

Fully-Dynamic Coresets

Monika Henzinger 

University of Vienna, Faculty of Computer Science, Vienna, Austria

<https://homepage.univie.ac.at/monika.henzinger/>

monika.henzinger@univie.ac.at

Sagar Kale

University of Vienna, Faculty of Computer Science, Vienna, Austria

<https://sagark4.github.io/>

sagar.kale@univie.ac.at

Abstract

With input sizes becoming massive, coresets—small yet representative summary of the input—are relevant more than ever. A weighted set C_w that is a subset of the input is an ε -coreset if the cost of any feasible solution S with respect to C_w is within $[1 \pm \varepsilon]$ of the cost of S with respect to the original input. We give a very general technique to compute coresets in the fully-dynamic setting where input points can be added or deleted. Given a static (i.e., not dynamic) ε -coreset-construction algorithm that runs in time $t(n, \varepsilon, \lambda)$ and computes a coreset of size $s(n, \varepsilon, \lambda)$, where n is the number of input points and $1 - \lambda$ is the success probability, we give a fully-dynamic algorithm that computes an ε -coreset with worst-case update time $O((\log n) \cdot t(s(n, \varepsilon / \log n, \lambda / n), \varepsilon / \log n, \lambda / n))$ (this bound is stated informally), where the success probability is $1 - \lambda$. Our technique is a fully-dynamic analog of the merge-and-reduce technique, which is due to Har-Peled and Mazumdar [17] and is based on a technique of Bentley and Saxe [3], that applies to the insertion-only setting where points can only be added. Although, our space usage is $O(n)$, our technique works in the presence of an adaptive adversary, and we show that $\Omega(n)$ space is required when adversary is adaptive.

As a concrete implication of our technique, using the result of Braverman et al. [4], we get fully-dynamic ε -coreset-construction algorithms for k -median and k -means with worst-case update time $O(\varepsilon^{-2} k^2 \log^5 n \log^3 k)$ and coreset size $O(\varepsilon^{-2} k \log n \log^2 k)$ ignoring $\log \log n$ and $\log(1/\varepsilon)$ factors and assuming that $\varepsilon = \Omega(1/\text{poly}(n))$ and $\lambda = \Omega(1/\text{poly}(n))$ (which are very weak assumptions made only to make these bounds easy to parse). This results in the first fully-dynamic constant-approximation algorithms for k -median and k -means with update times $O(\text{poly}(k, \log n, \varepsilon^{-1}))$. Specifically, the dependence on k is only quadratic, and the bounds are worst-case. The best previous bound for both problems was *amortized* $O(n \log n)$ by Cohen-Addad et al. [10] via randomized $O(1)$ -coresets in $O(n)$ space.

We also show that under the OMv conjecture [18], a fully-dynamic $(4 - \delta)$ -approximation algorithm for k -means must either have an amortized update time of $\Omega(k^{1-\gamma})$ or amortized query time of $\Omega(k^{2-\gamma})$, where $\gamma > 0$ is a constant.

2012 ACM Subject Classification Theory of computation \rightarrow Facility location and clustering

Keywords and phrases Clustering, Coresets, Dynamic Algorithms

Digital Object Identifier 10.4230/LIPIcs.ESA.2020.36

Related Version A full version of the paper is available at <https://arxiv.org/abs/2004.14891>.

Funding *Monika Henzinger*: The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 340506.

Sagar Kale: Fully supported by Vienna Science and Technology Fund (WWTF) through project ICT15-003.

1 Introduction

Clustering is an ubiquitous notion that one encounters in computer-science areas such as data mining, machine learning, image analysis, bioinformatics, data compression, and computer graphics, and also in the fields of medicine, social science, marketing, etc. Today, when the input data has become massive, one would rather run an algorithm on a small but representative summary of the input, and for clustering problems, a *coreset* serves that function perfectly. The concept of a coreset was defined first in computational geometry as a small subset of a point set that approximates the shape of the point set. The word coreset has now evolved to mean an appropriately weighted subset of the input that approximates the original input with respect to solving a computational problem.

Let P be a problem for which the input is a weighted subset¹ $X_w \subseteq U$; think of U as in a metric space (U, d) , so U is unweighted and d is the distance function. Let $n := |X_w|$ and $W := \sum_{x \in X_w} w(x)$. We also refer to elements of U as points. The goal in the problem P then is to output S^* that belongs to the feasible-solution space (or *query* space) Q such that the cost $c(S^*, X_w)$ is minimized. For example, in the *k-median* (respectively, *k-means*) problem, Q is the set of all (unweighted) subsets of X_w of cardinality at most k and $c(S, X_w) := \sum_{x \in X_w} w(x) \min_{s \in S} d(x, s)$ (respectively, $\sum_{x \in X_w} w(x) \min_{s \in S} (d(x, s))^2$). Then, for the problem P , a weighted set C_w such that $C_w \subseteq X_w$ is an ε -coreset if, for any feasible solution $S \in Q$, we have that $c(S, X_w) \in [1 \pm \varepsilon]c(S, C_w)$; we sometimes say that the *quality* of coreset C_w is ε . For many problems, fast coreset-construction algorithms exist; e.g., for *k-median* and *k-means*, $\tilde{O}(nk)$ -time² algorithms for computing ε -coresets of size $O(\varepsilon^{-2}k \text{ polylog}(n))$ exist.

Throughout the paper, we assume that the cost function c for the problem P is *linear*: for any weighted subsets $Y_w^1, Y_w^2 \subseteq U$ with disjoint supports and any $S \in Q$, we have that $c(S, Y_w^1 \cup Y_w^2) = c(S, Y_w^1) + c(S, Y_w^2)$, where the union $Y_w^1 \cup Y_w^2$ is the weighted union. It is easy to see that *k-median* and *k-means* cost functions are linear.

Our goal in this paper is to give *dynamic algorithms* for computing a coreset. In the dynamic setting, the input changes over time. A dynamic algorithm is a data structure that supports three types of operations: *Insert*(p, w), which inserts a point p with weight w into X_w ; *Delete*(p), which removes point p from X_w ; and *Query*(\cdot), which outputs a coreset of X_w . Weight updates can be simulated by deleting and re-inserting a given point, or the data structure may support a separate weight-changing operation. This is known as the *fully* dynamic model as opposed to the *insertion-only* setting where a point can only be inserted. At any time instant, a coreset is maintained by the algorithm, and the complexity measure of interest is the *update* time, i.e., how fast the solution can be updated after receiving a point update, and also the *size* of the coreset, which determines the *query* time. Suppose there is a dynamic coreset-construction algorithm, say ALG_D , for a problem P . Then a solution for the problem P can be maintained dynamically by running ALG_D , and on query, a solution is computed by querying ALG_D and running a static (i.e., not dynamic) algorithm for P on the returned coreset. In this paper, we give a very general technique on how to maintain a coreset in the fully-dynamic setting: given a *static* coreset-construction algorithm for any problem P , we show how to turn it into a dynamic coreset-construction algorithm for P .

Intuitively, our technique is to the fully-dynamic setting as the merge-and-reduce technique

¹ When using a set operation such as union or notation such as \subseteq with one or more weighted sets, we mean it for the underlying unweighted sets. Also, all weights are nonnegative.

² Logarithmic factors are hidden in the \tilde{O} notation.

is to the insertion-only setting. The *merge-and-reduce* technique, which is based on a technique of Bentley and Saxe [3], is due to Har-Peled and Mazumdar [17] and is a fundamental technique to obtain an insertion-only coreset-construction algorithm using a static coreset-construction algorithm, say ALG_S , as a black box. Loosely speaking, it is as follows. At any time instant, the algorithm maintains up to $\lceil \log n \rceil$ buckets. For $i \in \{1, 2, \dots, \lceil \log n \rceil\}$, the bucket B_i has capacity 2^{i-1} , each bucket can be either *full*, (i.e. at capacity 2^{i-1}) or *empty*, and each point goes in exactly one bucket. Then at any time-instant, the current number of points uniquely determines the states of the buckets. Whenever a point is inserted, the states of the buckets change like a binary counter. That is, the new point goes into B_i , where B_i is the smallest-index empty bucket, and all the points in $\cup_{j=1}^{i-1} B_j$ are moved to B_i (*merge*). Note that this creates a full bucket B_i . Then a coreset is computed on B_i by running ALG_S on it (*reduce*). The overall coreset is then just union of all non-empty buckets.

We show that a similar result can be achieved in the fully-dynamic setting. Our main result is the following theorem (stated slightly informally).

► **Theorem 1.** *Assume that there is a static coreset construction algorithm for a problem P with linear cost function that **a)** runs in time $t_P(n_s, \varepsilon_s, \lambda_s, W_s)$, **b)** always outputs a set of cardinality at most $s_P(\varepsilon_s, \lambda_s, W_s)$ and total weight at most $(1+\delta)W_s$, and **c)** has the guarantee that the output is an ε_s -coreset with probability at least $1-\lambda_s$, where n_s is the number of integer-weighted input points and W_s is the total weight of points.*

*Then there is a fully-dynamic coreset-construction algorithm for P that, with rational-weighted input points, **a)** always maintains an output set of cardinality at most $s_P(\varepsilon, \lambda, W)$, **b)** has the guarantee that the output is an ε -coreset with probability at least $1-\lambda$, and **c)** has worst-case update time*

$$O\left((\log n) \cdot t_P\left(s_P^*, \frac{\varepsilon}{\log n}, \frac{\lambda}{n}, W\right)\right),$$

where $W = O((1+\delta)^{\lceil \log n \rceil} \text{poly}(n))$, $s_P^* = s_P\left(\frac{\varepsilon}{\log n}, \frac{\lambda}{n}, W\right)$, and n is the current number of points.

We mention below a concrete implication of the above theorem for k -median and k -means using the result of Braverman et al. [4].

