to achieve the same approximation ratio in expectation with optimal $O(n)$ query complexity. Afterwards, Kuhnle (2021) obtained a deterministic algorithm that achieves ratio $1/4$ in exactly n queries, and the first deterministic algorithm with $O(n)$ query complexity and $(1 - 1/e - \varepsilon)$ approximation ratio. Since our goal is to develop practical algorithms in this work, we develop parallelizable and distributed algorithms with nearly linear query complexity (*i.e.* within a polylogarithmic factor of linear). Further, we develop the first linear-time algorithm with constant number of MR rounds, although it doesn't parallelize well.

Non-monotone Algorithms. If the submodular function is no longer assumed to be monotone, the SMCC problem becomes more difficult. Here, we highlight a few works in each of the categories of distributed, parallelizable, and query efficient. Lee et al. (2009) provided the first constant-factor approximation algorithm with an approximation ratio of $(1/4 - \varepsilon)$ and a query complexity of $O(nk^5 \log(k))$. Building upon that, Gupta et al. (2010) reduced the query complexity to $O(nk)$ by replacing the local search procedure with a greedy approach that resulted in a slightly worse approximation ratio. Subsequently, Buchbinder et al. (2014) incorporated randomness to achieve an expected $1/e$ approximation ratio with $O(nk)$ query complexity. Furthermore, Buchbinder et al. (2017) improved it to $O(n)$ query complexity while maintaining the same approximation ratio in expectation. For parallelizable algorithms in the adaptive complexity model, Ene and Nguyen (2020) achieved the current best $(1/e - \varepsilon)$ approximation ratio with $O(\log(n))$ adaptivity and $O(nk^2 \log^2(n))$ queries to the continuous oracle (a continuous relaxation of the original value oracle). Later, Chen and Kuhnle (2024) achieved the current best sublinear adaptivity and nearly linear query complexity with a $(1/6 - \varepsilon)$ approximation ratio. In terms of distributed algorithms in the MapReduce model, RANDGREEDI proposed by Barbosa et al. (2015) achieved a 0.1 approximation ratio within 2 MR rounds and $O(nk)$ query complexity. Simultaneously, Mirrokni and Zadimoghaddam (2015) also presented a 2-round algorithm with a 0.14 approximation ratio and $O(nk)$ query complexity. Afterwards, Barbosa et al. (2016) proposed a 2-round continuous algorithm with a $0.232(1 - \frac{1}{\ell})$ approximation ratio.

Algorithm 1 Low-Adaptive Threshold Algorithm (THRESHSEQMOD)

```

1: procedure THRESHSEQMOD( $f, X, k, \delta, \varepsilon, \tau, \mathbf{q}$ )
2:   Input: evaluation oracle  $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ , subset  $X \subseteq \mathcal{N}$ , constraint  $k$ , confidence  $\delta$ ,
   error  $\varepsilon$ , threshold  $\tau$ , a finite set of sequences of random bits  $\mathbf{q} \leftarrow \{\sigma_1, \dots, \sigma_{M+1}\}$ 
3:   Initialize  $S_0 \leftarrow \emptyset$ ,  $R_0 \leftarrow \emptyset$ ,  $V_0 \leftarrow X$ ,  $M \leftarrow \left\lceil 4 \left(1 + \frac{1}{\beta\varepsilon}\right) \log\left(\frac{n}{\delta}\right) \right\rceil$ ,  $\beta \leftarrow$ 
    $\varepsilon \left(16 \log\left(\frac{4}{1-e^{-\varepsilon/2}}\right)\right)^{-1}$ 
4:   for  $j \leftarrow 1$  to  $M + 1$  do
5:     Update  $V_j \leftarrow \{x \in V_{j-1} : \Delta(x | S_{j-1}) \geq \tau\}$ 
6:     if  $|V_j| = 0$  or  $|S_{j-1}| = k$  then
7:       return  $S_{j-1}, R_{j-1}$ , success
8:      $V_j \leftarrow \mathbf{random-permutation}(V_j, \sigma_j)$ 
9:      $s \leftarrow \min\{k - |S_{j-1}|, |V_j|\}$ 
10:     $\Lambda \leftarrow \{1, 2, \dots, \min\{s, \lceil \frac{1}{\varepsilon} \rceil\}\} \cup \{\lfloor (1 + \varepsilon)^u \rfloor : 1 \leq \lfloor (1 + \varepsilon)^u \rfloor \leq s, u \in \mathbb{N}\} \cup \{s\}$ 
11:     $B[\lambda] \leftarrow \mathbf{False}, \forall \lambda \in \Lambda$ 
12:    for  $\lambda \in \Lambda$  in parallel do
13:       $T_\lambda \leftarrow \{v_1, v_2, \dots, v_\lambda\}$ 
14:      if  $\Delta(T_\lambda | S_{j-1}) / |T_\lambda| \geq (1 - \varepsilon)\tau$  then
15:         $B[\lambda] \leftarrow \mathbf{True}$ 
16:       $\lambda'_j \leftarrow \min\{\lambda \in \Lambda : B[\lambda] = \mathbf{False} \text{ OR } \lambda = s\}$ 
17:       $R_j \leftarrow R_{j-1} \cup T_{\lambda'_j}$ 
18:      if  $\lambda'_j \leq \lceil \frac{1}{\varepsilon} \rceil$  then  $\lambda_j^* \leftarrow \lambda'_j - 1$  else  $\lambda_j^* \leftarrow \lambda'_j$ 
19:       $S_j \leftarrow S_{j-1} \cup T_{\lambda_j^*}$ 
20:   return  $S_{M+1}, R_{M+1}$ , failure

```

2. Low-Adaptive Algorithms That Satisfy the RCP

In this section, we analyze two low-adaptive procedures, THRESHSEQMOD (Alg. 1), LAG (Alg. 2), variants of low-adaptive procedures proposed in Chen et al. (2021). This analysis enables their use in the distributed, MapReduce setting. For convenience, we regard the randomness of the algorithms to be determined by a sequence of random bits \mathbf{q} , which is an input to each algorithm.

Observe that the randomness of both of THRESHSEQMOD and LAG only comes from the random permutations of V_j on Line 8 of THRESHSEQMOD, since LAG employs THRESHSEQMOD as a subroutine. Consider an equivalent version of these algorithms in which the entire ground set \mathcal{N} is permuted randomly, from which the permutation of V_j is extracted. That is, if σ is the permutation of \mathcal{N} , the permutation of V is given by $v < w$ iff $\sigma(v) < \sigma(w)$, for $v, w \in V$. Then, the random vector \mathbf{q} specifies a sequence of permutations of \mathcal{N} : $(\sigma_1, \sigma_2, \dots)$, which completely determines the behavior of both algorithms.

2.1 Low-Adaptive Threshold (ThreshSeqMod) Algorithm

This section presents the analysis of the low-adaptive threshold algorithm, THRESHSEQMOD (Alg. 1; a variant of THRESHOLDSEQ from Chen et al. (2021)). In THRESHSEQMOD, the

randomness is explicitly dependent on a random vector \mathbf{q} . This modification ensures the consistency property in the distributed setting. Besides the addition of randomness \mathbf{q} , THRESHSEQMOD employs an alternative strategy of prefix selection within the for loop. Instead of identifying the *failure point* (the average marginal gain is under the threshold) after the final *success point* (the average marginal gain is above the threshold) in Chen et al. (2021), THRESHSEQMOD determines the first failure point. This adjustment not only addresses the consistency problem but also preserves the theoretical guarantees below.

Theorem 1. *Suppose THRESHSEQMOD is run with input $(f, X, k, \delta, \varepsilon, \tau, \mathbf{q})$. Then, the algorithm has adaptive complexity $O(\log(n/\delta)/\varepsilon^3)$ and outputs $S, R \subseteq \mathcal{N}$, where S is the solution set with $|S| \leq k$ and R provides additional information with $|R| = O(k)$. The following properties hold: 1) The algorithm succeeds with probability at least $1 - \delta/n$. 2) There are $O(n/\varepsilon^3)$ oracle queries in expectation. 3) It holds that $f(S)/|S| \geq (1 - 2\varepsilon)\tau/(1 + \varepsilon)$. 4) If $|S| < k$, then $\Delta(x|S) < \tau$ for all $x \in \mathcal{N}$.*

2.1.1 ANALYSIS OF CONSISTENCY

Lemma 1. THRESHSEQMOD *satisfies the randomized consistency property (Property 1).*

Proof. Consider that the algorithm runs for all $M + 1$ iterations of the outer **for** loop; if the algorithm would have returned at iteration j , the values S_i and V_i , for $i > j$ keep their values from when the algorithm would have returned. The proof relies upon the fact that every call to THRESHSEQMOD(\cdot, \mathbf{q}) uses same sequences of permutations of \mathcal{N} : $\{\sigma_1, \sigma_2, \dots, \sigma_{M+1}\}$. We refer to iterations of the outer **for** loop on Line 4 of Alg. 1 simply as iterations. Since randomness only happens with the random permutation of \mathcal{N} at each iteration in Alg. 1, the randomness of THRESHSEQMOD(\cdot, \mathbf{q}) is determined by \mathbf{q} , which satisfies the hypothesis of Property 1.

For the two sets returned by THRESHSEQMOD(\mathcal{N}, \mathbf{q}), $S_{M+1} = \text{THRESHSEQMODSOL}(\mathcal{N}, \mathbf{q})$ represents the feasible solution, and $R_{M+1} = \text{THRESHSEQMODREL}(\mathcal{N}, \mathbf{q})$ is the set that provides additional information. We consider the runs of (1) THRESHSEQMOD(A, \mathbf{q}), (2) THRESHSEQMOD($A \cup \{b\}, \mathbf{q}$), and (3) THRESHSEQMOD($A \cup B, \mathbf{q}$) together. Variables of (1) are given the notation defined in the pseudocode; variables of (2) are given the superscript b ; and variables of (3) are given the superscript $'$. Suppose for each $b \in B$, $R_{M+1}^b = R_{M+1}$. The lemma holds if $S'_{M+1} = S_{M+1}$ and THRESHSEQMOD($A \cup B, \mathbf{q}$) terminates successfully.

Let $P(i)$ be the statement for iteration i of the outer for loop on Line 4 that

- (i) $S_i = S'_i$, and
- (ii) for all $b \in B$, $S_i = S_i^b$, $R_i = R_i^b$, and
- (iii) $V_i = V'_i \setminus B$, and
- (iv) for all $b \in B$, $V_i = V_i^b \setminus \{b\}$.

If $P(M + 1)$ is true, and THRESHSEQMOD(A, \mathbf{q}) and THRESHSEQMOD($A \cup B, \mathbf{q}$) terminate at the same iteration, implying that THRESHSEQMOD($A \cup B, \mathbf{q}$) also terminates successfully, then the lemma holds immediately. In the following, we prove these two statements by induction.

For $i = 0$, it is clear that $S_0 = S'_0 = R_0 = R_0^b = \emptyset$, and $V_0 = V'_0 \setminus B = V_0^b \setminus \{b\} = \mathcal{N} \setminus B$ for all $b \in B$. Therefore, $P(0)$ is true. Next, suppose that $P(i-1)$ is true. We show that $P(i)$ is also true, and if $\text{THRESHSEQMOD}(A, \mathbf{q})$ terminates at iteration i , $\text{THRESHSEQMOD}(A \cup B, \mathbf{q})$ also terminates at iteration i .

Firstly, we show that (iii) and (iv) of $P(i)$ hold. Since $P(i-1)$ holds, it indicates that $S_{i-1} = S'_{i-1} = S_{i-1}^b$ for any $b \in B$ by (i) and (ii) of $P(i-1)$. So, (iii) and (iv) of $P(i)$ clearly hold since the sets $S_{i-1}, S'_{i-1}, S_{i-1}^b$ involved in updating V_i, V'_i, V_i^b on Line 5 are equal (after B is removed).

Secondly, we prove that (i) and (ii) of $P(i)$ hold and that if $\text{THRESHSEQMOD}(A, \mathbf{q})$ terminates at iteration i , $\text{THRESHSEQMOD}(A \cup B, \mathbf{q})$ also terminates at iteration i .

If $\text{THRESHSEQMOD}(A, \mathbf{q})$ terminates at iteration i because $|S_{i-1}| = k$, then the other two runs of THRESHSEQMOD also terminate at iteration i , since $|S'_{i-1}| = |S_{i-1}^b| = |S_{i-1}| = k$ by (i) and (ii) of $P(i-1)$. Moreover, (i) and (ii) of $P(i)$ are true since S'_i, S_i^b, S_i, R_i^b , and R_i do not update.

If $\text{THRESHSEQMOD}(A, \mathbf{q})$ terminates at iteration i because $V_i = \emptyset$, then by (iii) and (iv) of $P(i)$, it holds that $V_i^b \setminus \{b\} = V'_i \setminus B = \emptyset$ for all $b \in B$. Thus, V_i^b must be empty; otherwise, $V_i^b = \{b\}$ and b will be added to R_{i-1}^b at iteration i which contradicts the fact that $b \notin R_{M+1} = R_{M+1}^b$. Therefore, b has been filtered out before iteration i or will be filtered out at iteration i in both $\text{THRESHSEQMOD}(A \cup \{b\}, \mathbf{q})$ and $\text{THRESHSEQMOD}(A \cup B, \mathbf{q})$ since the sets S_{i-1}^b and S_{i-1}^B involved in updating V_i^b and V_i^B are the same. Consequently, V'_i is also an empty set and $\text{THRESHSEQMOD}(A \cup B, \mathbf{q})$ terminates at iteration i . Furthermore, (i) and (ii) of $P(i)$ are true since S'_i, S_i^b, S_i, R_i^b , and R_i do not update.

Finally, consider the case that $\text{THRESHSEQMOD}(A, \mathbf{q})$ does not terminate at iteration i , where $V_i \neq \emptyset$ and $|S_{i-1}| < k$. By (iii) and (iv), it holds that $V_i = V'_i \setminus B = V_i^b \setminus \{b\}$, for any $b \in B$. Let $b_1 \in B$ be the first element appeared in the sequence of V'_i , $\lambda^\#$ be the index of b_1 in the order of the permutation of V'_i . If such b_1 does not exist (*i.e.* $V_i = V'_i = V_i^b$, for any $b \in B$), let $\lambda^\# = |V_i| + 1$. Fig. 1 depicts the relationship between three sequences $V_i, V_i^{b_1}$ and V'_i . Let $T_{\lambda'_i}, T_{\lambda'_i}^{b_1}, T_{\lambda'_i}$ be sets added to $R_{i-1}, R_{i-1}^{b_1}, R'_{i-1}$ at iteration i respectively. Since $V_i^{b_1}$ and V'_i are subsets of the same random permutation of \mathcal{N} , and $V'_i \setminus B = V_i^{b_1} \setminus \{b_1\}$, $\lambda^\#$ is also the index of b_1 in $V_i^{b_1}$. Also, it holds that

$$T_x = T_x^{b_1} = T'_x, \forall x < \lambda^\#. \quad (2)$$

Since $b_1 \notin R_{M+1}^{b_1}$, it holds that $b_1 \notin T_{\lambda'_i}^{b_1}$, which indicates that $\lambda'_i < \lambda^\#$. By the definition of λ'_i on Line 16,

$$\begin{cases} B^{b_1}[x] = \mathbf{True}, \forall x < \lambda'_i, \\ B^{b_1}[\lambda'_i] = \mathbf{False}. \end{cases}$$

Follows from Equality 2, and $\lambda'_i < \lambda^\#$, it holds that

$$\begin{cases} B[x] = B'[x] = \mathbf{True}, \forall x < \lambda'_i, \\ B[\lambda_i^*] = B'[\lambda'_i] = \mathbf{False}, \end{cases}$$

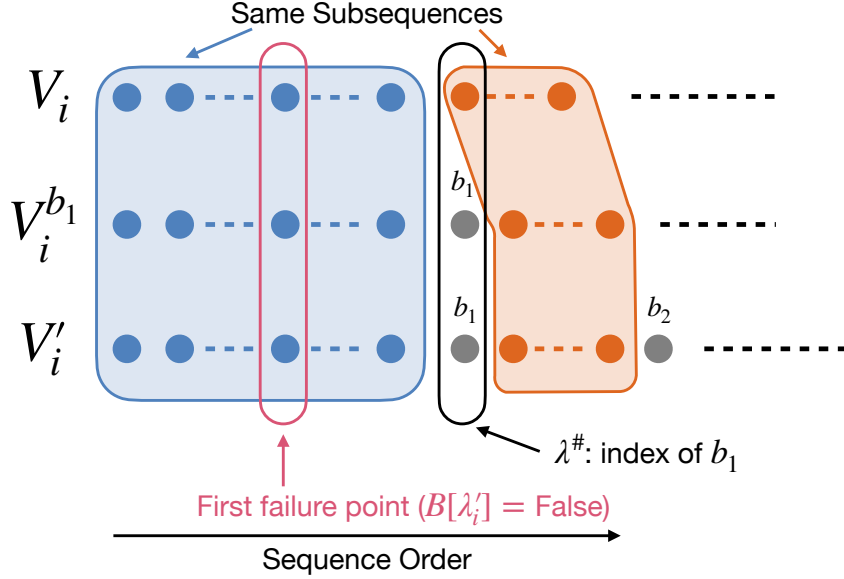


Figure 1: This figure depicts the relationship between V_i , $V_i^{b_1}$ and V_i' in the circumstance that $|S_{i-1}| < k$, $V_i \neq \emptyset$.

Therefore, $T_{\lambda_i^*} = T_{\lambda_i^*}^b = T_{\lambda_i^*}'$, and even further $T_{\lambda_i} = T_{\lambda_i}^b = T_{\lambda_i}'$. So, Property (i) and (ii) hold.

