

GREEDYML: A PARALLEL ALGORITHM FOR MAXIMIZING SUBMODULAR FUNCTIONS

SHIVARAM GOPAL*, S M FERDOUS†, HEMANTA K. MAJI*, AND ALEX POTHEN*

Abstract. We describe a parallel approximation algorithm for maximizing monotone submodular functions subject to hereditary constraints on distributed memory multiprocessors. Our work is motivated by the need to solve submodular optimization problems on massive data sets, for practical applications in areas such as data summarization, machine learning, and graph sparsification.

Our work builds on the randomized distributed RANDGREEDI algorithm, proposed by Barbosa, Ene, Nguyen, and Ward (2015). This algorithm computes a distributed solution by randomly partitioning the data among all the processors and then employing a single accumulation step in which all processors send their partial solutions to one processor. However, for large problems, the accumulation step could exceed the memory available on a processor, and the processor which performs the accumulation could become a computational bottleneck.

Here, we propose a generalization of the RANDGREEDI algorithm that employs multiple accumulation steps to reduce the memory required. We analyze the approximation ratio and the time complexity of the algorithm (in the BSP model). We evaluate the new GREEDYML algorithm on three classes of problems, and report results from massive data sets with millions of elements. The results show that the GREEDYML algorithm can solve problems where the sequential GREEDY and distributed RANDGREEDI algorithms fail due to memory constraints. For certain computationally intensive problems, the GREEDYML algorithm can be faster than the RANDGREEDI algorithm. The observed approximation quality of the solutions computed by the GREEDYML algorithm closely matches those obtained by the RANDGREEDI algorithm on these problems.

Keywords: Combinatorial optimization, submodular functions, parallel algorithms, approximation algorithms, data summarization.

Mathematics Subject Classification: 90C27, 68W10.

1. Introduction. We describe a scalable parallel approximation algorithm for maximizing monotone submodular functions subject to hereditary constraints on distributed memory multiprocessors. We build on an earlier distributed approximation algorithm which has limited parallelism and higher memory requirements. Although this problem is NP-hard (the objective function is nonlinear), a GREEDY algorithm that maximizes the marginal gain (defined later) at each step is $(1 - 1/e) \approx 0.63$ -approximate for cardinality constraints and $1/2$ -approximate for matroid constraints; here e is Euler’s number.

Combinatorial optimization with a submodular objective function (rather than a linear objective function) leads to diversity in the computed solution, since at each step the algorithm chooses an element with the least properties in common with the current solution set. A broad collection of optimization problems could be modeled using submodular functions, including data and document summarization [24], load balancing parallel computations in quantum chemistry [9], sensor selection [6], resource allocation [27], active learning [11], interpretability of neural networks [7], influence maximization in social networks [13], diverse recommendation [5] etc. Submodular optimization problems often have efficient approximation algorithms to solve them, since submodular functions have properties that make them discrete analogs of both convex and concave continuous functions. Surveys discussing submodular optimization formulations, algorithms, and computational experiments include Tohidi et al. [28] and Krause and Golovin [14].

Our algorithm builds on the RANDGREEDI framework [2], a state-of-the-art randomized distributed algorithm for monotone submodular function maximization un-

*Purdue University.

†Pacific Northwest National Lab

49 der hereditary constraints, which has an approximation ratio half that of the GREEDY
 50 algorithm. The RANDGREEDI algorithm randomly partitions the data among all the
 51 processors, runs the standard GREEDY algorithm on each partition independently in
 52 parallel, and then executes a *single accumulation step* in which all processors send
 53 their partial solutions to one processor. However, this step could exceed the mem-
 54 ory available on a processor when the memory is small relative to the size of the
 55 data, or when solutions are large. Additionally, this merging step serializes both the
 56 computation and communication and is a bottleneck when scaled to more machines.

57 The new GREEDYML algorithm brings additional parallelism to this step and can
 58 lower the memory and time required to solve the problem. We randomly partition
 59 the data among all the processors, which constitute the leaves of an *accumulation*
 60 *tree*, and then merge partial solutions at multiple levels in the tree. We prove that
 61 the GREEDYML algorithm has a worst-case approximation guarantee of $1/(L + 1)$
 62 of the serial GREEDY algorithm, where L is the total number of accumulation levels
 63 in the accumulation tree. Using the BSP model, we analyze the time complexity of
 64 both computation and communication steps in the GREEDYML and RANDGREEDI
 65 algorithms, and show that the former has lower computation and communication costs
 66 than the latter.

67 We evaluate the parallel algorithms on the maximum k -set cover problem, the
 68 maximum k -vertex dominating set in graphs, and exemplar-based clustering (modeled
 69 by the k -medoid problem); all of these problems arise in data reduction or summa-
 70 rization. We experiment on massive data sets with millions of elements that exceed
 71 the memory constraints (a few GBs) on a single processor.

72 We demonstrate how solutions may be computed using the parallel algorithm
 73 by organizing the accumulation tree to have more levels to adapt to the memory
 74 available on a processor. This strategy also enables us to solve for larger values of
 75 the parameter k in the problems discussed above, which corresponds to the size of
 76 the solution sought. We show that the number of function evaluations on the critical
 77 path of the accumulation tree, and hence the run time, could be reduced when the
 78 parallel algorithm is employed. Also, we do not observe the deterioration in objective
 79 function values expected from the worst-case approximation ratio of the GREEDYML
 80 algorithm, and the observed approximation quality of the computed solutions closely
 81 matches those obtained by the RANDGREEDI algorithm on these problems.

82 2. Background and Related Work.

2.1. Submodular functions. A set function $f: 2^W \rightarrow \mathbb{R}^+$ defined on the power
 set of a ground set W is *submodular* if it satisfies the *diminishing marginal gain*
 property. That is,

$$f(X \cup \{w\}) - f(X) \geq f(Y \cup \{w\}) - f(Y), \text{ for all } X \subseteq Y \subseteq W \text{ and } w \in W \setminus Y.$$

A submodular function f is *monotone* if for every $X \subseteq Y \subseteq W$, we have $f(X) \leq f(Y)$.
 The *constrained submodular maximization* problem maximizes a submodular function
 subject to certain constraints:

$$\max f(S) \text{ subject to } S \in \mathcal{C}, \text{ where } \mathcal{C} \subseteq 2^W \text{ is the family of feasible solutions.}$$

83 We consider *hereditary constraints*: i.e., for every set $S \in \mathcal{C}$, every subset of S is
 84 also in \mathcal{C} . The hereditary family of constraints includes various common ones such
 85 as cardinality constraints ($\mathcal{C} = \{A \subseteq W : |A| \leq k\}$) and matroid constraints (\mathcal{C}
 86 corresponds to the collection of independent sets of a matroid).

2.2. Lovász extension. For the analysis of our algorithm, we use the *Lovász extension* [20], a relaxation of submodular functions. A submodular function f can be viewed as a function defined over the vertices of the unit hypercube, $f : \{0, 1\}^n \rightarrow \mathbb{R}^+$, by identifying sets $V \subseteq W$ with binary vectors of length $n = |W|$ in which the i^{th} component is 1 if $i \in V$, and 0 otherwise. The *Lovász extension* [20] $\widehat{f} : [0, 1]^n \rightarrow \mathbb{R}^+$ is a convex extension that extends f over the entire hypercube, which is given by

$$\widehat{f}(x) = \mathbb{E}_{\theta \in \mathcal{U}[0,1]} [f(\{i : x_i \geq \theta\})].$$

87 Here, θ is uniformly random in $[0, 1]$. For any submodular function f , the Lovász
88 extension \widehat{f} satisfies the following properties [20]:

- 89 1. $\widehat{f}(1_S) = f(S)$, for all $S \subseteq V$ where $1_S \in [0, 1]^n$ is a vector containing 1 for
90 the elements in S and 0 otherwise,
- 91 2. $\widehat{f}(x)$ is convex, and
- 92 3. $\widehat{f}(c \cdot x) \geq c \cdot \widehat{f}(x)$, for any $c \in [0, 1]$.

93 An α -*approximation* algorithm (where $\alpha \in [0, 1)$) for maximizing a submodular
94 function $f : 2^W \rightarrow \mathbb{R}^+$ subject to a hereditary constraint \mathcal{C} produces a solution $S \subseteq W$
95 with $S \in \mathcal{C}$, satisfying $f(S) \geq \alpha \cdot f(S^*)$, where S^* is an optimal solution.

96 **2.3. The GREEDI and RANDGREEDI Algorithms.** The GREEDI algorithm
97 (shown in Algorithm 2.1) for maximizing submodular functions subject to constraints
98 is an iterative algorithm that starts with an empty solution. Given the current so-
99 lution, an element is *feasible* if it can be added to the solution without violating the
100 constraints. In each iteration, the GREEDI algorithm chooses a feasible element $e \in V$
101 that maximizes the marginal gain, $f(S \cup \{e\}) - f(S)$, w.r.t. the current solution S .
102 The algorithm terminates when the maximum marginal gain is zero or all elements
103 in the ground set have been considered.