► **Theorem 2.** *For the k -median and k -means problems, there is a fully-dynamic algorithm that maintains a set of cardinality $O(\varepsilon^{-2}k(\log n \log k \log(k\varepsilon^{-1} \log n) + \log(1/\lambda)))$, that is an ε -coreset with probability at least $1-\lambda$, and has worst-case update time*

$$O\left(\varepsilon^{-2}k^2 \log^5 n \log^3 k \log^2(\varepsilon^{-1})(\log \log n)^3\right),$$

assuming that $\varepsilon = \Omega(1/\text{poly}(n))$ and $\lambda = \Omega(1/\text{poly}(n))$.³

Ignoring $\log \log n$ and $\log(1/\varepsilon)$ above, the coreset cardinality is $O(\varepsilon^{-2}k \log n \log^2 k)$ and worst-case update time is $O(\varepsilon^{-2}k^2 \log^5 n \log^3 k)$. It can be easily proved that running an α -approximation algorithm for k -median on an ε -coreset gives a $2\alpha(1+\varepsilon)$ -approximation whereas that for k -means gives a $4\alpha(1+\varepsilon)$ -approximation. Any such polynomial-time static algorithm—say, e.g., $(5 + \varepsilon')$ -approximation algorithm for k -median by Arya et al. [2] and 16-approximation algorithm for k -means by Gupta and Tangwongsan [15]—can be run on

³ We make these very weak assumptions to simplify some extremely unhandy factors involving ε and λ in the expression for the update time.

127 our output coreset in $O(\text{poly}(k, \log n, \varepsilon^{-1}))$ time to obtain a constant approximation. This is
 128 the first fully-dynamic constant-approximation algorithm for k -median and k -means whose
 129 *worst-case* time per operation is polynomial in k , $\log n$, and ε^{-1} . The best previous result
 130 was a randomized algorithm with *amortized* $O(n \log n)$ update time and $O(n)$ space by
 131 Cohen-Addad et al. [10].

132 With a simple reduction, we also show a conditional lower bound on the time per operation
 133 for k -means. The following theorem is proved as Theorem 20 in Section 4.

134 ► **Theorem 3.** *Let $\gamma > 0$ be a constant. Under the OMv conjecture [18], for any $\delta > 0$, there*
 135 *does not exist a fully-dynamic algorithm that maintains a $(4 - \delta)$ -approximation for k -means*
 136 *with amortized update time $O(k^{1-\gamma})$ and query time $O(k^{2-\gamma})$ such that over a polynomial*
 137 *number of updates, the error probability is at most $1/3$.*

138 Our technique

139 At the core, our technique is simple. We always maintain a balanced binary tree of depth
 140 $\lceil \log n \rceil$ containing exactly n leaf nodes (recall that n is the current number of points). Each
 141 node corresponds to a subset of X_w , the current input: each leaf node corresponds to a
 142 singleton (hence n leaf nodes), and an internal node corresponds to the weighted union of the
 143 sets represented by its children. If the cardinality of the union exceeds a certain threshold,
 144 then we use the static coreset-construction algorithm to compute its coreset. The root gives
 145 a coreset of the whole input.

146 We next explain how we handle updates in this data structure. Point insertions are
 147 straightforward: create a new leaf node and run all the static-algorithm instances at the
 148 nodes on the leaf-to-root path. The way we handle point deletions is similar in spirit to
 149 the way delete-min works in a min-heap data structure: whenever a point at leaf-node ℓ_d is
 150 deleted, we swap contents of ℓ_d with those of the rightmost leaf-node, say ℓ_r , and delete ℓ_r ,
 151 thus maintaining the balance of the tree. Then we run all the static-algorithm instances at
 152 the nodes on the two affected leaf-to-root paths.

153 However, there are some complications that require new techniques to make it work in
 154 *worst-case* time. To maintain guarantees for the output coreset quality and overall success
 155 probability, we need to adapt the parameters ε_s and λ_s used for the static algorithm at the
 156 internal nodes. The problem is that both depend on n , which changes over time and thus
 157 might become outdated. To show an *amortized* update-time bound, we can simply rerun
 158 the static algorithms at all internal nodes whenever n has changed by a constant factor. To
 159 achieve our worst-case bound, we use two *refresh pointers* that point at leaf nodes, and after
 160 each update operation, we rerun using the new values of ε_s and λ_s all the static-algorithm
 161 instances at the nodes on the leaf-to-root path from the leaf nodes pointed to by the refresh
 162 pointers. This keeps the outputs of the static-algorithm instances at the internal nodes
 163 always fresh. After every update, we move these pointers to the right so that they point to
 164 the next leaf nodes.

165 Further complications are caused by fractional weights at the leaf nodes and fractional
 166 intermediate-output weights. A problem arises when the weights in X_w are fractional, and
 167 the static algorithm expects integer-weighted input [9]. Even if the static algorithm can
 168 handle fractional weights [12, 4], there can be a problem. The output of the static algorithm
 169 at an internal node is the input for the static algorithm at its parent. Naïvely feeding these
 170 output fractional weights directly to the static algorithm at the parent may result in numbers
 171 exponential in n near the root, thus prohibitively increasing the update time. To deal with
 172 these problems, rounding is needed for the input, i.e., at the leaf nodes, as well as for each

173 intermediate-output at an internal node. Thus, we propose a more sophisticated rounding
 174 scheme and show that the rounding errors accumulated by our rounding are not too high.

175 We note that our balanced binary-tree data structure may be used to get dynamic
 176 algorithms in the following situations. Let $f : \mathbb{R}^{\dim} \rightarrow \mathbb{R}^{\dim}$ be a *multi-valued* function.
 177 Suppose for any u and v with disjoint supports and for any $f_u \in f(u)$ and $f_v \in f(v)$, we
 178 have $f_u + f_v \in f(u + v)$. Also suppose that $f(f(v)) \subseteq f(v)$ for any v . Now, given input v ,
 179 we want to compute some vector in $f(v)$. If there is a static algorithm for this, then using
 180 our technique, we can maintain some vector in $f(v)$ for a dynamically changing vector v .
 181 The allowed dynamic operation on v is “add a to the i th coordinate of v ,” where $a \in \mathbb{R}$.
 182 The resulting dynamic algorithm is fast if the static algorithm always outputs a “small”
 183 vector; this is true for coresets because coresets are small by nature. Thinking about coresets
 184 in the above language, each point is an identity vector in $\mathbb{R}_+^{|U|}$, and then each weighted
 185 set of points naturally identifies with a vector. An ε -coreset reduces the number of points
 186 drastically. Union of coresets of two disjoint sets is a coreset of the union of those two sets (see
 187 Lemma 5). Although an ε -coreset of an ε -coreset is not an ε -coreset, it is a $(2\varepsilon + \varepsilon^2)$ -coreset
 188 (see Lemma 6).

189 Space

190 In the merge-and-reduce technique, a bucket B_i will not actually contain 2^{i-1} points but
 191 just a coreset of 2^{i-1} points that would have been there otherwise at any time instant. Thus,
 192 using space just $\lceil \log n \rceil$ times the coreset size for a bucket, one can get a coreset of the whole
 193 input [17]. This makes it also applicable in the more restricted streaming model, where the
 194 input points arrive in a sequence and the goal is to compute a coreset using sublinear space.
 195 In the fully-dynamic setting, deletions also need to be handled, and hence no deterministic
 196 or randomized algorithm against an adaptive adversary that stores only a coreset is possible:
 197 the adversary generating the input could simply ask a query and then delete all points in
 198 the returned coreset. Hence, an algorithm that does not store any information about the
 199 non-coreset points would not be able to maintain a valid coreset. Even though we store
 200 all the points in our fully-dynamic technique, i.e., its space usage is $O(n)$, it works against
 201 an adaptive adversary because we never make any assumption about the next update and
 202 perform each update independently of all previous updates. By a straightforward reduction
 203 from the communication problem of INDEX, we show that $\Omega(n)$ space is required in the
 204 presence of an adaptive adversary. The proof of the following theorem appears in Section 4.

205 ► **Theorem 4.** *A fully-dynamic algorithm that obtains any bounded approximation for 1-
 206 median or 1-means that works in the presence of an adaptive adversary and has success
 207 probability $1 - 1/(8n^2)$ must use $\Omega(n)$ space, where n is the current number of points.*

208 Comparison with the sparsification technique

209 Our technique is close to the sparsification technique of Eppstein et al. [11] that is used to
 210 speed up dynamic graph algorithms. There, one has to assume that the number of vertices
 211 in the input graph, say n_v , does not change, but the edge set changes dynamically, and the
 212 bounds are obtained in terms of n_v and m , the current number of edges. Their dynamic
 213 edge-tree structure is based on a fixed vertex-partition tree. In the vertex-partition tree, a
 214 node at level i corresponds to a vertex-set of cardinality $n_v/2^i$, and a vertex-set at a node is
 215 a union of its children’s vertex sets (cf. our technique). To start using the edge tree, the
 216 vertex-partition tree has to be built first and hence the knowledge of n_v is necessary. Neither
 217 do we need such a fixed structure nor any information about the number of points. Also,

218 in the sparsification technique, there is no analog of weight handling/rounding. Another
 219 crucial difference is that they do not use a routine analogous to our refresh-pointers routine
 220 because their internal-node guarantees are always fresh. As we discussed before, these refresh
 221 pointers are critical for us also in making sure that the error introduced by the unavoidable
 222 rounding of output weights of the static-algorithm instances is kept in check.

223 1.1 Related Work

224 The most related work is by Cohen-Addad et al. [10] who give an $O(1)$ -coreset for k -median
 225 and k -means in *amortized* update time of $O(n \log n)$.

226 For k -median and k -means, the first coreset-construction algorithms were by Har-Peled
 227 and Mazumdar [17] for Euclidean metrics and by Chen [9] for general metrics. Improved
 228 algorithms computing smaller coresets were later obtained by Har-Peled and Kushal [16]
 229 and by Feldman and Langberg [12]. The current known best is by Braverman et al. [4]:
 230 $O(\varepsilon^{-2} k \log k \log n)$ -size coresets in $\tilde{O}(nk)$ time, who also give an excellent summary of the
 231 literature on coresets that we highly recommend. Note that by merge-and-reduce technique,
 232 each improvement also gave rise to better (insertion-only) streaming coreset constructions.
 233 For k -median and k -means, Frahling and Sohler [14] gave the first coreset-construction
 234 algorithm in the dynamic-streaming setting where points can be added or removed. It uses
 235 space and update time of $O(\text{poly}(\varepsilon^{-1}, \log m, \log \Delta))$ for constant k and dim when the points
 236 lie in the discrete Euclidean metric space $\{1, \dots, \Delta\}^{\text{dim}}$; for k -median, this was recently
 237 improved to $O(\varepsilon^{-2} k \text{poly}(\text{dim}, \log \Delta))$ space and update time of $O(\text{poly}(\varepsilon^{-1}, k, \text{dim}, \log \Delta))$
 238 by Braverman et al. [5]. Coreset constructions with improvements in certain parameters in
 239 the Euclidean settings have been obtained [13, 24].