□

2.1.2 ANALYSIS OF GUARANTEES

There are three adjustments in THRESHSEQMOD compared with THRESHOLDSEQ in Chen et al. (2021) altogether. First, we made the randomness \mathbf{q} explicit allowing us to still consider each permutation on Line 8 as a random step. Therefore, the analysis of the theoretical guarantees is not influenced by adopting the randomness \mathbf{q} . Second, the prefix selection step is changed from identifying the failure point after the final success point to identifying the initial failure point. Note that, this change is necessary for the algorithm to satisfy RCP (Property1). However, by making this change, fewer elements could be filtered out at the next iteration. Fortunately, we are still able to filter out a constant fraction of the candidate set. Third, the algorithm returns one more set that provides extra information. Since the additional set does not affect the functioning of the algorithm, it does not influence the theoretical guarantees. In the following, we will provide the detailed analysis for the theoretical guarantees.

As demonstrated by Lemma 2 below, an increase in the number of selected elements results in a corresponding increase in the number of filtered elements. Consequently, there exists a point, say t , such that a given constant fraction of V_j can be filtered out if we are adding more than t elements. The proof of this lemma is quite similar to the proof of Lemma 12 in Chen et al. (2021) which can be found in Appendix C.

Lemma 2. *At an iteration j , let $A_i = \{x \in V_j : \Delta(x | T_i \cup S_{j-1}) < \tau\}$ after Line 5. It holds that $|A_0| = 0$, $|A_{|V_j|}| = |V_j|$ and $|A_i| \leq |A_{i+1}|$.*

Although we use a smaller prefix based on Line 16 and 18 compared to THRESHOLDSEQ, it is still possible that with probability at least $1/2$, a constant fraction of V_j will be filtered out at the beginning of iteration $j + 1$, or the algorithm terminates with $|S_j| = k$; given as Lemma 3 in the following. Intuitively, if there are enough such iterations, THRESHSEQMOD will succeed. Also, the calculation of query complexity comes from Lemma 3 directly. Since the analyses of success probability, adaptivity and query complexity simply follows the analysis in Chen et al. (2021), we provide these analyses in Appendix C. Next, we prove that Lemma 3 always holds.

Lemma 3. *At any iteration $j + 1$, with probability at least $1/2$, either $|V_{j+1}| \leq (1 - \beta\varepsilon)|V_j|$ after Line 5, or the algorithm terminates with $|S_j| = k$ on Line 7.*

Proof. By the definition of A_i in Lemma 2, $A_{\lambda_j^*}$ will be filtered out from V_j at the next iteration. Equivalently, $A_{\lambda_j^*} = V_j \setminus V_{j+1}$. After Line 8 at iteration $j + 1$, by Lemma 2, there exists a t such that $t = \min\{i \in \mathbb{N} : |A_i| \geq \beta\varepsilon|V_j|\}$. We say an iteration j succeeds, if $\lambda_j^* \geq \min\{s, t\}$. In this case, it holds that either $|V_j \setminus V_{j+1}| \geq \beta\varepsilon|V_j|$, or $|S_j| = k$. In other words, iteration j succeeds if there does not exist $\lambda \leq \min\{s, t\}$ such that $B[\lambda] = \mathbf{False}$ by the selection of λ_j^* on Line 18. Instead of directly calculating the success probability, we consider the failure probability. Define an element v_i *bad* if $\Delta(v_i | S_{j-1} \cup T_{i-1}) < \tau$, and *good* otherwise. Consider the random permutation of V_j as a sequence of dependent Bernoulli trials, with *success* if the element is bad and *failure* otherwise. When $i \leq \min\{s, t\}$, the probability of v_i is bad is less than $\beta\varepsilon$. If $B[\lambda] = \mathbf{False}$, there are at least $\varepsilon\lambda$ bad elements in T_λ . Let $\{Y_i\}_{i=1}^\infty$ be a sequence of independent and identically distributed Bernoulli trials, each with success probability $\beta\varepsilon$. Next, we consider the following sequence of dependent Bernoulli trials $\{X_i\}_{i=1}^\infty$, where $X_i = 1_{v_i \text{ is bad}}$ when $i \leq t$, and $X_i = Y_i$ otherwise. Define $B'[\lambda] = \mathbf{False}$ if there are at least $\varepsilon\lambda$ bad elements in $\{X_i\}_{i=1}^\lambda$. Then, for any $\lambda \leq \min\{s, t\}$, it holds that $Pr(B[\lambda] = \mathbf{False}) = Pr(B'[\lambda] = \mathbf{False})$. In the following, we bound the probability that there are more than ε -fraction 1s in $\{X_i\}_{i=1}^\lambda$ for any $\lambda \leq s$,

$$\begin{aligned}
 Pr(B'[\lambda] = \mathbf{False}) &\leq Pr\left(\sum_{i=1}^{\lambda} X_i \geq \varepsilon\lambda\right) \\
 &\leq Pr\left(\sum_{i=1}^{\lambda} Y_i \geq \varepsilon\lambda\right) && \text{(Lemma 9)} \\
 &\leq \min\left\{\beta, e^{-\frac{(1-\beta)^2}{1+\beta}\varepsilon\lambda}\right\} && \text{(Markov's Inequality, Lemma 10)} \\
 &\leq \min\left\{\beta, e^{-\varepsilon\lambda/2}\right\} && (\beta < 1/16)
 \end{aligned}$$

Subsequently, the probability of failure at iteration j is calculated below,

$$\begin{aligned}
 Pr(\text{iteration } j \text{ fails}) &= Pr(\exists \lambda \leq \min\{s, t\}, \text{ s.t.}, B[\lambda] = \mathbf{False}) \\
 &\leq Pr(\exists \lambda \leq s, \text{ s.t.}, B'[\lambda] = \mathbf{False})
 \end{aligned}$$

Algorithm 2 Low-Adaptive Greedy (LAG)

1: **Input:** Evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$, subset C , constraint k , fixed sequence of random bits $(\mathbf{q}_1, \mathbf{q}_2, \dots)$, accuracy parameter ε , constant α , value Γ such that $\Gamma \leq f(O) \leq \Gamma/\alpha$
 2: **if** $\Gamma = \emptyset$ or $\alpha = \emptyset$ **then**
 3: $\Gamma \leftarrow \max_{x \in \mathcal{N}} f(x)$
 4: $\alpha \leftarrow 1/k$
 5: Initialize $\tau \leftarrow \Gamma/\alpha k$, $\delta \leftarrow 1/(\log_{1-\varepsilon}(\alpha/3) + 1)$, $S \leftarrow \emptyset, R \leftarrow \emptyset, i \leftarrow 0$
 6: **for** $i \leftarrow 0$ to $\log_{1-\varepsilon}(\alpha/3)$ **do**
 7: $\tau \leftarrow \tau(1 - \varepsilon)^i$, $g(\cdot) \leftarrow f(S \cup \cdot)$
 8: $T_1, T_2 \leftarrow \text{THRESHSEQMOD}(g, C, k - |S|, \delta, \varepsilon/3, \tau, \mathbf{q}_i)$
 9: $S \leftarrow S \cup T_1$, $R \leftarrow R \cup T_2$
 10: **if** $|S| = k$ **then**
 11: **return** S, R
 12: **return** S, R

$$\begin{aligned}
 &\leq \sum_{\lambda \leq s} Pr(B'[\lambda] = \mathbf{False}) \\
 &\leq \sum_{\lambda=1}^{\infty} \min \left\{ \beta, e^{-\varepsilon\lambda/2} \right\} \\
 &\leq a\beta + \sum_{\lambda=a+1}^{\infty} e^{-\varepsilon\lambda/2}, \text{ where } a = \left\lfloor \frac{1}{4\beta} \right\rfloor = \left\lfloor \frac{2}{\varepsilon} \log \left(\frac{4}{1 - e^{-\varepsilon/2}} \right) \right\rfloor \\
 &\leq a\beta + \frac{e^{-\varepsilon(a+1)/2}}{1 - e^{-\varepsilon/2}} \leq \frac{1}{4} + \frac{1}{4} = \frac{1}{2}
 \end{aligned}$$

□

As for the Properties (3) and (4) in Theorem 1, since the structure of the solution set is the same as the solution of THRESHOLDSEQ. So, the analysis is similar to what in Chen et al. (2021). We provide these proofs in Appendix C.

2.2 Low-Adaptive Greedy (LAG) Algorithm

Another building block for our distributed algorithms is a simple, low-adaptive greedy algorithm LAG (Alg. 2). This algorithm is an instantiation of the PARALLELGREEDYBOOST framework of Chen et al. (2021), and it relies heavily on the low-adaptive procedure THRESHSEQMOD (Alg. 1).

Guarantees. Following the analysis of Theorem 3 in Chen et al. (2021), LAG achieves ratio $1 - 1/e - \varepsilon$ in $O\left(\frac{\log(\alpha^{-1}) \log(n \log(\alpha^{-1})/\varepsilon)}{\varepsilon^4}\right)$ adaptive rounds and $O\left(\frac{n \log(\alpha^{-1})}{\varepsilon^4}\right)$ total queries with probability at least $1 - 1/n$. Setting $\Gamma = \max_{x \in \mathcal{N}} f(\{x\})$ and $\alpha = 1/k$, LAG has adaptive complexity of $O\left(\frac{\log n \log k}{\varepsilon^4}\right)$ and a query complexity of $O\left(\frac{n \log k}{\varepsilon^4}\right)$.

Lemma 4. LAG satisfies the randomized consistency property (Property 1).

Proof of Lemma 4. Observe that the only randomness in LAG is from the calls to THRESHSEQMOD. Since $\text{LAG}(A, \mathbf{q})$ succeeds, every call to THRESHSEQMOD must succeed as well. Moreover, considering that \mathbf{q} is used to permute the underlying ground set \mathcal{N} , changing the set argument A of LAG does not change the sequence received by each call to THRESHSEQMOD. Order the calls to THRESHSEQMOD: $\text{THRESHSEQMOD}(\cdot, \mathbf{q}_1)$, $\text{THRESHSEQMOD}(\cdot, \mathbf{q}_2)$, \dots , $\text{THRESHSEQMOD}(\cdot, \mathbf{q}_m)$. Since $\text{LAGREL}(A \cup \{b\}, \mathbf{q}) = \text{LAGREL}(A, \mathbf{q})$, it holds that $b \notin \text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q}_i)$, for each $b \in B$ and each i . If $b \notin \text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q}_i)$ is equivalent to $\text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q}_i) = \text{THRESHSEQMODREL}(A, \mathbf{q}_i)$ (as shown in Claim 1 in the following), then by application of Lemma 1, $\text{THRESHSEQMOD}(A \cup B, \mathbf{q}_i)$ terminates successfully and $\text{THRESHSEQMODSOL}(A \cup B, \mathbf{q}_i) = \text{THRESHSEQMODSOL}(A, \mathbf{q}_i)$ for each i . Therefore, $\text{LAGSOL}(A \cup B, \mathbf{q}) = \text{LAGSOL}(A, \mathbf{q})$ and the former call terminates successfully. \square

Next, we provide Claim 1 and its analysis below.

Claim 1. *Let \mathbf{q} be a fixed sequence of random bits, $A \subseteq \mathcal{N}$, and $b \in \mathcal{N} \setminus A$. Then, $b \notin \text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q})$ if and only if $\text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q}) = \text{THRESHSEQMODREL}(A, \mathbf{q})$.*

Proof. It is obvious that if $\text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q}_i) = \text{THRESHSEQMODREL}(A, \mathbf{q}_i)$, then $b \notin \text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q}_i)$. In the following, we prove the reverse statement.

Following the analysis of RCP for THRESHSEQMOD in Section 2.1.1, we consider the runs of (1) $\text{THRESHSEQMOD}(A, \mathbf{q})$, and (2) $\text{THRESHSEQMOD}(A \cup \{b\}, \mathbf{q})$. Variables of (1) are given the notation defined in the pseudocode; variables of (2) are given the superscript b . We analyze the following statement $P(i)$ for each iteration i of the outer for loop on Line 4 in Alg. 1.

- (i) for all $b \in B$, $S_i = S_i^b$, $R_i = R_i^b$, and
- (ii) for all $b \in B$, $V_i = V_i^b \setminus \{b\}$.

Since $b \notin \text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q}) = R_{M+1}^b$, then $b \notin R_i^b$ for each i . Thus, the analysis in Section 2.1.1 also holds in this case. According to the inductive method, we are able to prove that $P(M+1)$ is true for each i which indicates $\text{THRESHSEQMODREL}(A \cup \{b\}, \mathbf{q}) = \text{THRESHSEQMODREL}(A, \mathbf{q})$. \square

3. Consistent Linear Time (Linear-TimeCardinality) Algorithm

In this section, we present the analysis of LINEAR-TIMECARDINALITY (Alg. 3, LTC), a consistent linear time algorithm which is an extension of the highly adaptive linear-time algorithm (Alg. 3) in Chen et al. (2021). Notably, to ensure consistency within a distributed setting and bound the solution size, LTC incorporates the randomness \mathbf{q} , and initializes the solution set with the maximum singleton. The algorithm facilitates the creation of linear-time MapReduce algorithms, enhancing overall computational efficiency beyond the capabilities of current state-of-the-art superlinear algorithms.

Algorithm 3 LINEAR-TIMECARDINALITY (LTC)

```

1: procedure LTC( $f, \mathcal{N}, k, \mathbf{q}$ )
2:   Input: evaluation oracle  $f$ , constraint  $k$ , fixed sequence of random bits  $\mathbf{q} \leftarrow (\sigma)$ .
3:    $\mathcal{N} \leftarrow$  random-permutation( $\mathcal{N}, \sigma$ )
4:   Initialize  $a \leftarrow$  first element in  $\arg \max_{x \in \mathcal{N}} f(x)$ ,  $S \leftarrow \{a\}$ 
5:   for  $x \in \mathcal{N}$  do
6:     if  $\Delta(x|S) \geq \frac{f(S)}{k}$  then  $S \leftarrow S \cup \{x\}$ 
7:   return  $S$ 
    
```

Theorem 2. *Let (f, k) be an instance of SM. The algorithm LINEAR-TIMECARDINALITY outputs $S \subseteq \mathcal{N}$ such that the following properties hold: 1) There are $O(n)$ oracle queries and $O(n)$ adaptive rounds. 2) Let S' be the last k elements in S . It holds that $f(S') \geq \frac{1}{2}f(S) \geq \frac{1}{4}f(O)$, where O is an optimal solution to the instance (f, k) . 3) The size of S is limited to $O(k \log(n))$.*

3.1 Analysis of Consistency

The highly adaptive linear-time algorithm (Alg. 3) outlined in Chen et al. (2021) commences with an empty set and adds elements to it iff. $\Delta(x|S) \geq f(S)/k$. Without introducing any randomness, this algorithm can be deterministic only if the order of the ground set is fixed. Additionally, to limit the solution size, we initialize the solution set with the maximum singleton, a choice that also impacts the algorithm's consistency. However, by selecting the first element that maximizes the objective value, the algorithm can maintain its consistency. In the following, we provide the analysis of randomized consistency property of LTC.

Lemma 5. *LTC satisfies the randomized consistency property (Property 1).*

Proof. Consider the runs of (1) LTC(A, \mathbf{q}), (2) LTC($A \cup \{b\}, \mathbf{q}$), and (3) LTC($A \cup B, \mathbf{q}$). With the same sequence of random bits \mathbf{q} , after the random permutation, A , $A \cup \{b\}$ and $A \cup B$ are the subsets of the same sequence. For any $x \in \mathcal{N}$, let i_x be the index of x . Then, \mathcal{N}_{i_x} are the elements before x (including x). Define S_{i_x} be the intermediate solution of (1) after we consider element x . If $x \notin A$, define $S_{i_x} = S_{i_x-1}$. Similarly, define S'_{i_x} and $S^b_{i_x}$ be the intermediate solution of (2) and (3). If $S_{i_x} = S^b_{i_x} = S'_{i_x}$, for any $x \in \mathcal{N}$ and $b \in B$, LTC($A \cup B, \mathbf{q}$) = LTC(A, \mathbf{q}). Next, we prove the above statement holds.

For any $b \in B$, since LTC(A, \mathbf{q}) = LTC($A \cup \{b\}, \mathbf{q}$), it holds that either $f(b) < \max_{x \in A} f(x)$, or $f(b) = \max_{x \in A} f(x)$, and $i_a < i_b$, where a is the first element in $\arg \max_{x \in A} f(x)$. Therefore, a is also the first element in $\arg \max_{x \in A \cup B} f(x)$. Furthermore, it holds that $S_0 = S_0^b = S'_0 = \{a\}$.

Suppose that $S_{i_x-1} = S^b_{i_x-1} = S'_{i_x-1}$. If $x \in A$ or $x \notin A \cup B$, naturally, $S_{i_x} = S^b_{i_x} = S'_{i_x}$. If $x \in B$, $S_{i_x} = S_{i_x-1}$ immediately. Since LTC(A, \mathbf{q}) = LTC($A \cup \{b\}, \mathbf{q}$), it holds that $\Delta(x|S^b_{i_x-1}) < f(S^b_{i_x-1})/k$ and $S^b_{i_x} = S^b_{i_x-1}$. So, $\Delta(x|S'_{i_x-1}) < f(S'_{i_x-1})/k$. Then, $S'_{i_x} = S'_{i_x-1}$. Thus, $S_{i_x} = S^b_{i_x} = S'_{i_x}$. \square

3.2 Analysis of Guarantees

Query Complexity and Adaptivity. Alg. 3 queries on Line 4 and 6, where there are n oracle queries on Line 4 and 1 oracle query for each element received on Line 6. Therefore, the query complexity and the adaptivity are both $O(n)$.