Algorithm 2.1 GREEDI Algorithm

```

1: procedure GREEDI ( $V$ : Dataset)
2:    $S \leftarrow \emptyset$ 
3:   while True do
4:      $E \leftarrow \{e \in V \setminus S : S \cup \{e\} \in \mathcal{C}\}$ 
5:      $e' \leftarrow \arg \max_{e \in E} f(S \cup \{e\})$ 
6:     if  $f(S \cup \{e'\}) = f(S)$  or  $E = \emptyset$  then
7:       break
8:     end if
9:      $S \leftarrow S \cup \{e'\}$ 
10:  end while
11:  return  $S$ 
12: end procedure

```

104 We now discuss the GREEDI and RANDGREEDI, which are the state-of-the-art
105 distributed algorithms for constrained submodular maximization. The GREEDI algo-
106 rithm [24] partitions the data *arbitrarily* on available machines, and on each partition,
107 it runs the GREEDI algorithm in parallel to compute a *local* solution. These solutions
108 are then sent to a single *global* machine, where they are accumulated. The GREEDI
109 algorithm is then again executed on the accumulated data to get a global solution.
110 The final solution is the best solution among all the local and global solutions. For
111 a cardinality constraint, where k is the solution size, the GREEDI algorithm has a

Algorithm 2.2 RANDGREEDI framework for maximizing constrained submodular function

```

1: procedure RANDGREEDI( $V$ : Dataset,  $m$ : number of machines)
2:    $S \leftarrow \emptyset$ 
3:   Let  $\{P_0, P_1, \dots, P_{m-1}\}$  be an uniform random partition of  $V$ .
4:   Run GREEDY( $P_i$ ) on each machine  $i \in [0, m - 1]$  to compute the solution  $S_i$ 
5:   Place  $S = \bigcup_i S_i$  on machine 0
6:   Run GREEDY( $S$ ) to compute the solution  $T$  on machine 0
7:   return  $\arg \max \{f(T), f(S_1), f(S_2), \dots, f(S_{m-1})\}$ 
8: end procedure

```

112 worst-case approximation guarantee of $1/\Theta(\min(\sqrt{k}, m))$, where m is the number of
 113 machines.

114 Although the GREEDI algorithm performs well in practice [24], its approximation
 115 ratio is not a constant but depends on k . To improve the approximation guarantee
 116 of GREEDI algorithm, Barbosa et al. proposed the RANDGREEDI algorithm [2]. By
 117 partitioning the data uniformly at random on machines, RANDGREEDI achieves an
 118 expected approximation guarantee of $\frac{1}{2}(1 - 1/e)$ for cardinality and $1/4$ for matroid
 119 constraints. In general, it has an approximation ratio of $\alpha/2$ where α is the approx-
 120 imation ratio of the GREEDY algorithm used at the local and global machines. We
 121 present the pseudocode of RANDGREEDI framework in Algorithm 2.2. Note that for
 122 a cardinality constraint, both GREEDI and RANDGREEDI perform $O(nk(k + m))$ calls
 123 to the objective function and has $O(mk)$ elements communicated to the single central
 124 machine where n is the number of elements in the ground set, m is the number of
 125 machines, and k is solution size.

126 Both GREEDI and RANDGREEDI require a single global accumulation from the
 127 solutions generated in local machines. This single accumulation step can quickly
 128 become dominating since the runtime, memory, and complexity of this global aggre-
 129 gation grows linearly with the number of machines. We propose to alleviate this by
 130 introducing a hierarchical aggregation strategy that maintains an accumulation tree.
 131 Our GREEDYML framework generalizes the RANDGREEDI from a single accumula-
 132 tion to a multi-level accumulation. The number of partial solutions to be aggregated
 133 depends on the branching factor of the tree, which can be a constant. Thus, the num-
 134 ber of accumulation levels grows logarithmically with the number of machines, and
 135 the total aggregation is not likely to become a memory, runtime, and communication
 136 bottleneck with the increase in the number of machines.

137 **2.4. Other Related Work.** Kumar et al. [17] have developed the sample and
 138 prune algorithm which achieves an expected approximation ratio of $1/(2 + \epsilon)$ for
 139 k -cardinality constraints, using $O(1/\delta)$ rounds, when the memory per machine is
 140 $O(kn^\delta \log n)$, where $\delta > 0$ is a parameter, and n is the number of elements in the
 141 ground set. Barbosa et al. [2] have compared their RANDGREEDI algorithm with this
 142 one and show that the former performs better than the latter for the practical quality
 143 of the computed approximate solution. They observed this even though the sample
 144 and prune algorithm has a better worst-case approximation ratio in expectation.

145 More recent work on distributed submodular maximization uses the multi-linear
 146 extension to map the submodular optimization problem into a continuous domain.
 147 This line of work [4, 25, 26] typically performs a gradient ascent on each local ma-
 148 chine and builds a consensus solution in each round, which improves the approxima-




Fig. 1: An accumulation tree with $L = 2$ levels, $m = b^2$ machines, and a branching factor b . Each node has a label of the form (ℓ, id) . Here there are b nodes as children at each level, but when there are fewer than b^L leaf nodes, then the number of children at levels closer to the root may be fewer than b .

149 tion factor to $(1 - 1/e)$. However, we believe these latter papers represent primarily
 150 a theoretical contribution rather than one that leads to practical algorithms. The
 151 reason is the high (exponential) cost of computing a single gradient by sampling
 152 many points; even randomized approximations of gradient computations are expen-
 153 sive. Most of these algorithms are not implemented and the ones with implementation
 154 solve problems with only a hundred elements in the data set [26].

155 **3. Description of Our Algorithm.** We describe and analyze our algorithm
 156 that generalizes the RANDGREEDI algorithm from a single accumulation step to mul-
 157 tiple accumulation steps. Each accumulation step corresponds to a level in an *accu-*
 158 *mutation tree* which we describe next. We assume that there are m machines identified
 159 by the set of ids: $\{0, 1, \dots, m - 1\}$.

160 **Accumulation Tree.** An accumulation tree (T) is defined by the number of machines
 161 (m), and branching factor (b). It has the same structure as a complete b -ary tree with
 162 m leaves which means all the leaves are at the same depth. The tree nodes correspond
 163 to processors along with the corresponding subset of data accessible to them. The
 164 edges of the tree determine the accumulation pattern of the intermediate solutions.
 165 The final solution is generated on the root node of T . Thus, the branching factor b
 166 of the tree indicates the maximum number of nodes that send data to its parent. For
 167 each internal node of the tree, we attempt to have exactly b children. Note that since
 168 we plan to construct a complete b -ary tree, in the case where m is not multiple of b , in
 169 each level of the tree, there could be at most one node whose arity is less than b . The
 170 number of accumulation levels, L (i.e., the height of the tree minus 1) is $\lceil \log_b m \rceil$.

171 To uniquely identify a node in the tree, we will assign an identifier (ℓ, id) to each
 172 node of T , where ℓ represents the accumulation level of the node and id represents
 173 the machine id corresponding to the node. The id for each leaf node is the id of
 174 the machine that the leaf node corresponds to. All the leaf nodes are at level 0.
 175 Each internal node receives the lowest id of its children, i.e., any node (l, i) has node
 176 $(l + 1, \lfloor i/b^{l+1} \rfloor * b^{l+1})$ as the parent. Therefore the root node will always have level L
 177 with id value 0. Also, we characterize an accumulation tree T by the triple $T(m, L, b)$,
 178 where m is the number of leaves (machines), L is the number of levels, and b is the
 179 branching factor.

180 Figure 1 shows an example of a generic accumulation tree with b^2 leaves and
 181 branching factor b . The number of accumulation levels is the level of the root. Here

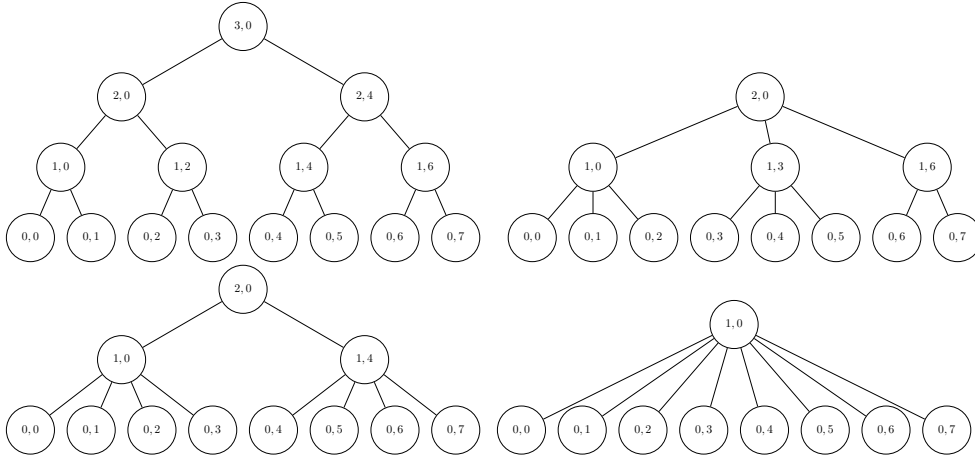


Fig. 2: Accumulation tree with 8 machines and branching factors 2 (top-left), 3 (top-right), 4 (bottom-left), and 8 (bottom-right). The level inside a node represents the identification of the node.

182 we have $L = \lceil \log_b b^2 \rceil = 2$. Figure 2 shows accumulation trees with 8 leaves and
 183 with branching factors 2, 3, 4, and 8. The trees with branching factors 2 and 8 have
 184 the same branching factor for every internal node as these trees have b^L nodes. But
 185 the tree with branching factor 3 has the last node in level 1 with only 2 children.
 186 Similarly, the tree with branching factor 4 has the last node in level 2 i.e. the root
 187 with 2 children. Observe that the id parameter remains the same in multiple nodes
 188 that are involved in computations at multiple levels. For this paper, we show analysis
 189 by keeping the branching factor constant across all levels.

190 **Data Accessibility.** We use P_{id} to denote the elements assigned to machine id . To
 191 indicate the data *accessible* to a particular node in the tree, we describe a set for the
 192 input data set as $V_{\ell, id}$. It corresponds to all the data used to compute the solution at
 193 the node (ℓ, id) and consists of all the elements assigned to its descendants:

$$194 \quad V_{\ell, id} = \bigcup_{i=0}^{\min(b^\ell - 1, m - id)} P_{id+i}.$$

195

196 **Randomness.** The randomness in the algorithm is *only* in the initial placement of the
 197 data on the machines, and we use a random tape to encapsulate this. The random
 198 tape r_W has a randomized entry for each element in W to indicate the machine
 199 containing that element. Any expectation results proved henceforth are over the
 200 choice of this random tape. Moreover, if the data accessible to a node is V , we
 201 consider the randomness over just r_V . Whenever the expectation is over r_V , we
 202 denote the expectation as \mathbb{E}_V .

