

Competitive Data-Structure Dynamization*

CLAIRE MATHIEU, CNRS, France

RAJMOHAN RAJARAMAN[†], Northeastern University, USA

NEAL E. YOUNG[‡], University of California Riverside, USA

ARMAN YOUSEFI, Google, USA

Data-structure dynamization is a general approach for making static data structures dynamic. It is used extensively in geometric settings and in the guise of so-called *merge (or compaction)* policies in big-data databases such as LevelDB and Google Bigtable. Previous theoretical work is based on worst-case analyses for uniform inputs – insertions of one item at a time and non-varying read rate. In practice, merge policies must not only handle batch insertions and varying read/write ratios, they can take advantage of such non-uniformity to reduce cost on a per-input basis.

To model this, we initiate the study of data-structure dynamization through the lens of competitive analysis, via two new online set-cover problems. For each, the input is a sequence of disjoint sets of weighted items. The sets are revealed one at a time. The algorithm must respond to each with a set cover that covers all items revealed so far. It obtains the cover incrementally from the previous cover by adding one or more sets and optionally removing existing sets. For each new set the algorithm incurs *build cost* equal to the weight of the items in the set. In the first problem the objective is to minimize total build cost plus total *query cost*, where the algorithm incurs a query cost at each time t equal to the current cover size. In the second problem, the objective is to minimize the build cost while keeping the query cost from exceeding k (a given parameter) at any time. We give deterministic online algorithms for both variants, with competitive ratios of $\Theta(\log^* n)$ and k , respectively. The latter ratio is optimal for the second variant.

CCS Concepts: • **Theory of computation** → **Online algorithms**; • **Applied computing** → *Enterprise computing infrastructures*; • **Information systems** → *Data management systems*.

Additional Key Words and Phrases: online algorithms, competitive analysis, data-structure dynamization, log-structured merge-tree, compaction

1 Introduction

1.1 Background

A *static* data structure is built once to hold a fixed set of items, queried any number of times, and then destroyed, without changing throughout its lifespan. *Dynamization* is a generic technique for transforming any static container data structure into a *dynamic* one that supports insertions and queries intermixed arbitrarily. (Deletions and updates can be supported as described in Section 1.3.) The dynamic structure holds all items inserted so far in a collection of static containers. Insertions are supported by adding new static containers and deleting old ones. Queries are supported by querying all (current) static containers. Static containers are called *components*. Dynamization has been applied in computational geometry [1, 2, 18, 22, 30], in geometric streaming algorithms [7, 31, 34], and to design external-memory dictionaries [3, 6, 11, 52].

Bentley’s *binary transform* [12, 13], later called the *logarithmic method* [45, 51], is a widely used example. It maintains the invariant that the number of items in each component is a distinct power of two. Each insert operation mimics a binary increment: it destroys the components of size

*A conference version of this paper appeared in SODA 2021 [42]. The journal version was published in ACM Transactions on Algorithms, <https://doi.org/10.1145/3672614>.

[†]Supported by NSF grants 1535929 and 1909363.

[‡]Supported by Google Research Award & NSF grant 1619463.




Fig. 1. Steps 1–11 of the binary transform [12, 13]. Each cell \boxed{i} is a component holding i items, where i is a distinct power of two. In each step one item is inserted and held in the new (top, bolded) component.

$2^0, 2^1, 2^2, \dots, 2^{j-1}$, where $j \geq 0$ is the minimum such that there is no component of size 2^j , and builds one new component of size 2^j , holding the contents of the destroyed components and the inserted item. (See Figure 1.) Meanwhile, each query operation queries all current components, combining the results appropriately for the data type. During n insertions, whenever an item is incorporated into a new component, the item’s new component is at least twice as large as its previous component, so the item is in at most $\log_2 n$ component builds. That is, the worst-case *write amplification* is at most $\log_2 n$. Meanwhile, the number of components never exceeds $\log_2 n$, so each query examines at most $\log_2 n$ components. That is, the worst-case *read amplification* is at most $\log_2 n$.

Bentley and Saxe’s *k-binomial* transform is a variant of the binary transform [13]. It maintains k components at all times, of respective sizes $\binom{i_1}{1}, \binom{i_2}{2}, \dots, \binom{i_k}{k}$ such that $0 \leq i_1 < i_2 < \dots < i_k$. (This decomposition is guaranteed to exist and to be unique. Figure 2 gives an example.) It thus ensures read amplification at most k , independent of n , but its write amplification is at most $(k! n)^{1/k}$, about $\frac{k}{e} n^{1/k}$ for large k . This tradeoff between *worst-case* read amplification and *worst-case* write amplification is optimal up to lower-order terms, as is the tradeoff achieved by the binary transform (see Section 1.4).