Solution Size. Given that f is a monotone function, it holds that $\arg \max_{S \in \mathcal{N}} f(S) = \mathcal{N}$. Furthermore, since $f(\{a\}) = \max_{x \in \mathcal{N}} f(\{x\})$, the submodularity property implies that $f(\{a\}) \geq \frac{1}{n} \sum_{x \in \mathcal{N}} f(\{x\}) \geq \frac{1}{n} f(\mathcal{N})$. Each element added to S contributes to an increase in the solution value by a factor of $(1 + \frac{1}{k})$. Therefore,

$$\begin{aligned} f(\mathcal{N}) &\geq f(S) \geq \left(1 + \frac{1}{k}\right)^{|S|} f(a) \geq \left(1 + \frac{1}{k}\right)^{|S|} \frac{f(\mathcal{N})}{n} \\ \Rightarrow |S| &\leq \log_{1+\frac{1}{k}}(n) \leq (k+1) \log(n) \qquad (\log(x) \geq 1 - \frac{1}{x}) \end{aligned}$$

Objective Value. For any $x \in \mathcal{N}$, let S_x be the intermediate solution before we process x . It holds that $\Delta(o | S_o) < f(S_o)/k$ for any $o \in O \setminus S$. Then,

$$f(O) - f(S) \stackrel{(1)}{\leq} \sum_{o \in O \setminus S} \Delta(o | S) \stackrel{(2)}{\leq} \sum_{o \in O \setminus S} \Delta(o | S_o) < \sum_{o \in O \setminus S} f(S_o)/k \leq f(S),$$

where Inequality (1) follows from monotonicity and submodularity, Inequality (2) follows from submodularity.

For any $x \in S'$, it holds that $\Delta(x | S_x) \geq f(S_x)/k$. Then,

$$\begin{aligned} f(S) - f(S \setminus S') &= \sum_{x \in S'} \Delta(x | S_x) \geq \sum_{x \in S'} f(S_x)/k \stackrel{(1)}{\geq} \sum_{x \in S'} f(S \setminus S')/k = f(S \setminus S') \\ \Rightarrow f(S') &\stackrel{(2)}{\geq} f(S) - f(S \setminus S') \geq \frac{1}{2} f(S) \end{aligned}$$

where Inequalities (1) and (2) follow from submodularity.

4. Low-Adaptive Algorithms with Constant MR Rounds (Greedy-DASH and Randomized-DASH)

Once we have the randomized consistency property of LAG, we can almost immediately use it to obtain parallelizable MapReduce algorithms.

4.1 Randomized-DASH

R-DASH (Alg. 4) is a two MR-rounds algorithm obtained by plugging LAG into the RANDGREEDI algorithm of Mirzasoleiman et al. (2013), Barbosa et al. (2015). R-DASH runs in two MapReduce rounds, $O(\log(k) \log(n))$ adaptive rounds, and guarantees the ratio $\frac{1}{2}(1 - 1/e - \varepsilon)$ ($\simeq 0.316$).

Description. The ground set is initially distributed at random by R-DASH across all machines M . In its first MR round, R-DASH runs LAG on every machine to obtain

Algorithm 4 RANDOMIZED-DASH (R-DASH)

1: **Input:** Evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, constraint k , error ε , available machines $M \leftarrow \{1, 2, \dots, \ell\}$
2: **for** $e \in \mathcal{N}$ **do**
3: Assign e to machine chosen uniformly at random
4: **for** $i \in M$ **do**
5: ▷ On machine i
6: Let \mathcal{N}_i be the elements assigned to machine i
7: $S_i, R_i \leftarrow \text{LAG}(f, \mathcal{N}_i, k, \varepsilon)$
8: Send S_i, R_i to primary machine
9: ▷ On primary machine
10: Gather $R \leftarrow \bigcup_{i=1}^{\ell} R_i$
11: $T \leftarrow \text{LAG}(f, R, k, \varepsilon)$ on machine m_1
12: **return** $V \leftarrow \arg \max \{f(T), f(S_1)\}$

S_i, R_i in $O(\log(k) \log(|\mathcal{N}_i|))$ adaptive rounds. The solutions from every machine are then returned to the primary machine, where LAG selects the output solution that guarantees $\frac{1}{2}(1 - 1/e - \varepsilon)$ approximation in $O(\log(k) \log(|R|))$ adaptive rounds as stated in Corollary 1. First, we provide Theorem 3 and its analysis by plugging any randomized algorithm ALG which follows RCP (Property 1) into the framework of RANDGREEDI. The proof is a minor modification of the proof of Barbosa et al. (2015) to incorporate randomized consistency. Then, we can get Corollary 1 for R-DASH immediately.

Theorem 3. *Let (f, k) be an instance of SM where $k = O(\psi/\ell)$. ALG is a randomized algorithm which satisfies the Randomized Consistency Property with α approximation ratio, $\Phi(n)$ query complexity, and $\Psi(n)$ adaptivity. By replacing LAG with ALG, R-DASH returns set V with two MR rounds, $2\Psi(n/\ell)$ adaptive rounds, $(\ell + 1)\Phi(n/\ell)$ total queries, $O(n)$ communication complexity such that*

$$\mathbb{E}[f(V)] \geq \frac{\alpha}{2} \text{OPT}.$$

Proof. Query Complexity and Adaptivity. R-DASH runs with two MR rounds. In the first MR round, ALG is invoked ℓ times in parallel, each with a ground set size of n/ℓ . So, during the first MR round, the number of queries is $\ell\Phi(n/\ell)$ and the number of adaptive rounds is $\Psi(n/\ell)$. Then, the second MR round calls ALG once, handling at most n/ℓ elements. Consequently, the total number of queries of R-DASH is $(\ell + 1)\Phi(n/\ell)$, and the total number of adaptive rounds is $2\Psi(n/\ell)$.

Approximation Ratio. Let R-DASH be run with input (f, k, ε, M) . Since $\ell \leq n^{1-c}$, $\mathbb{E}[|\mathcal{N}_i|] = n/\ell \geq n^c$. By an application of Chernoff’s bound to show that the size $|\mathcal{N}_i|$ is concentrated.

Let $\mathcal{N}(1/\ell)$ denote the random distribution over subsets of \mathcal{N} where each element is included independently with probability $1/\ell$. For $x \in \mathcal{N}$, let

$$p_x = \begin{cases} \Pr_{X \sim \mathcal{N}(1/\ell), \mathbf{q}} [\text{ALGREL}(X \cup \{x\}, \mathbf{q}) = \text{ALGREL}(X, \mathbf{q})], & \text{if } x \in O \\ 0 & , \text{ otherwise} \end{cases} .$$

Consider $S_1 = \text{ALGSOL}(\mathcal{N}_1, \mathbf{q})$ and $R_1 = \text{ALGREL}(\mathcal{N}_1, \mathbf{q})$ on machine m_1 . Let $O_1 = \{o \in O : \text{ALGREL}(\mathcal{N}_1 \cup \{o\}, \mathbf{q}) = \text{ALGREL}(\mathcal{N}_1, \mathbf{q})\}$. Since ALG follows the Randomized Consistency Property, $\text{ALGSOL}(\mathcal{N}_1 \cup O_1, \mathbf{q}) = \text{ALGSOL}(\mathcal{N}_1, \mathbf{q}) = S_1$. Therefore, $f(S_1) \geq \alpha f(O_1)$. Next, let $O_2 = O \cap R$, where $R = \bigcup_{i=1}^{\ell} R_i$. It holds that $f(T) \geq \alpha f(O_2)$. Let $o \in O$; o is assigned to \mathcal{N}_c on some machine m_c . It holds that,

$$\begin{aligned} \Pr[o \in O_2] &= \Pr[o \in \text{ALGREL}(\mathcal{N}_c) | o \in \mathcal{N}_c] \\ &= \Pr[o \in \text{ALGREL}(\mathcal{N}_c \cup \{o\})] \\ &= 1 - p_o. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E}[f(V)] &\geq \frac{1}{2} (\mathbb{E}[f(S_1)] + \mathbb{E}[f(T)]) \\ &\geq \frac{\alpha}{2} (\mathbb{E}[f(O_1)] + \mathbb{E}[f(O_2)]) \\ &\geq \frac{\alpha}{2} (F(\mathbf{p}) + F(\mathbf{1}_O - \mathbf{p})) \\ &\geq \frac{\alpha}{2} F(\mathbf{1}_O), \end{aligned} \tag{3}$$

where Inequality 3 follows from Lemma 8 and F is convex. \square

Corollary 1. *Let (f, k) be an instance of SM where $k = O(\psi/\ell)$. R-DASH returns set V with two MR rounds, $O\left(\frac{1}{\varepsilon^4} \log(k) \log(n)\right)$ adaptive rounds, $O\left(\frac{n \log(k)}{\varepsilon^4}\right)$ total queries, $O(n)$ communication complexity, and probability at least $1 - n^{1-2c}$ such that*

$$\mathbb{E}[f(V)] \geq \frac{1 - 1/e - \varepsilon}{2} \text{OPT}.$$

4.2 Greedy-DASH

Next, we obtain the nearly the optimal ratio by applying LAG and randomized consistency to DISTRIBUTEDDISTORTED proposed by Kazemi et al. (2021). DISTRIBUTEDDISTORTED is a distributed algorithm for regularized submodular maximization that relies upon (a distorted version of) the standard greedy algorithm. For SMCC, the distorted greedy algorithm reduces to the standard greedy algorithm GREEDY. In the following, we show that any algorithm that satisfies randomized consistency can be used in place of GREEDY as stated in Theorem 4. By introducing LAG into DISTRIBUTEDDISTORTED, G-DASH achieves the near optimal $(1 - 1/e - \varepsilon)$ expected ratio in $\lceil \frac{1}{\varepsilon} \rceil$ MapReduce rounds, $O(\log(n) \log(k))$ adaptive rounds, and $O(n \log(k))$ total queries. We also generalize this to any algorithm that satisfies randomized consistency.

Algorithm 5 GREEDY-DASH(G-DASH)

Input: evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, constraint k , error ε , available machines $M \leftarrow \{1, 2, \dots, \ell\}$
 2: $S \leftarrow \emptyset, C_0 \leftarrow \emptyset$
 for $r \leftarrow 1$ to $\lceil \frac{1}{\varepsilon} \rceil$ **do**
 4: $X_{r,i} \leftarrow$ Elements assigned to machine i chosen uniformly at random in round r
 $\mathcal{N}_{r,i} \leftarrow X_{r,i} \cup C_{r-1}$
 6: **for** $i \in M$ in parallel **do**
 $S_{r,i}, R_{r,i} \leftarrow \text{LAG}(f, \mathcal{N}_{r,i}, k, \varepsilon)$
 8: Send $S_{r,i}, R_{r,i}$ to each machine
 $S \leftarrow \arg \max \{f(S), f(S_{r,1}), \dots, f(S_{r,\ell})\}$
 10: $C_r \leftarrow \bigcup_{i=1}^{\ell} R_{r,i} \cup C_{r-1}$
return S

The Framework of Kazemi et al. (2021). DISTRIBUTEDDISTORTED has $\lceil 1/\varepsilon \rceil$ MR rounds where each round r works as follows: First, it distributes the ground set into m machines uniformly at random. Then, each machine i runs GREEDY (when modular term $\ell(\cdot) = 0$) on the data $\mathcal{N}_{r,i}$ that combines the elements distributed before each round $X_{r,i}$ and the elements forwarded from the previous rounds C_{r-1} to get the solution $S_{r,i}$. At the end, the final solution, which is the best among $S_{r,1}$ and all the previous solutions, is returned. To improve the adaptive rounds of DISTRIBUTEDDISTORTED, we replace standard greedy with LAG to get G-DASH.

Theorem 4. Let (f, k) be an instance of SM where $k = O(\varepsilon\psi/\ell)$. ALG is a randomized algorithm which satisfies the Randomized Consistency Property with α approximation ratio, $\Phi(n)$ query complexity, and $\Psi(n)$ adaptivity. By replacing LAG with ALG, G-DASH returns set S with $O(1/\varepsilon)$ MR rounds and $\frac{1}{\varepsilon}\Psi\left(\frac{n}{\ell}\right)$ adaptive rounds, $\frac{\ell}{\varepsilon}\Phi\left(\frac{n}{\ell}\right)$ total queries, $O\left(\frac{n}{\varepsilon}\right)$ communication complexity such that,

$$\mathbb{E}[f(S)] \geq (\alpha - \varepsilon)\text{OPT}.$$

Proof. Query Complexity and Adaptivity. G-DASH operates with $\lceil \frac{1}{\varepsilon} \rceil$ MR rounds, where each MR round calls ALG ℓ times in parallel. As each call of ALG works on the ground set with a size of at most $\frac{n}{\ell}$, the total number of queries for G-DASH is $\frac{\ell}{\varepsilon}\Phi\left(\frac{n}{\ell}\right)$, and the total number of adaptive rounds is $\frac{1}{\varepsilon}\Psi\left(\frac{n}{\ell}\right)$.

Approximation ratio. Let G-DASH be run with input (f, k, ε, M) . Similar to the analysis of Theorem 3, The size of \mathcal{N}_i is concentrated. Let O be the optimal solution. For any $x \in \mathcal{N}$, define that,

$$p_x^r = \begin{cases} Pr_{X \sim \mathcal{N}(1/\ell), \mathbf{q}} [x \notin C_{r-1} \text{ and} \\ x \in \text{ALGREL}(X \cup C_{r-1} \cup \{x\}, \mathbf{q})] & , \text{ if } x \in O \\ 0 & , \text{ otherwise} \end{cases}.$$

Then we provide the following lemma.

Lemma 6. For any $x \in O$ and $1 \leq r \leq 1/\varepsilon$, $Pr(x \in C_r) = \sum_{r'=1}^r p_x^{r'}$.

Proof.

$$\begin{aligned}
 Pr(x \in C_r) &= \sum_{r'=1}^r Pr(x \in C_{r'} \setminus C_{r'-1}) \\
 &= \sum_{r'=1}^r Pr\left(x \in \bigcup_{i=1}^{\ell} R_{r',i} \setminus C_{r'-1}\right) \\
 &= \sum_{r'=1}^r \sum_{i=1}^{\ell} \frac{1}{\ell} Pr(x \in R_{r',i} \setminus C_{r'-1} | x \in \mathcal{N}_{r',i}) \\
 &= \sum_{r'=1}^r \sum_{i=1}^{\ell} \frac{1}{\ell} Pr(x \notin C_{r'-1} \text{ and} \\
 &\quad x \in \text{ALGREL}(X_{r',i} \cup C_{r'-1} \cup \{x\}, \mathbf{q})) \\
 &= \sum_{r'=1}^r p_x^{r'}
 \end{aligned}$$

□

The rest of the proof bounds $f(S_{r,1})$ in the following two ways.

First, let $O_{r,1} = \{o \in O : o \notin \text{ALGREL}(X_{r,1} \cup C_{r-1} \cup \{o\}, \mathbf{q})\}$, and $O_{r,2} = (C_{r-1} \cap O) \cup O_{r,1}$. Since ALG follows the Randomized Consistency Property, it holds that $\text{ALGSOL}(X_{r,1} \cup C_{r-1}, \mathbf{q}) = \text{ALGSOL}(X_{r,1} \cup C_{r-1} \cup O_{r,1}, \mathbf{q}) = S_{r,1}$, and $f(S_{r,1}) \geq \alpha f(O_{r,2})$. So, for any $o \in O$,

$$Pr(o \in O_{r,2}) = Pr(o \in C_{r-1} \text{ or } o \notin \text{ALGREL}(X_{r,1} \cup C_{r-1} \cup \{o\}, \mathbf{q})) = 1 - p_o^r.$$

Therefore,

$$\mathbb{E}[f(S_{r,1})] \geq \alpha \mathbb{E}[f(O_{r,2})] \geq \alpha F(\mathbf{1}_O - \mathbf{p}^r). \tag{4}$$

Second, let $O_{r,3} = C_{r-1} \cap O$. Similarly, it holds that $f(S_{r,1}) \geq \alpha f(O_{r,3})$. And for any $o \in O$, by Lemma 6, it holds that,

$$Pr(o \in O_{r,3}) = Pr(o \in C_{r-1}) = \sum_{r'=1}^{r-1} p_o^{r'},$$

Therefore,

$$\mathbb{E}[f(S_{r,1})] \geq \alpha \mathbb{E}[f(O_{r,3})] \geq \alpha F\left(\sum_{r'=1}^{r-1} \mathbf{p}^{r'}\right). \tag{5}$$

By Inequalities 4 and 5, we bound the approximation ratio of G-DASH by the following,

$$\begin{aligned}
 \mathbb{E}[f(S)] &\geq \varepsilon \sum_{r=1}^{1/\varepsilon} \mathbb{E}[f(S_{r,1})] \\
 &\geq \varepsilon \alpha \cdot \left(F \left(\sum_{r'=1}^{1/\varepsilon-1} \mathbf{p}^{r'} \right) + \sum_{r=1}^{1/\varepsilon-1} F \left(\mathbf{1}_O - \mathbf{p}^{r'} \right) \right) \\
 &\stackrel{(a)}{\geq} \alpha(1 - \varepsilon)F(\mathbf{1}_O) \\
 &\geq (\alpha - \varepsilon)\text{OPT},
 \end{aligned}$$

where Inequality (a) follows from Lemma 8 and F is convex. □

Corollary 2. *Let (f, k) be an instance of SM where $k = O(\varepsilon\psi/\ell)$. G-DASH returns set S with $O(1/\varepsilon)$ MR round and $O\left(\frac{1}{\varepsilon^5} \log(k) \log(n)\right)$ adaptive rounds, $O\left(\frac{n \log(k)}{\varepsilon^5}\right)$ total queries, $O\left(\frac{n}{\varepsilon}\right)$ communication complexity, and probability at least $1 - \varepsilon^{-1}n^{1-2c}$ such that,*

$$\mathbb{E}[f(S)] \geq (1 - 1/e - \varepsilon)\text{OPT}.$$

5. Threshold-DASH: Two MR-Rounds Algorithm with Improved Theoretical Properties

In this section, we present a novel $(0.375 - \varepsilon)$ -approximate, two MR-rounds algorithm THRESHOLD-DASH (T-DASH) that achieves nearly optimal $O(\log(n))$ adaptivity in $O(n \log(k))$ total time complexity.