203 **Recurrence Relation.** Figure 3 shows a recurrence relation defined for every node
 204 in the accumulation tree and will be the basis for our multilevel distributed algorithm.
 205 At level 0 (leaves), the recurrence function returns the GREEDY solution of the random
 206 subset of data P_{id} assigned to it. At other levels (internal nodes), it returns the better
 207 among the GREEDY solution computed from the union of the received solution sets
 208 of its children and its solution from its previous level. It is undefined for (ℓ, id) tuples

$$\text{GREEDYML}(\ell, id) = \begin{cases} \text{GREEDY}(P_{id}) & \ell = 0 \\ \arg \max \begin{cases} \text{GREEDY} \left(\bigcup_{i \in \{0, 1, \dots, b-1\}} \text{GREEDYML}(\ell - 1, id + i \cdot b^{\ell-1}) \right) \\ \text{GREEDYML}(\ell - 1, id) \end{cases} & id \bmod b^\ell = 0 \\ \text{undefined} & \text{otherwise} \end{cases}$$

Fig. 3: The recurrence relation for the multilevel GREEDYML which is defined for each node in the accumulation tree. We denote the random subset assigned to machine id by P_{id} .

209 that do not correspond to nodes in the tree (at higher levels). We call our algorithm
 210 associated with the recurrence relation as the GREEDYML algorithm.

211 We can compare it with the RANDGREEDI algorithm by looking at the recurrence
 212 relation at level one. Our recurrence relation takes the $\arg \max$ for the accumulated
 213 solution and *one* solution from the previous level. However, the RANDGREEDI algo-
 214 rithm takes the $\arg \max$ of the accumulated solution and the *best* solution from the
 215 children. Our choice reduces the computation at the internal node. We show that this
 216 modification produces the same approximation ratio as the RANDGREEDI algorithm.
 217 **Pseudocode.** Algorithm 3.1 describes our multilevel distributed algorithm in two
 218 functions. The first function GREEDYML is a wrapper function that sets up the
 219 environment to run the distributed algorithm. The second function GREEDYML' is
 220 the iterative implementation of the recurrence relation that runs on each machine.
 221 The wrapper function partitions the data into m subsets and assigns them to the
 222 machines (Line 2). Then each machine runs the GREEDYML' function on the subset
 223 assigned to it (Line 5, Line 7). The wrapper function uses and returns the solution
 224 from machine 0 (Line 8) as it is the root of the accumulation tree.

225 The GREEDYML' procedure is an iterative implementation of the recurrence re-
 226 lation 3 that runs on every machine. Each machine checks whether it needs to be
 227 active at a particular level (Line 5) and decides whether it needs to receive from (Line
 228 11) or send to other machines (Line 6). The function returns the solution from the
 229 last level of the machine.

230 **4. Analysis of Our Algorithm.** In this section, we will derive the expected
 231 approximation ratio of the GREEDYML algorithm. We will then describe the three
 232 submodular functions we experiment with and derive their computation and commu-
 233 nication complexities.

234 **4.1. Expected Approximation Ratio.** This subsection proves the expected
 235 approximation ratio of our GREEDYML algorithm in Theorem 4.4. To do so, we need
 236 three Lemmas. The first Lemma characterizes elements that do not change the solu-
 237 tion computed by the GREEDYML algorithm. We need some preliminary notation.
 238 When we wish to indicate the data set that a node in the tree and its descendants work
 239 with, we add an argument to GREEDYML(ℓ, id), and write GREEDYML($V_{\ell, id}, \ell, id$).
 240 When we perform a union operation on this data set with some set B , and execute the
 241 GREEDYML algorithm on the union, i.e., GREEDYML($V_{\ell, id} \cup B, \ell, id$), then elements
 242 in B are assigned randomly to the leaves of the subtree rooted at node (ℓ, id) and the
 243 algorithm is run with the updated data sets. Lemma 4.1 compares executions of the
 244 algorithm when this union operation is performed for a special set B . It states that
 245 if adding an individual element of a set, B , to the input of the GREEDYML does not
 246 change the solution set then adding the whole set, B to the input will also have no
 247 effect on the solution. Here we consider executions that use the same random tape,
 248 number of machines, and branching factor. We use the same random tape to couple

Algorithm 3.1 Our Randomized Multi-level GREEDYML Algorithm

```

1: procedure GREEDYML( $V$ : Dataset,  $b$ : branching factor,  $m$ : number of machines,  $r$ :
   random tape)
2:   Let  $\{P_0, P_1, \dots, P_{m-1}\}$  be uniform random partition of  $V$  using  $r$ .
3:   for  $i = 1 \dots m - 1$  in parallel do       $\triangleright$  Run GREEDYML' on all machines except 0
4:      $\ell = \text{level}(i, b)$                        $\triangleright \text{level}(i, b) = \max\{l : id \bmod b^l \text{ is } 0\}$ 
5:     Run GREEDYML'( $V_i, \ell, b, i$ ) to obtain  $S_i$  on machine  $i$ 
6:   end for
7:   Run GREEDYML'( $V_0, \lceil \log_b m \rceil, b, 0$ ) to obtain  $S_0$  on machine 0
8:   return  $S_0$ 
9: end procedure

1: procedure GREEDYML'( $P$ : Partial Data-set,  $\ell$ : levels;  $b$ : branching factor,  $id$ : machine
   ID)
2:    $S = \text{GREEDY}(P)$ 
3:    $S_{prev} = S$ 
4:   for  $i = 1 \dots \ell$  do
5:     if  $id \neq \text{parent}(id, i)$  then
6:       Send  $S_{prev}$  to  $\text{parent}(id, i)$            $\triangleright \text{parent}(id, i) = b^i \cdot \lfloor id/b^i \rfloor$ 
7:       break
8:     end if
9:      $D = S_{prev}$                                  $\triangleright$  Prepare  $D$  for current iteration
10:    for  $j = 1 \dots b - 1$  do
11:      Receive  $D_j$  from  $\text{child}(id, i, j)$        $\triangleright \text{child}(id, i, j) = id + j \cdot b^{i-1}$ 
12:       $D = D \cup D_j$ 
13:    end for
14:    Run GREEDY( $D$ ) to obtain  $S$ 
15:     $S_{prev} = \arg \max\{f(S), f(S_{prev})\}$ 
16:  end for
17:  return  $S_{prev}$ 
18: end procedure

```

249 the executions. Therefore, the result of the Lemma is not over the expectation of the
250 random tape.

251 LEMMA 4.1. *Let $T(m, L, b)$ be an accumulation tree. Consider a universal set*
252 *W , and a random tape r_W that maps elements of W to the leaves of T . Let $V \subseteq W$*
253 *denote the set of elements accessible to a node (ℓ, id) , and consider adding elements of*
254 *$B \subseteq W$ to this node. If we have $\text{GREEDYML}(V \cup \{e\}, \ell, id) = \text{GREEDYML}(V, \ell, id)$,*
255 *for each element $e \in B$, then $\text{GREEDYML}(V \cup B, \ell, id) = \text{GREEDYML}(V, \ell, id)$.*

256 Note: Function calls in this analysis use the same random tape for assigning
257 elements; hence elements are assigned uniformly at random to the machines, but they
258 use the same random assignment in all runs involving V ; $V \cup \{e\}, \forall e \in B$; and $V \cup B$.

259 *Proof.* If possible, let $\text{GREEDYML}(V \cup B, \ell, id) \neq \text{GREEDYML}(V, \ell, id)$. Let e
260 be the first element of B to be selected by the GREEDY algorithm at the final level.
261 Let i be the level in which e was in the input in some machine but not selected in a
262 solution for the next level in $\text{GREEDYML}(V \cup \{e\}, \ell, id)$. Since e is the first element
263 of B that was selected by the GREEDY algorithm, the elements chosen before it at
264 level i in $\text{GREEDYML}(V \cup B, \ell, id)$ are the same ones chosen before it at level i in
265 $\text{GREEDYML}(V \cup \{e\}, \ell, id)$. Since it was not selected in $\text{GREEDYML}(V \cup \{e\}, \ell, id)$ it
266 will not be selected in $\text{GREEDYML}(V \cup B, \ell, id)$. This is a contradiction since e must

267 be selected at every level to be present in the final solution. \square

268 Now we turn to the two Lemmas that provide bounds on the quality of the
 269 computed solutions in terms of the optimal solution at an internal node in the accu-
 270 mulation tree.

271 Lemma 4.2 provides a lower bound on the expected function value of the solution
 272 of the GREEDYML algorithm from a child of the internal node. Lemma 4.3 provides
 273 a lower bound on the expected function value of the solution set from the GREEDY
 274 algorithm executed at each internal node on the union of the partial solutions from
 275 its children. These Lemmas depend on the probability distribution defined below.

276 Let $p_{\ell,id}: V_{\ell,id} \rightarrow [0, 1]$ be a probability distribution over the elements in $V_{\ell,id}$,
 277 which we shall define below. Here $A \sim V_{\ell,id}(1/b)$ denotes a random subset of $V_{\ell,id}$ such
 278 that each element is independently present in A with probability $1/b$. This probability
 279 corresponds to the distribution from the random tape because each element is present
 280 with the same likelihood from any child of the node. Let $OPT_{\ell,id}$ be an optimal
 281 solution of the constrained submodular maximization problem when the input data
 282 is $V_{\ell,id}$.