These worst-case guarantees on read- and write-amplification hold both for *uniform* inputs (where the inserted items have roughly the same sizes, and insertions and queries occur at uniform and balanced rates) and for non-uniform inputs. But a non-uniform input can be substantially easier, in that it admits a solution with average write amplification (over all inserted items) and average read amplification (over all queries) well below worst case, achieving lower total cost. (Roughly, this is achieved by trading build cost for query cost as the read/write ratio varies. For intuition consider a long sequence of insertions followed by a long sequence of queries.) Worst-case dynamization analyses do not capture this. Indeed, transforms such as those above do not adapt to non-uniformity. Their build and query costs are close to worst case even on non-uniform inputs.

We propose two new dynamization problems—*Min-Sum Dynamization* and *k-Component Dynamization*—that model non-uniform insertions and queries. We consider these as online problems and use competitive analysis to measure how well algorithms for them take advantage of non-uniformity. We introduce new algorithms that have substantially better competitive ratios than existing algorithms.

Relevance to industrial LSM systems. Dynamization algorithms underlie standard implementations of external-memory (i.e., disk-based) ordered dictionaries, where they are called *merge* (or *compaction*) policies [41]. Recently inserted key/value pairs are cached in RAM to the extent possible, while older pairs are stored in immutable (static) on-disk files (the components). Each query (if not resolved in cache) searches the current components for the queried key, using one disk




Fig. 2. Steps 1–11 of the 2-binomial transform [13]. At time t the top and bottom components hold $\binom{i_1}{1}$ and $\binom{i_2}{2}$ items where $0 \leq i_1 < i_2$ and $\binom{i_1}{1} + \binom{i_2}{2} = t$. For example at time $t = 8$, $i_1 = 2$ and $i_2 = 4$. If $i_1 = 0$ there is only one component, the bottom component.

access¹ per component. The components are managed using the merge policy: periodically, the cached pairs are flushed to disk in one batch, which is treated as an inserted item and incorporated by building and deleting components² according to the policy. The build cost captures the time building on-disk components, while the query cost captures the time responding to queries.

O’Neil et al’s seminal *log-structured merge* (LSM) architecture [44] (building on [47, 48]) was one of the first to adapt dynamization to external-memory dictionaries as described above. Its dynamization scheme can be viewed as a parameterized generalization of Bentley’s binary transform. As the parameter varies, the tradeoff it achieves between read amplification and write amplification is optimal (in some parameter regimes) among all external-memory structures [5, 16, 53].

Many subsequent and current industrial systems—including so-called NoSQL and NewSQL databases—have LSM architectures. These include Google’s Bigtable [21] (and Spanner [25]), Amazon’s Dynamo [27], Accumulo (by the NSA) [36], AsterixDB [4], Facebook’s Cassandra [38], HBase and Accordion (used by Yahoo! and others) [15, 32], LevelDB [28], and RocksDB [29].

Non-uniform inputs can be particularly important in production LSM systems, where the sizes of inserted batches can vary by orders of magnitude [15, §2] (see also [9, 10, 17]) and the query and insertion rates can vary substantially with time. As discussed previously, such non-uniform workloads can have optimal cost well below the worst-case cost. Industrial compaction policies do adapt to non-uniformity, but only heuristically. Bigtable’s default compaction policy (which, like the k -binomial transform, is configured by a single parameter k and maintains at most k components) is as follows: in response to each insert (cache flush), create a new component holding the inserted items; then, if there are more than k components, merge the i most-recently created components into one, where $i \geq 2$ is chosen minimally so that, for each remaining component S , the size of S in bytes exceeds the total size of all components newer than S [50]. Both the worst-case build cost and the competitive ratio of this algorithm are suboptimal.

1.2 Problem definitions

The definitions of the two dynamization problems below model insertions and queries. The end of the section gives generalizations that allow updates, deletions, and item expiration as implemented (lazily) in typical LSM systems.

Recall that a *set cover* of a given set S of items is a collection of subsets whose union is S .

¹Database servers are typically configured so that RAM size is 1–3% of disk size, even as RAM and disk sizes grow according to Moore’s law [33, p. 227]. A disk block typically holds at least thousands of items. Hence, an index for every disk component, storing the minimum item in each disk block in the component, fits easily in RAM. Then querying any component (a file storing its items in sorted order) for a given item requires accessing just one disk block, determined by checking the index [33, p. 232].

²Crucially, builds use sequential (as opposed to random) disk access. This is why these systems outperform B^+ trees on write-heavy workloads. See [41, § 2.2.1–2.2.2] for details.

Definitions 1.1. An *input* is a sequence $I = (I_1, I_2, \dots, I_n)$ of pairwise-disjoint sets of weighted items. Each item $x \in I_t$ is said to be “inserted at time t ”. The weight of each item x , denoted $\text{wt}(x)$, must be non-negative.