Description. The T-DASH algorithm (Alg. 6) is a two MR-rounds algorithm that runs THRESHSEQMOD concurrently on every machine with a specified threshold value of $\alpha\text{OPT}/k$. The primary machine then builds up its solution S_1 by adding elements with THRESHSEQMOD from the pool of solutions returned by the other machines. Notice that there is a small amount of data duplication as elements of the ground set are not randomly partitioned in the same way as in the other algorithms. This version of the algorithm requires to know the OPT value; in Appendix E we show how to remove this assumption. In Appendix I, we further discuss the similar two MR round algorithm of Liu and Vondrak (2019), that provides an improved 1/2-approximation but requires four times the data duplication of T-DASH.

Theorem 5. *Let (f, k) be an instance of SM where $k = O\left(\frac{\psi}{\ell}\right)$. T-DASH knowing OPT returns set T' with two MR rounds, $O\left(\frac{1}{\varepsilon^3} \log(n)\right)$ adaptive rounds, $O\left(\frac{n}{\varepsilon^3}\right)$ total queries, $O(n)$ communication complexity, and probability at least $1 - n^{-c}$ such that*

$$\mathbb{E}[f(T')] \geq \left(\frac{3}{8} - \varepsilon\right)\text{OPT}.$$

Algorithm 6 THRESHOLD-DASH Knowing OPT

- 1: **Input:** Evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, constraint k , error ε , available machines $M \leftarrow \{1, 2, \dots, \ell\}$, and OPT
 - 2: Initialize $\delta \leftarrow 1/(\ell + 1)$, $\mathbf{q} \leftarrow$ a fixed sequence of random bits.
 - 3: Set $\alpha \leftarrow \frac{3}{8}$, $(\mathbf{q}_1, \dots, \mathbf{q}_{\ell+1}) \leftarrow \mathbf{q}$
 - 4: **for** $e \in \mathcal{N}$ **do** **do**
 - 5: Assign e to each machine independently with probability $1/\ell$
 - 6: **for** $i \in M$ **do**
 - 7: ▷ On machine i
 - 8: Let \mathcal{N}_i be the elements assigned to machine i
 - 9: $S_i, R_i \leftarrow \text{THRESHSEQMOD}(f, \mathcal{N}_i, k, \delta, \varepsilon, (\alpha + \varepsilon)\text{OPT}/k, \mathbf{q}_i)$
 - 10: **if** $|S_i| = k$ **then** return $T' \leftarrow S_i$
 - 11: Send S_i, R_i to primary machine
 - 12: ▷ On primary machine
 - 13: $R \leftarrow \bigcup_{i=1}^{\ell} R_i$
 - 14: Let $g(\cdot) \leftarrow f(S_1 \cup \cdot) - f(S_1)$
 - 15: $T \leftarrow \text{THRESHSEQMOD}(g, R, k - |S_1|, \delta, \varepsilon, (\alpha + \varepsilon)\text{OPT}/k, \mathbf{q}_{\ell+1})$
 - 16: $T' \leftarrow S_1 \cup T$
 - 17: **return** T'
-

Proof. Query Complexity and Adaptivity. T-DASH runs with two MR rounds, where the first MR round invokes THRESHSEQMOD ℓ times in parallel, and the second MR round invokes THRESHSEQMOD once. By Theorem 1, each call of THRESHSEQMOD with a ground set of size at most $\frac{n}{\ell}$ queries $O\left(\frac{n}{\ell \varepsilon^3}\right)$ within $O\left(\frac{\log(n/\ell)}{\varepsilon^3}\right)$ adaptive rounds. The total number of queries is $O\left(\frac{n}{\varepsilon^3}\right)$, and the total number of adaptive rounds is $O\left(\frac{\log(n)}{\varepsilon^3}\right)$.

Approximation Ratio. In Algorithm 6, there are $\ell + 1$ independent calls of THRESHSEQMOD. With $|\mathcal{N}_i| \geq n^c$, the success probability of each call of THRESHSEQMOD is larger than $1 - \frac{1}{n^{c(\ell+1)}}$. Thus, Algorithm 6 succeeds with probability larger than $1 - n^{-c}$. For the remainder of the analysis, we condition on the event that all calls to THRESHSEQMOD succeed.

In the case that $|T'| = k$, by Theorem 1 in Section 2.1 and $\tau = \left(\frac{3}{8} + \varepsilon\right) \frac{\text{OPT}}{k}$, it holds that $f(T') \geq \frac{1-2\varepsilon}{1+\varepsilon} \tau \cdot k \geq \left(\frac{3}{8} - \varepsilon\right) \text{OPT}$. Otherwise, we consider the case that $|T'| < k$ in the following. Let $(\text{TSMSOL}(\mathcal{N}, \mathbf{q}), \text{TSMREL}(\mathcal{N}, \mathbf{q}))$ be the pair of sets returned by $\text{THRESHSEQMOD}(\mathcal{N}, \mathbf{q})$. For any $x \in \mathcal{N}$, let

$$p_x = \begin{cases} Pr_{X \sim \mathcal{N}(1/\ell), \mathbf{q}} [\text{TSMREL}(X \cup \{x\}, \mathbf{q}) \\ \quad = \text{TSMREL}(X, \mathbf{q})] & , \text{ if } x \in O \\ 0 & , \text{ otherwise} \end{cases} .$$

Let $O_1 = \{o \in O : o \notin \text{TSMREL}(N_1 \cup \{o\}, q)\}$, $O_2 = R \cap O$. By Randomized Consistency Property, it holds that $S_1 = \text{TSMSOL}(N_1, q) = \text{TSMSOL}(N_1 \cup O_1, q)$. For any $o \in O_1$, o is not selected in S_1 . Since, $|S_1| < k$, by Property (4) in Theorem 1, $\Delta(o | T') < \Delta(o | S_1) < \tau$.

Algorithm 7 LINEAR-TIME-DISTRIBUTED

- 1: **Input:** Evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, constraint k , error ε , available machines $M \leftarrow \{1, 2, \dots, \ell\}$
 - 2: **for** $e \in \mathcal{N}$ **do**
 - 3: Assign e to machine chosen uniformly at random
 - 4: **for** $i \in M$ **do**
 - 5: ▷ On machine i
 - 6: Let \mathcal{N}_i be the elements assigned to machine i
 - 7: $S_i \leftarrow \text{LTC}(f, \mathcal{N}_i, k, \mathbf{q})$
 - 8: Send S_i to primary machine
 - 9: ▷ On primary machine
 - 10: Gather $S \leftarrow \bigcup_{i=1}^{\ell} S_i$
 - 11: $T_1 \leftarrow \text{LTC}(f, S, k, \mathbf{q})$ on machine m_1
 - 12: $T'_1 \leftarrow$ last k elements in T_1
 - 13: $T_2 \leftarrow \text{THRESHOLDGREEDY}(f, T_1, k, \varepsilon, f(T'_1), 1/2)$ ▷ Post-Processing
 - 14: $S'_1 \leftarrow$ last k elements added to S_1
 - 15: **return** $V \leftarrow \arg \max \{f(S'_1), f(T'_1), f(T_2)\}$
-

Also, for any $o \in O_2 \setminus T$, o is not selected in T . Similarly, $\Delta(o|T') < \tau$. Then, we can get,

$$\begin{aligned}
 f(O_1 \cup O_2) - f(T') &\leq f(O_1 \cup O_2 \cup T') - f(T') \\
 &\leq \sum_{o \in O_1 \cup O_2 \setminus T'} \Delta(o|T') \leq k \cdot \tau = \left(\frac{3}{8} + \varepsilon\right) \text{OPT}. \tag{6}
 \end{aligned}$$

Next, we provide the following lemma to complete the rest of the proof.

Lemma 7. *For any $o \in O$, it holds that $\Pr(o \in O_1 \cup O_2) \geq 3/4$.*

Then, by Lemma 8 in Appendix A, $\mathbb{E}[f(O_1 \cup O_2)] \geq 3/4 \cdot F(\mathbf{1}_O) = 3/4 \cdot \text{OPT}$. From this and Inequality 6, we can bound the approximation ratio for Algorithm T-DASH knowing OPT by follows, $\mathbb{E}[f(T')] \geq (\frac{3}{8} - \varepsilon) \text{OPT}$. □

6. LinearTime-Distributed (L-Dist)

In this section, we introduce L-DIST, the first linear time MapReduce algorithm with constant approximation. We derive L-DIST from R-DASH by incorporating the analysis of LTC (Alg. 3), an algorithm with $O(n)$ query complexity, into the distributed setting. This integration allows L-DIST to operate within two MapReduce rounds, demonstrating an adaptive complexity of $O(\frac{n}{\ell})$ and a query complexity of $O(n)$. Also, as stated in Theorem 6, it guarantees an approximation ratio of $\frac{1}{8}$. Additionally, we improve the objective value by implementing THRESHOLDGREEDY (Alg. 9 in Appendix D) on the set returned at the second MapReduce round. The integration of LTC into LINEAR-TIME-DISTRIBUTED provides enhanced capabilities and improves the efficiency of the overall algorithm.

Description: Initially, L-DIST involves randomly distributing the ground set across all machines M . In the first MR round, LINEAR-TIME-DISTRIBUTED applies LTC on each machine to obtain S_i within $O(n/\ell)$ query calls. The solutions from all machines are then returned to the primary machine, where LTC selects the output solution T_1 . By selecting the best last k among all solutions, L-DIST ensures a $\frac{1}{8}$ -approximation. Besides returning the last k elements of the solution T_1 which is an $1/2$ -approximation of T_1 , we employ THRESHOLDGREEDY, a $(1 - 1/e - \varepsilon)$ approximation algorithm, to boost the objective value. We provide the theoretical guarantees of L-DIST as Theorem 6. The proof involves a minor modification of the proof provided in Theorem 3.

Theorem 6. *Let (f, k) be an instance of SM where $k < \frac{\psi}{\ell \log(\psi)}$. L-DIST returns a set V with two MR rounds, $O(\frac{n}{\ell})$ adaptive rounds, $O(n)$ total queries, $O(n)$ communication complexity such that*

$$\mathbb{E}[f(V)] \geq \frac{1}{8} \text{OPT}.$$

6.1 Analysis of Query Complexity and Adaptivity

L-DIST runs with two MR rounds. In the first MR round, LTC is invoked ℓ times in parallel, each with $O(n/\ell)$ queries and $O(n/\ell)$ adaptive rounds by Theorem 2. So, during the first MR round, the number of queries is $O(n)$ and the number of adaptive rounds is $O(n/\ell)$. Then, the second MR round calls LTC and THRESHOLDGREEDY once respectively, handling at most n/ℓ elements. By Theorem 2 and 8, the number of queries is $O(n/\ell)$ and the number of adaptive rounds is $O(n/\ell)$. Consequently, the total number of queries of L-DIST is $O(n)$, and the total number of adaptive rounds is $O(n/\ell)$.

6.2 Analysis of Approximation Ratio

Let L-DIST be executed with input (f, k, ε, M) . Since $\ell \leq n^{1-c}$, it follows that the expected size of each subset $|\mathcal{N}_i|$ satisfies $\mathbb{E}[|\mathcal{N}_i|] = n/\ell \geq n^c$. To ensure that the size $|\mathcal{N}_i|$ is concentrated, we apply Chernoff's bound. Let $\mathcal{N}(1/\ell)$ denote the random distribution over subsets of \mathcal{N} where each element is included independently with probability $1/\ell$. Let $\mathbf{p} \in [0, 1]^n$ be the following vector. For $x \in \mathcal{N}$, let

$$p_x = \begin{cases} \Pr_{X \sim \mathcal{N}(1/\ell), \mathbf{q}} [\text{LTC}(X \cup \{x\}, \mathbf{q}) = \text{LTC}(X, \mathbf{q})], & \text{if } x \in O \\ 0 & \text{, otherwise} \end{cases}.$$

Consider $S_1 = \text{LTC}(\mathcal{N}_1, \mathbf{q})$ on machine m_1 . Let $O_1 = \{o \in O : \text{LTC}(\mathcal{N}_1 \cup \{o\}, \mathbf{q}) = \text{LTC}(\mathcal{N}_1, \mathbf{q})\}$. By Lemma 5, $\text{LTC}(\mathcal{N}_1 \cup O_1, \mathbf{q}) = S_1$. Therefore, by Theorem 2, it holds that $f(S'_1) \geq f(O_1)/4$. Next, let $O_2 = O \cap S$, where $S = \bigcup_{i=1}^{\ell} S_i$. Similarly, by Theorem 2, it holds that $f(T'_1) \geq f(O_2)/4$. Let $o \in O$; o is assigned to \mathcal{N}_c on some machine m_c . It holds that,

$$\begin{aligned} \Pr[o \in O_2] &= \Pr[o \in \text{LTC}(\mathcal{N}_c) | o \in \mathcal{N}_c] \\ &= \Pr[o \in \text{LTC}(\mathcal{N}_c \cup \{o\})] \\ &= 1 - p_o. \end{aligned}$$

Algorithm 8 MED($f, k, M, \text{ALG}, \mathbf{q}$)

- 1: **Input:** evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, constraint k , available machines $M \leftarrow \{1, 2, \dots, \ell\}$, MR algorithm ALG, random bits \mathbf{q}
 - 2: $n \leftarrow |\mathcal{N}|$, $\Psi \leftarrow$ Memory capacity (# elements) of primary machine
 - 3: Choose $k' \leftarrow \max\{k' \in \mathbb{N} : k' \leq \frac{\Psi}{\ell}\}$
 - 4: $m \leftarrow \lceil k/k' \rceil$, $(\mathbf{q}_1, \dots, \mathbf{q}_m) \leftarrow \mathbf{q}$
 - 5: **for** $i \leftarrow 1$ **to** m **do**
 - 6: Let $g(\cdot) \leftarrow f(\cdot \cup S_i) - f(S_i)$
 - 7: $A_i \leftarrow \text{ALG}(g, M, \min\{k', k - |S_i|\}, \mathbf{q}_i)$
 - 8: $S_{i+1} \leftarrow S_i \cup A_i$
 - 9: **return** S_m
-

Therefore,

$$\begin{aligned}
 \mathbb{E}[f(V)] &\geq \frac{1}{2} (\mathbb{E}[f(S'_1)] + \mathbb{E}[f(T'_1)]) \\
 &\geq \frac{1}{2} (\mathbb{E}[f(O_1)/4] + \mathbb{E}[f(O_2)/4]) \\
 &\geq \frac{1}{8} (F(\mathbf{p}) + F(\mathbf{1}_O - \mathbf{p})) \\
 &\geq \frac{1}{8} F(\mathbf{1}_O) = \frac{1}{8} \text{OPT},
 \end{aligned} \tag{7}$$

where inequality 7 follows from Lemma 8 and F is convex.

6.3 Post-Processing

In this algorithm, we employ a simple post-processing procedure. Since T'_1 is an $1/2$ -approximation of T_1 with size k , T'_1 is also an $1/2$ -approximation of the instance (f, k) on the ground set T_1 . By running any linear time algorithm that has a better approximation ratio on T_1 , we are able to boost the objective value returned by the algorithm with the same theoretical guarantees. By Theorem 8, THRESHOLDGREEDY achieves $(1 - 1/e - \varepsilon)$ -approximation in linear time with a guess on the optimal solution. Therefore, with input $\alpha = 1/2$ and $\Gamma = f(T'_1)$, utilizing THRESHOLDGREEDY can effectively enhance the objective value.

7. Towards Larger k : A Memory-Efficient, Distributed Framework (MED)

In this section, we propose a general-purpose plug-in framework for distributed algorithms, MEMORYEFFICIENTDISTRIBUTED (MED, Alg. 8). MED increases the largest possible constraint value from $O(n/\ell^2)$ to $O(n/\ell)$ in the value oracle model. Under some additional assumptions, we remove all restrictions on the constraint value.

As discussed in Section 1, the k value is limited to a fraction of the machine memory for MapReduce algorithms: $k \leq O(n/\ell^2)$, since those algorithms need to merge a group of solutions and pass it to a single machine. MED works around this limitation as follows: MED can be thought of as a greedy algorithm that uses an approximate greedy selection through ALG. One machine manages the partial solutions $\{S_i\}$, built up over m iterations.

In this way, each call to ALG is within the constraint restriction of ALG, *i.e.* $O(n/\ell^2)$, but a larger solution of up to size $O(n/\ell)$ can be constructed.

The restriction on k of MED of $O(n/\ell)$ comes from passing the data of the current solution to the next round. Intuitively, if we can send some auxiliary information about the function instead of the elements selected, the k value can be unrestricted.

Assumption 1. *Let f be a set function with ground set \mathcal{N} of size n . If for all $S \subseteq \mathcal{N}$, there exists a bitvector \mathbf{v}_S , such that the function $g(X) = f(S \cup X) - f(S)$ can be computed from X and \mathbf{v}_S , then f satisfies Assumption 1.*

We show in Appendix F that all four applications evaluated in Section 8 satisfy Assumption 1. As an example, consider MaxCover, which can be expressed as $f(S) = \sum_{i \in \mathcal{N}} f_i(S)$, where $f_i(S) = \mathbf{1}_{\{i \text{ is covered by } S\}}$. Let $g_i(X) = f_i(S \cup X) - f_i(S)$, and $\mathbf{v}_S = (f_1(S), \dots, f_n(S))$. Then, $f_i(S \cup X) = \mathbf{1}_{\{i \text{ is covered by } S\}} \vee \mathbf{1}_{\{i \text{ is covered by } X\}}$, where the first term is given by \mathbf{v}_S and the second term is calculated by X . Therefore, since $g(X) = \sum_{i \in \mathcal{N}} g_i(X)$, $g(X)$ can be computed from X and \mathbf{v}_S .