283 The probability $p_{\ell,id}$ is defined as follows:

$$284 \quad p_{\ell,id}(e) = \begin{cases} \Pr_{A \sim V_{\ell,id}(1/b)} [e \in \text{GREEDYML}(A \cup \{e\}, \ell - 1, id)], & \text{if } e \in OPT_{\ell,id}; \\ 0, & \text{otherwise.} \end{cases}$$

285 For any internal node (ℓ, id) , the distribution $p_{\ell,id}$ defines the probability that each
 286 element of $OPT_{\ell,id}$

287 is in the solution of the GREEDYML algorithm of a child when it is accessible to
 288 the child node.

289 Next, we state and prove Lemma 4.2 that relates the expected solution of the
 290 GREEDYML algorithm at a child node with the optimal solution at the node when
 291 the approximation ratio of the GREEDYML algorithm at the child is β .

LEMMA 4.2. *Let $c = (\ell - 1, id_c)$ be a child of an internal node $n = (\ell, id)$ of the
 accumulation tree. Let S_c be the solution computed from child c , and $V_c \subset V_n$ denote
 the elements considered in forming S_c . If $\mathbb{E}_{V_c}[f(S_c)] \geq \beta \cdot f(OPT_{\ell-1, id_c})$, then*

$$\mathbb{E}_{V_n}[f(S_c)] \geq \beta \cdot \widehat{f}(1_{OPT_{\ell, id}} - p_{\ell, id}).$$

Proof. We first construct a subset of $OPT_{\ell, id}$ that contains all the elements that
 do not appear in S_c when added to some leaf node in the subtree rooted at child c .
 Let O_c be the rejected set that can be added to V_c without changing S_c ; i.e.,

$$O_c = \{e \in OPT_{\ell, id} : e \notin \text{GREEDYML}(V_c \cup \{e\}, \ell', id)\}.$$

292 To clarify further, O_c is a randomized set dependent on the tape $r_{V_{\ell, id}}$. Since the
 293 distribution of V_c is the same as $V_{\ell, id}(1/b)$ for each element e in $OPT_{\ell, id}$,

$$294 \quad (4.1) \quad \Pr[e \in O_c] = 1 - \Pr[e \notin O_c] = 1 - p_{\ell, id}(e).$$

295 From Lemma 4.1, we know that $\text{GREEDYML}(V_c \cup O_c, \ell - 1, id_c) = \text{GREEDYML}(V_c, \ell -$
 296 $1, id_c)$. Since the rejected set $O_c \subseteq OPT_{\ell, id}$ and the constraints are hereditary, $O_c \in \mathcal{C}$
 297 (i.e. O_c is a feasible solution of child node c).

$$298 \quad (4.2) \quad f(OPT_{\ell-1, id_c}) \geq f(O_c).$$

299 Then from the condition of Lemma 4.2, we have

$$\begin{aligned}
300 \quad \mathbb{E}_{V_n}[\mathbb{E}_{V_c} f(S_c)] &\geq \beta \cdot \mathbb{E}_{V_n}[f(OPT_{\ell-1, id_c})] \\
301 \quad \mathbb{E}_{V_n}[f(S_c)] &\geq \beta \cdot \mathbb{E}_{V_n}[f(OPT_{\ell-1, id_c})] && [V_c \subset V_n] \\
302 \quad &\geq \beta \cdot \widehat{f}(\mathbb{E}_{V_n}[1_{O_c}]) && [\text{Eqn. (4.2)}] \\
303 \quad &\geq \beta \cdot \mathbb{E}_{V_n}[f(O_c)] && [\text{Lovász (2), 2.2}] \\
304 \quad &\geq \beta \cdot f^-(\mathbb{E}_{V_n}[1_{O_c}]) = \beta \cdot \widehat{f}(1_{OPT_{\ell, id}} - p_{\ell, id}) && [\text{Eqn. (4.1)}. \quad \square]
\end{aligned}$$

306 Now we show how the solution of the GREEDY algorithm that runs at each internal
307 node compares with the optimal solution at the internal node.

308 LEMMA 4.3. *For an internal node $n = (\ell, id)$, let D be the union of all the
309 solutions computed by the children of node n in the accumulation tree. Let $S =$
310 GREEDY(D) be the solution from the Greedy algorithm on the set D . If GREEDY is
311 an α -approximate algorithm, then*

$$312 \quad \mathbb{E}_{V_n}[f(S)] \geq \alpha \cdot f^-(p_{\ell, id}).$$

Proof. We first show a preliminary result on the union set D . Consider an element
 $e \in D \cap OPT_{\ell, id}$ present in some solution S_c from a child c . Then,

$$\Pr[e \in S_c | e \in V_c] = \Pr[e \in \text{GREEDYML}(V_c, \ell - 1, id) | e \in V_c].$$

Since the distribution of $V_c \sim V_{\ell, id}(1/b)$ conditioned on $e \in V_c$ is identical to the
distribution of $B \cup \{e\}$, where $B \sim V_{\ell, id}(1/b)$, we have,

$$\Pr[e \in S_c | e \in V_c] = \Pr_{B \sim V_{\ell, id}(1/b)}[e \in \text{GREEDYML}(B \cup \{e\}, \ell - 1, id)] = p_{\ell, id}(e).$$

313 Since this result holds for every child c , and each subset V_c is disjoint from the
314 corresponding subsets mapped to the other children, we have

$$315 \quad (4.3) \quad \Pr(D \cap OPT_{\ell, id}) = p_{\ell, id}.$$

316 Now, we are ready to prove the Lemma. The subset $D \cap OPT_{\ell, id} \in \mathcal{C}$, since it is
317 a subset of $OPT_{\ell, id}$ and the constraints are hereditary. Further, since the GREEDY
318 algorithm is α -approximate, we have

$$\begin{aligned}
319 \quad &f(S) \geq \alpha \cdot f(D \cap OPT_{\ell, id}) \\
320 \quad \mathbb{E}_{V_n}[f(S)] &\geq \mathbb{E}_{V_n}[\alpha \cdot f(D \cap OPT_{\ell, id})] \\
321 \quad &\geq \alpha \cdot \mathbb{E}_{V_n}[f(D \cap OPT_{\ell, id})] \\
322 \quad &\geq \alpha \cdot \widehat{f}(\mathbb{E}_{V_n}[1_{D \cap OPT_{\ell, id}}]) && [\text{Lovász Ext. (2), 2.2}] \\
323 \quad (4.4) \quad &= \alpha \cdot \widehat{f}(p_{\ell, id}). && [\text{Eqn. 4.3}. \quad \square]
\end{aligned}$$

THEOREM 4.4. *Let $T(m, L, b)$ be an accumulation tree. For a universal set W and
random tape r_W that maps elements of W to the leaves of the tree T , let $V_{\ell, id} \subseteq W$
denote the subset of W accessible to a node (ℓ, id) . Let $OPT_{\ell, id}$ be an optimal solution
computed from the subset $V_{\ell, id}$ for the submodular function f with constraints \mathcal{C} . If
GREEDY is an α -approximate algorithm, then*

$$\mathbb{E}_{V_{\ell, id}}[f(\text{GREEDYML}(V_{\ell, id}, \ell, id))] \geq \frac{\alpha}{(\ell + 1)} f(OPT_{\ell, id}).$$

325 *Proof.* We prove this theorem by induction on the level ℓ .

326 **Base case**, $\ell = 0$: Here, there is no accumulation step, and we obtain the solution
 327 from a single node. Thus we run the GREEDY algorithm on $V_{\ell, id}$. The result follows
 328 since the GREEDY algorithm has the approximation ratio α .

329 **Inductive case**, $\ell = \ell' + 1$: We first obtain a relation for the quality of the
 330 solutions at level ℓ' compared to the quality of an optimal solution.

331 For each child c , let S_c be a solution computed by the GREEDYML algorithm
 332 from the data $V_c \subset V_{\ell, id}$;

333 From the **induction hypothesis** applied to child $c = (\ell', id)$, the approximation
 334 ratio obtained as a result of the computation $\text{GREEDYML}(V_c, \ell', id)$ is $\alpha/(\ell'+1) = \alpha/\ell$.
 335 This implies that $\mathbb{E}_{V_{\ell', id}}[f(S_c)] \geq \frac{\alpha}{\ell} \cdot f^-(1_{OPT_{\ell', id}})$. Therefore we can apply Lemma 4.2
 336 to get

$$337 \quad (4.5) \quad \mathbb{E}_{V_{\ell, id}}[f(S_c)] \geq \frac{\alpha}{\ell} \cdot \widehat{f}(1_{OPT_{\ell, id}} - p_{\ell, id}).$$

338 After obtaining the solutions from the children, we get the solution S computed
 339 by the GREEDY algorithm on the union of these solution sets. From Lemma 4.3, we
 340 have

$$341 \quad (4.6) \quad \mathbb{E}_{V_{\ell, id}}[f(S)] \geq \alpha \cdot \widehat{f}(p_{\ell, id}).$$

342 Now we obtain the relation between the solution at level $\ell' + 1$ and the optimal
 343 solution. Let the solution set at level $\ell' + 1$ be T . We have $T = \arg \max\{f(S), f(S_c)\}$.
 344 Then, we can use the lower bounds calculated earlier in Eqn. 4.5 and Eqn. 4.6 to find
 345 lower bounds for T .

$$346 \quad \mathbb{E}_{V_{\ell, id}}[f(T)] \geq \alpha \cdot \widehat{f}(p_{\ell, id}) \text{ and } \mathbb{E}_{V_{\ell, id}}[f(T)] \geq \frac{\alpha}{\ell} \cdot \widehat{f}(1_{OPT_{\ell, id}} - p_{\ell, id}).$$

347 By multiplying the second inequality by ℓ and then adding it to the first, we get

$$348 \quad (\ell + 1)\mathbb{E}_{V_{\ell, id}}[f(T)] \geq \alpha \cdot (\widehat{f}(1_{OPT_{\ell, id}} - p_{\ell, id}) + \widehat{f}(p_{\ell, id}))$$

$$349 \quad = \alpha \cdot \widehat{f}(1_{OPT_{\ell, id}}). \quad [\text{Lovász Ext. (2), 2.2}]$$

351 Dividing by $\ell + 1$, and substituting from Lovász Ext. (1), 2.2 we conclude that the
 352 algorithm is $\alpha/(\ell + 1)$ -approximate. \square

353 **4.2. Submodular Functions and Complexity.** Here, we describe three sub-
 354 modular functions that we consider in our experiments and then discuss their com-
 355 putational and communication complexities.