A *solution* is a sequence $C = (C_1, C_2, \dots, C_n)$, where each C_t is a set cover for the items inserted by time t . That is, $\bigcup_{S \in C_t} S = U_t$, where $U_t = \bigcup_{i=1}^t I_i$. The sets in each C_t are called *components*. The *build cost* at time t is the total weight in *new* sets, that is $\sum_{S \in C_t \setminus C_{t-1}} \text{wt}(S)$, where $\text{wt}(S)$ denotes $\sum_{x \in S} \text{wt}(x)$. (For time $t = 1$ we define C_0 to be the empty set.)

The *query cost* at time t is $|C_t|$, that is, the number of components in the current cover, C_t .

Given an input, the objective of the *Min-Sum Dynamization* problem is to find a solution of minimum *total cost* (the sum of all build costs and query costs over time).

The objective of the *k-Component Dynamization* problem is to find a solution having minimum total build cost, among solutions with maximum query cost k (that is, $\max_t |C_t| \leq k$).

An algorithm is *online* if for every input I it outputs a solution C such that at each time t its cover C_t is independent of $I_{t+1}, I_{t+2}, \dots, I_n$, all build costs $\text{wt}_{t'}(S)$ at times $t' > t$, and n .

An algorithm’s *competitive ratio*, $c^*(m)$, is the supremum, over all inputs with m non-empty insertions, of the cost of the algorithm’s solution divided by the optimum cost for the input.

An algorithm is $c(m)$ -*competitive* if its competitive ratio $c^*(m)$ is at most $c(m)$ for all m .

Remarks for Min-Sum Dynamization. The definition of total read cost (as $\sum_{t=1}^n |C_t|$) models, a-priori, exactly one query per insert. This keeps the problem statement relatively simple. However, applications can have any number of queries per insert. This can be modeled by reduction. To model consecutive queries with no intervening insertions, separate the consecutive queries by artificial insertions with $I_t = \emptyset$ (inserting an empty set). To model consecutive insertions with no intervening query, aggregate the consecutive insertions into a single insertion.

Note that uniformly scaling item weights changes build cost relative to query cost. In LSM applications, each unit of query cost represents the time for one *random-access* disk read, whereas each unit of build cost represents the (much smaller, amortized) time per byte during *sequential* disk reads and writes. To model these costs, take the weight of each item x to be its size in bytes, times the time per byte for a sequential disk read and write, divided by the (much larger) time for one random-access disk read.

Remark for k-Component Dynamization. Among well-studied problems, *Dynamic TCP Acknowledgment* [20, 35], a generalization of the classic ski-rental problem, is perhaps technically closest to k -Component Dynamization. TCP Acknowledgment can be viewed as a continuous-time variant of 2-Component Dynamization in which building a new component that contains all items inserted so far (corresponding to a “TCP-ack”) has cost 1 (regardless of the component weight).

Deletions, updates, and expiration. The problems as defined above model queries and insertions. Next we extend the definitions to allow modelling updates, deletions, and item expiration as they typically happen (lazily) in LSM dictionaries.

In this context we assume each item is a weighted key/value pair, timestamped by insertion time, and possibly having an expiration time. (The item weight is typically proportional to the size in bytes of the key/value pair.) Updates and deletions are lazy (“out of place” [41, §2], [40]): update just inserts an item with the given key/value pair (as usual), while delete inserts an item for the given key with a so-called *tombstone* (a.k.a. *antimatter*) value. Multiple items with the same key may be stored (possibly in multiple components), but only the newest matters: a query, given a key, returns the newest item inserted for that key, or “none” if that item is a tombstone or has expired. When a component S is built, it is “garbage collected”: for each key, among the items in S with that key, only the newest is written to disk—all others are discarded.

To model this, we define three generalizations of the problems. To keep the definitions clean, in each variant the input sets must still be disjoint and the current cover must still contain all items inserted so far. To model updates, deletions, and expirations we only redefine the build cost.

Definitions 1.2. Decreasing Weights. Each item $x \in I_t$ has weights $\text{wt}_t(x) \geq \text{wt}_{t+1}(x) \geq \dots \geq \text{wt}_n(x)$. The cost of building a component $S \subseteq U_t$ at time t is redefined as $\text{wt}_t(S) = \sum_{x \in S} \text{wt}_t(x)$. We use this variant as a stepping stone to the LSM variant, next.

LSM. Each item is a timestamped key/value pair with an expiration time. Given a subset S of items, the set of *non-redundant* items in S , denoted $\text{nonred}(S)$, consists of those that have no newer item in S with the same key. The cost of building a component S at time t , denoted $\text{wt}_t(S)$, is redefined as the sum, over all non-redundant items x in S , of the item weight $\text{wt}(x)$, or the weight of the tombstone item for x if x has expired. The latter weight must be at most $\text{wt}(x)$. Items with the same key may have different weights, must have distinct timestamps, and can occur in different components. For any two items $x \in I_t$ and $x' \in I_{t'}$ with $t < t'$, the timestamp of x must be less than the timestamp of x' . This variant applies to LSM systems.³