Theorem 7. *Let (f, k) be an instance of SMCC distributed over ℓ machines. For generic objectives, where the data of current solution need to be passed to the next round, $k \leq \min\{\frac{\ell\psi-n}{\ell-1} + 1, \psi - \ell + 1\}$; for special objectives, where only one piece of compressed data need to be passed to the next round, $k \leq n$. Let S_m be the set returned by MED. Then $\mathbb{E}[f(S_m)] \geq (1 - e^{-\gamma})\text{OPT}$, where γ is the expected approximation of ALG.*

Proof. Let O be an optimal and O_1, O_2, \dots, O_m be a partition of O into m pieces, each of size $\leq k'$. Also $\forall_i \hat{f}_i$ is monotone SMCC since f is monotone SMCC.

$$\begin{aligned} \mathbb{E}[f(S_{i+1}) - f(S_i)|S_i] &= \mathbb{E}[f(S_i \cup A_i) - f(S_i)|S_i] \\ &= \mathbb{E}\left[\hat{f}_i(A_i)|S_i\right] \\ &\stackrel{(a)}{\geq} \frac{1}{m} \sum_{j=1}^m \gamma \hat{f}_i(O_j) \\ &\geq \frac{\gamma}{m} \hat{f}_i(O) \\ &= \frac{\gamma}{m} (f(S_i \cup O) - f(S_i)) \\ &\geq \frac{\gamma}{m} (\text{OPT} - f(S_i)), \end{aligned}$$

where Inequality (a) follows from \hat{f}_i is SMCC; ALG is γ -approximation. Unfix S_i , it holds that

$$\begin{aligned} \text{OPT} - \mathbb{E}[f(S_{i+1})] &\leq \left(1 - \frac{\gamma}{m}\right) [\text{OPT} - \mathbb{E}[f(S_i)]] \\ &\leq \left(1 - \frac{\gamma}{m}\right)^i [\text{OPT} - f(\emptyset)] \\ &= \left(1 - \frac{\gamma}{m}\right)^i \cdot \text{OPT} \end{aligned}$$

Table 2: Environment Setup

Experiment	Number of Machines (ℓ)*	Number of Threads Per Machine (\mathcal{T})	Dataset Size (n)
Experiment 1 (Fig. 2)	8	4	10K - 100K
Experiment 2 (Fig. 3)	64	32	1M - 5M
Experiment 3 (Fig. 4(a))	8	4	3M
Experiment 3 (Fig. 4(b))	8	4	100K
Experiment 4 (Fig. 5)	32	1	100K

* Prior MR (greedy) algorithms were assessed using $\ell \cdot \mathcal{T}$ machines, utilizing each thread as an individual core. Thus, all algorithms were configured to use all available cores.

‡ We kept a variant of R-DASH evaluated on 8 machines (with 4 cores each) as a baseline. All other algorithms use 2 machines (with 4 cores each).

$$\begin{aligned} \therefore \mathbb{E}[f(S_m)] &\geq \left[1 - \left(1 - \frac{\gamma}{m}\right)^m\right] \cdot OPT \\ &\geq (1 - e^{-\gamma}) \cdot OPT \end{aligned}$$

To analyze the memory requirements of MED, consider the following. To compute $g(\cdot) \leftarrow f(\cdot \cup S_i) - f(S_i)$, the current solution S_i need to be passed to each machine in M for the next call of ALG. Suppose $|S| = k - x$ where $1 \leq x \leq k$. The size of data stored on any non-primary machine in cluster M , as well as on the primary machine of M can be bounded as follows:

$$\begin{aligned} k - x + (n - k + x)/\ell \leq \Psi &\Rightarrow k \leq (\ell\Psi - n)/(\ell - 1) + 1 \\ k - x + (k\ell)/m \leq \Psi &\Rightarrow k \leq \Psi - \ell + 1 \end{aligned}$$

Therefore, if $\Psi \geq 2n/\ell$, MED can run since $\ell \leq n/\ell$ in our MapReduce model. Under the alternative assumption, it is clear that MED can run for all $k \leq n$. \square

8. Empirical Evaluation

This section presents empirical comparison of R-DASH, T-DASH, G-DASH, L-DIST to the state-of-the-art distributed algorithms. Distributed algorithms include RANDGREEDI (RG) of Barbosa et al. (2015) and DISTRIBUTEDDISTORTED (DDIST) of Kazemi et al. (2021).

Environment. Our experiments utilized diverse computational setups. Experiments 1, and 3 employed an eight-machine cluster, each with four CPU cores, totaling 32 cores. Experiment 2 was conducted on a larger 64-machine cluster, each featuring 32 CPU cores. Experiment 4 operated on a cluster of 32 single-core machines. Notably, in Experiment 1, our algorithms used $\ell = 8$ machines, while prior MapReduce (MR) algorithms employed $\ell = 32$, fully utilizing 32 cores. MPICH version 3.3a2 was installed on every machine, and we used the python library *mpi4py* for implementing and parallelizing all algorithms with the Message Passing Interface (MPI). These algorithms were executed using the *mpirun* command, with runtime tracked using *mpi4py.MPI.Wtime()* at the algorithm’s start and completion.

Datasets. Table 2 presents dataset sizes for our experiments. In Experiment 1 and 3 (Fig. 4(b)), we assess algorithm performance on small datasets, varying from $n = 10,000$ to 100,000. These sizes enable evaluation of computationally intensive algorithms, such as DDIST, T-DASH, and G-DASH. Experiment 2, 3 (Fig. 4(a)), and 4 focus on larger datasets, ranging from $n = 50,000$ to 5 million.

Applications. We evaluated image summarization (ImageSumm), influence maximization (InfMax), revenue maximization (RevMax) and maximum coverage (MaxCover). Details are provided in Appendix F.1.

Experiment Objectives. The primary objective of our experiment set is to comprehensively evaluate the practicality of the distributed algorithms across varying cluster and dataset sizes. The specific objectives of each experiment are outlined below:

- Experiment 1¹: Baseline experiment aimed at assessing the performance of the algorithms using small datasets and cluster setup.
- Experiment 2: Assess the performance of the algorithms on large datasets and cluster setup.
- Experiment 3*: Investigate the influence of the number of nodes in a cluster on algorithm performance.
- Experiment 4: Examine the impact of increasing cardinality constraints on the performance of MED.

8.1 Experiment 1 - Comparative Analysis on Small Datasets

The results of Experiment 1 (Fig. 2) show that all algorithms provide similar solution values (with T-DASH being a little worse than the others). However, there is a large difference in runtime, with R-DASH the fastest by orders of magnitude. The availability of only 4 threads per machine severely limits the parallelization of T-DASH, resulting in longer runtime; access to $\log_{1+\epsilon}(k)$ threads per machine should result in faster runtime than R-DASH.

8.2 Experiment 2 - Performance Analysis of R-DASH, RandGreeDI and L-Dist on a Large Cluster

This section presents a comparative analysis of R-DASH, L-DIST, and RANDGREEDI on a large 64-node cluster, each node equipped with 32 cores. We assess two versions of L-DIST and RANDGREEDI, one with $\ell = 64$ and the other with $\ell = 2048$, with a predetermined time limit of 3 hours for each application. The plotted results depict the instances completed by each algorithm within this time constraint.

In terms of solution quality, as depicted in Figure 3(a), statistically indistinguishable results, while L-DIST exhibits an average drop of 9% in solution quality across completed instances. Examining runtime (Figure 3(d)), R-DASH consistently outperforms L-DIST and RANDGREEDI by orders of magnitude. As discussed in Experiment 3, for both L-DIST and RANDGREEDI, instances with fewer distributions ($\ell = 64$) display faster runtimes compared

¹Results published in Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023

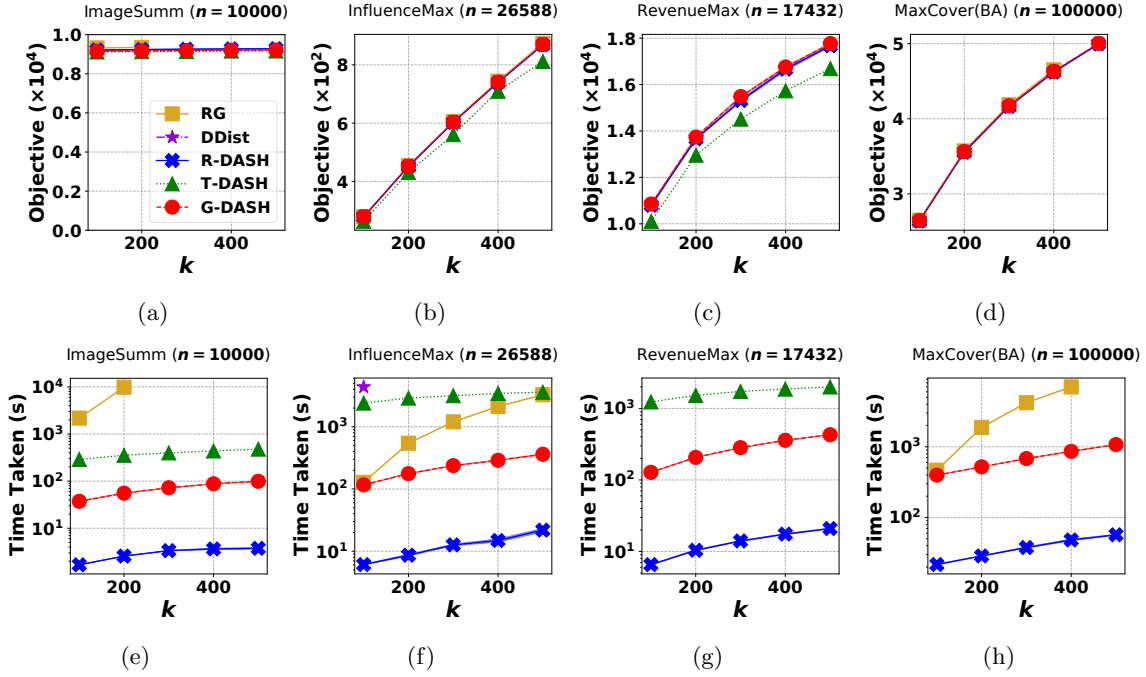


Figure 2: Performance comparison of distributed algorithms on ImageSumm, InfluenceMax, RevenueMax and MaxCover; RANDGREEDI (RG) is run with GREEDY as the algorithm ALG (Barbosa et al., 2015) to ensure the $\frac{1}{2}(1 - 1/e)$ ratio. All GREEDY implementations used lazy greedy to improve the runtime. *Timeout* for each application: 6 hours per algorithm.

to instances using a larger setup ($\ell = 2048$), attributed to k values surpassing $\frac{n}{\ell^2}$ ($\simeq 0.27$ for InfMax). The subsequent experiment demonstrates how MED effectively addresses this challenge. Comparing the fastest variant of each algorithm, we observe that for the smallest k values, RANDGREEDI outperforms L-DIST. However, as k increases, RANDGREEDI’s runtime linearly grows, while L-DIST’s runtime remains unaffected due to its linear time complexity, enabling L-DIST to outperform RANDGREEDI for larger constraints.

8.3 Experiment 3 - Scalability Assessment

Figure 4(a) illustrates a linear speedup for R-DASH as the number of machines ℓ increases. Figure 4(b) highlights an intriguing observation that, despite having sufficient available memory, increasing ℓ can result in inferior performance when $k > \frac{n}{\ell^2}$. Specifically, as depicted in Figure 4(b), we initially witness the expected faster execution of RANDGREEDI with $\ell = 32$ compared to RANDGREEDI with $\ell = 8$. However, once $k > \frac{n}{32^2}$, the relative performance of RANDGREEDI with $\ell = 32$ rapidly deteriorates. This decline can be attributed to the degradation of RANDGREEDI’s optimal performance beyond $k = \frac{n}{\ell^2}$.

When running RANDGREEDI with a single thread on each machine, the total running time on ℓ machines can be computed based on two components. First, the running time for one machine in the first MR round, which is proportional to $(n/\ell - (k - 1)/2)k$. Second, the running time for the primary machine in the second MR round (post-processing step), which

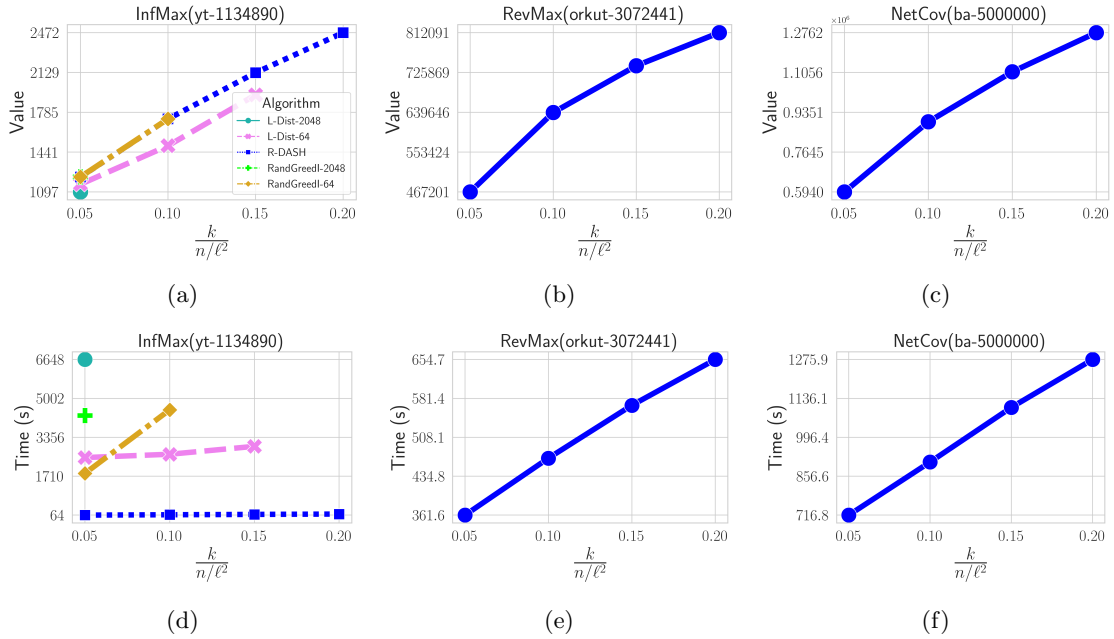


Figure 3: Empirical comparison of R-DASH and L-DIST. The plotted metrics are solution value (Fig. 3(a)-3(c)) and runtime (Fig. 3(d)-3(f)).

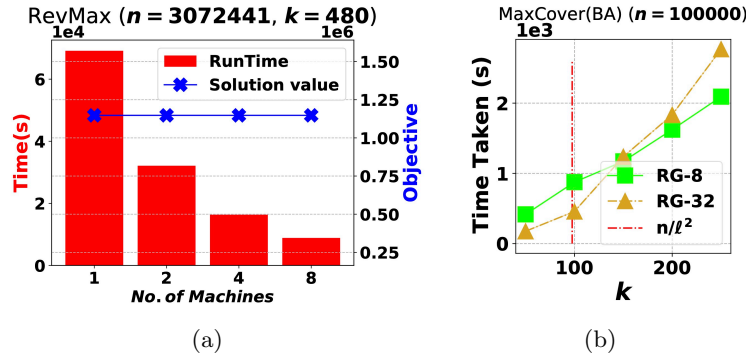


Figure 4: (a): Scalability of R-DASH vs. ℓ (b): RANDGREEDI with $\ell = 8$ Vs. $\ell = 32$.

is proportional to $(k\ell - (k - 1)/2)k$. Consequently, the total running time is proportional to $nk/\ell + \ell k^2 - k(k - 1)$. Optimal performance is achieved when $\ell = \sqrt{n/k}$, which justifies the preference for parallelization within a machine to maintain a lower ℓ rather than distributing the data across separate processors.

Furthermore, in Experiment 5 (Section 8.4), we demonstrate that utilizing MED enables MR algorithms to produce solutions much faster with no compromise in solution value, particularly when solving for $k > \frac{n}{\ell^2}$. These results provide further support for the advantage of incorporating MED in achieving efficient and effective parallelization in MR algorithms.

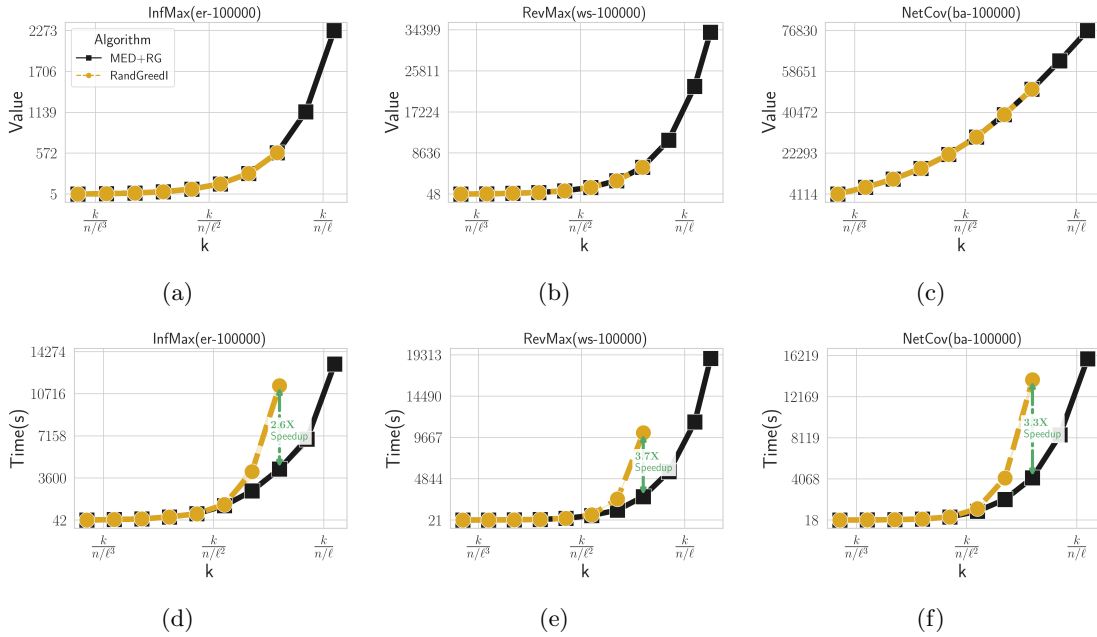


Figure 5: Empirical comparison of MED+RG to RANDGREEDI. The plotted metrics are solution value (Fig. 5(a)-5(c)) and runtime (Fig. 5(d)-5(f)).