356 Our algorithm can handle any hereditary constraint, but in our experiments, we
 357 consider only cardinality constraints to keep the computations simple. (More general
 358 constraints involve additional computations to check if an element can be added to the
 359 current solution set and satisfy the constraints.) Cardinality constraints are widely
 360 used in various applications such as sensor placement [16], text, image, and document
 361 summarization [18, 19], and information gathering [15]. The problem of maximizing
 362 a submodular function under cardinality constraints can be expressed as follows.

$$363 \quad \max_{S \subseteq V} f(S)$$

$$364 \quad \text{s.t.} \quad |S| \leq k.$$

366 Here V is the ground set, f is a non-negative monotone submodular function, and k
 367 is the size of the solution set S .

368 In our experiments, We have considered the following three submodular functions.

369 **k -cover.** The first problem we consider is the k -cover. Given a ground set B ,
 370 a collection of subsets $V \subseteq 2^B$, and an integer k , the goal is to select a set $S \subseteq V$
 371 containing k of these subsets to maximize $f(S) = |\bigcup_{S_i \in S} S_i|$.

372 **k -dominating set.** The k -dominating set problem is a special case of the k -cover
 373 problem defined on graphs with the ground set V as the set of vertices. We say a vertex
 374 $u \in V$ dominates all its adjacent vertices (denoted by $\delta(u)$). Our goal is to select a
 375 set S of k vertices to dominate as many vertices as possible, i.e., $f(S) = |\bigcup_{u \in S} \delta(u)|$.
 376 The marginal gain of any vertex is the number of vertices in its neighborhood that
 377 are not yet dominated. Therefore the problem shows diminishing marginal gains and
 378 is submodular.

379 **k -medoid problem.** The k -medoid problem [12] is used to compute exemplar-
 380 based clustering, where we want to select a set of exemplars (cluster centers) that are
 381 representatives of a large dataset. Given a collection of elements in a ground set V ,
 382 and a dissimilarity $d(u, v)$, we define a loss function (denoted by L) as the average
 383 pairwise dissimilarity between the exemplars (S) and the elements of the data set,
 384 i.e., $L(S) = \frac{1}{|V|} \sum_{u \in V} \min_{v \in S} d(u, v)$. Following [24], we turn this loss minimization
 385 to a submodular maximization problem by setting $f(S) = L(\{e_0\}) - L(S \cup \{e_0\})$, where
 386 e_0 is an auxiliary element specific to the dataset. The goal is to select a subset $S \subseteq V$
 387 of size k representing the exemplars that maximize $f(S)$.

388 Next, we will analyze the computational and communication complexity of our
 389 GREEDYML algorithm using the bulk synchronous parallel (BSP) model of parallel
 390 computation [29]. For the analysis, we will denote the number of elements in the
 391 ground set by $n = |V|$, the solution size by k , the number of machines by m , and the
 392 number of levels in the accumulation tree by L .

393 **Computational Complexity.** The number of objective function calls by the se-
 394 quential GREEDY algorithm (shown in Algorithm 2.1) is $O(nk)$, since k elements are
 395 selected to be in the solution, and we may need to compute $O(n)$ marginal gains for
 396 each of them. Each machine in RANDGREEDI algorithm makes $O(k(n/m + mk))$ func-
 397 tion calls, where the second term comes from the accumulation step. Each machine
 398 of the GREEDYML algorithm with branching factor b makes $O(k(n/m + Lbk))$ calls.
 399 Recall that $L = \lceil \log_b m \rceil$.

400 We note that the time complexity of a function call depends on the specific func-
 401 tion being computed. For example, in the k -coverage and the k -dominating set prob-
 402 lems, computing a function costs $O(\delta)$, where δ is the size of the largest itemset for
 403 k -coverage, and the maximum degree of a vertex for the vertex dominating set. In
 404 both cases, the runtime complexity is $O(\delta k(n/m + mk))$ for the RANDGREEDI, and
 405 $O(\delta k(n/m + Lbk))$ for the GREEDYML algorithm. The k -medoid problem computes
 406 a local objective function value and has a complexity of $O(n'\delta)$ where δ is the number
 407 of features, and n' is the number of elements present in the machine. For the leaves of
 408 the accumulation tree, $n' = n/m$, and for interior nodes, $n' = bk$. Therefore its com-
 409 plexity is $O(k\delta((n/m)^2 + (mk)^2))$ for the RANDGREEDI, and $O(k\delta((n/m)^2 + L(bk)^2))$
 410 for the GREEDYML algorithm.

411 **Communication Complexity.** Each edge in the accumulation tree represents com-
 412 munication from a machine at a lower level to one at a higher level and contains four
 413 messages. They are the indices of the selected elements of size k , the size of the data
 414 associated with each selection (proportional to the size of each adjacency list ($\leq \delta$), the
 415 total size of the data elements, and the data associated with each selection. Therefore

416 the total volume of communication is $O(k\delta)$ per child. Since at each level, a parent
 417 node receives messages from b children, the communication complexity is $O(k\delta Lb)$
 418 for each parent. Therefore the communication complexity for the RANDGREEDI algo-
 419 rithm is $O(k\delta m)$ and for the GREEDYML algorithm is $O(k\delta L \lceil m^{1/L} \rceil)$. We summarize
 420 these results in Table 1.

Algorithms	Metric	GREEDY	RANDGREEDI	GREEDYML
All	Elements per leaf node	n	n/m	n/m
	Calls per leaf node	nk	nk/m	nk/m
	Elements per interior node	0	km	$k \lceil m^{1/L} \rceil$
	Calls per interior node	0	k^2m	$k^2 \lceil m^{1/L} \rceil$
	Total Function Calls	kn	$k(n/m + km)$	$k(n/m + Lk \lceil m^{1/L} \rceil)$
k -cover / k -dominating set	Cost Per call	δ	δ	δ
	Computational complexity	δkn	$\delta k(n/m + km)$	$\delta k(n/m + Lk \lceil m^{1/L} \rceil)$
	Communication cost	0	δkm	$\delta kL \lceil m^{1/L} \rceil$
k -medoid	Cost Per call in Leaf node	δn	$\delta n/m$	$\delta n/m$
	Cost Per call in interior node	0	δkm	$\delta k \lceil m^{1/L} \rceil$
	Computational complexity	δkn^2	$\delta k((n/m)^2 + (km)^2)$	$\delta k((n/m)^2 + L(k \lceil m^{1/L} \rceil)^2)$
	Communication cost	0	δkm	$\delta kL \lceil m^{1/L} \rceil$

Table 1: Complexity Results of the submodular functions for different algorithms. The number of elements in the ground set is n , the selection size is k , the number of machines is m , and the number of levels in the accumulation tree is L .

421 **5. Experimental setup.** We conduct experiments to evaluate our algorithms
 422 using different accumulation tree structures and compare them with GREEDY and
 423 RANDGREEDI to assess the quality, runtime, and memory footprints of these algo-
 424 rithms. All the algorithms are executed on the Bell community cluster [22] of Purdue
 425 University with 448 nodes, each of which is an AMD EPYC 7662 node with 256 GB of
 426 total memory shared by the 128 cores. Each core operates at 2.0 GHz frequency. The
 427 cores on a node are organized hierarchically: four cores form a core complex, two core
 428 complexes form a core complex die, eight core complex dies form a socket, and two
 429 sockets constitute a node. Unfortunately, there are only 16 memory controllers for the
 430 128 cores, and hence in this NUMA architecture, memory contention is an issue on
 431 cores within a node. To simulate a completely distributed environment on this cluster
 432 we needed to ensure that the memory is not shared between nodes. Therefore, in what
 433 follows, a machine will denote one node with just one core assigned for computation,
 434 but having access to all 256 GB of memory. We also found that this made the run
 435 time results more reproducible.

436 For our experimental evaluation, we report the *runtime* and *quality* of the algo-
 437 rithms being compared. For runtime, we exclude the file reading time in each machine,
 438 and for the quality, we show the objective function value of the corresponding submod-
 439 ular function. Since the RANDGREEDI and GREEDYML are distributed algorithms,
 440 we also report the *number of function calls in the critical path* of the computational
 441 tree, which represents the parallel runtime of the algorithm. Given an accumulation
 442 tree, the number of function calls in the critical path refers to the maximum number
 443 of function calls the algorithm makes along a path from the leaf to the root. In our
 444 implementation, this quantity can be captured by the number of function calls made
 445 by nodes of the accumulation tree with $mid = 0$ since this node participates in the
 446 function calls from all levels of the tree.

Function	Dataset	$n = V $	$\sum_u \delta(u)$	avg. $\delta(u)$
k -dominating set	Friendster	65,608,366	1,806,067,135	27.52
	road_usa	23,947,347	57,708,624	2.41
	road_central	14,081,816	33,866,826	2.41
	belgium_osm	1,441,295	3,099,940	2.14
k -cover	webdocs	1,692,082	299,887,139	177.22
	kosarak	990,002	8,018,988	8.09
	retail	88,162	908,576	10.31
k -medoid	Tiny ImageNet	100,000	1,228,800,000	12,288

Table 2: Properties of Datasets used in the experiments. $\delta(u)$ is the number of neighbors of vertex u for the k -dominating set problem, the cardinality of the subset u for the k -cover problem, and the size of the vector representation of the pixels of image u for the k -medoid problem.