General. Instead of weighting the items, build costs are specified directly for sets. At each time t a *build-cost function* $\text{wt}_t : 2^{U_t} \rightarrow \mathbb{R}_+$ is revealed (along with I_t), directly specifying the build cost $\text{wt}_t(S)$ for every possible component $S \subseteq U_t$. The build-cost function must obey the following restrictions, for all times $i \leq t$ and sets $S, S' \subseteq U_t$:

- (R1) sub-additivity:** $\text{wt}_t(S \cup S') \leq \text{wt}_t(S) + \text{wt}_t(S')$. (The cost of building a component holding the union of two sets is at most the combined cost of building two components that hold the respective sets.)
- (R2) suffix monotonicity:** if $S \neq U_t$, then $\text{wt}_t(S \setminus U_i) \leq \text{wt}_t(S)$, (The cost of building a component holding a set S of items is at least the cost of building a component holding just those items in S that were inserted after time i . The exception for $S = U_t$ allows modeling full removal of tombstone items during full merges.)
- (R3) temporal monotonicity:** $\text{wt}_i(S) \geq \text{wt}_t(S)$ (The cost of building a component to hold S does not increase over time. Note, for example, that in the LSM model item expirations can cause the cost to decrease over time.)

We chose Restrictions (R1)–(R3) so that the resulting problem has several competing properties: it should be relatively simple, sufficiently general to model practical LSM systems, and sufficiently restricted to allow competitive online algorithms. The build costs implicit in the LSM and Decreasing Weights variants do obey (R1)–(R3).⁴ The restrictions would also hold, for example, if each item had a fixed weight and $\text{wt}_t(S) = \max_{x \in S} \text{wt}(x)$.

1.3 Statement of results

Min-Sum Dynamization. Recall that the *iterated logarithm* (base e) is the slowly growing function defined inductively by $\log_e^* m = 1 + \log_e^* \log_e m$, with the base case $\log_e^* m = 0$ for $m \leq 1$. (Our analysis will use base $\sqrt{2}$ instead of e . Note that $\log_{\sqrt{2}}^* m = \Theta(\log_e^* m)$, so inside O-notation the base is omitted.)

THEOREM 2.1 (SECTION 2). *For Min-Sum Dynamization, the online algorithm Adaptive-Binary (Figure 3) has competitive ratio $\Theta(\log^* m)$, where $m \leq n$ is the number of non-empty insertions.*

³LSM systems delete tombstone items during full merges (i.e., when building a component $S = U_t$ at time t). This is not captured by the LSM model here, but is captured by the following *general* model. See Section 5.2.

⁴The LSM build cost obeys (R1) because $\text{nonred}(S \cup S') \subseteq \text{nonred}(S) \cup \text{nonred}(S')$. It obeys (R2) because $\text{nonred}(S \setminus U_i) \subseteq \text{nonred}(S)$. It obeys (R3) because the tombstone weight for each item x is at most $\text{wt}(x)$.

Roughly speaking, every 2^j time steps ($j \in \{0, 1, 2, \dots\}$), the algorithm merges all components of weight 2^j or less into one. Figure 5 illustrates one execution of the algorithm. The bound in the theorem is tight for the algorithm.

In contrast, consider the naive adaptation of Bentley’s binary transform (i.e., treat each insertion I_t as a size-1 item, then apply the transform). On inputs with $\text{wt}(I_t) = 1$ for all t the algorithms produce the same (optimal) solution. But, as we show in Lemma A.1 in the appendix, the competitive ratio of the naive adaptation is $\Omega(\log n)$.

Min-Sum Dynamization is a special case of *Set Cover with Service Costs*, for which Buchbinder et al. give a randomized online algorithm [19]. For Min-Sum Dynamization, their bound on the algorithm’s competitive ratio simplifies to $O(\log^2 n)$.

K-Component Dynamization and its generalizations.

THEOREM 3.1 (SECTION 3.1). *For k -Component Dynamization (and consequently for its generalizations) no deterministic online algorithm has ratio competitive ratio less than k .*

THEOREM 3.2 (SECTION 3.2). *For k -Component Dynamization with decreasing weights (and plain k -Component Dynamization) the deterministic online algorithm in Figure 8 is k -competitive.*

For comparison, consider the naive adaptation of Bentley and Saxe’s k -binomial transform to k -Component Dynamization (treat each insertion I_t as one size-1 item, then apply the transform). On inputs with $\text{wt}(I_t) = 1$ for all t , the two algorithms produce essentially the same optimal solution. But, as we show in Lemma A.2 in the appendix, the competitive ratio of the naive adaptation is $\Omega(kn^{1/k})$ for any $k \geq 2$.