240 The k -median and k -means problems have received significant attention in the algorithms
 241 community [8, 20, 19, 7, 2, 23, 21, 15, 22, 1, 6]. The best approximation ratio for k -median
 242 is 2.675 by Byrka et al. [6] and that for k -means is $9 + \varepsilon$ by Ahmadian et al. [1].

243 2 Preliminaries

244 Let us fix a problem P with the input X_w , the set of feasible solutions Q , and the linear
 245 cost function $c : Q \times \mathcal{W} \rightarrow \mathbb{R}_+$, where \mathcal{W} is the set of all weighted subsets⁴ of X_w . All the
 246 numbers encountered are nonnegative.

247 The computational model

248 The input set X_w is a weighted set of n points having rational weights whose numerators
 249 and denominators are bounded by $O(\text{poly}(n))$. The algorithm works in the random access
 250 machine model with word size $O(\log n)$. Each memory word can be accessed in constant
 251 time. With each update, a new point is inserted, an existing point is deleted, or the weight
 252 of an existing point is modified by adding or subtracting a nonnegative number. The net
 253 weight of each point always stays nonnegative with its numerator and denominator always
 254 bounded by $O(\text{poly}(n))$.

255 We will prove some basic lemmas about coresets. Using these, we can take weighted
 256 union of two coresets without any loss (Lemma 5) and take a coreset of a coreset without
 257 much loss (Lemma 6).

⁴ To be precise: denote unweighted version of X_w by X' , then \mathcal{W} is essentially $\mathbb{R}_+^{X'}$.

36:6 Fully-Dynamic Coresets

258 ► **Lemma 5.** *If C_w^1 and C_w^2 are ε -coresets of X_w^1 and X_w^2 , respectively, with respect to a*
 259 *linear cost function c such that $X_w^1 \cap X_w^2 = \emptyset$, then $C_w^1 \cup C_w^2$ is an ε -coreset of $X_w^1 \cup X_w^2$.*

260 **Proof.** By linearity of c : for any $S \in Q$,

$$261 \quad c(S, X_w^1 \cup X_w^2) = c(S, X_w^1) + c(S, X_w^2) \in [1 \pm \varepsilon] (c(S, C_w^1) + c(S, C_w^2)) = [1 \pm \varepsilon] c(S, C_w^1 \cup C_w^2),$$

262 where, recall that, $C_w^1 \cup C_w^2$ is a weighted union. ◀

263 ► **Lemma 6.** *If C'_w is an ε -coreset of C_w , and C''_w is a δ -coreset of C'_w , both with respect to*
 264 *c , then C''_w is an $(\varepsilon + \delta + \varepsilon\delta)$ -coreset of C_w with respect to c .*

265 **Proof.** For any $S \in Q$, we have $c(S, C_w) \in [1 \pm \varepsilon] c(S, C'_w)$ and $c(S, C'_w) \in [1 \pm \delta] c(S, C''_w)$. So,

$$266 \quad c(S, C_w) \geq (1 - \varepsilon) c(S, C'_w) \geq (1 - \varepsilon)(1 - \delta) c(S, C''_w) \geq (1 - \varepsilon - \delta - \varepsilon\delta) c(S, C''_w),$$

$$267 \quad \text{and } c(S, C_w) \leq (1 + \varepsilon) c(S, C'_w) \leq (1 + \varepsilon)(1 + \delta) c(S, C''_w) = (1 + \varepsilon + \delta + \varepsilon\delta) c(S, C''_w). \quad \blacktriangleleft$$

268 Let C_w^1 be an ε -coreset of C_w and C_w^2 be an ε -coreset of C_w^1 . Then we say that C_w^1 and
 269 C_w^2 are, respectively, 1-level and 2-level ε -coresets of C_w . Extending this notion, we define
 270 an i -level ε -coreset to be an ε -coreset of an $(i - 1)$ -level ε -coreset.

271 ► **Lemma 7.** *If C_w^ℓ is an ℓ -level ε -coreset of C_w , then C_w^ℓ is a $\left(\sum_{i=1}^{\ell} \binom{\ell}{i} \varepsilon^i\right)$ -coreset of C_w .*

272 **Proof.** The proof is by induction on ℓ . Base case is when $\ell = 1$, and by definition, a 1-level
 273 coreset is an ε -coreset. By induction hypothesis, we have that $C_w^{\ell-1}$ is a $\left(\sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \varepsilon^i\right)$ -
 274 coreset of C_w . Now, C_w^ℓ is an ε -coreset of $C_w^{\ell-1}$, hence C_w^ℓ is an $\left(\varepsilon + (1 + \varepsilon) \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \varepsilon^i\right)$ -
 275 coreset of C_w by Lemma 6. Now, use Lemma 8, which appears below, with $\alpha = \varepsilon$ to finish
 276 the proof. ◀

277 We prove two basic lemmas.

278 ► **Lemma 8.** *For any positive integer ℓ and $\alpha \in \mathbb{R}_+$, we have $\alpha + (1 + \alpha) \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^i =$
 279 $\sum_{i=1}^{\ell} \binom{\ell}{i} \alpha^i$.*

280 **Proof idea.** The proof is provided in Appendix A and uses elementary identities involving
 281 binomial coefficients and algebraic manipulations. ◀

282 ► **Lemma 9.** *For any positive integer ℓ and $\alpha \in [0, 1]$, we have $\sum_{i=1}^{\ell} \binom{\ell}{i} \left(\frac{\alpha}{2\ell}\right)^i \leq \alpha$.*

283 **Proof.** $\sum_{i=1}^{\ell} \binom{\ell}{i} \left(\frac{\alpha}{2\ell}\right)^i \leq \sum_{i=1}^{\ell} \ell^i \frac{\alpha^i}{2^i \ell^i} = \sum_{i=1}^{\ell} \frac{\alpha^i}{2^i} \leq \sum_{i=1}^{\ell} \frac{\alpha}{2^i} \leq \alpha$. ◀

284 Now, as a corollary to Lemma 7, we get the following using Lemma 9.

285 ► **Corollary 10.** *If C_w^ℓ is an ℓ -level $(\varepsilon/(2\ell))$ -coreset of C_w , then C_w^ℓ is an ε -coreset of C_w .*

286 As we discussed earlier, rounding of the weights at internal nodes is needed in our dynamic
 287 algorithm to achieve the desired worst-case update time. Towards that, we need two lemmas.

288 In the next lemma, think of a/b as the original weight of the point, c/d as the weight that
 289 we want to approximate a/b with, and D as the cost of this point with respect to a feasible
 290 solution in Q . So the lemma says that by rounding, the cost of the point stays within $1 \pm b/d$
 291 of the original cost.

292 ► **Lemma 11.** *For positive integers a, b , and d , let $c = \lfloor ad/b \rfloor$. Then $cD/d \in [1 \pm b/d] aD/b$
 293 for any nonnegative real D .*

294 **Proof.** By the definition of c , we have that $c/d \leq a/b \leq c/d + 1/d$, and $1/d \leq a/d$ because
 295 $a \geq 1$; hence $a/b \geq c/d \geq a/b - a/d$, which implies that $aD/b \geq cD/d \geq aD/b - aD/d =$
 296 $(1 - b/d)aD/b$. ◀


297 The proof of the following lemma is very similar. Here, think that we approximate the
 298 weight r of a point by $\lfloor r \rfloor + c/d$ and the cost of the point stays within $1 \pm 1/d$ of the original
 299 cost.

300 ▶ **Lemma 12.** *Let $r \geq 1$ be a rational number, a and b be positive integers such that*
 301 *$a/b = r - \lfloor r \rfloor$, d be any positive integer, and $c = \lfloor ad/b \rfloor$. Then $(\lfloor r \rfloor + c/d)D \in (1 \pm 1/d)rD$*
 302 *for any nonnegative real D .*

303 **Proof.** By the definition of c and using $r \geq 1$, we get that $a/b \geq c/d \geq a/b - r/d$; adding
 304 $\lfloor r \rfloor$ and multiplying by D finishes the proof. ◀

305 **3 A Dynamic Coreset**

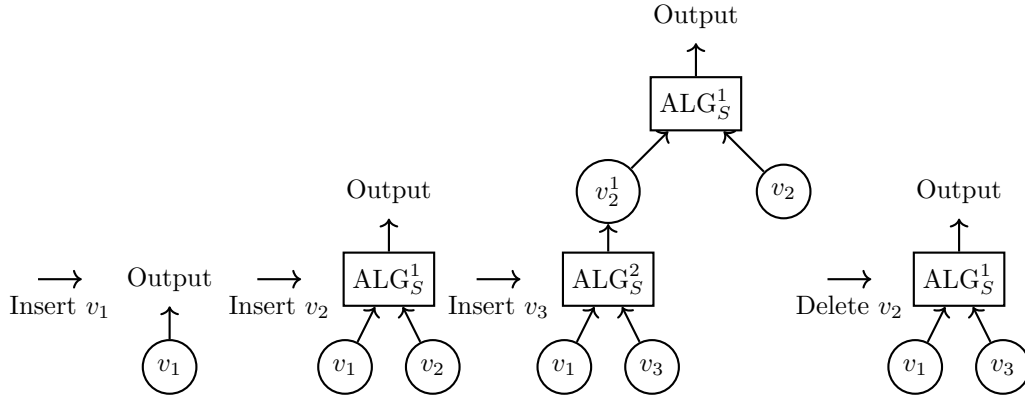
306 We describe our dynamic algorithm for maintaining an ε -coreset for a problem P with query
 space Q that uses a static coreset algorithm, say, ALG_S .



■ **Figure 1** An ALG_S node takes input from two point-nodes. If the union of the sets has cardinality greater than s' , then the ALG_S node computes a coreset of cardinality at most s' and passes it on to the point-node above it (its parent). The number of leaf nodes is always n , and the number of levels is always $O(\log n)$, where n is the current number of points.