447 **Datasets.** In this paper, we limit our experiments to cardinality constraints using
 448 three different submodular functions described in detail in Section 4.2.

449 Our benchmark dataset is shown in Table 2. They are grouped based on the
 450 objective function and are sorted by the size of the dataset in each group. For the
 451 k -dominating set, our testbed consists of the Friendster social network graph [30] and
 452 a collection of road networks from **DIMACS10** dataset. We chose these graphs since
 453 they have relatively small average vertex degrees, leading to large vertex-dominating
 454 sets. For the k -cover objective, our datasets come from the **Frequent Itemset**
 455 **Mining Dataset Repository** [10] which contains popular benchmarks for set covers.
 456 We choose *webdocs*[21], *retail* [3], and *kosarak*. For the k -medoid problem, we use the
 457 **Tiny ImageNet** dataset [8], which contains 100,000 images with 200 different classes
 458 and 500 images from each class. Each image is 64×64 pixels in size.

459 **MPI Implementation.** Our codes are implemented using C++11, and compiled with
 460 g++9.3.0, using the O3 optimization flag. Our implementation of the GREEDY algo-
 461 rithm uses the Lazy Greedy [23] variant that has the same approximation guarantee as
 462 the GREEDY but is faster in practice since it potentially reduces the number of func-
 463 tion evaluations needed to choose the next element (by using the monotone decreasing
 464 gain property of submodular functions). Our implementation of the GREEDYML algo-
 465 rithm uses Open MPI implementation for the inter-node communication. We use the
 466 MPI_Gather and MPI_Gatherv primitives to receive all the solution sets from the chil-
 467 dren (Line 11 in Algorithm 3.1). We generated custom MPI_Comm communicators to
 468 enable this communication using MPI_Group primitives. Customized communicators
 469 are required since every machine has different children at each level. Additionally, we
 470 use the MPI_Barrier primitive to synchronize all the computations at each level.

471 **6. Experimental Results.** The experiments are executed with different accu-
 472 mulation trees that vary in the number of machines (m) and the number of levels (L)
 473 and branching factors (b) to assess their performance. We repeat each experiment
 474 six times and report the geometric mean of the results. Unless otherwise stated, a
 475 machine in our experiments represents a node in the cluster with only one core as-
 476 signed for computation as stated in Section 5. Whenever memory constraints allow,
 477 we compare our results with the sequential GREEDY algorithm that achieves $(1 - 1/e)$
 478 approximation guarantee.

479 Recall that our GREEDYML algorithm generalizes the RANDGREEDI algorithm

480 by allowing multiple levels in the accumulation tree, thus removing the bottleneck of
 481 a single aggregation. In the following, we verify this through a series of experiments.

482 In Section 6.1, we assess the performance of our algorithm using different ac-
 483 cumulation tree structures. We fix the number of machines and construct the best
 484 parameters of the accumulation tree for our dataset. Additionally, the experiment
 485 also demonstrates that the number of function calls in the critical path is a good
 486 estimate of the parallel runtime. In Section 6.2, we show the memory benefit of our
 487 GREEDYML w.r.t RANDGREEDI with two experiments. In Section 6.2.1, we impose
 488 a limit of 100 MB for each node and vary k , the selection size. This also simulates
 489 how the new algorithm can find applications in the *edge-computing* context. In Sec-
 490 tion 6.2.2, we fix the k value and vary the memory limits, necessitating different
 491 numbers of nodes to fit the data in the leaves. We observe the quality and runtime
 492 of different accumulation tree structures in these two experiments. Both these ex-
 493 periments are designed to show that the RANDGREEDI algorithm quickly runs out of
 494 memory with increasing m and k , and by choosing an appropriate accumulation tree,
 495 our GREEDYML algorithm can solve this problem with negligible drop in accuracy.
 496 For these experiments, we will choose the computational tree with the lowest depth
 497 that can be used with the memory limit and k values.

498 In Section 6.3, we perform a scaling experiment by varying the number of machines
 499 and using the tallest tree by setting a branching factor of two for the accumulation
 500 tree. We specifically show that even though the RANDGREEDI algorithm has a low
 501 asymptotic communication cost, it can become a bottleneck when scaled to a large
 502 number of machines. We also show how our algorithm alleviates this bottleneck.
 503 Finally, in Section 6.4, we perform experiments for the k -medoid objective function
 504 and show that we can provide a significant speedup by using taller accumulation trees
 505 without loss in quality. The k -medoid function is extensively used in machine learning
 506 as a solution to exemplar-based clustering problems.

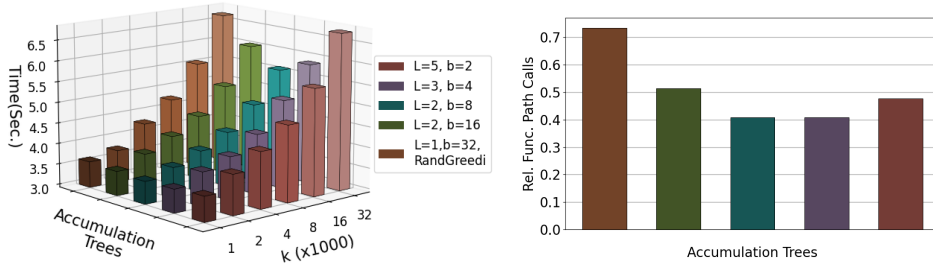


Fig. 4: Geometric means of results from GREEDYML for k -dominating set on different road datasets and k -cover on different set cover benchmark datasets on 32 machines. The first subfigure shows the execution times for different k values and accumulation trees. The second subfigure shows the Geometric mean values of the number of function calls in the critical path relative to the GREEDY algorithm for $k = 32,000$.

507 **6.1. Accumulation tree parameter selection.** Our first experiment explores
 508 the effect of choosing different branching factors and different accumulation levels in
 509 the accumulation tree for a fixed number of machines. In this experiment, we vary
 510 the selection set sizes k for each of these accumulation trees. We obtain results for the
 511 six datasets for k -dominating set and k -coverage detailed in Table 2. In Figure 4, we
 512 provide summary results on the relative number of function evaluations in the critical

513 path relative to the GREEDY algorithm and the running times by taking a geometric
 514 mean over all the datasets.

515 The first subfigure shows the execution time in seconds for the GREEDYML and
 516 RANDGREEDI algorithms, as the number of levels and the parameter k are varied.
 517 When k is small, there is less variation in the execution time, since a significant
 518 amount of work is performed at the leaves. As k increases, we can observe that the
 519 GREEDYML algorithm runs faster relative to the RANDGREEDI algorithm ($L = 1, b =$
 520 32). Note that, although we present in Figure 4 the geometric mean results over all
 521 the six datasets, the runtime and the function values for the individual datasets follow
 522 the same trend. The belgium_osm dataset has the largest reduction in run time with
 523 a reduction of around 22% and the smallest reduction in runtime is in the kosarak
 524 dataset with a reduction of 1% across all k values.

525 The second subfigure chooses $k = 32,000$ and plots the number of function calls in
 526 the critical path of the accumulation tree relative to the GREEDY algorithm for differ-
 527 ent (L, b) pairs. We observe that the relative number function calls for RANDGREEDI
 528 is around 70% of GREEDY, whereas the GREEDYML (with $L = 2$ and $b = 8$) cuts
 529 down the time by 15 percent. From Table 1, we can see that the function calls at a
 530 leaf node is $O(nk/m)$ whereas the function calls at an accumulation node is $O(mk^2)$
 531 for the RANDGREEDI algorithm. The accumulation node dominates the computation
 532 since it has a quadratic dependence on k , becoming a bottleneck for large k values.

533 This plot shows that the number of calls is a good indicator of the run time of
 534 the algorithm and that the cost of function evaluations dominates the time taken by
 535 the algorithm. On the other hand, the communication costs are small but, for the
 536 GREEDYML, they do grow with the number of levels when k is very large.

537 We additionally note (not in figure), that the objective function values obtained
 538 by the GREEDYML algorithm are not sensitive to the choice of the number of levels
 539 and the branching factors of the accumulation tree and differ by less than 1% from the
 540 values of the RANDGREEDI algorithm. For the webdocs k -coverage problem however,
 541 GREEDY obtains objective function values that are about 20% higher than both the
 542 RANDGREEDI and GREEDYML algorithms.




Fig. 5: Results from GREEDYML for the k -dominating set problem on the road_usa dataset on 16 nodes with varying k . The tuple (L, b) shows the number of levels and branching factors chosen for specific k values. The function values are relative to the GREEDY algorithm. Note that the leftmost bars in both plots represent the RANDGREEDI results.

544 **6.2.1. Varying k .** For this experiment, we use 16 machines with a limit on
 545 available memory of 100 MB per machine and vary k from 128,000 to 1,024,000.
 546 We consider the k -dominating set problem on the road_usa [1] dataset and large k
 547 values are chosen since the graph has an even larger maximum dominating set. Note
 548 that the k -values other than 128,000 cause the RANDGREEDI algorithm to run out of
 549 memory in accumulating all the solutions in the root node. We note that the small
 550 memory limit in this experiment can also be motivated from *edge computing* context.

551 The left plot in Figure 5 shows the number of function calls with varying values
 552 of k for the GREEDY and GREEDYML algorithms. For the GREEDYML (and the
 553 RANDGREEDI), we are interested in the number of function calls in the critical path
 554 since it represents the parallel runtime of the algorithm. With our memory limits,
 555 only $k = 128,000$ instance can be solved using the RANDGREEDI algorithm, which is
 556 shown in the leftmost blue bar of the plot.

557 As we increase k , we were able to generate solutions using our GREEDYML with
 558 different accumulation trees. The corresponding lowest-depth accumulation tree with
 559 the number of levels and branching factor is shown on top of the blue bars.