Bigtable’s default algorithm (Section 1.1) solves k -Component Dynamization, but its competitive ratio is $\Omega(n)$. For example, with $k = 2$, given an instance with $\text{wt}(I_1) = 3$, $\text{wt}(I_2) = 1$, and $\text{wt}(I_t) = 0$ for $t \geq 3$, it pays $n + 2$, while the optimum is 4. (In fact, the algorithm is memoryless — each C_t is determined by C_{t-1} and I_t . No deterministic memoryless algorithm has competitive ratio independent of n .) Even for uniform instances ($\text{wt}(I_t) = 1$ for all t), Bigtable’s default incurs cost quadratic in n , whereas the optimum is $\Theta(kn^{1+1/k})$.

Bentley and Saxe showed that their solutions were optimal (for uniform inputs) among a restricted class of solutions that they called *arboreal transforms* [13]. Here we call such solutions *newest-first*:

Definition 1.6. A solution C is *newest-first* if at each time t , if $I_t = \emptyset$ it creates no new components, and otherwise it creates one new component, by merging I_t with some $i \geq 0$ newest components into a single component (destroying the merged components). Likewise, C is *lightest-first* if, at each time t with $I_t \neq \emptyset$, it merges I_t with some $i \geq 0$ lightest components. An *algorithm* is *newest-first* (*lightest-first*) if it produces only newest-first (*lightest-first*) solutions.

The Min-Sum Dynamization algorithm Adaptive-Binary (Figure 3) is *lightest-first*. The k -Component Dynamization algorithm Greedy-Dual (Figure 8) is *newest-first*. In a *newest-first* solution, every cover C_t partitions the set U_t of current items into components of the form $\bigcup_{h=i}^j I_h$ for some i, j .

Any *newest-first* algorithm for the decreasing-weights variant of either problem can be “bootstrapped” into an equally good algorithm for the LSM variant:

THEOREM 3.5 (SECTION 3.3). *Any newest-first online algorithm A for k -Component (or Min-Sum) dynamization with decreasing weights can be converted into an equally competitive online algorithm A' for the LSM variant.*

With Theorem 3.2 this gives the following corollary:

COROLLARY 1.8 (SECTION 3.3). *There is a deterministic online algorithm for LSM k -Component Dynamization with competitive ratio k .*

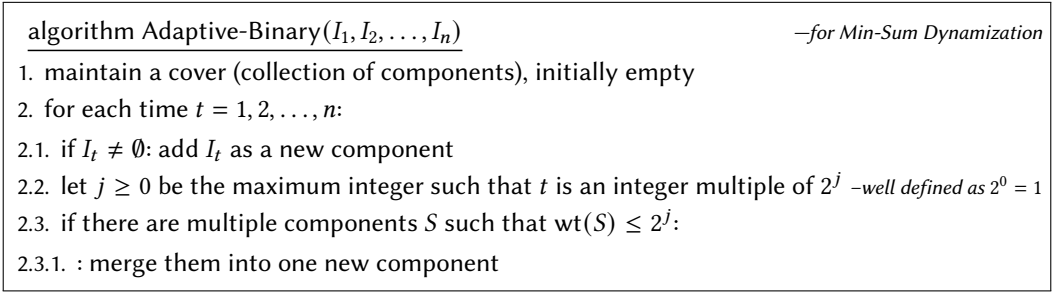


Fig. 3. A $\Theta(\log^* m)$ -competitive algorithm for Min-Sum Dynamization (Theorem 2.1).

Finally we give an algorithm for the general variant:

THEOREM 3.9 (SECTION 3.4). *For general k -Component Dynamization, the deterministic online algorithm B_k in Figure 9 is k -competitive.*

The algorithm B_k partitions the input sequence into phases. Before the start of each phase, it has just one component in its cover, called the current “root”, containing all items inserted before the start of the phase. During the phase, B_k recursively simulates B_{k-1} to handle the insertions occurring during the phase, and uses the cover that consists of the root component together with the (at most $k - 1$) components currently used by B_{k-1} . At the end of the phase, B_k does a *full merge* – it merges all components into one new component, which becomes the new root. It extends the phase maximally subject to the constraint that the cost incurred by B_{k-1} during the phase does not exceed $k - 1$ times the cost of the full merge that ends the phase.

1.4 Properties of optimal offline solutions

Bentley and Saxe showed that, among newest-first solutions (which they called *arboreal*), their various transforms were near-optimal for uniform inputs [12, 13]. Mehlhorn showed (also for uniform inputs) that the best newest-first solutions have cost at most a constant times optimum [43]. We generalize and strengthen Mehlhorn’s result:

THEOREM 4.1 (SECTION 4). *Every instance of k -Component or Min-Sum Dynamization has an optimal solution that is newest-first and lightest-first.*

One consequence is that Bentley and Saxe’s transforms give optimal solutions (up to lower-order terms) for uniform inputs. Another is that, for Min-Sum and k -Component Dynamization, optimal solutions can be computed in time $O(n^3)$ and $O(kn^3)$, respectively, because optimal newest-first solutions can be computed in these time bounds via natural dynamic programs.