307 The main idea is described in Figure 1 using a tree with a special structure. Each node
 308 is of one of the two types: a point-node representing a weighted set of points or an alg-node
 309 representing an instance of ALG_S . We sometimes use a point-node to denote the point set it
 310 represents and an alg-node to denote the ALG_S instance it represents. Each level contains
 311 either only point-nodes or only alg-nodes. All leaf nodes are point-nodes and represent a
 312 weighted singleton with an input point. Each alg-node gets as input the weighted union of
 313 its children, and its output is represented by its parent node (which is a point-node). When
 314 running ALG_S at an alg-node A , if the union of its children has cardinality larger than s' ,
 315 then A would compute a coreset of cardinality at most s' otherwise it would just output the
 316 weighted union. We will later fix this threshold s' for computing a coreset. An example of
 317

318 how insertions and deletions are handled is shown in Figure 2 (where all weights are assumed
 319 to be one). For the ease of description, from now onwards, we will think of this tree with
 320 alg-nodes being collapsed into their parent nodes. Then each leaf node would contain a
 321 weighted singleton and each internal node would contain the output of the ALG_S instance
 322 run on the weighted union of its children's sets.



■ **Figure 2** An example of how insertions and deletions are handled. We start with an empty tree. The first point that is inserted is represented by v_1 . We use a point and the node that represents it interchangeably. Then v_2 is inserted followed by v_3 . Next, if v_4 is inserted, we get exactly the tree shown in Figure 1, and if v_2 is deleted, then we get the last tree.

323 We guarantee that the resulting tree then will always be a *complete* binary tree, i.e.,
 324 every level except possibly the lowest is completely filled, and the nodes at the lowest level
 325 are packed to the left. To describe the updates briefly, let ℓ_r denote the rightmost leaf
 326 node at the lowest level; for simplicity, assume that the lowest level is not full. Insertion is
 327 straightforward: the new point goes in a new leaf node to the right of ℓ_r . For deletion of a
 328 point at leaf node ℓ_d , if $\ell_d \neq \ell_r$, then we replace contents of ℓ_d with those of ℓ_r and delete ℓ_r .
 329 See Section 3.1 for details of these operations. For weight update, the tree does not change.

330 ► **Remark 13.** Since a coreset will not be computed until a node has more than s' points,
 331 the tree can be modified so that each leaf node corresponds to a set of $\Theta(s')$ points. Then
 332 the number of nodes in the tree is $\Theta(n/s')$. This reduces the additional space used for
 333 maintaining this tree. This is important when the number of points is very large. See
 334 Section 3.2 for further details. This is essentially the same idea as used for asymmetric
 335 sparsification in Section 3.4 in Eppstein et al. [11].

336 We call the leaf nodes at the same level as that of the leftmost leaf node to be at level 0.
 337 We increment these level numbers naturally as we move upwards in the tree. Since we
 338 maintain a complete binary tree, the root, which is at the highest level, is on level $\lceil \log n \rceil$.
 339 After a point insertion, deletion, or weight update, we recompute all the nodes that are
 340 affected by running ALG_S from scratch. Once we update a leaf node, all the nodes on its
 341 leaf-to-root path are affected. Since at most two leaf nodes are updated after every point
 342 update, we run at most $2\lceil \log n \rceil$ instances of ALG_S . Finally, to reduce the cardinality of
 343 our output coreset, we run another *outer* instance of ALG_S with $\varepsilon_s = \varepsilon/3$ and $\lambda_s = \lambda/2$
 344 with input as the output of the root. Here, ε_s and λ_s are parameters for ALG_S as described
 345 below, and our goal is to compute an ε -coreset with probability at least $1-\lambda$. The outer
 346 instance is run after every update.

347 The static coreset algorithm ALG_S takes as input an integer weighted set of n_s points

with total weight W_s and always returns a weighted set of cardinality at most $s(\varepsilon_s, \lambda_s, W_s)$; this set is an ε_s -coreset with probability at least $1-\lambda_s$. Let the running time of ALG_S be $t(n_s, \varepsilon_s, \lambda_s, W_s)$. We assume that the functions t and s are nondecreasing in W_s and nonincreasing in ε_s and λ_s , and also that t is nondecreasing in n_s . We call such functions t and s *well-behaved*.

We note that t and s implicitly depend on the query space Q as well. In particular, for k -median and k -means, they depend on k and the dimension or the cardinality of the universe from which a solution is allowed to be picked. Also, assume that the total weight of ALG_S 's output is at most $1+\delta$ times the total input weight and it outputs a coreset of points with integer weights. For the dynamic algorithm, n denotes the current number of points, and we assume that any input weight is a rational number with numerator and denominator bounded by n^c , for a fixed constant c .

► **Theorem 14.** *Assume that there is a static algorithm ALG_S that takes as input an integer-weighted set of n_s points with total weight W_s and always returns an integer-weighted set of cardinality at most $s(\varepsilon_s, \lambda_s, W_s)$ with total weight at most $(1+\delta)W_s$, and this set is an ε_s -coreset with probability at least $1-\lambda_s$. Let the running time of ALG_S be $t(n_s, \varepsilon_s, \lambda_s, W_s)$, and assume that both s and t are well-behaved. Then there is a fully-dynamic algorithm that, on rational-weighted input points, always maintains an $s(\frac{\varepsilon}{3}, \frac{\lambda}{2}, W_p)$ -cardinality weighted set. This set is an ε -coreset with probability at least $1-\lambda$. Its worst-case update time is*

$$O\left(t\left(2s^*, \frac{\varepsilon}{6\lceil\log n_p\rceil}, \frac{\lambda}{2n_p}, W_p\right) \cdot \left(1 + \log(1+\delta) + \frac{\log \varepsilon^{-1}}{\log n}\right) \cdot \log n\right),$$

where $W_p = (1+\delta)^{\lceil\log n_p\rceil} n_p^{c'} \lceil 1/\varepsilon \rceil$, c' is a constant, $s^* = s\left(\frac{\varepsilon}{6\lceil\log 2n_p\rceil}, \frac{\lambda}{4n_p}, W_p\right)$, and $8n/3 \leq n_p \leq 8n$.

Proof. We first prove that the output of the algorithm is an ε -coreset if every non-outer ALG_S instance outputs an ε_s -coreset of its input for some $\varepsilon_s \leq \varepsilon/(6\lceil\log n\rceil)$ and the outer ALG_S instance outputs an $(\varepsilon/3)$ -coreset of its input. We prove the following by induction on level number: every node at level ℓ contains a $(\sum_{i=1}^{\ell} \binom{\ell}{i} \varepsilon_s^i)$ -coreset of the leaf nodes in its subtree. In the base case, a node at level 1 contains an ε_s -coreset of its input trivially. An ALG_S instance A at level i gets as input two sets, say C'_w and C''_w , each of which is a $(\sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \varepsilon_s^i)$ -coreset for the leaf nodes in their respective nodes' subtrees. Hence, $C'_w \cup C''_w$ is a $(\sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \varepsilon_s^i)$ -coreset for leaf nodes in the subtree rooted at A by Lemma 5. Now, A outputs an ε_s -coreset of $C'_w \cup C''_w$, hence by Lemma 6, its output is an $(\varepsilon_s + (1+\varepsilon_s) \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \varepsilon_s^i)$ -coreset of the leaf nodes in its subtree, which, by Lemma 8, means a $(\sum_{i=1}^{\ell} \binom{\ell}{i} \varepsilon_s^i)$ -coreset. This completes the induction step. Hence, the root node, which is at level $\lceil\log n\rceil$, contains $(\sum_{i=1}^{\lceil\log n\rceil} \binom{\lceil\log n\rceil}{i} \varepsilon_s^i)$ -coreset. Now, since $\varepsilon_s \leq \varepsilon/(6\lceil\log n\rceil)$, by Lemma 9, the output at the root is an $(\varepsilon/3)$ -coreset. The outer ALG_S instance outputs an $(\varepsilon/3)$ -coreset of this, hence, by Lemma 6, the final output is an $(2\varepsilon/3 + \varepsilon^2/9)$ -coreset, which is an ε -coreset of all points.

Recall that the running time of ALG_S is $t(n_s, \varepsilon_s, \lambda_s, W_s)$ to compute an ε_s -coreset with probability at least $1-\lambda_s$, where n_s is the number of points in the input. Our output success probability will depend on λ_s , and ε depends on ε_s as proved in the previous paragraph. We will need $\varepsilon_s \leq \varepsilon/(6\lceil\log n\rceil)$ and $\lambda_s \leq \lambda/(2n)$, so these depend on n , which can change a lot over time. We now show how to maintain these guarantees for ε_s and λ_s after each update.

Towards this, we need a little tweak to our algorithm and an additional maintenance routine that we call the *refresher*. The algorithm works in phases. The refresher routine maintains two *refresh* pointers that always point to consecutive leaf nodes, say r_1 and r_2 .

36:10 Fully-Dynamic Coresets

393 The refresh pointers are reset after the end of a phase as follows. If the number of leaf nodes
394 is a power of 2, then r_1 and r_2 point to the two leftmost leaf nodes, otherwise they point to
395 the two leftmost leaf nodes at the level above the lowest level. Assume, for completeness,
396 that the very first phase ends after receiving two points, so the tree is just two leaf nodes
397 and their parent as the root.

398 For each subsequent phase, let n_0 be the value of n at the beginning of the phase. Each
399 phase ends after $n_0/2$ updates, and we set $n_p = 4n_0$. This guarantees that n_p is greater
400 than n throughout the whole phase and even the next phase (details appear below). After
401 receiving an update, we rerun all the ALG_S instances on the leaf-to-root path starting at
402 r_1 and r_2 (at most $2\lceil\log n\rceil$ such instances). This is the refresher routine. Then we move
403 the refresh pointers to the next two leaf nodes on the right. If we reach the right end, then
404 we go to the next level if it exists, otherwise we stop. If we stop, then we achieved the
405 goal of (re-)running all the ALG_S instances that are present at the end of the phase at
406 least once in this phase (this will become clearer below). After the refresher routine, we
407 execute the update which affects at most two leaf nodes. We rerun all the ALG_S instances
408 that are affected by this update, again, at most $2\lceil\log n\rceil$ such instances. So in total, at
409 most $4\lceil\log n\rceil$ of non-outer ALG_S instances are run after an update and one outer instance,
410 which explains the $\log n$ factor in the update time. We now explain the parameters used
411 in the ALG_S instances. For all the non-outer ALG_S instances, we use $\varepsilon_s = \varepsilon/(6\lceil\log n_p\rceil)$
412 and $\lambda_s = \lambda/(2n_p)$. (This explains the ε_s and λ_s parameters of the functions t and s in the
413 theorem statement.) Note here that the running time of the outer instance is going to be
414 less than any non-outer instance because t is non-increasing in ε_s and λ_s .