8.4 Experiment 4 - Performance Analysis of MED

This section presents experimental results comparing the performance of MED+RG and the vanilla RANDGREEDI algorithm. In terms of solution value, MED+RG consistently provides nearly identical solutions to the vanilla RANDGREEDI algorithm across all instances for the three applications studied. Regarding runtime, the following observations are made: Initially, both algorithms exhibit similar execution times up to a threshold of $k = \frac{n}{\ell^2}$. However, beyond this threshold, the performance gap between MED+RG and RANDGREEDI widens linearly. MED+RG achieves notable average speedup factors of 1.8, 2.2, and 2.3 over RANDGREEDI for the respective applications. Moreover, beyond $k = \frac{n}{\ell^2}$, MED+RG outperforms RANDGREEDI in terms of completing instances within a 12-hour timeout, completing 77% more instances of k across all three applications. These findings highlight the promising performance of MED+RG, demonstrating comparable solution values while significantly improving runtime efficiency compared to the vanilla RANDGREEDI algorithm.

The empirical findings from this experiment provide insights into the capabilities of the MED+ALG framework. Our results indicate that MED+ALG achieves solution quality that is practically indistinguishable from the vanilla ALG algorithm, even for significantly larger values of k . Additionally, even with sufficient available memory to run ALG beyond $k = \frac{n}{\ell^2}$, MED+ALG demonstrates noteworthy computational efficiency, surpassing that of the ALG algorithm. This combination of comparable solution quality and improved computational efficiency positions MED+ALG as a highly promising framework. These findings have significant implications and underscore the potential of MED+ALG to address complex problems in a distributed setting with large-scale values of k .

9. Discussion and Conclusion

Prior to this work, no MR algorithms for SMCC could parallelize within a machine; these algorithms require many sequential queries. Moreover, increasing the number of machines to the number of threads available in the cluster may actually harm the performance, as we showed empirically; intuitively, this is because the size of the set on the primary machine scales linearly with the number of machines ℓ . In this paper, we have addressed this limitation by introducing a suite of algorithms that are both parallelizable and distributed. Specifically, we have presented R-DASH, T-DASH, and G-DASH, which are the first MR algorithms with sublinear adaptive complexity (highly parallelizable). Moreover, our algorithms make have nearly linear query complexity over the entire computation. We also provide the first distributed algorithm with $O(n)$ total query complexity, improving on the $O(npolylog(n))$ of our other algorithms and the algorithm of Liu and Vondrak (2019).

When RANDGREEDI was introduced by Mirzasoleiman et al. (2013), the empirical performance of the algorithm was emphasized, with theoretical guarantees unproven until the work of Barbosa et al. (2015). Since that time, RANDGREEDI has remained the most practical algorithm for distributed SMCC. Our R-DASH algorithm may be regarded as a version of RANDGREEDI that is 1) parallelized; 2) nearly linear time in total (RANDGREEDI is quadratic); and 3) empirically orders of magnitude faster in parallel wall time in our evaluation.

R-DASH achieves approximation ratio of $(1 - 1/e)/2$ in two MR rounds; the first round merely being to distribute the data. We provide G-DASH to close the gap to $(1 - 1/e)$ ratio in a constant number of rounds. However, MR rounds are expensive (as shown by our Experiment 1). The current best ratio achieved in two rounds is the 0.545-approximation algorithm of Mirrokni and Zadimoghaddam (2015). Therefore, a natural question for future research is what is the best ratio achievable in two MR rounds. Moreover, non-monotone or partially monotone objective functions and more sophisticated constraint systems are directions for future research.

Acknowledgements

Yixin Chen and Tonmoy Dey contributed equally to this work. The work of Tonmoy Dey was partially supported by Florida State University; the work of Yixin Chen and Alan Kuhnle was partially supported by Texas A & M University. The authors have received no third-party funding in direct support of this work. The authors have no additional revenues from other sources related to this work.

Appendix A. Lovász Extension of Submodular Function.

Given submodular function f , the Lovász extension F of f is defined as follows: For $\mathbf{z} \in [0, 1]^{|N|}$,

$$F(\mathbf{z} = (z_i)_{i \in N}) = \mathbb{E}_{\lambda \sim \mathcal{U}[0,1]}[f(\{i : z_i \geq \lambda\})].$$

The Lovász extension satisfies the following properties: (1) F is convex; (2) $F(c\mathbf{z}) \geq cF(\mathbf{z})$ for any $c \in (0, 1)$. Moreover, we will require the following simple lemma:

Lemma 8 (Barbosa et al. (2015)). *Let S be a random set, and suppose $\mathbb{E}[\mathbf{1}_S] = c \cdot \mathbf{z}$, for $c \in (0, 1)$. Then $\mathbb{E}[f(S)] \geq c \cdot F(\mathbf{z})$.*

Appendix B. Probability Lemma and Concentration Bounds

Lemma 9. (Chen et al. (2021)) *Suppose there is a sequence of n Bernoulli trials: X_1, X_2, \dots, X_n , where the success probability of X_i depends on the results of the preceding trials X_1, \dots, X_{i-1} . Suppose it holds that*

$$\Pr(X_i = 1 | X_1 = x_1, X_2 = x_2, \dots, X_{i-1} = x_{i-1}) \geq \eta,$$

where $\eta > 0$ is a constant and x_1, \dots, x_{i-1} are arbitrary.

Then, if Y_1, \dots, Y_n are independent Bernoulli trials, each with probability η of success, then

$$\Pr\left(\sum_{i=1}^n X_i \leq b\right) \leq \Pr\left(\sum_{i=1}^n Y_i \leq b\right),$$

where b is an arbitrary integer.

Moreover, let A be the first occurrence of success in sequence X_i . Then,

$$\mathbb{E}[A] \leq 1/\eta.$$

Lemma 10 (Chernoff bounds (Mitzenmacher & Upfal, 2017)). *Suppose X_1, \dots, X_n are independent binary random variables such that $\Pr(X_i = 1) = p_i$. Let $\mu = \sum_{i=1}^n p_i$, and $X = \sum_{i=1}^n X_i$. Then for any $\delta \geq 0$, we have*

$$\Pr(X \geq (1 + \delta)\mu) \leq e^{-\frac{\delta^2 \mu}{2 + \delta}}. \quad (8)$$

Moreover, for any $0 \leq \delta \leq 1$, we have

$$\Pr(X \leq (1 - \delta)\mu) \leq e^{-\frac{\delta^2 \mu}{2}}. \quad (9)$$

Appendix C. Omitted Proof of ThreshSeqMod

Lemma 2. *At an iteration j , let $A_i = \{x \in V_j : \Delta(x | T_i \cup S_{j-1}) < \tau\}$ after Line 5. It holds that $|A_0| = 0$, $|A_{|V_j|}| = |V_j|$ and $|A_i| \leq |A_{i+1}|$.*

Proof. After filtration on Line 5, it holds that, for any $x \in V_j$, $\Delta(x | S_{j-1}) \geq \tau$ and $\Delta(x | V_j \cup S_{j-1}) = 0$. Therefore,

$$\begin{aligned} |A_0| &= |\{x \in V_j : \Delta(x | S_{j-1}) < \tau\}| = 0, \\ |A_{|V_j|}| &= |\{x \in V_j : \Delta(x | V_j \cup S_{j-1}) < \tau\}| = |V_j|. \end{aligned}$$

For any $x \in A_i$, it holds that $\Delta(x | T_i \cup S_{j-1}) < \tau$. Due to submodularity, $\Delta(x | T_{i+1} \cup S_{j-1}) \leq \Delta(x | T_i \cup S_{j-1}) < \tau$. Therefore, A_i is subset of A_{i+1} , which indicates that $|A_i| \leq |A_{i+1}|$. \square

Proof of Success Probability. The algorithm successfully terminates if, at some point, $|V_j| = 0$ or $|S_j| = k$. To analyze the success probability, we consider a variant of the algorithm in which it does not terminate once $|V_j| = 0$ or $|S_j| = k$. In this case, the algorithm keeps running with $s = 0$ and selecting empty set after inner for loop in the following iterations. Thus, with probability 1, either $|V_{j+1}| = 0 \leq (1 - \beta\varepsilon)|V_j| = 0$, or $|S_j| = k$. Lemma 3 holds for all $M + 1$ iterations of the outer for loop.

If the algorithm fails, there should be no more than $m = \lceil \log_{1-\beta\varepsilon}(1/n) \rceil$ successful iterations where $\lambda_j^* \geq \min\{s, t\}$. Otherwise, by Lemma 3, there exists an iteration j such that $|S_j| = k$, or $|V_{M+1}| < n(1 - \beta\varepsilon)^m \leq 1$. Let X be the number of successful iterations. Then, X is the sum of $M + 1$ dependent Bernoulli random variables, where each variable has a success probability of more than $1/2$. Let Y be the sum of $M + 1$ independent Bernoulli random variables with success probability $1/2$. By probability lemmata, the failure probability of Alg. 1 is calculated as follows,

$$\begin{aligned}
 \Pr(\text{Alg. 1 fails}) &\leq \Pr(X \leq m) \\
 &\leq \Pr(Y \leq m) && \text{(Lemma 9)} \\
 &\leq \Pr\left(Y \leq \log\left(\frac{n}{\delta}\right) / (\beta\varepsilon)\right) && (\log(x) \geq 1 - \frac{1}{x}, \forall x > 0) \\
 &\leq e^{-\frac{1}{2}\left(1 - \frac{1}{2(1+\beta\varepsilon)}\right)^2 \cdot 2(1+1/\beta\varepsilon) \log\left(\frac{n}{\delta}\right)} && \text{(Lemma 10)} \\
 &= \left(\frac{\delta}{n}\right)^{\frac{(2\beta\varepsilon+1)^2}{4\beta\varepsilon(\beta\varepsilon+1)}} \leq \frac{\delta}{n} && \square
 \end{aligned}$$

Proof of Adaptivity and Query Complexity. In Alg. 1, oracle queries incurred on Line 5 and 14, and can be done in parallel. Thus, there are constant number of adaptive rounds within one iteration of the outer for loop. The adaptivity is $O(M) = O(\log(n/\delta)/\varepsilon^3)$.

Consider an iteration j , there are no more than $|V_{j-1}| + 1$ queries on Line 5 and $\log_{1+\varepsilon}(|V_j|) + 2$ queries on Line 14. Let Y_i be the i -th successful iteration. Then, for any $Y_{i-1} < j \leq Y_i$, it holds that $|V_j| \leq n(1 - \beta\varepsilon)^{i-1}$. By Lemma 9, for any $i \geq 1$, $\mathbb{E}[Y_i - Y_{i-1}] \leq 2$. Then, the expected number of oracle queries is as follows,

$$\begin{aligned}
 \mathbb{E}[\#\text{queries}] &\leq \sum_{j=1}^{M+1} \mathbb{E}[|V_{j-1}| + \log_{1+\varepsilon}(|V_j|) + 3] \\
 &\leq n + \sum_{i:n(1-\beta\varepsilon)^{i-1} \geq 1} \mathbb{E}[Y_i - Y_{i-1}] (n(1 - \beta\varepsilon)^{i-1} + \log_{1+\varepsilon}(n(1 - \beta\varepsilon)^{i-1})) + 3(M + 1) \\
 &\leq n + 2n \sum_{i \geq 1} (1 - \beta\varepsilon)^{i-1} + 2 \log_{1+\varepsilon}(n) \cdot \log_{1-\beta\varepsilon}(1/n) + 3 \log(n/\delta)/\varepsilon^3 \\
 &\leq \left(1 + \frac{2}{\beta\varepsilon}\right) n + 2 \log_{1+\varepsilon}(n) \cdot \log_{1-\beta\varepsilon}(1/n) + 3 \log(n/\delta)/\varepsilon^3 \\
 &\leq O(n/\varepsilon^3) && \square
 \end{aligned}$$

Algorithm 9 THRESHOLDGREEDY

- 1: **Input:** Evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, constraint k , constant α , value Γ s.t. $\Gamma \leq f(O) \leq \Gamma/\alpha$, error ε
 - 2: Initialize $\tau \leftarrow \Gamma/(\alpha k)$, $S \leftarrow \emptyset$
 - 3: **while** $\tau \geq \varepsilon\Gamma/k$ **do**
 - 4: **for** $e \in \mathcal{N}$ **do**
 - 5: **if** $\Delta(e|S) \geq \tau$ **then** $S \leftarrow S \cup \{e\}$
 - 6: **if** $|S| = k$ **then return** S
 - 7: $\tau \leftarrow \tau(1 - \varepsilon)$
 - 8: **return** S
-

Proof of Property (3). For each iteration j , consider two cases of λ_j^* by the choice of λ_j^* on Line 18. Firstly, if $\lambda_j^* \leq \lceil \frac{1}{\varepsilon} \rceil$, clearly,

$$\Delta(T_{\lambda_j^*} | S_{j-1}) \geq (1 - \varepsilon)\tau\lambda_j^*.$$

Secondly, if $\lambda_j^* > \lceil \frac{1}{\varepsilon} \rceil$, let $\lambda_j = \max\{\lambda \in \Lambda : \lambda < \lambda_j^*\}$. Then, $B[\lambda_j] = \mathbf{True}$, and

$$\Delta(T_{\lambda_j} | S_{j-1}) \geq (1 - \varepsilon)\tau\lambda_j.$$

Let $\lambda_j = \lfloor (1 + \varepsilon)^u \rfloor$, and $\lambda_j^* = \lfloor (1 + \varepsilon)^{u+1} \rfloor$. Then,

$$\frac{|T_{j,\lambda_j}|}{|T_{j,\lambda_j^*}|} = \frac{\lambda_j}{\lambda_j^*} = \frac{\lfloor (1 + \varepsilon)^u \rfloor}{\lfloor (1 + \varepsilon)^{u+1} \rfloor} \geq \frac{(1 + \varepsilon)^u - 1}{(1 + \varepsilon)^{u+1}} \geq \frac{1}{1 + \varepsilon} - \varepsilon,$$

where the last inequality follows from $(1 + \varepsilon)^{u+1} \geq \lambda_j^* > \lceil \frac{1}{\varepsilon} \rceil \geq \frac{1}{\varepsilon}$. Therefore, by above two inequalities and monotonicity of f ,

$$\Delta(T_{j,\lambda_j^*} | S_{j-1}) \geq \Delta(T_{j,\lambda_j} | S_{j-1}) \geq (1 - \varepsilon) \left(\frac{1}{1 + \varepsilon} - \varepsilon \right) \tau\lambda_j^* \geq (1 - 2\varepsilon)\tau\lambda_j^*/(1 + \varepsilon).$$

The objective value of the returned solution can be bounded as follows,

$$f(S) = \sum_j \Delta(T_{\lambda_j^*} | S_{j-1}) \geq \sum_j (1 - 2\varepsilon)\tau\lambda_j^*/(1 + \varepsilon) = (1 - 2\varepsilon)\tau|S|/(1 + \varepsilon). \quad \square$$

Proof of Property (4). If the output set S follows that $|S| < k$, for any $x \in \mathcal{N}$, it is filtered out at some point. Let x be discarded at iteration j_x . Then,

$$\Delta(x|S) \leq \Delta(x|S_{j_x-1}) < \tau. \quad \square$$

Appendix D. ThresholdGreedy

In this section, we introduce a variant of near-linear time algorithm (Alg. 1) in Badanidiyuru and Vondrak (2014), THRESHOLDGREEDY (Alg. 9), which requires two more parameters, constant α and value Γ . With the assumption that Γ is an α -approximation solution, THRESHOLDGREEDY ensures a $(1 - 1/e - \varepsilon)$ -approximation with $O(n/\varepsilon)$ query calls.

Theorem 8. *Let (f, k) be an instance of SM, and O is the optimal solution to (f, k) . With input α, Γ and ε , where $\Gamma \leq f(O) \leq \Gamma/\alpha$, THRESHOLDGREEDY outputs solution S with $O(n\varepsilon^{-1} \log(\alpha\varepsilon)^{-1})$ queries such that $f(S) \geq (1 - 1/e - \varepsilon)f(O)$.*

Proof. Query Complexity. The while loop in Alg. 9 has at most $\log_{1-\varepsilon}(\alpha\varepsilon) \leq \varepsilon^{-1} \log(\alpha\varepsilon)^{-1}$ iterations. For each iteration of the while loop, there are $O(n)$ queries. Therefore, the query complexity of THRESHOLDGREEDY is $O(n\varepsilon^{-1} \log(\alpha\varepsilon)^{-1})$.

Approximation Ratio. If $|S| < k$ at termination, for any $e \in \mathcal{N} \setminus S$, it holds that $\Delta(e|S) < \varepsilon\Gamma/k \leq \varepsilon f(O)/k$. Then,

$$\begin{aligned} f(O) - f(S) &\stackrel{(1)}{\leq} \sum_{o \in O \setminus S} \Delta(o|S) \leq \varepsilon f(O) \\ \Rightarrow f(S) &\geq (1 - \varepsilon)f(O) \end{aligned}$$

where Inequality (1) follows from monotonicity and submodularity.