The body of the paper gives the proofs of Theorems 2.1–4.1.

2 Min-Sum Dynamization (Theorem 2.1)

THEOREM 2.1. *For Min-Sum Dynamization, the online algorithm Adaptive-Binary (Figure 3) has competitive ratio $\Theta(\log^* m)$, where $m \leq n$ is the number of non-empty insertions.*

We prove the theorem in two parts:

- (i) The competitive ratio is $O(\log^* m)$ (proof in Section 2.1).
- (ii) The competitive ratio is $\Omega(\log^* m)$ (proof in Section 2.2).

2.1 Part (i): the competitive ratio is $O(\log^* m)$

Fix an input $I = (I_1, I_2, \dots, I_n)$ with $m \leq n$ non-empty sets. Let C be the algorithm's solution. Let C^* be an optimal solution, of cost OPT . For any time t , call the 2^j chosen in Line 2.2 the *capacity* $\mu(t)$ of time t , and let S_t be the newly created component (if any) in Line 2.3.

It is convenient to over-count the algorithm's build cost as follows. In Line 2.3, if there is exactly one component S with $\text{wt}(S) \leq 2^j$, the algorithm as stated does not change the current cover, but we pretend for the analysis that it does — specifically, that it destroys and rebuilds S , paying its build cost $\text{wt}(S)$ again at time t . This allows a clean statement of the next lemma. In the remainder of the proof, the “build cost” of the algorithm refers to this over-counted build cost.

We first bound the total query cost, $\sum_t |C_t|$, of C .

LEMMA 2.2. *The total query cost of C is at most twice the (over-counted) build cost of C , plus OPT .*

PROOF. Consider the components with weight less than 1. By inspection of the algorithm each cover C_t has at most one such component — the component S_t created at time t . Therefore, the query cost from the components with weight less than 1 is at most n .

It remains to consider the components with weight at least 1. Let S be any component in C of weight $\text{wt}(S) \geq 1$. Each new occurrence of S in C contributes at most $2 \text{wt}(S)$ to C 's query cost. Indeed, let $2^j \geq \text{wt}(S)$ be the next larger power of 2. Times with capacity 2^j or more occur every 2^j time steps. So, after C creates S , C destroys S within $2^j \leq 2 \text{wt}(S)$ time steps; note that we are using here the over-counted build cost. So C 's query cost from such components is at most twice the build cost of C .

Thus, the total query cost from all components is at most twice the build cost of C plus n . The lemma follows since the query cost of C^* is at least n , so $n \leq \text{OPT}$. \square

Define Δ to be the maximum number of components merged by the algorithm in response to any query. Note that $\Delta \leq m$ simply because there are at most m components at any given time in C . (Only Line 2.1 increases the number of components, and it does so only if I_t is non-empty.) The remainder of the section bounds the build cost of C by $O(\log^*(\Delta) \text{OPT})$. By Lemma 2.2, this will imply prove Part (i) of the theorem.

The total weight of all components I_t that the algorithm creates in Line 2.1 is $\sum_t \text{wt}(I_t)$, which is at most OPT because every $x \in I_t$ is in at least one new component in C^* (at time t). To finish, we bound the (over-counted) build cost of the components that the algorithm builds in Line 2.3, i.e., $\sum_t \text{wt}(S_t)$.

OBSERVATION 2.3. *The difference between any two distinct times t and t' is at least $\min\{\mu(t), \mu(t')\}$.*

(This holds because t and t' are distinct integer multiples of $\min\{\mu(t), \mu(t')\}$. See Figure 4.)

Charging scheme. For each time t at which Line 2.3 creates a new component S_t , have S_t charge to each item $x \in S_t$ the weight $\text{wt}(x)$ of x . Have x in turn charge $\text{wt}(x)$ to each optimal component $S^* \in C_t^*$ that contains x at time t . The entire build cost $\sum_t \text{wt}(S_t)$ is charged to components in C^* . To finish, we show that each component S^* in C^* is charged $O(\log^* \Delta)$ times S^* 's contribution (via its build and query costs) to OPT .

Throughout, given integer times t and t' , let $[t, t']$ denote the (*time*) *interval* $\{t, t+1, \dots, t'\}$. (This is non-standard notation.)

Fix any such S^* . Define $[t_1, t_2]$ to be the interval of S^* in C^* . That is, C^* adds S^* to its cover at time t_1 , where it remains through time t_2 , so its contribution to OPT is $t_2 - t_1 + 1 + \text{wt}(S^*)$. At each (integer) time $t \in [t_1, t_2]$, component S^* is charged $\text{wt}(S^* \cap S_t)$. To finish, we show $\sum_{t=t_1}^{t_2} \text{wt}(S^* \cap S_t) = O(t_2 - t_1 + \log^*(\Delta) \text{wt}(S^*))$.




Fig. 4. The capacities $\mu(t)$ as a function of t .