415 As we use $n_p = 4n_0$ and there could be at most $n_0/2$ insertions in a phase, the final value
416 of n is at most $3n_0/2$, and, thus, n_p is always greater than n . In fact, crucially, n_p is an upper
417 bound on n for even the next phase; in the next phase, $n \leq n_0 + n_0/2 + (n_0 + n_0/2)/2 =$
418 $9n_0/4 \leq n_p$. Also, in the current phase, $n_0/2 \leq n \leq 3n_0/2$, hence $8n/3 \leq n_p \leq 8n$, as
419 required (cf. the theorem statement).

420 We now prove that any non-outer ALG_S instance uses $\varepsilon_s \leq \varepsilon/(6\lceil\log n\rceil)$ and $\lambda_s \leq \lambda/(2n)$
421 at any time instant. Let L be the set of leaf nodes at the beginning of the phase; therefore,
422 $|L| = n_0$. An ALG_S instance that exists at the end of the phase is either on the leaf-to-root
423 path for some leaf in L or it was created/updated in this phase. At the end of the phase, the
424 refresh pointers will hit all surviving leaf nodes in L ; the argument is as follows. Each phase
425 lasts for $n_0/2$ updates, $|L| = n_0$, and we move the two refresh pointers to the right on next
426 two leaf nodes after each update. Importantly, new leaf nodes are added only to the right of
427 the rightmost leaf node at the lowest level, and hence, the refresher routine will have hit all
428 surviving leaf nodes in L before hitting a newly created leaf node.

429 This shows that, in any case (being either hit by a point update or by the refresher
430 routine), each ALG_S instance is run with $n_p = 4n_0$, setting up these instances for the next
431 phase. This means that at any time instant, each ALG_S instance was created/updated in the
432 current phase or created/updated in the previous phase, thus showing that $\varepsilon_s \leq \varepsilon/(6\lceil\log n\rceil)$
433 and $\lambda_s \leq \lambda/(2n)$ for all ALG_S instances at all times.

434 At any time instant, there are at most n non-outer instances of ALG_S , each with success
435 probability at least $1 - \lambda/(2n)$, and the outer ALG_S instance has success probability at least
436 $1 - \lambda/2$. Hence, the final success probability is at least $1 - \lambda$ by the union bound over these
437 $n + 1$ instances.

438 **How to handle weights**

439 We will need one further tweak to argue that each weight ever encountered by the algorithm
 440 can be stored using $O(1 + \log(1+\delta) + \log(1/\varepsilon)/\log n)$ words, which also explains that factor
 441 in the update time. By assumption, an insertion or weight update comes with a weight that is
 442 a fraction with the numerator and the denominator bounded by n^c for some fixed constant c .
 443 After receiving such an update, we approximate the weight by a fraction that has numerator
 444 bounded by $n_p^{c'} \lceil 1/\varepsilon \rceil$, where $c' = 2c + 1$ is also a fixed constant, and the denominator is
 445 equal to $n_p^{c+1} \lceil 1/\varepsilon \rceil^5$. The change in the cost due to this approximation is at most ε/n_p times
 446 the original cost; hence, by the linearity of the cost function, the output coreset quality is
 447 affected by at most an additive factor of $O(\varepsilon/n)$. More formally, the following claim holds
 448 by Lemma 11 and using $b/d \leq \varepsilon/n_p$ below (think of D below as cost).

449 \triangleright **Claim 15.** Let $d = n_p^{c+1} \lceil 1/\varepsilon \rceil$. Given a rational number a/b , where a and b are integers,
 450 $a \leq n_p^c$ and $b \leq n_p^c$, let $f = \lceil ad/b \rceil$. Then $f \leq n_p^{2c+1} \lceil 1/\varepsilon \rceil$ and $(f/d)D \in [1 \pm \varepsilon/n_p](a/b)D$ for
 451 any nonnegative real D .

452 Recall that due to the refresher routine, at any time instant, the denominator of the weight
 453 at any leaf node can be one of the two: $n_p^{c+1} \lceil 1/\varepsilon \rceil$ or $n_{pp}^{c+1} \lceil 1/\varepsilon \rceil$, where n_{pp} is the value of n_p
 454 for the previous phase. When the two children of an internal node use different denominators,
 455 this complicates our rounding scheme. Thus, when taking a union of the children's sets at
 456 an internal node, for each weight, we make its numerator an integer and the denominator
 457 equal to $(n_p n_{pp})^{c+1} \lceil 1/\varepsilon \rceil$, which is a common multiple of $n_p^{c+1} \lceil 1/\varepsilon \rceil$ and $n_{pp}^{c+1} \lceil 1/\varepsilon \rceil$ —the only
 458 possible denominators of an input weight after rounding. Next, we run the ALG_S instance
 459 with integer weights as given by the numerator, then (implicitly) dividing the output weights
 460 by the denominator $(n_p n_{pp})^{c+1} \lceil 1/\varepsilon \rceil$ afterwards. Since each ALG_S instance can increase the
 461 total weight by at most a factor of $1+\delta$, the sum of the numerators of all weights at level i is
 462 always bounded by $n(1+\delta)^i (n_p n_{pp})^{c'} \lceil 1/\varepsilon \rceil$. Since $i \leq \lceil \log n \rceil$ and $n_{pp} = \Theta(n_p)$, there exists
 463 a constant c'' , such that the sum of the numerators of all weights at any level i and all the
 464 possible numerators and denominators are bounded by $(1+\delta)^{\lceil \log n_p \rceil} n_p^{c''} \lceil 1/\varepsilon \rceil =: W_p$, and
 465 hence, can be stored in $O(1 + \log(1+\delta) + \log(1/\varepsilon)/\log n)$ words as desired (see the beginning
 466 of the paragraph before Claim 15). This also justifies the W_s parameters of the functions t
 467 and s in the theorem statement.

468 Now we put everything together. The outer ALG_S instance outputs a weighted set of
 469 size at most $s(\frac{\varepsilon}{3}, \frac{\lambda}{2}, W_p)$. This set is an ε -coreset with probability at least $1-\lambda$, which
 470 we proved by a union bound over all ALG_S instances. We set $s' = s(\frac{\varepsilon}{6^{\lceil \log n_p \rceil}}, \frac{\lambda}{2n_p}, W_p)$,
 471 which is the threshold for computing a coreset at each internal node, i.e., (recall that) if the
 472 number of points at an internal node is greater than s' , then we run ALG_S to compute a
 473 coreset. An upper bound on the threshold for the current phase and the previous phase is
 474 $s^* = s(\frac{\varepsilon}{6^{\lceil \log 2n_p \rceil}}, \frac{\lambda}{4n_p}, W_p)$ because the n_p value for the previous phase can be at most twice
 475 that of the current phase. Then the worst-case update time is dominated by the non-outer

⁵ The static algorithm ALG_S expects integer-weighted input and outputs integer-weighted points, whereas our dynamic algorithm handles fractional weights. If fractional weights are naively stored in our dynamic algorithm, then at internal nodes, combining two fractions may result in larger magnitude numbers. E.g., naively handling two points with weights a/b and c/d so as to be used in ALG_S results in weights $ad/(bd)$ and $bc/(bd)$. Thus, at level i , the numerators and denominators may be as large $(\text{poly}(n))^{2^i}$. Note that some rounding would be needed even if ALG_S can handle rational weights, because its output may be points with rational weights having much larger magnitude; e.g., even if the output magnitude is about only quadratic in that of the input, the blowup near the root in our dynamic algorithm would be n th power of the input. In fact, we do this rounding in the proof of Theorem 2.

36:12 Fully-Dynamic Coresets

476 ALG_S instances, each running in time $t\left(2s^*, \frac{\varepsilon}{2^{\lceil \log n_p \rceil}}, \frac{\lambda}{2n_p}, W_p\right)$, and we run $O(\log n)$ of
 477 these after receiving an update. An additional factor of $1 + \log(1+\delta) + \log(1/\varepsilon)/\log n$ appears
 478 because each weight may need memory worth $O(1 + \log(1+\delta) + \log(1/\varepsilon)/\log n)$ words, and
 479 we need constant time to access each memory word. ◀

480 Before proving the concrete bounds for k -median and k -means that are stated in Theorem 2,
 481 we prove a weaker theorem that is a direct consequence of Theorem 14 using the static
 482 algorithm of Chen [9].

483 ► **Theorem 16.** *For the k -median and k -means problems, there is a fully-dynamic algorithm
 484 that maintains a set of cardinality $O(\varepsilon^{-2}k \log^2(n/\varepsilon)(k \log n + \log(1/\lambda)))$, that is an ε -coreset
 485 with probability at least $1-\lambda$, and has worst-case update time*

$$486 \quad O\left(\varepsilon^{-2}k^2 \log^3 n \log^2 \frac{n}{\varepsilon} \log \frac{n}{\lambda} \left(k \log n + \log \frac{n}{\lambda}\right) \log \log \frac{n}{\varepsilon} \left(1 + \frac{\log \varepsilon^{-1}}{\log n}\right)\right).$$

487 *Ignoring the $\log \log n$ factors, for $\lambda = \Omega(1/\text{poly}(n))$ and $\varepsilon = \Omega(1/\text{poly}(n))$, the coreset
 488 cardinality is $O(\varepsilon^{-2}k^2 \log^3 n)$, and the worst-case update time is $O(\varepsilon^{-2}k^3 \log^7 n)$.*

489 **Proof.** Chen's algorithm takes in an integer weighted set and outputs also an integer weighted
 490 set. Its output has the same total weight as the input, so $\delta = 0$ (see Theorem 14). Also,
 491 for Chen's algorithm, $s(\varepsilon_s, \lambda_s, W_s) = O(\varepsilon_s^{-2}k(k \log n + \log(1/\lambda_s)) \log^2 W_s)$ and the running
 492 time $t(n_s, \varepsilon_s, \lambda_s, W_s) = O(n_s k \log(1/\lambda_s) \log \log W_s)$ (see Theorems 3.6 and 5.5 in Chen [9]),
 493 which is dominated by the computation of a bicriteria approximation. Note that both s and
 494 t are well-behaved. Using $W_p = O(\text{poly}(n)/\varepsilon)$,

$$495 \quad s^* = O\left(\varepsilon^{-2}k \log^2 n \log^2 \frac{n}{\varepsilon} \left(k \log n + \log \frac{n}{\lambda}\right)\right),$$

496 and $\delta = 0$ in Theorem 14 gives the desired bounds using the functions t and s above. ◀

497 Now we use the result of Braverman et al. [4] to get better bounds as stated in Theorem 2
 498 in the introduction section. Unfortunately, we cannot use Theorem 14 as a complete black
 499 box for this because in this case, on integer weighted input, ALG_S does not output an integer
 500 weighted coreset. The proof of the following theorem is thus an extension of the proof of
 501 Theorem 14.