560 For each k value, we also executed the GREEDY algorithm shown in the orange
 561 bars. The result shows that the number of function evaluations on the critical path in
 562 the GREEDYML algorithm is smaller than the number of function evaluations in the
 563 sequential GREEDY algorithm. While the number of calls for accumulation trees with
 564 smaller b values is larger than RANDGREEDI, we can see that GREEDYML can solve
 565 the problems with larger k values in the same machine setup, which was not possible
 566 with RANDGREEDI. But it comes with a trade-off on parallel runtime. We observe
 567 that as we make the branching factor smaller our number of function calls in the
 568 critical path increases. That suggests that it is sufficient to choose the accumulation
 569 trees with the largest branching factor (thus the lowest depth tree) whenever the
 570 memory allows it.

571 The right plot of Figure 5 shows the relative objective function value, i.e., the
 572 relative number of vertices covered by the dominating set compared to the GREEDY
 573 algorithm, with varying k . The figure shows that the RANDGREEDI and GREEDYML
 574 algorithms attain quality at most 6% lesser than the serial GREEDY algorithm. Similar
 575 trends can be observed for other datasets in the summary of results shown in Figure
 576 4.

577 **6.2.2. Varying Memory Limits.** This experiment demonstrates the capabil-
 578 ity of the GREEDYML algorithm to solve a problem with a fixed k value on parallel
 579 machines when the memory is insufficient for the RANDGREEDI and GREEDY algo-
 580 rithms. Unlike the experiment in Section 6.2.1, where we selected the accumulation
 581 trees based on the k value for the problem, here, we fix k and choose accumulation
 582 trees based on the size of memory available on the machines. We consider the k -
 583 dominating set problem on graphs, and first report results on the Friendster dataset
 584 [30]. We set the cardinality constraint k so that the k -dominating set requires 512
 585 MB, roughly a factor of 64 smaller than the original graph. The RANDGREEDI algo-
 586 rithm can execute this problem only on 8 machines, each with 4 GB of memory, since
 587 in the accumulation step, one machine receives solutions of size 512 MB each from
 588 8 machines. The GREEDYML algorithm having multiple levels of accumulation can
 589 run on 16 machines with only 2 GB memory, using $L = 2$ and $b = 4$. Furthermore, it
 590 can also run on 32 machines with only 1 GB memory, using $L = 5$ and $b = 2$.

591 We show results from these three machine configurations in Table 3. We report the
 592 number of function calls on a critical path and the objective function values normalized

Dataset	Alg.	Mem. Limit	m	b	L	Rel. Func. Val.(%)	Time (sec.)
Friendster	RG	4GB	8	8	1	96.294	69.81
	GML	2GB	16	4	2	96.232	82.92
	GML	1GB	32	2	5	96.175	112.17
road_usa	RG		8	8	1	99.034	1.25
	GML		16	4	2	99.005	1.63
	GML		32	2	5	99.027	3.56
webdocs	RG		8	8	1	79.948	4.50
	GML		16	4	2	78.723	4.72
	GML		32	2	5	79.743	8.59

Table 3: Results from GREEDYML (GML) for k -dominating set on the Friendster, road_usa and webdocs datasets. The memory size per machine is varied for the Friendster dataset. The number of machines m and the accumulation tree are selected based on the size of the data and the size of the solutions to get three different machine organizations. We report the function values relative to the GREEDY algorithm and the execution time in seconds. Note that the 4GB entries run with $L = 1$ and correspond to the RANDGREEDI (RG) algorithm. We use the same three machine organizations for the road_usa and webdocs datasets to show they follow similar trends in solution quality and execution time.

593 by those obtained from the serial GREEDY algorithm. Our results show that objective
594 function values computed by the GREEDYML algorithm (the 2 and 1 GB results) are
595 insensitive to the number of levels in the tree. Similar trends are observed for the
596 webdocs [21] and road_usa [1] datasets when we used the same number of machines
597 and accumulation trees. As we increase the number of machines and levels in the
598 accumulation tree, the execution times (in seconds) increase for this problem due to
599 the communication and synchronization costs involved. However, the larger numbers
600 of machines enable us to solve large problems by overcoming memory constraints. So,
601 in this scenario, it is sufficient to select the number of machines depending on the size
602 of the dataset and then select the branching factor such that the accumulation step
603 does not exceed the memory limits. We also notice that the RANDGREEDI algorithm
604 has an inherently serial accumulation step, and the GREEDYML algorithm provides
605 a mechanism to parallelize it.

606 **6.3. Strong Scaling.** Next, we show how the GREEDYML algorithm serves as
607 a solution to the scaling bottlenecks that arise in the RANDGREEDI algorithm. For
608 the scaling experiment, we consider the k -dominating set problem on the Friendster
609 dataset. We set the branching factor $b = 2$ for the GREEDYML algorithm since
610 this has the highest number of levels and, thus, the lowest approximation guarantee.
611 We compare this with the RANDGREEDI algorithm starting from 8 machines to 128
612 machines with $k = 50$. We compare the total execution time, communication time,
613 and computation time for the GREEDYML and the RANDGREEDI algorithms.

614 In Figure 6, we plot the total execution time by stacking communication and
615 computation time for the two algorithms. We observe that the communication cost
616 does not scale for the RANDGREEDI algorithm. From Table 1, we can see that the time
617 spent by the central node collecting the solutions is $O(km)$ and, therefore, increases
618 linearly with the number of machines. In contrast, for GREEDYML algorithm (with




Fig. 6: Strong scaling results of the RANDGREEDI and GREEDYML algorithms for $k = 50$ on Friendster dataset for k -dominating set problem. We set $b = 2$ for the GREEDYML algorithm.

L	b	Local Obj.		Added Images	
		Rel. Func. Val. (%)	Speedup	Rel. Func. Val. (%)	Speedup
5	2	92.22	2.00	93.69	2.01
3	4	92.21	1.96	92.70	1.94
2	8	92.73	1.95	92.77	1.93
2	16	92.22	1.49	93.34	1.44

Table 4: Results from GREEDYML for the k -medoid function on the **Tiny ImageNet** data set using different accumulation trees. The table shows the relative function values and speedup compared to the RANDGREEDI algorithm using two different local objective values computation schemes executed on 32 nodes. For both, higher values are better. Here L and b are the number of levels and branching factor, respectively.

619 a constant branching factor, $b = 2, L = \log_2 m$), the communication cost $O(k \log m)$
 620 which grows logarithmically in the number of machines. The total communication
 621 times of the GREEDYML algorithm across different machines are consistently around
 622 0.25 seconds, whereas the RANDGREEDI increases from 0.05 second to 2 seconds
 623 linearly. We observe that computation times for both RANDGREEDI and GREEDYML
 624 changes similarly with m , indicating that the majority of the computation work is
 625 performed at the leaf nodes. For computation time, we observe a slightly worse scaling
 626 of RANDGREEDI compared to GREEDYML, again because the central node becomes
 627 a computational bottleneck as m increases. Similar to other experiments, we see an
 628 almost identical quality in the solutions where the GREEDYML solution has a quality
 629 reducing by less than 1% from the solution of the RANDGREEDI algorithm.

630 **6.4. The k -Medoid Problem.** In our final experiment, we consider the k -
 631 medoid function that solves the exemplar-based clustering problem. Our dataset
 632 consists of the **Tiny ImageNet** dataset [8], which contains 100,000 images with 200
 633 different classes and 500 images from each class. Each of the images is 64×64 pixels.
 634 We flatten each image into a vector of 12,288 length. We then subtracted the mean
 635 value and normalized the vector. We compute the dissimilarity between two images
 636 as the Euclidean distance between the normalized vector representations. Here, the




Fig. 7: Results from GREEDYML for the k -medoid problem on the Tiny ImageNet dataset on 32 nodes with $k = 200$ with no images added at each accumulation step. The subfigure on the left shows the first 16 image results for one of the runs for the GREEDYML algorithm with branching factor $b = 2$, and the subfigure on the right shows the top 16 image results for one of the runs for the RANDGREEDI algorithm.

637 auxiliary image e_0 is a pixel vector of all zeros. Note that, unlike the other two
 638 functions, the k -medoid function, requires access to the full dataset for computing
 639 the functional value. Since the dataset is distributed, this poses an issue in the
 640 experiment. To overcome this, following [24, 2], we calculate the objective function
 641 value using only the images available *locally* on each machine. This means the ground
 642 set for each machine is just the images present in that machine. This is motivated
 643 by an analysis from Mirzasoleiman et al. (Theorem 10, [24]) showing that computing
 644 $f(S)$ with the ground set as some subset $D \subseteq V$ chosen uniformly at random provides
 645 a high-probability additive approximation to the function value $f(S)$ evaluated with
 646 ground set V . Additionally, they have also added subsets of randomly chosen images
 647 to the central machine to provide practical quality improvement. We have followed
 648 these techniques (local only and local with additional images) in the experiments for
 649 our multilevel GREEDYML algorithm.

650 In our experiments, we fix the number of machines ($m = 32$) and vary the ac-
 651 cumulation trees by choosing different L and b . We set the solution size k to 200
 652 images. For the variant with additional images, we add 1,000 random images from
 653 the original dataset to each accumulation step.

654 In Table 4, we show the relative objective function values and speedup for different
 655 accumulation trees relative to the RANDGREEDI algorithm. We observe that the ob-
 656 jective function values for GREEDYML algorithm are almost similar to RANDGREEDI.
 657 Our results show that the GREEDYML algorithm becomes gradually faster as we in-
 658 crease the number of levels with runtime improvement ranging from $1.45 - 2.01 \times$.
 659 This is because the k -medoid function is compute-intensive, where computation cost
 660 increases quadratically with the number of images (Table 1). With $k = 200$ and
 661 $m = 32$, the RANDGREEDI algorithm has $km = 6,400$ images at the root node but
 662 only $n/m = 313$ images at the leaves, thus the computation at the root node domi-
 663 nates in cost. On the other hand, as we decrease the branching factor (from $b = 16$
 664 to 2), the number of images (kb) in the interior nodes decreases from 3,200 to 400 for

665 the GREEDYML algorithm. This gradual decrease in compute time is reflected in the
666 total time, and also in the observed speedup.