By Observation 2.3, there can be at most one time $t' \in [t_1, t_2]$ with capacity $\mu(t') > t_2 - t_1 + 1$. If there is such a time t' , the charge received then, i.e. $\text{wt}(S^* \cap S_{t'})$, is at most $\text{wt}(S^*)$. To finish, we bound the charges at the times $t \in [t_1, t_2] \setminus \{t'\}$, with $\mu(t) \leq t_2 - t_1 + 1$.

Definition 2.4 (dominant). Classify each such time t and C 's component S_t as *dominant* if the capacity $\mu(t)$ strictly exceeds the capacity $\mu(i)$ of every earlier time $i \in [t_1, t-1]$ ($\mu(t) > \max_{i=t_1}^{t-1} \mu(i)$) in S^* 's interval $[t_1, t_2]$. Otherwise t and S_t are *non-dominant*.

LEMMA 2.5 (NON-DOMINANT TIMES). *The net charge to S^* at non-dominant times is at most $t_2 - t_1$.*

PROOF. Let τ_1 be any dominant time. Let $\tau_2 > \tau_1$ be the next larger dominant time step, if any, else $t_2 + 1$. Consider the charge to S^* during the open interval (τ_1, τ_2) . We show that this charge is at most $\tau_2 - \tau_1 - 1$.

Component S^* is built at time $t_1 \leq \tau_1$, so $S^* \subseteq U_{\tau_1}$. At time τ_1 , every item x that can charge S^* (that is, $x \in S^*$) is in some component S in C_{τ_1} . By the definition of dominant, each time in $t \in (\tau_1, \tau_2)$ has capacity $\mu(t) \leq \mu(\tau_1)$, since otherwise τ_2 would not be the *next* dominant time. So, the components S in C_{τ_1} that have weight $\text{wt}(S) > \mu(\tau_1)$ remain unchanged in C throughout (τ_1, τ_2) , and the items in them do not charge S^* during (τ_1, τ_2) . So we need only consider items in components S in C_{τ_1} with $\text{wt}(S) \leq \mu(\tau_1)$. Assume there are such components. By inspection of the algorithm, there can only be one: the component S_{τ_1} built at time τ_1 . All charges in (τ_1, τ_2) come from items $x \in S_{\tau_1} \cap S^*$.

Let $\tau_1 = t'_1 < t'_2 < \dots < t'_\ell$ be the times in (τ_1, τ_2) when these items are put in a new component. These are the times in (τ_1, τ_2) when S^* is charged, and, at each, the charge is $\text{wt}(S^* \cap S_{t'_i}) \leq \text{wt}(S_{\tau_1})$, so the total charge to S^* during (τ_1, τ_2) is at most $(\ell - 1) \text{wt}(S_{\tau_1})$.

At each time t'_i with $i \geq 2$ the previous component $S_{t'_{i-1}}$, of weight at least $\text{wt}(S_{\tau_1})$, is merged. So each time t'_i has capacity $\mu(t'_i) \geq \text{wt}(S_{\tau_1})$. By Observation 2.3, the difference between each time t'_i and the next t'_{i+1} is at least $\text{wt}(S_{\tau_1})$. So $(\ell - 1) \text{wt}(S_{\tau_1}) \leq t'_\ell - t'_1 \leq \tau_2 - \tau_1 - 1$.

By the two previous paragraphs the charge to S^* during (τ_1, τ_2) is at most $\tau_2 - \tau_1 - 1$. Summing over the dominant times τ_1 in $[t_1, t_2]$ proves the lemma. \square

Let D be the set of dominant times. For the rest of the proof all times that we consider are dominant. Note that all times that are congested or uncongested (as defined next) are dominant.

Definition 2.6 (congestion). For any time $t \in D$ and component S_t , define the *congestion* of t and S_t to be $\text{wt}(S_t \cap S^*)/\mu(t)$, the amount S_t charges S^* , divided by the capacity $\mu(t)$. Call t and

S_t *congested* if this congestion exceeds κ , and *uncongested* otherwise ($\kappa > 0$ is a constant that is specified later).

LEMMA 2.7 (UNCONGESTED TIMES). *The total charge to S^* at uncongested times is $O(t_2 - t_1)$.*

PROOF. The charge to S^* at any uncongested time t is at most $\kappa\mu(t)$, so the total charge to C^* during such times is at most $\kappa \sum_{t \in D} \mu(t)$. By definition of *dominant*, the capacity $\mu(t)$ for each $t \in D$ is a distinct power of 2 no larger than $t_2 - t_1 + 1$. So $\sum_{t \in D} \mu(t)$ is at most $2(t_2 - t_1 + 1)$, and the total charge to C^* during uncongested times is $O(t_2 - t_1)$. \square

LEMMA 2.8 (CONGESTED TIMES). *The total charge to S^* at congested times is $O(\text{wt}(S^*) \log^* \Delta)$.*

PROOF. Let Z denote the set of congested times. For each item $x \in S^*$, let $W(x)$ be the collection of congested components that contain x and charge S^* . The total charge to S^* at congested times is $\sum_{x \in S^*} |W(x)| \text{wt}(x)$.