If $|S| = k$ at termination, let S_j be the intermediate solution after iteration j of the while loop, and τ_j is the corresponding threshold. Suppose that $S_j \setminus S_{j-1} \neq \emptyset$. Then, each element added during iteration j provides a minimum incremental benefit of τ ,

$$f(S_j) - f(S_{j-1}) \geq |S_j \setminus S_{j-1}| \tau_j.$$

Next, we consider iteration $j - 1$. If $j > 1$, since the algorithm does not terminate during iteration $j - 1$, each element in \mathcal{N} is either added to the solution or omitted due to small marginal gain with respect to the current solution. Therefore, for any $o \in O \setminus S_{j-1}$, it holds that $\Delta(o|S_{j-1}) < \tau_{j-1}$ by submodularity. Then,

$$f(O) - f(S_{j-1}) \leq \sum_{o \in O \setminus S_{j-1}} \Delta(o|S_{j-1}) < k\tau_{j-1} = k\tau_j/(1 - \varepsilon).$$

If $j = 1$, $\tau_j = \Gamma/(\alpha k)$ and $S_{j-1} = \emptyset$.

$$f(O) - f(S_{j-1}) = f(O) \leq \Gamma/\alpha = k\tau_j < k\tau_j/(1 - \varepsilon).$$

Therefore, for any iteration j that has elements being added to the solution, it holds that

$$\begin{aligned} f(O) - f(S_j) &\leq \left(1 - \frac{1 - \varepsilon}{k} |S_j \setminus S_{j-1}|\right) (f(O) - f(S_{j-1})) \\ \Rightarrow f(O) - f(S) &\leq \prod_j \left(1 - \frac{1 - \varepsilon}{k} |S_j \setminus S_{j-1}|\right) f(O) \end{aligned}$$

Table 3: Small and Large Data

Application	Small Data		Large Data	
	n	Edges	n	Edges
ImageSumm	10,000	$\simeq 1.0 \times 10^8$	50,000	$\simeq 2.5 \times 10^9$
InfluenceMax	26,588	$\simeq 1.0 \times 10^5$	1,134,890	$\simeq 6.0 \times 10^6$
RevenueMax	17,432	$\simeq 1.8 \times 10^5$	3,072,441	$\simeq 2.3 \times 10^8$
MaxCover (BA)	100,000	$\simeq 5.0 \times 10^5$	1,000,000	$\simeq 1.0 \times 10^9$

$$\leq \prod_j e^{-\frac{1-\varepsilon}{k}|S_j \setminus S_{j-1}|} f(O) \quad (x+1 \leq e^x)$$

$$= e^{-1+\varepsilon} f(O) \leq (1/e + \varepsilon) f(O) \quad (0 < \varepsilon < 1)$$

$$\Rightarrow f(S) \geq (1 - 1/e - \varepsilon) f(O) \quad \square$$

Appendix E. Threshold-DASH with no Knowledge of OPT

Lemma 7. *For any $o \in O$, it holds that $\Pr(o \in O_1 \cup O_2) \geq 3/4$.*

Proof. By Definition (p_x in Section 5), it holds that $\Pr(o \in O_1) = p_o$. Since o is assigned to each machine randomly with probability $1/\ell$,

$$\begin{aligned} \Pr(o \in O_2) &= \sum_{i=1}^{\ell} \Pr(o \in S_i) \\ &= \ell \cdot \Pr(o \in S_1 | o \in \mathcal{N}_1) \cdot \Pr(o \in \mathcal{N}_1) \\ &= \Pr(o \in \text{TSMREL}(\mathcal{N}_1) | o \in \mathcal{N}_1) \\ &= \Pr(o \in \text{TSMREL}(\mathcal{N}_1 \cup \{o\})) \\ &= 1 - p_o. \end{aligned}$$

Moreover, we know that any two machines selects elements independently. So,

$$\begin{aligned} \Pr(o \in O_2 | o \in O_1) &= \Pr(o \in O_2 | o \notin \text{TSMREL}(\mathcal{N}_1 \cup \{o\}, q)) \\ &= 1 - \Pr(o \notin O_2 | o \notin \text{TSMREL}(\mathcal{N}_1 \cup \{o\}, q)) \\ &= 1 - \prod_{i=1}^{\ell} \Pr(o \notin S_i | o \notin \text{TSMREL}(\mathcal{N}_1 \cup \{o\}, q)) \\ &\stackrel{(a)}{=} 1 - \prod_{i=2}^{\ell} \Pr(o \notin S_i) \\ &\leq 1 - \Pr(o \notin O_2) \\ &= \Pr(o \in O_2) = 1 - p_o, \end{aligned}$$

where Inequality (a) follows from $\Pr(o \notin \text{TSMREL}(\mathcal{N}_1) | o \notin \text{TSMREL}(\mathcal{N}_1 \cup \{o\})) = 1$.

Thus, we can bound the probability by the following,

$$\begin{aligned}
 \Pr(o \in O_1 \cup O_2) &= \Pr(o \in O_1) + \Pr(o \in O_2) \\
 &\quad - \Pr(o \in O_1 \cap O_2) \\
 &\geq p_o + 1 - p_o - p_o(1 - p_o) \\
 &\geq 3/4.
 \end{aligned}$$

□

Description. The T-DASH algorithm described in Alg. 10 is a two MR-rounds algorithm using the AFD approach that runs THRESHSEQMOD concurrently on every machine for $\log_{1+\varepsilon}(k)$ different guesses of threshold $\tau_{i,j}$ in the range $[\frac{\alpha\Delta_i^*}{k}, \alpha\Delta_i^*]$; where α is the approximation of T-DASH and Δ_i^* is the maximum singleton in \mathcal{N}_i . Every solution returned to the primary machine are placed into bins based on their corresponding threshold guess $\tau_{i,j}$ such that $\frac{\Delta^*}{k}(1+\varepsilon)^x \leq \tau_{i,j} \leq \frac{\Delta^*}{k}(1+\varepsilon)^{x+1}$; where Δ^* is the maximum singleton in \mathcal{N} . Since there must exist a threshold τ^* that is close enough to $\alpha\text{OPT}/k$; running THRESHSEQMOD (Line 21) on every bin in the range $[\frac{\alpha\Delta^*}{k}, \alpha\Delta^*]$ and selecting the best solution guarantees the α -approximation of T-DASH in $O(\log(n))$ adaptive rounds.

Theorem 9. *Let (f, k) be an instance of SM where $k \log(k) < \frac{\varepsilon\psi}{\ell}$. T-DASH with no knowledge of OPT returns set T' with two MR rounds, $O(\frac{1}{\varepsilon^3} \log(n))$ adaptive rounds, $O(\frac{n \log(k)}{\varepsilon^4})$ total queries, $O(n)$ communication complexity, and probability at least $1 - n^{-c}$ such that*

$$\mathbb{E}[f(T')] \geq \left(\frac{3}{8} - \varepsilon\right) \text{OPT}.$$

Overview of Proof. Alg. 10 is inspired by Alg. 6 in Section 5, which is a version of the algorithm that knows that optimal solution value. With $\Delta^* = \max\{f(e) : e \in \mathcal{N}\}$, there exists an x_0 such that $\tau_{x_0} \leq \alpha\text{OPT}(1+\varepsilon)/k \leq \tau_{x_0+1}$ with $O(\log(k)/\varepsilon)$ guesses. Then, on each machine i , we only consider sets $S_{i,j}$ and $R_{i,j}$ such that $\tau_{x_0} \leq \tau_{i,j} \leq \tau_{x_0+1}$. If this $\tau_{i,j}$ does exist, $(S_{i,j}, R_{i,j})$ works like (S_i, R_i) in Alg. 6. If this $\tau_{i,j}$ does not exist, then for any $e \in \mathcal{N}_i$, it holds that $f(e) < \alpha\text{OPT}/k$, which means THRESHSEQMOD(\mathcal{N}_i) with $\tau = \alpha\text{OPT}/k$ will return an empty set. Since each call of THRESHSEQMOD with different guesses on τ is executed in parallel, the adaptivity remains the same and the query complexity increases by a factor of $\log(k)/\varepsilon$.

Proof. First, for $x = 0$, $\tau_x = (\alpha + \varepsilon)\Delta^*/k \leq (\alpha + \varepsilon)\text{OPT}/k$; and for $x = \lceil \log_{1+\varepsilon}(k) \rceil + 1$,

$$\tau_x \geq (\alpha + \varepsilon)(1 + \varepsilon)\Delta^* \geq (\alpha + \varepsilon)(1 + \varepsilon) \sum_{o \in O} f(o)/k \geq (\alpha + \varepsilon)(1 + \varepsilon)\text{OPT}/k.$$

Therefore, there exists an x_0 such that $\tau_{x_0} \leq (\alpha + \varepsilon)(1 + \varepsilon)\text{OPT}/k \leq \tau_{x_0+1}$. Since $\tau_{x_0+1} = \tau_{x_0}(1 + \varepsilon)$, it holds that $\tau_{x_0} \geq (\alpha + \varepsilon)\text{OPT}/k$ and $\tau_{x_0+1} \leq (\alpha + \varepsilon)(1 + \varepsilon)^2\text{OPT}/k$.

Then, we only consider T'_{x_0} . If $|T'_{x_0}| = k$, by Theorem 1, it holds that,

$$f(T) \geq f(T'_{x_0}) \geq \frac{1 - 2\varepsilon}{1 + \varepsilon} \tau_{x_0} k \geq \left(\frac{3}{8} - \varepsilon\right) \text{OPT}.$$

Algorithm 10 Threshold-DASH with no knowledge of OPT (T-DASH)

- 1: **Input:** Evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, constraint k , error ε , available machines $M \leftarrow \{1, 2, \dots, \ell\}$
 - 2: Initialize $\delta \leftarrow 1/(\ell + 1)$, $\mathbf{q} \leftarrow$ a fixed sequence of random bits.
 - 3: Set $\alpha \leftarrow \frac{3}{8}$, $(\mathbf{q}_{i,j})_{i \in [\ell+1], j \in [\log_{1+\varepsilon}(k)]} \leftarrow \mathbf{q}$
 - 4: **for** $e \in \mathcal{N}$ **do** **do**
 - 5: Assign e to each machine independently with probability $1/\ell$
 - 6: **for** $i \in M$ **do**
 - 7: ▷ On machine i
 - 8: Let \mathcal{N}_i be the elements assigned to machine i
 - 9: Set $\Delta_i^* \leftarrow \max\{f(e) : e \in \mathcal{N}_i\}$
 - 10: **for** $j \leftarrow 0$ to $\log_{1+\varepsilon}(k)$ in parallel **do**
 - 11: $\tau_{i,j} \leftarrow \frac{(\alpha+\varepsilon)\Delta_i^*}{k}(1+\varepsilon)^j$
 - 12: $S_{i,j}, R_{i,j} \leftarrow \text{THRESHSEQMOD}(f, \mathcal{N}_i, k, \delta, \varepsilon, \tau_j, \mathbf{q}_{i,j})$
 - 13: Send Δ_i^* and all $(\tau_{i,j}, S_{i,j}, R_{i,j})$ to primary machine
 - 14: ▷ On primary machine
 - 15: Set $\Delta^* \leftarrow \max\{\Delta_i^* : 1 \leq i \leq \ell\}$
 - 16: **for** $x \leftarrow 0$ to $\lceil \log_{1+\varepsilon}(k) \rceil + 1$ in parallel **do**
 - 17: $\tau_x \leftarrow \frac{(\alpha+\varepsilon)\Delta^*}{k}(1+\varepsilon)^x$
 - 18: Let $R_x \leftarrow \{\bigcup R_{i,j} : \tau_x \leq \tau_{i,j} \leq \tau_{x+1}\}$
 - 19: Let $A_x \leftarrow \{\text{Sample a solution } S_{i,j} : \tau_x \leq \tau_{i,j} \leq \tau_{x+1}\}$
 - 20: Let $g_x(\cdot) \leftarrow f(A_x \cup \cdot) - f(A_x)$
 - 21: $T_x \leftarrow \text{THRESHSEQMOD}(g_x, R_x, k - |A_x|, \delta, \varepsilon, \tau_x, \mathbf{q}_{\ell+1,x})$
 - 22: $T'_x \leftarrow A_x \cup T_x$
 - 23: $T \leftarrow \arg \max\{f(T'_x) : 0 \leq x \leq \log_{1+\varepsilon}(k)\}$
 - 24: **return** T
-

Otherwise, in the case that $|T'_{x_0}| < k$, let $A_{x_0} = S_{i_0, j_0}$, $O_1 = \{o \in O : o \notin \text{TSMREL}(N_{i_0} \cup \{o\}, q)\}$. Also, let $\tau_{i, (j_0)}$ be returned by machine i that $\tau_{x_0} \leq \tau_{i, (j_0)} \leq \tau_{x_0+1}$, define

$$(S'_{i, (j_0)}, R'_{i, (j_0)}) = \begin{cases} (S_{i, (j_0)}, R_{i, (j_0)}) & , \text{ if machine } i \text{ returned a } \tau_{i, (j_0)} \\ (\emptyset, \emptyset) & , \text{ otherwise} \end{cases}.$$

On the machine which does not return a $\tau_{i, (j_0)}$, we consider it runs $\text{THRESHSEQMOD}(\mathcal{N}_i, \tau_{x_0+1})$ and returns two empty sets, and hence $\max\{f(e) : e \in \mathcal{N}_i\} < \tau_{x_0}$. Let $R'_{x_0} = \bigcup_{i \in M} R'_{i, (j_0)}$, $O_2 = R'_{x_0} \cap O$. Then, Lemma 7 in Appendix E still holds in this case. We can calculate the approximation ratio as follows with $\varepsilon \leq 2/3$,

$$\begin{aligned} \mathbb{E}[T] &\geq \mathbb{E}[T'_{x_0}] \geq \mathbb{E}[f(O_1 \cup O_2)] - k \cdot \tau_{x_0+1} \\ &\geq \left(\frac{3}{8} - \varepsilon\right) \text{OPT}. \end{aligned}$$

□

Appendix F. Experiment Setup

F.1 Applications

Given a constraint k , the objectives of the applications are defined as follows:

F.1.1 MAX COVER

Maximize the number of nodes covered by choosing a set S of maximum size k , such that the number of nodes having at least one neighbour in the set S . The application is run on synthetic random BA graphs of groundset size 100,000, 1,000,000 and 5,000,000 generated using Barabási–Albert (BA) models for the centralized and distributed experiments respectively.

F.1.2 IMAGE SUMMARIZATION ON CIFAR-10 DATA

Given large collection of images, find a subset of maximum size k which is representative of the entire collection. The objective used for the experiments is a monotone variant of the image summarization from Fahrback, Mirrokni, and Zadimoghaddam (2019a). For a groundset with N images, it is defined as follows:

$$f(S) = \sum_{i \in N} \max_{j \in S} s_{i,j}$$

where $s_{i,j}$ is the cosine similarity of the pixel values between image i and image j . The data for the image summarization experiments contains 10,000 and 50,000 CIFAR-10 (Krizhevsky, Hinton, et al., 2009) color images respectively for the centralized and distributed experiments.

F.1.3 INFLUENCE MAXIMIZATION ON A SOCIAL NETWORK.

Maximise the aggregate influence to promote a topic by selecting a set of social network influencers of maximum size k . The probability that a random user i will be influenced by the set of influencers in S is given by:

$$\begin{aligned} f_i(S) &= 1 \quad , \text{ for } i \in S \\ f_i(S) &= 1 - (1 - p)^{|N_S(i)|} \quad , \text{ for } i \notin S \end{aligned}$$

where $|N_S(i)|$ is the number of neighbors of node i in S . We use the Epinions data set consisting of 27,000 users from Rossi and Ahmed (2015) for the centralized data experiments and the Youtube online social network data Yang and Leskovec (2012) consisting more than 1 million users for distributed data experiments. The value of p is set to 0.01.

F.1.4 REVENUE MAXIMIZATION ON YOUTUBE.

Maximise revenue of a product by selecting set of users S of maximum size k , where the network neighbors will be advertised a different product by the set of users S . It is based on the objective function from Mirzasoleiman, Badanidiyuru, and Karbasi (2016). For a given set of users X and $w_{i,j}$ as the influence between user i and j , the objective function can be

defined by:

$$f(S) = \sum_{i \in X} V \left(\sum_{j \in S} w_{i,j} \right)$$

$$V(y) = y^\alpha$$

where $V(S)$, the expected revenue from an user is a function of the sum of influences from neighbors who are in S and $\alpha : 0 < \alpha < 1$ is a rate of diminishing returns parameter for increased cover.

We use the Youtube data set from Mirzasoleiman et al. (2016) consisting of 18,000 users for centralized data experiments. For the distributed data experiments we perform empirical evaluation on the Orkut online social network data from Yang and Leskovec (2012) consisting more than 3 million users. The value of α is set to 0.3

Appendix G. Replicating the Experiment Results

Our experiments can be replicated by running the following scripts:

- Install **MPICH** version **3.3a2** (DO NOT install OpenMPI and ensure *mpirun* utilizes mpich using the command *mpirun -version* (Ubuntu))
- Install **pandas**, **mpi4py**, **scipy**, **networkx**
- Set up an MPI cluster using the following tutorial:
<https://mpitutorial.com/tutorials/running-an-mpi-cluster-within-a-lan/>
- Create and update the host file *../nodesFileIPnew* to store the ip addresses of all the connected MPI machines before running any experiments (First machine being the primary machine)
 - NOTE: Please place *nodesFileIPnew* inside the MPI shared repository; *"cloud/"* in this case (at the same level as the code base directory). DO NOT place it inside the code base *"DASH-Distributed_SMCC-python/"* directory.
- Clone the *DASH-Distributed_SMCC-python* repository inside the MPI shared repository (*/cloud* in the case using the given tutorial)
 - NOTE: Please clone the *"DASH-Distributed_SMCC-python"* repository and execute the following commands on a machine with **sufficient memory** (RAM); capable of generating the large datasets. This repository NEED NOT be the primary repository (*"/cloud/DASH-Distributed_SMCC-python/"*) on the shared memory of the cluster; that will be used for the experiments.
- Additional Datasets For Experiment 1 : Please download the Image Similarity Matrix file **"images_10K_mat.csv"** (<https://drive.google.com/file/d/1s9PzUhV-C5dW8iL4tZPVjSRX4PBhrsiJ/v>) and place it in the *data/data_exp1/* directory.
- To generate the decentralized data for **Experiment 2** and **3** : Please follow the below steps:

- Execute `bash GenerateDistributedData.bash nThreads nNodes`
- The previous command should generate `nNodes` directories in `loading_data/` directory (with names `machine<nodeNo>Data`)
- Copy the `data_exp2_split/` and `data_exp3_split/` directories within each `machine<i>Data` directory to the corresponding machine M_i and place the directories outside `/cloud` (directory created after setting up an MPI cluster using the given tutorial)).