502 ► **Theorem 2.** *For the k -median and k -means problems, there is a fully-dynamic algorithm
 503 that maintains a set of cardinality $O(\varepsilon^{-2}k(\log n \log k \log(k\varepsilon^{-1} \log n) + \log(1/\lambda)))$, that is an
 504 ε -coreset with probability at least $1-\lambda$, and has worst-case update time*

$$505 \quad O\left(\varepsilon^{-2}k^2 \log^5 n \log^3 k \log^2(\varepsilon^{-1})(\log \log n)^3\right),$$

506 *assuming that $\varepsilon = \Omega(1/\text{poly}(n))$ and $\lambda = \Omega(1/\text{poly}(n))$.*

507 **Proof.** Our dynamic algorithm expects to have at its disposal a static algorithm ALG_S that
 508 takes integer-weighted input and outputs an integer-weighted coreset. Since the algorithm
 509 of Braverman et al. that we use as ALG_S outputs on *integer* weighted input a coreset with
 510 *fractional* weights, we need some modifications. Hence, before ALG_S is ready to be used in
 511 the dynamic algorithm, we round its output to turn it into integers.

512 **Weight-Rounding Modifications for ALG_S**

513

- 514 ■ Let input to ALG_S be Y_w which is a set of n_s points with integer weights $w(1), \dots, w(n_s)$.
- 515 ■ We run ALG_S on the same points with scaled weights $s'w(1), \dots, s'w(n_s)$, where s' is
- 516 the desired cardinality of the output coresets (which is the same as the threshold for
- 517 computing a coresets at an internal node in this case). We set s' later in a such a way that
- 518 it can be computed by our dynamic algorithm. This step of multiplying input weights by
- 519 s' is done to make sure that each of the fractional weights output by ALG_S is at least 1
- 520 (see Line 6 of Algorithm 2 in Braverman et al. [4]).
- 521 ■ Let the output C_w of ALG_S be a weighted set of s' points with fractional weights
- 522 $w_o(1), \dots, w_o(s')$. Using the rounding strategy of Lemma 12, round these fractional
- 523 weights to have an integer numerator and the denominator equal to $\lceil (\log n_p)/\varepsilon \rceil$ to get
- 524 weights $\tilde{w}(1), \dots, \tilde{w}(s')$, where n_p is as defined in the proof of Theorem 14. Formally, for
- 525 $i \in \{1, \dots, s'\}$:

$$526 \quad \tilde{w}(i) = \lfloor w_o(i) \rfloor + \frac{\lfloor (w_o(i) - \lfloor w_o(i) \rfloor) \lceil \frac{\log n_p}{\varepsilon} \rceil \rfloor}{\lceil \frac{\log n_p}{\varepsilon} \rceil}.$$

527

Since $w_o(i) \geq 1$, by Lemma 12, for any real $D \geq 0$, we have $\tilde{w}(i)D \in [1 \pm \varepsilon / \log n_p]w_o(i)D$.

528

- 528 ■ Hence, by the linearity of the cost function, C_w with weights $\tilde{w}(1)/s', \dots, \tilde{w}(s')/s'$
- 529 is an $(\varepsilon_s + 2\varepsilon/\log n_p)$ -coresets of Y_w with weights $w(1), \dots, w(n_s)$ if C_w with weights
- 530 $w_o(1), \dots, w_o(s')$ is an ε_s -coresets of Y_w with weights $s'w(1), \dots, s'w(n_s)$. Note that
- 531 $\tilde{w}(i)/s'$ can be represented as a fraction with an integer numerator and denominator
- 532 equal to $s' \lceil (\log n_p)/\varepsilon \rceil$.

533

- 533 ■ The additive loss of $2\varepsilon/\log n_p$ in the coresets quality due to this rounding is tolerable
- 534 because every non-outer ALG_S instance will be run with $\varepsilon_s = O(\varepsilon/\log n_p)^6$. Hence, the
- 535 coresets quality at internal nodes will always be $O(\varepsilon_s + \varepsilon/\log n_p) = O(\varepsilon/\log n)$, as desired.

536

- 536 ■ This rounding ensures that on integer-weighted input with total weight W , the output
- 537 weights of ALG_S are fractions with integer numerator bounded by $(1+\delta)W s' \lceil (\log n_p)/\varepsilon \rceil$
- 538 and integer denominator equal to $s' \lceil (\log n_p)/\varepsilon \rceil$. Here, $1+\delta$ is the factor by which ALG_S
- 539 can increase the total weight.

540

540 To handle rational weights in the dynamic algorithm, we first proceed as described in the

541 paragraph on how to handle weights in the proof of Theorem 14. Recall that we assume that

542 each insertion or weight update by the adversary comes with a weight that is a fraction with

543 the numerator and the denominator bounded by n^c for some fixed constant c , and we set

544 $c' = 2c+1$. Also, each leaf node was created/updated in the current phase or created/updated

545 in the previous phase and thus uses the value either n_p or n_{pp} , where n_{pp} is the value of n_p

546 for the previous phase. We then showed the following. At any time instant, the weight of

547 the point at a leaf node is rounded in such a way that the numerator is bounded by $n_p^{c'} \lceil 1/\varepsilon \rceil$

548 and the denominator is equal to $n_p^{c'+1} \lceil 1/\varepsilon \rceil$, or the numerator is bounded by $n_{pp}^{c'} \lceil 1/\varepsilon \rceil$ and

549 the denominator is equal to $n_{pp}^{c'+1} \lceil 1/\varepsilon \rceil$. Due to this rounding, the output coresets quality is

550 affected by at most an additive factor of $\max\{2\varepsilon/n_p, 2\varepsilon/n_{pp}\} = O(\varepsilon/n)$. We now prove the

551 following more general statement towards the current proof.

⁶ If we go for smaller additive loss, say ε/n_p , the denominators of resulting numbers due to this rounding would become exponential in n_p . And if we go for a larger additive loss, it would worsen the coresets quality at non-outer instances to $\omega(\varepsilon/\log n_p)$ resulting in the quality of the output coresets worse than ε .

36:14 Fully-Dynamic Coresets

552 ► **Lemma 17.** *At any time instant, every weight at a node at level i has an integer numerator*
 553 *and a denominator that is a factor of $(n_p n_{pp})^{c+1} \lceil 1/\varepsilon \rceil (s'_p s'_{pp} \lceil (\log n_p)/\varepsilon \rceil \lceil (\log n_{pp})/\varepsilon \rceil)^i =:$*
 554 *$D(i)$, where s'_p and s'_{pp} are values of the threshold s' in the current and the previous phase,*
 555 *respectively.*

556 **Proof.** We prove this statement by induction over the sequence of nodes updated by the
 557 algorithm.

558 In the base case, the first ever node update will be due to creation of a leaf node, and the
 559 weight will have denominator $n_p^{c+1} \lceil 1/\varepsilon \rceil$. Next we discuss the induction step. Let the update
 560 be on a node at level i , so we run the *modified* ALG_S instance with all weights having a
 561 denominator that is a factor of $D(i-1)$, which is true by induction hypothesis. Then, since
 562 the modified ALG_S adds a factor of $s'_p \lceil (\log n_p)/\varepsilon \rceil$ to the denominator, all resulting output
 563 weights have a denominator that is a factor of $D(i-1) s'_p \lceil (\log n_p)/\varepsilon \rceil$, which is a factor of
 564 $D(i)$. This finishes the induction step for the case when the node update is not the last of
 565 the phase. When the node being updated is the last of the phase, we have to be careful. In
 566 this case, we need to show that for all weights in all nodes, n_{pp} or s'_{pp} do not appear in the
 567 denominator, as this will set these denominators for the next phase. Towards this, we need
 568 the following claim.

569 ▷ **Claim 18.** Let u be a node at level i . Fix a time instant. Suppose, in the current phase,
 570 all nodes in the subtree rooted at u were updated and u was updated after the update of the
 571 last-updated leaf node in the subtree. Then the denominator of the weights at u is a factor
 572 of $n_p^{c+1} \lceil 1/\varepsilon \rceil (s'_p \lceil (\log n_p)/\varepsilon \rceil)^i$ at the fixed time instant.

573 We omit the proof of this claim as it can be proved easily by induction on the level number
 574 at any fixed time instant.

575 After the last node update of the phase, every node in the tree has been updated in the
 576 current phase and the premise of Claim 18 holds due to the refresher routine. Hence, by
 577 Claim 18, after the last node update of the phase, i.e., just before the new phase begins, all
 578 denominators at level i are a factor of $n_p^{c+1} \lceil 1/\varepsilon \rceil (s'_p \lceil (\log n_p)/\varepsilon \rceil)^i$. Since n_p and s'_p of this
 579 phase will become n_{pp} and s'_{pp} in the next phase, the induction hypothesis stays true for the
 580 next phase as well. This finishes the proof of Lemma 17. ◀

581 Since an ALG_S instance may increase the total weight by at most a factor of $1+\delta$, the
 582 sum of the numerators of weights at any level i is at most

$$583 \quad n_p (1+\delta)^i (n_p n_{pp})^{c'} \left\lceil \frac{1}{\varepsilon} \right\rceil \left(s'_p s'_{pp} \left\lceil \frac{\log n_p}{\varepsilon} \right\rceil \left\lceil \frac{\log n_{pp}}{\varepsilon} \right\rceil \right)^i ;$$

584 this can be seen by an easy induction on the level number. Using this bound, we set the
 585 threshold s' in a way similar to that in the proof of Theorem 14: we set

$$586 \quad s'_p = s \left(\frac{\varepsilon}{6 \lceil \log n_p \rceil}, \frac{\lambda}{2n_p}, W_p \right),$$

587 where

$$588 \quad W_p = (1+\delta)^{\lceil \log n_p \rceil} n_p^{c_1} \left(k \left\lceil \frac{\log n_p}{\varepsilon} \right\rceil \right)^{c_2 \lceil \log n_p \rceil},$$

589 and c_1 and c_2 are chosen to be large enough constants so that W_p upper bounds the sum of the
 590 numerators of all weights at any level. From now onwards, we assume that $\lambda = \Omega(1/\text{poly}(n))$.