667 Finally, in Fig. 7, we show 16 out of the 200 images determined to be cluster
668 centers by the GREEDYML and RANDGREEDI algorithms. We can draw the conclusion
669 that the submodular k -medoid function is able to generate a diverse set of exemplar
670 images for this clustering problem.

671 **7. Conclusion and Future work.** We have developed a new distributed algo-
672 rithm that generalizes the existing state-of-the-art algorithm for monotone submodu-
673 lar maximization subject to hereditary constraints. We prove that the new algorithm
674 is $\alpha/(L + 1)$ approximate and showed its quality doesn't degrade for the k -cover, k -
675 dominating set, and k -medoid problems. We showed how this new algorithm reduces
676 the inherent serial computation and communication bottlenecks of the RANDGREEDI
677 algorithm. We also reduce the memory required to solve the problem enabling sub-
678 modular maximization to be solved in an edge computation context and with larger
679 k values. Finally, We showed a significant speedup in solving the popular exemplar-
680 based clustering problem.

681 As part of our future work, we plan to run experiments for other hereditary
682 constraints, such as matroid and p -system constraints. We will also explore how this
683 generalization technique can be applied to other classes of NP-Hard problems such as
684 non-monotone submodular functions and weakly-submodular functions.

685

REFERENCES

- 686 [1] D. A. BADER, H. MEYERHENKE, P. SANDERS, AND D. WAGNER, *Graph partitioning and graph*
687 *clustering*, in Contemporary Mathematics, vol. 588, American Mathematical Society, 2012.
688 10th DIMACS Implementation Challenge Workshop.
- 689 [2] R. D. BARBOSA, A. ENE, H. L. NGUYEN, AND J. WARD, *The power of randomization: Distribu-*
690 *tuted submodular maximization on massive datasets*, in Proceedings of the 32nd Interna-
691 tional Conference on Machine Learning, JMLR.org, 2015, pp. 1236–1244.
- 692 [3] T. BRIJS, G. SWINNEN, K. VANHOOF, AND G. WETS, *Using association rules for product as-*
693 *ortment decisions: A case study*, in Proceedings of the Fifth ACM SIGKDD International
694 Conference on Knowledge Discovery and Data Mining, ACM, 1999, pp. 254–260.
- 695 [4] C. CHEKURI AND K. QUANRUD, *Submodular function maximization in parallel via the multilin-*
696 *ear relaxation*, in Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete
697 Algorithms, Society for Industrial and Applied Mathematics, 2019, p. 303–322.
- 698 [5] L. CHEN, G. ZHANG, AND E. ZHOU, *Fast greedy MAP inference for determinantal point process*
699 *to improve recommendation diversity*, in Advances in Neural Information Processing Sys-
700 tems, 2018, pp. 5627–5638.
- 701 [6] M. COUTINO, S. P. CHEPURI, AND G. LEUS, *Submodular sparse sensing for Gaussian de-*
702 *tection with correlated observations*, IEEE Transactions on Signal Processing, 66 (2018),
703 p. 4025–4039.
- 704 [7] E. ELENBERG, A. G. DIMAKIS, M. FELDMAN, AND A. KARBASI, *Streaming weak submodularity:*
705 *Interpreting neural networks on the fly*, in Advances in Neural Information Processing
706 Systems, 2017, p. 4044–4054.
- 707 [8] A. K. FEI-FEI LI, *Tiny imagenet challenge*. <http://cs231n.stanford.edu/tiny-imagenet-200.zip>,
708 2017. [Online; last accessed 13-Mar-2024].
- 709 [9] S. M. FERDOUS, A. POTHEN, A. KHAN, A. PANYALA, AND M. HALAPPANAVAR, *A parallel*
710 *approximation algorithm for maximizing submodular b-matching*, in SIAM Conference on
711 Applied and Computational Discrete Algorithms (ACDA), 2021, pp. 45–56.
- 712 [10] FIMI, *Frequent itemset mining dataset repository*. <http://fimi.uantwerpen.be/data/>, 2003.
- 713 [11] D. GOLOVIN AND A. KRAUSE, *Adaptive submodularity: Theory and applications in active*
714 *learning and stochastic optimization*, Journal of Artificial Intelligence Research, 42 (2011),
715 p. 427–486.
- 716 [12] L. KAUFMAN AND P. J. ROUSSEEUW, *Finding Groups in Data: An Introduction to Cluster*
717 *Analysis.*, John Wiley, 1990.
- 718 [13] D. KEMPE, J. KLEINBERG, AND É. TARDOS, *Maximizing the spread of influence through a social*

- 719 *network*, in Proceedings of the ninth ACM SIGKDD international conference on Knowledge
720 discovery and data mining, 2003, pp. 137–146.
- 721 [14] A. KRAUSE AND D. GOLOVIN, *Submodular function maximization*, in Tractability: Practical
722 Approaches to Hard Problems, Cambridge University Press, 2014, pp. 71 – 104.
- 723 [15] A. KRAUSE AND C. GUESTRIN, *Near-optimal observation selection using submodular functions*,
724 in Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, vol. 7,
725 2007, pp. 1650–1654.
- 726 [16] A. KRAUSE, J. LESKOVEC, C. GUESTRIN, J. VANBRIESEN, AND C. FALOUTSOS, *Efficient sensor
727 placement optimization for securing large water distribution networks*, Journal of Water
728 Resources Planning and Management, 134 (2008), pp. 516–526.
- 729 [17] R. KUMAR, B. MOSELEY, S. VASSILVITSKII, AND A. VATTANI, *Fast greedy algorithms in mapre-
730 duce and streaming*, ACM Trans. Parallel Comput., 2 (2015), pp. 14:1–14:22.
- 731 [18] H. LIN AND J. BILMES, *Multi-document summarization via budgeted maximization of submodu-
732 lar functions*, in Human Language Technologies: The 2010 Annual Conference of the North
733 American Chapter of the Association for Computational Linguistics, 2010, pp. 912–920.
- 734 [19] H. LIN AND J. BILMES, *A class of submodular functions for document summarization*, in Pro-
735 ceedings of the 49th Annual Meeting of the Association for Computational Linguistics:
736 Human Language Technologies, 2011, pp. 510–520.
- 737 [20] L. LOVÁSZ, *Submodular functions and convexity*, in Mathematical Programming The State
738 of the Art: Bonn 1982, A. Bachem, B. Korte, and M. Grötschel, eds., Springer Berlin
739 Heidelberg, 1983, pp. 235–257, https://doi.org/10.1007/978-3-642-68874-4_10.
- 740 [21] C. LUCCHESI, S. ORLANDO, R. PEREGO, AND F. SILVESTRI, *WebDocs: A real-life huge trans-
741 actional dataset*, in Proceedings of the IEEE ICDM Workshop on Frequent Itemset Mining
742 Implementations (FIMI 04), 2004.
- 743 [22] G. MCCARTNEY, T. HACKER, AND B. YANG, *Empowering Faculty: A Campus Cyberinfrastructure
744 Strategy for Research Communities*, Educause Review, (2014), [https://er.educause.
745 edu/articles/2014/7/empowering-faculty-a-campus-cyberinfrastructure-strategy-for-re-
746 search-communities](https://er.educause.edu/articles/2014/7/empowering-faculty-a-campus-cyberinfrastructure-strategy-for-research-communities).
- 747 [23] M. MINOUX, *Accelerated greedy algorithms for maximizing submodular set functions*, in Opti-
748 mization Techniques, J. Stoer, ed., Springer, 1978, pp. 234–243.
- 749 [24] B. MIRZASOLEIMAN, A. KARBASI, R. SARKAR, AND A. KRAUSE, *Distributed submodular max-
750 imization: Identifying representative elements in massive data*, in Advances in Neural
751 Information Processing Systems, vol. 26, 2013, pp. 2049–2057.
- 752 [25] A. MOKHTARI, H. HASSANI, AND A. KARBASI, *Decentralized submodular maximization: Bridg-
753 ing discrete and continuous settings*, in Proceedings of the 35th International Conference
754 on Machine Learning, 2018, pp. 3616–3625.
- 755 [26] A. ROBAY, A. ADIBI, B. SCHLOTFELDT, H. HASSANI, AND G. J. PAPPAS, *Optimal algorithms for
756 submodular maximization with distributed constraints*, in Proceedings of the 3rd Conference
757 on Learning for Dynamics and Control, 2021, pp. 150–162.
- 758 [27] K. THEKUMPARAMPIL, A. THANGARAJ, AND R. VAZE, *Combinatorial resource allocation using
759 submodularity of waterfilling*, IEEE Transactions on Wireless Communications, 15 (2016),
760 p. 206–216.
- 761 [28] E. TOHIDI, R. AMIRI, M. COUTINO, D. GESBERT, G. LEUS, AND A. KARBASI, *Submodularity in
762 action: From machine learning to signal processing applications*, IEEE Signal Processing
763 Magazine, 37 (2020), pp. 120–133, <https://doi.org/10.1109/MSP.2020.3003836>.
- 764 [29] L. G. VALIANT, *A bridging model for parallel computation*, Commun. ACM, 33 (1990),
765 p. 103–111, <https://doi.org/10.1145/79173.79181>.
- 766 [30] J. YANG AND J. LESKOVEC, *Defining and evaluating network communities based on ground-
767 truth*, in Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics, 2012,
768 pp. 1–8.