To run all experiments in the apper

Please read the `README.md` file in the `"DASH-Distributed.SMCC-python"` (Code/Data Appendix) for detailed information.

Appendix H. The Greedy Algorithm (Nemhauser and Wolsey (1978))

The standard GREEDY algorithm starts with an empty set and then proceeds by adding elements to the set over k iterations. In each iteration the algorithm selects the element with the maximum marginal gain $\Delta(e|A_{i-1})$ where $\Delta(i|A) = f(A \cup \{i\}) - f(A)$. The Algorithm 11 is a formal statement of the standard GREEDY algorithm. The intermediate solution set A_i represents the solution after iteration i . The $(1 - 1/e \approx 0.632)$ -approximation of the GREEDY algorithm is the optimal ratio possible for monotone objectives.

Algorithm 11 GREEDY(f, \mathcal{N}, k)

- 1: **Input:** evaluation oracle $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, constraint k
 - 2: Let $A_i \leftarrow \emptyset$
 - 3: **for** $i = 1$ to k **do**
 - 4: $u_i \leftarrow \arg \max\{f(e|A_{i-1}); \text{ where } e \in \mathcal{N} \setminus A_{i-1}\}$
 - 5: $A_i \leftarrow A_{i-1} \cup u_i$
 - 6: **return** A_k
-

Appendix I. Discussion of the 1/2-approximate Algorithm in Liu and Vondrak (2019)

Algorithm 12 THRESHOLDGREEDY(S, G, τ)

- 1: **Input:** An input set S , a partial greedy solution G with $|G| \leq k$, and a threshold τ
 - 2: $G' \leftarrow G$
 - 3: **for** $e \in S$ **do**
 - 4: **if** $\Delta(e|G_0) \geq \tau$ and $|G_0| < k$ **then**
 - 5: $G_0 \leftarrow G_0 \cup \{e\}$
 - 6: **return** G'
-

The simple two-round distributed algorithm proposed by Liu and Vondrak (2019) achieves a deterministic $1/2 - \epsilon$ approximation ratio. The pseudocode of this algorithm, assuming that we know OPT, is provided in Alg. 14. The guessing OPT version runs $O(\log(k))$ copies of Alg. 14 on each machine, which has $O(k)$ adaptivity and $O(nk \log(k))$ query complexity.

Algorithm 13 THRESHOLDFILTER(S, G, τ)

```

1: Input: An input set  $S$ , a partial greedy solution  $G$  with  $G \leq k$ , and a threshold  $\tau$ 
2:  $S' \leftarrow S$ 
3: for  $e \in S$  do
4:   if  $\Delta(e|G) < \tau$  then  $S' \leftarrow S' \setminus \{e\}$ 
5: return  $S'$ 

```

Algorithm 14 A simple 2-round 1/2 approximation, assuming OPT is known

```

1: round 1:
2:  $S \leftarrow$  sample each  $e \in \mathcal{N}$  with probability  $p = 4\sqrt{k/n}$ 
3: send  $S$  to each machine and the central machine  $C$ 
4: partition  $\mathcal{N}$  randomly into sets  $V_1, V_2, \dots, V_m$ 
5: send  $V_i$  to machine  $i$  for each  $i \in [m]$ 
6: for each machine  $M_i$  (in parallel) do
7:    $\tau \leftarrow \frac{\text{OPT}}{2k}$ 
8:    $G_0 \leftarrow$  THRESHOLDGREEDY( $S, \emptyset, \tau$ )
9:   if  $|G_0| < k$  then
10:     $R_i \leftarrow$  THRESHOLDFILTER( $V_i, G_0, \tau$ )
11:   else
12:     $R_i = \emptyset$ 
13:   send  $R_i$  to the central machine  $C$ 

14: round 2 (only on  $C$ ):
15: compute  $G_0$  from  $S$  as in first round
16: for  $e \in \cup_i R_i$  do
17:   if  $\Delta(e|G) \geq \tau$  then  $G \leftarrow G \cup \{e\}$ 
18: return  $G$ 

```

Intuitively, the result of Alg. 14 is equivalent to a single run of THRESHOLDGREEDY on the whole ground set. Within the first round, it calls THRESHOLDGREEDY on S to get an intermediate solution G_0 . After filtering out the elements with gains less than τ in V_i , the rest of the elements in V_i are sent to the primary machine. By careful analysis, there is a high probability that the total number of elements returned within the first round is less than \sqrt{nk} .

Compared to the 2-round algorithms in Table 1, Alg. 14 is the only algorithm that requires data duplication, with four times more elements distributed to each machine in the first round. As a result, distributed setups have a more rigid memory restriction when running this algorithm. For example, when the dataset exactly fits into the distributed setups which follows that $\Psi = n/\ell$, there is no more space on each machine to store the random subset S . In the meantime, other 2-round algorithms can still run with $k \leq n/\ell^2$.

Additionally, since the approximation analysis of Alg. 14 does not take randomization into consideration, we may be able to substitute its THRESHOLDGREEDY subroutine with any threshold algorithm (such as THRESHOLDSEQ in Chen et al. (2021)) that has optimal

logarithmic adaptivity and linear query complexity. However, given its analysis of the success probability considers S as $3k$ blocks where each block has random access to the whole ground set, it may not be easy to fit parallelizable threshold algorithms into this framework.

References

- Badanidiyuru, A., & Vondrak, J. (2014). Fast algorithms for maximizing submodular functions. In Chekuri, C. (Ed.), *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pp. 1497–1514. SIAM.
- Balakrishnan, R., Li, T., Zhou, T., Himayat, N., Smith, V., & Bilmes, J. A. (2022). Diverse client selection for federated learning via submodular maximization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Balkanski, E., Rubinfeld, A., & Singer, Y. (2019). An exponential speedup in parallel running time for submodular maximization without loss in approximation. In Chan, T. M. (Ed.), *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pp. 283–302. SIAM.
- Balkanski, E., & Singer, Y. (2018). The adaptive complexity of maximizing a submodular function. In Diakonikolas, I., Kempe, D., & Henzinger, M. (Eds.), *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pp. 1138–1151. ACM.
- Barbosa, R. d. P., Ene, A., Nguyen, H. L., & Ward, J. (2015). The power of randomization: Distributed submodular maximization on massive datasets. In Bach, F. R., & Blei, D. M. (Eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, Vol. 37 of *JMLR Workshop and Conference Proceedings*, pp. 1236–1244. JMLR.org.
- Barbosa, R. d. P., Ene, A., Nguyen, H. L., & Ward, J. (2016). A new framework for distributed submodular maximization. In Dinur, I. (Ed.), *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pp. 645–654. IEEE Computer Society.
- Blelloch, G. E., Simhadri, H. V., & Tangwongsan, K. (2012). Parallel and I/O efficient set covering algorithms. In Blelloch, G. E., & Herlihy, M. (Eds.), *24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '12, Pittsburgh, PA, USA, June 25-27, 2012*, pp. 82–90. ACM.
- Breuer, A., Balkanski, E., & Singer, Y. (2020). The FAST algorithm for submodular maximization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Vol. 119 of *Proceedings of Machine Learning Research*, pp. 1134–1143. PMLR.

- Buchbinder, N., Feldman, M., Naor, J., & Schwartz, R. (2014). Submodular maximization with cardinality constraints. In Chekuri, C. (Ed.), *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pp. 1433–1452. SIAM.
- Buchbinder, N., Feldman, M., & Schwartz, R. (2017). Comparing apples and oranges: Query trade-off in submodular maximization. *Math. Oper. Res.*, 42(2), 308–329.
- Calinescu, G., Chekuri, C., Pal, M., & Vondrak, J. (2007). Maximizing a submodular set function subject to a matroid constraint (extended abstract). In Fischetti, M., & Williamson, D. P. (Eds.), *Integer Programming and Combinatorial Optimization, 12th International IPCO Conference, Ithaca, NY, USA, June 25-27, 2007, Proceedings*, Vol. 4513 of *Lecture Notes in Computer Science*, pp. 182–196. Springer.
- Chekuri, C., & Quanrud, K. (2019). Submodular function maximization in parallel via the multilinear relaxation. In Chan, T. M. (Ed.), *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pp. 303–322. SIAM.
- Chen, Y., Dey, T., & Kuhnle, A. (2021). Best of both worlds: Practical and theoretically optimal submodular maximization in parallel. In Ranzato, M., Beygelzimer, A., Dauphin, Y. N., Liang, P., & Vaughan, J. W. (Eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 25528–25539.
- Chen, Y., & Kuhnle, A. (2024). Practical and parallelizable algorithms for non-monotone submodular maximization with size constraint. *J. Artif. Intell. Res.*, 79, 599–637.
- Chierichetti, F., Kumar, R., & Tomkins, A. (2010). Max-cover in map-reduce. In Rappa, M., Jones, P., Freire, J., & Chakrabarti, S. (Eds.), *Proceedings of the 19th International Conference on World Wide Web, WWW 2010, Raleigh, North Carolina, USA, April 26-30, 2010*, pp. 231–240. ACM.
- Dey, T., Chen, Y., & Kuhnle, A. (2023). DASH: A distributed and parallelizable algorithm for size-constrained submodular maximization. In Williams, B., Chen, Y., & Neville, J. (Eds.), *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pp. 3941–3948. AAAI Press.
- Ene, A., & Nguyen, H. L. (2019). Submodular maximization with nearly-optimal approximation and adaptivity in nearly-linear time. In Chan, T. M. (Ed.), *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pp. 274–282. SIAM.
- Ene, A., & Nguyen, H. L. (2020). Parallel algorithm for non-monotone dr-submodular maximization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Vol. 119 of *Proceedings of Machine Learning Research*, pp. 2902–2911. PMLR.

- Epasto, A., Mirrokni, V. S., & Zadimoghaddam, M. (2017). Bicriteria distributed submodular maximization in a few rounds. In Scheideler, C., & Hajiaghayi, M. T. (Eds.), *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2017, Washington DC, USA, July 24-26, 2017*, pp. 25–33. ACM.
- Ettinger, S., Cheng, S., Caine, B., Liu, C., Zhao, H., Pradhan, S., Chai, Y., Sapp, B., Qi, C. R., Zhou, Y., Yang, Z., Chouard, A., Sun, P., Ngiam, J., Vasudevan, V., McCauley, A., Shlens, J., & Anguelov, D. (2021). Large scale interactive motion forecasting for autonomous driving : The waymo open motion dataset. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pp. 9690–9699. IEEE.
- Fahrbach, M., Mirrokni, V. S., & Zadimoghaddam, M. (2019a). Non-monotone submodular maximization with nearly optimal adaptivity and query complexity. In Chaudhuri, K., & Salakhutdinov, R. (Eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, Vol. 97 of *Proceedings of Machine Learning Research*, pp. 1833–1842. PMLR.
- Fahrbach, M., Mirrokni, V. S., & Zadimoghaddam, M. (2019b). Submodular maximization with nearly optimal approximation, adaptivity and query complexity. In Chan, T. M. (Ed.), *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pp. 255–273. SIAM.
- Gupta, A., Roth, A., Schoenebeck, G., & Talwar, K. (2010). Constrained non-monotone submodular maximization: Offline and secretary algorithms. In Saberi, A. (Ed.), *Internet and Network Economics - 6th International Workshop, WINE 2010, Stanford, CA, USA, December 13-17, 2010. Proceedings*, Vol. 6484 of *Lecture Notes in Computer Science*, pp. 246–257. Springer.
- Joseph, K. J., R, V. T., Singh, K., & Balasubramanian, V. N. (2019). Submodular batch selection for training deep neural networks. In Kraus, S. (Ed.), *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pp. 2677–2683. ijcai.org.
- Kazemi, E., Minaee, S., Feldman, M., & Karbasi, A. (2021). Regularized submodular maximization at scale. In Meila, M., & Zhang, T. (Eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, Vol. 139 of *Proceedings of Machine Learning Research*, pp. 5356–5366. PMLR.
- Kazemi, E., Mitrovic, M., Zadimoghaddam, M., Lattanzi, S., & Karbasi, A. (2019). Submodular streaming in all its glory: Tight approximation, minimum memory and low adaptive complexity. In Chaudhuri, K., & Salakhutdinov, R. (Eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, Vol. 97 of *Proceedings of Machine Learning Research*, pp. 3311–3320. PMLR.
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images. In *University of Toronto, Canada*. Citeseer.

- Kuhnle, A. (2021). Quick streaming algorithms for maximization of monotone submodular functions in linear time. In Banerjee, A., & Fukumizu, K. (Eds.), *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, Vol. 130 of *Proceedings of Machine Learning Research*, pp. 1360–1368. PMLR.
- Kumar, R., Moseley, B., Vassilvitskii, S., & Vattani, A. (2013). Fast greedy algorithms in mapreduce and streaming. In Blelloch, G. E., & Vöcking, B. (Eds.), *25th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '13, Montreal, QC, Canada - July 23 - 25, 2013*, pp. 1–10. ACM.
- Lee, J., Mirrokni, V. S., Nagarajan, V., & Sviridenko, M. (2009). Non-monotone submodular maximization under matroid and knapsack constraints. In Mitzenmacher, M. (Ed.), *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pp. 323–332. ACM.
- Liu, P., & Vondrak, J. (2019). Submodular optimization in the mapreduce model. In Fineman, J. T., & Mitzenmacher, M. (Eds.), *2nd Symposium on Simplicity in Algorithms, SOSA 2019, January 8-9, 2019, San Diego, CA, USA*, Vol. 69 of *OASiCs*, pp. 18:1–18:10. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Liu, S. (2020). A review for submodular optimization on machine scheduling problems. In Du, D., & Wang, J. (Eds.), *Complexity and Approximation - In Memory of Ker-I Ko*, Vol. 12000 of *Lecture Notes in Computer Science*, pp. 252–267. Springer.
- Liu, Y., Chong, E. K. P., Pezeshki, A., & Zhang, Z. (2020). Submodular optimization problems and greedy strategies: A survey. *Discret. Event Dyn. Syst.*, 30(3), 381–412.
- Mao, J., Niu, M., Jiang, C., Liang, H., Chen, J., Liang, X., Li, Y., Ye, C., Zhang, W., Li, Z., Yu, J., Xu, C., & Xu, H. (2021). One million scenes for autonomous driving: ONCE dataset. In Vanschoren, J., & Yeung, S. (Eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Mirroknii, V. S., & Zadimoghaddam, M. (2015). Randomized composable core-sets for distributed submodular maximization. In Servedio, R. A., & Rubinfeld, R. (Eds.), *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pp. 153–162. ACM.
- Mirzasoleiman, B., Badanidiyuru, A., & Karbasi, A. (2016). Fast constrained submodular maximization: Personalized data summarization. In Balcan, M., & Weinberger, K. Q. (Eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, Vol. 48 of *JMLR Workshop and Conference Proceedings*, pp. 1358–1367. JMLR.org.
- Mirzasoleiman, B., Badanidiyuru, A., Karbasi, A., Vondrak, J., & Krause, A. (2015). Lazier than lazy greedy. In Bonet, B., & Koenig, S. (Eds.), *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pp. 1812–1818. AAAI Press.

- Mirzasoleiman, B., Jegelka, S., & Krause, A. (2018). Streaming non-monotone submodular maximization: Personalized video summarization on the fly. In McIlraith, S. A., & Weinberger, K. Q. (Eds.), *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pp. 1379–1386. AAAI Press.
- Mirzasoleiman, B., Karbasi, A., Sarkar, R., & Krause, A. (2013). Distributed submodular maximization: Identifying representative elements in massive data. In Burges, C. J. C., Bottou, L., Ghahramani, Z., & Weinberger, K. Q. (Eds.), *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pp. 2049–2057.
- Mitzenmacher, M., & Upfal, E. (2017). *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press.
- Nemhauser, G. L., & Wolsey, L. A. (1978). Best algorithms for approximating the maximum of a submodular set function. *Math. Oper. Res.*, 3(3), 177–188.
- Nemhauser, G. L., Wolsey, L. A., & Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions - I. *Math. Program.*, 14(1), 265–294.
- Rangwani, H., Jain, A., Aithal, S. K., & Babu, R. V. (2021). S³vaada: Submodular subset selection for virtual adversarial active domain adaptation. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pp. 7496–7505. IEEE.
- Rossi, R. A., & Ahmed, N. K. (2015). The network data repository with interactive graph analytics and visualization. In Bonet, B., & Koenig, S. (Eds.), *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pp. 4292–4293. AAAI Press.
- Yang, J., & Leskovec, J. (2012). Defining and evaluating network communities based on ground-truth. In Zaki, M. J., Siebes, A., Yu, J. X., Goethals, B., Webb, G. I., & Wu, X. (Eds.), *12th IEEE International Conference on Data Mining, ICDM 2012, Brussels, Belgium, December 10-13, 2012*, pp. 745–754. IEEE Computer Society.