591 For ALG_S , the function s is $s(\varepsilon_s, \lambda_s, W_s) = O(\varepsilon_s^{-2} k (\log k \log W_s + \log(1/\lambda_s)))$ and $\delta = O(\varepsilon)$.
 592 Then, using $n_{pp} = \Theta(n_p)$, we get that both s'_p and s'_{pp} are $O\left(\left(k \left\lceil \frac{\log n_p}{\varepsilon} \right\rceil\right)^{c_3}\right)$, where c_3 is a
 593 fixed constant (so, independent of c_1 and c_2). Observe that W_p and thus s'_p are determined
 594 by the phase and hence can be computed by our algorithm. More concretely, we get that
 595 both s'_p and s'_{pp} are

$$596 \quad O\left(\varepsilon^{-2} k \log^3 n \log k \log \frac{k \log n}{\varepsilon}\right).$$

597 All possible numerators and denominators encountered by the algorithm are bounded by

$$598 \quad N := O\left(\text{poly}(n) \left(\frac{k \log n}{\varepsilon}\right)^{O(\log n)}\right),$$

599 so, can be stored in $m := (\log N) / \log n = O(\log((k \log n)/\varepsilon))$ words.

600 The running time of ALG_S is $t(n_s, \varepsilon_s, \lambda_s, W_s) = O(n_s k \log(1/\lambda_s) \log \log W_s)$, which,
 601 similar to Chen's algorithm, is dominated by computation of a bicriteria approximation. At
 602 a non-outer ALG_S instance, $n_s = O(s'_p)$, $\varepsilon_s = O(\varepsilon/\log n_p)$, $\lambda_s = O(\lambda/n_p)$, and $W_s \leq W_p$.
 603 With every update, $O(\log n)$ instances of ALG_S are run, and an additional m factor appears
 604 because a weight may need up to m words. Hence, the worst-case update time assuming
 605 $\varepsilon = \Omega(1/\text{poly}(n))$ and $\lambda = \Omega(1/\text{poly}(n))$ is

$$606 \quad O\left(t\left(s'_p, \frac{\varepsilon}{\log n}, \frac{\lambda}{n}, W_p\right) m \log n\right) = O\left(\varepsilon^{-2} k^2 \log^5 n \log k \log^2\left(\frac{k \log n}{\varepsilon}\right) \log \log\left(\frac{k \log n}{\varepsilon}\right)\right)$$

607 and a looser, easier to parse, bound is $O(\varepsilon^{-2} k^2 \log^5 n \log^3 k \log^2(1/\varepsilon)(\log \log n)^3)$. The
 608 output coreset cardinality is

$$609 \quad s\left(\frac{\varepsilon}{3}, \frac{\lambda}{2}, W_p\right) = O\left(\varepsilon^{-2} k \left(\log n \log k \log\left(\frac{k \log n}{\varepsilon}\right) + \log \frac{1}{\lambda}\right)\right).$$

610 This finishes the proof of Theorem 2. ◀

611 3.1 The Binary-Tree Structure

612 We describe the tree structure in more detail, especially, how insertions and deletions are
 613 handled. We always maintain a complete binary tree, in which every level except possibly the
 614 lowest is completely filled, and the nodes in the lowest level are packed to the left. We also
 615 maintain the property that each internal node has exactly two children. Our data structure
 616 behaves somewhat like a heap, though a crucial difference is that we do not have keys. This
 617 structure supports insertion and deletion of a leaf node. Insertion of a new leaf-node ℓ works
 618 as follows.

- 619 ■ If the current number of leaf nodes is a power of 2, then let v be the leftmost leaf node,
- 620 ■ Else let v be the leftmost leaf node in the level above the lowest level.
- 621 ■ Let p be v 's parent.
- 622 ■ Create a new node u .
- 623 ■ Make p to be u 's parent; u replaces v , so if v was p 's right (respectively, left) child, then
 624 u is now p 's right (respectively, left) child.
- 625 ■ Make v to be u 's left child and ℓ to be u 's right child. This way, ℓ the rightmost leaf node
 626 at the lowest level.

627 Deletion of a leaf-node ℓ works as follows. Let v be the rightmost leaf node at the lowest
 628 level, p be v 's parent, and v' be v 's sibling. Replace ℓ 's contents by v 's contents and replace
 629 p 's contents by the contents of v' . Delete v and v' .

630 3.2 Reducing the Number of Nodes

631 The tree can be modified to have each leaf node correspond to a set of $\Theta(s')$ points to reduce
 632 the additional space used for maintaining this tree (pointers and such). Recall that s' is the
 633 threshold for computing a coreset. To reduce the number of nodes in the tree this way, we
 634 maintain the invariant that each leaf node, except possibly one, contains a set of size s_ℓ with
 635 $s'/2 \leq s_\ell \leq s'$. To maintain this invariant, we use a pointer p_s that points to a leaf node
 636 with less than $s'/2$ elements if such a leaf node exists.

637 Whenever a point is inserted, we add it to the leaf node, say ℓ_e pointed to by p_s . If ℓ_e
 638 now contains at least $s'/2$ points, then we make p_s a null pointer. If p_s was a null pointer
 639 already, then we create a new leaf node, say ℓ_n , insert the new point in ℓ_n , and make p_s
 640 point to ℓ_n . The new leaf node ℓ_n is inserted in the tree as described in Section 3.1.

641 Whenever a point is deleted, we check if the leaf node, say ℓ_d that contains it now contains
 642 less than $s'/2$ points. If ℓ_d contains less than $s'/2$ points, and p_s points to some leaf node,
 643 say ℓ_e , then we move points in ℓ_d into ℓ_e and delete ℓ_d . (Deletion of a leaf node is handled as
 644 described in Section 3.1.) If p_s does not point to any leaf node, then we make it point to ℓ_d .

645 As usual, we recompute all nodes on the affected leaf-to-root path.

646 4 Lower Bounds

647 In this section, we show lower bounds. We first see a space lower bound and then a conditional
 648 lower bound on the time per operation.

649 4.1 Space Lower Bound

650 We show a simple and very general space lower bound. Consider any problem that on
 651 input X has to output a feasible solution that is a subset of X . Moreover, if X non-empty,
 652 then all feasible solutions are also non-empty. Call such a problem *compliant*. Clearly,
 653 computing any bounded approximation for k -median and k -means and the problem of
 654 constructing any coreset with bounded quality are compliant. To get a linear space lower
 655 bound for fully-dynamic algorithms that solve a compliant problem, we use the communication
 656 problem of INDEX. In INDEX $_N$, Alice's input is an N -bit string and Bob's input is an index
 657 $I \in \{1, 2, \dots, N\}$. Alice sends one message to Bob, and he needs to correctly output the bit
 658 at position I . By a well-known communication complexity lower bound, Alice must send
 659 a message of size $(1 - H_2(3/4))N \geq 2N/11$ bits so that Bob can correctly output with a
 660 success probability of $3/4$; here H_2 is the binary entropy function.

661 ► **Theorem 19.** *A fully-dynamic algorithm for a compliant problem that works in the presence*
 662 *of an adaptive adversary and has success probability $1 - 1/(8n^2)$ must use space $\Omega(n)$, where*
 663 *n is the current input size.*

664 **Proof.** We describe the reduction for any compliant problem in a metric space, such as
 665 1-median or 1-means, but it can be naturally generalized to any compliant problem. Alice
 666 defines

$$667 \quad X = \{j : j\text{th bit in her string} = 1\},$$

668 and distance between any two points of X to be 1. She runs the fully-dynamic algorithm on
 669 X and sends the memory snapshot to Bob. Bob queries for a solution and if X is nonempty,
 670 a nonempty solution S_1 would be returned. He deletes the points in S_1 and queries again to
 671 get S_2 , and so on until \emptyset is returned. There would be at most N such queries. Note that

672 this works because the algorithm works under an adaptive adversary. If one of the S_ℓ s in
 673 this process contains I , which is Bob's input for the index problem, then Bob outputs 1, else
 674 he outputs 0. In the worst case, Bob makes N queries, where query number i would have
 675 failure probability at most $1/(8(N - i + 1)^2)$. So overall failure probability by the union
 676 bound is at most

$$677 \quad \sum_{i=1}^N \frac{1}{8(N - i + 1)^2} \leq \frac{1}{8} \sum_{i=1}^{\infty} \frac{1}{i^2} = \frac{1}{8} \frac{\pi^2}{6} \leq \frac{1}{4}.$$

678 Alice communicated as many bits as the space usage of the dynamic algorithm. Then, by the
 679 INDEX_N lower bound, the space usage of the algorithm is at least $2N/11 \geq 2n/11$ bits. ◀

680 4.2 Conditional Lower Bounds on the Time Per Operation

681 Now, we show conditional lower bounds on the time per update and query for fully-dynamic
 682 k -means algorithms. They are based on the *OMv-conjecture* [18]: You are given an $N \times N$
 683 Boolean matrix M that can be preprocessed in polynomial time. Then, an online sequence
 684 of N -dimensional Boolean vectors v^1, \dots, v^N is presented and the task is to compute each
 685 Mv^i (using Boolean matrix-vector multiplication) before seeing the next vector v^{i+1} . The
 686 conjecture is that finding all the N answers takes time $\Omega(N^{3-\gamma})$ for any constant $\gamma > 0$.
 687 In [18] also the following *OuMv problem* was presented: You are given an $N \times N$ Boolean
 688 matrix M that can be preprocessed in polynomial time and an online sequence of Boolean
 689 vector pairs $(u^1, v^1), \dots, (u^N, v^N)$ with the goal to compute each $(u^i)^T Mv^i$ (using Boolean
 690 matrix-vector multiplication) before seeing the next vector pair (u^{i+1}, v^{i+1}) . Under the OMv
 691 conjecture, finding N answers for the OuMv problem such that the error probability is at
 692 most $1/3$ takes time $\Omega(N^{3-\gamma})$ for any constant $\gamma > 0$. We will show a reduction from the
 693 latter problem to prove the following result.

694 ▶ **Theorem 20.** *Let $\gamma > 0$ be a constant. Under the OMv conjecture, for any $\delta > 0$, there*
 695 *does not exist a fully-dynamic algorithm that maintains a $(4 - \delta)$ -approximation for k -means*
 696 *with amortized update time $O(k^{1-\gamma})$ and query time $O(k^{2-\gamma})$ such that over a polynomial*
 697 *number of updates the error probability is at most $1/3$.*

698 **Proof.** For the ease of presentation, we assume that k is even; if k is odd, the construction
 699 can be easily adapted. We set $N = k/2$. Given an OuMv instance with $N \times N$ matrix M ,
 700 we construct the following metric space with distance function d from it:

701 The metric space U consists of $4N$ points numbered from 1 to $4N$. For any $1 \leq i < j \leq N$
 702 and $N + 1 \leq i < j \leq 2N$, the distance $d(i, j) = 2$. Furthermore, for $1 \leq i \leq N$ and
 703 $N + 1 \leq j \leq 2N$, the distance $d(i, j) = 1$ if $M_{i, j-N} = 1$, and $d(i, j) = 2$ otherwise.
 704 Additionally, all $2N$ points $2N + 1, \dots, 4N$ are at distance 100 from each other and from all
 705 the other points.

706 We use a k -means data structure to solve a $u^T Mv$ computation as follows: Initially the set
 707 X is empty. When given a vector pair (u, v) , let p be the number of ones in v and in u . Note
 708 that $p \leq 2N$. We insert the points i such that $u_i = 1$ and the points j such that $v_{j-N} = 1$
 709 into X and additionally $2N + 1 - p$ of the points ℓ with $\ell > 2N$. Thus $|X| = 2N + 1 = k + 1$.
 710 Then we ask a k -means query. Afterwards, we delete the inserted points.

711 If $u^T Mv = 1$, then there exist indices i and j such that $u_i = 1, M_{i, j} = 1$, and $v_j = 1$.
 712 Consider the optimal solution that consists of all points in X except for point i . Note that
 713 the cost of this solution for the k -means problem is 1.

714 If $u^T Mv = 0$, then any optimal solution must also consist of $2N + 1 - p$ of the points
 715 ℓ with $\ell > 2N$, and all but one of the other points in X . But as none of the points in X

716 has distance smaller than 2 to any other point in X , the cost of the solution is at least 4 for
 717 k -means. Thus, any $(4 - \delta)$ -approximation for k -means can distinguish between the cases
 718 $u^T Mv = 1$ and $u^T Mv = 0$. Hence, the OMv conjecture implies that it takes at least time
 719 $\Omega(N^{2-\gamma})$ time to execute the above $2N$ update operations and 1 query operation. This
 720 implies the claimed lower bound. ◀

721 — References —

- 722 1 Sara Ahmadian, Ashkan Norouzi-Fard, Ola Svensson, and Justin Ward. Better guarantees for
 723 k -means and euclidean k -median by primal-dual algorithms. In *58th IEEE Annual Symposium*
 724 *on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*,
 725 pages 61–72. IEEE Computer Society, 2017.
- 726 2 Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and
 727 Vinayaka Pandit. Local search heuristics for k -median and facility location problems. *SIAM*
 728 *J. Comput.*, 33(3):544–562, 2004.
- 729 3 Jon Louis Bentley and James B. Saxe. Decomposable searching problems I: Static-to-dynamic
 730 transformation. *J. Algorithms*, pages 301–358, 1980.
- 731 4 Vladimir Braverman, Dan Feldman, and Harry Lang. New Frameworks for Offline and
 732 Streaming Coreset Constructions. *arXiv e-prints*, 2016.
- 733 5 Vladimir Braverman, Gereon Frahling, Harry Lang, Christian Sohler, and Lin F. Yang.
 734 Clustering high dimensional dynamic data streams. In *Proceedings of the 34th International*
 735 *Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*,
 736 pages 576–585, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- 737 6 Jaroslaw Byrka, Thomas W. Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh.
 738 An improved approximation for k -median and positive correlation in budgeted optimization.
 739 *ACM Trans. Algorithms*, 13(2):23:1–23:31, 2017.
- 740 7 Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for facility location
 741 problems. *SIAM J. Comput.*, 34(4):803–824, 2005.
- 742 8 Moses Charikar, Sudipto Guha, Éva Tardos, and David B. Shmoys. A constant-factor
 743 approximation algorithm for the k -median problem. *J. Comput. Syst. Sci.*, 65(1):129–149,
 744 2002.
- 745 9 Ke Chen. On coresets for k -median and k -means clustering in metric and euclidean
 746 spaces and their applications. *SIAM Journal on Computing*, 39(3):923–947, 2009. URL:
 747 <https://doi.org/10.1137/070699007>, arXiv:<https://doi.org/10.1137/070699007>, doi:
 748 [10.1137/070699007](https://doi.org/10.1137/070699007).
- 749 10 Vincent Cohen-Addad, Niklas Oskar D Hjuler, Nikos Parotsidis, David Saulpic, and Chris
 750 Schwiegelshohn. Fully dynamic consistent facility location. In *Advances in Neural Information*
 751 *Processing Systems 32*, pages 3255–3265. Curran Associates, Inc., 2019. URL: [http://papers.](http://papers.nips.cc/paper/8588-fully-dynamic-consistent-facility-location.pdf)
 752 [nips.cc/paper/8588-fully-dynamic-consistent-facility-location.pdf](http://papers.nips.cc/paper/8588-fully-dynamic-consistent-facility-location.pdf).
- 753 11 David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Amnon Nissenzweig. Sparsification—
 754 a technique for speeding up dynamic graph algorithms. *J. ACM*, 44(5):669–696, 1997.
- 755 12 Dan Feldman and Michael Langberg. A unified framework for approximating and clustering
 756 data. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San*
 757 *Jose, CA, USA, 6-8 June 2011*, pages 569–578, 2011.
- 758 13 Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning big data into tiny data:
 759 Constant-size coresets for k -means, PCA and projective clustering. In *Proceedings of the*
 760 *Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New*
 761 *Orleans, Louisiana, USA, January 6-8, 2013*, pages 1434–1453. SIAM, 2013.
- 762 14 Gereon Frahling and Christian Sohler. Coresets in dynamic geometric data streams. In
 763 *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD,*
 764 *USA, May 22-24, 2005*, pages 209–217. ACM, 2005.

- 765 15 Anupam Gupta and Kanat Tangwongsan. Simpler analyses of local search algorithms for
766 facility location, 2008. [arXiv:0809.2554](https://arxiv.org/abs/0809.2554).
- 767 16 Sariel Har-Peled and Akash Kushal. Smaller coresets for k -median and k -means clustering.
768 *Discret. Comput. Geom.*, 37(1):3–19, 2007.
- 769 17 Sariel Har-Peled and Soham Mazumdar. On coresets for k -means and k -median clustering. In
770 *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '04,
771 page 291–300. Association for Computing Machinery, 2004.
- 772 18 Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak.
773 Unifying and strengthening hardness for dynamic problems via the online matrix-vector
774 multiplication conjecture. In *Symposium on Theory of Computing (STOC)*, pages 21–30, 2015.
775 doi:10.1145/2746539.2746609.
- 776 19 Kamal Jain, Mohammad Mahdian, and Amin Saberi. A new greedy approach for facility
777 location problems. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing*,
778 *May 19-21, 2002, Montréal, Québec, Canada*, pages 731–740. ACM, 2002.
- 779 20 Kamal Jain and Vijay V. Vazirani. Approximation algorithms for metric facility location
780 and k -median problems using the primal-dual schema and lagrangian relaxation. *J. ACM*,
781 48(2):274–296, 2001.
- 782 21 Tapas Kanungo, David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman,
783 and Angela Y. Wu. A local search approximation algorithm for k -means clustering. *Comput.*
784 *Geom.*, 28(2-3):89–112, 2004.
- 785 22 Shi Li and Ola Svensson. Approximating k -median via pseudo-approximation. *SIAM J.*
786 *Comput.*, 45(2):530–547, 2016.
- 787 23 Ramgopal R. Mettu and C. Greg Plaxton. Optimal time bounds for approximate clustering.
788 *Mach. Learn.*, 56(1-3):35–60, 2004.
- 789 24 Christian Sohler and David P. Woodruff. Strong coresets for k -median and subspace approx-
790 imation: Goodbye dimension. In *59th IEEE Annual Symposium on Foundations of Computer*
791 *Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 802–813. IEEE Computer
792 Society, 2018.

793 **A** Proof of Lemma 8

794 ► **Lemma 8.** *For any positive integer ℓ and $\alpha \in \mathbb{R}_+$, we have*

$$795 \quad \alpha + (1 + \alpha) \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^i = \sum_{i=1}^{\ell} \binom{\ell}{i} \alpha^i.$$

Proof.

$$796 \quad \alpha + (1 + \alpha) \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^i = \alpha + \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^i + \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^{i+1}$$

$$797 \quad = \binom{\ell-1}{0} \alpha + \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^i + \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^{i+1}$$

$$798 \quad \quad \quad \text{using the fact } \binom{\ell-1}{0} = 1$$

$$799 \quad = \binom{\ell-1}{0} \alpha + \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^i + \sum_{i=2}^{\ell} \binom{\ell-1}{i-1} \alpha^i$$

$$800 \quad \quad \quad \text{change of index in the second summation}$$

$$801 \quad = \sum_{i=1}^{\ell-1} \binom{\ell-1}{i} \alpha^i + \sum_{i=1}^{\ell} \binom{\ell-1}{i-1} \alpha^i$$

36:20 Fully-Dynamic Coresets

802

incorporating first term in second summation

803

$$= \sum_{i=1}^{\ell} \binom{\ell-1}{i} \alpha^i + \sum_{i=1}^{\ell} \binom{\ell-1}{i-1} \alpha^i$$

804

using the fact $\binom{\ell-1}{\ell} = 0$

805

$$= \sum_{i=1}^{\ell} \left(\binom{\ell-1}{i} + \binom{\ell-1}{i-1} \right) \alpha^i$$

806

$$= \sum_{i=1}^{\ell} \binom{\ell}{i} \alpha^i,$$

807

808 where we use $\binom{\ell}{i} = \binom{\ell-1}{i} + \binom{\ell-1}{i-1}$ in the last step. ◀