To bound this, we use a random experiment that starts by choosing a random item X in S^* , where each item x has probability proportional to $\text{wt}(x)$ of being chosen: $\Pr[X = x] = \text{wt}(x)/\text{wt}(S^*)$.

We will show that $\mathbb{E}_X[|W(X)|]$ is $O(\log^* \Delta)$. Since $\mathbb{E}_X[|W(X)|] = \sum_{x \in S^*} |W(x)| \text{wt}(x)/\text{wt}(S^*)$, this will imply that the total charge is $O(\log^* \Delta) \text{wt}(S^*)$, proving the lemma.

The merge forest for S^ .* Define the following *merge forest*. There is a leaf $\{x\}$ for each item $x \in S^*$. There is a non-leaf node S_t for each congested component S_t . The parent of each leaf $\{x\}$ is the first congested component S_t that contains x (that is, $t = \min\{i \in Z : x \in S_i\}$, if any). The parent of each node S_t is the next congested component $S_{t'}$ that contains all items in S_t (that is, $t' = \min\{i \in Z : i > t, S_t \subseteq S_i\}$, if any). Parentless nodes are roots.

The random walk starts at the root of the tree that holds leaf $\{X\}$, then steps along the path to that leaf in the tree. In this way it traces (in reverse) the sequence $W(X) = \{S_i : X \in S_i\}$ of congested components that X entered during $[t_1, t_2]$. The number of steps is $|W(X)|$. To finish, we show that the expected number of steps is $O(\log^* \Delta)$.

Each non-leaf node S_t in the tree has congestion $\text{wt}(S_t \cap S^*)/\mu(t)$, which is at least κ and at most Δ . For the proof, define the congestion of each leaf x to be 2^Δ . To finish, we argue that *with each step of the random walk, the iterated logarithm of the current node's congestion increases in expectation by at least $1/5$.*

A step in the random walk. Fix any non-leaf node S_t . Let $\alpha_t = \text{wt}(S_t \cap S^*)/\mu(t)$ be its congestion. The walk visits S_t with probability $\text{wt}(S^* \cap S_t)/\text{wt}(S^*)$. Condition on this event (that is, $X \in S_t$). Let random variable α' be the congestion of the child of S_t next visited.

SUBLEMMA 1. *For any $\beta \in [\alpha_t, 2^\Delta)$, $\Pr[\alpha' > \beta | X \in S_t]$ is at least $1 - \alpha_t^{-1}(2 + \log_2 \beta)$.*

PROOF. Consider any child $S_{t'}$ of S_t with $\alpha_{t'} \leq \beta$. We will bound the probability that $S_{t'}$ is visited next (i.e., $X \in S_{t'}$). Node $S_{t'}$ is not a leaf, as $\alpha_{t'} < 2^\Delta$. Define $j(t')$ so that its capacity $\mu(t')$ equals $\mu(t)/2^{j(t')}$. (That is, $j(t') = \log_2(\mu(t)/\mu(t'))$.) The definitions and $\alpha_{t'} \leq \beta$ imply

$$\Pr[X \in S_{t'} | X \in S_t] = \frac{\text{wt}(S_{t'} \cap S^*)}{\text{wt}(S_t \cap S^*)} = \frac{\alpha_{t'} \mu(t')}{\alpha_t \mu(t)} \leq \frac{\beta \mu(t)/2^{j(t')}}{\alpha_t \mu(t)} = \frac{\beta}{\alpha_t 2^{j(t')}}. \quad (1)$$

Also, the algorithm merged a component containing $S_{t'}$ at time t , so $\text{wt}(S_{t'}) \leq \mu(t)$, so

$$\Pr[X \in S_{t'} | X \in S_t] = \frac{\text{wt}(S_{t'} \cap S^*)}{\text{wt}(S_t \cap S^*)} = \frac{\text{wt}(S_{t'} \cap S^*)}{\alpha_t \mu(t)} \leq \frac{\text{wt}(S_{t'})}{\alpha_t \mu(t)} \leq \frac{1}{\alpha_t}. \quad (2)$$

Combining Bounds (1) and (2), $\Pr[X \in S_{t'} | X \in S_t]$ is at most $\alpha_t^{-1} \min(1, \beta 2^{-j(t')})$. Summing this bound over all children $S_{t'}$ of S_t with congestion $\alpha_{t'} \leq \beta$, and using that each $j(t')$ is a distinct

positive integer, the probability that $\alpha' \leq \beta$ is at most

$$\alpha_t^{-1} \sum_{j=1}^{\infty} \min(1, \beta 2^{-j}) \leq \alpha_t^{-1} \int_0^{\infty} \min(1, \beta 2^{-j}) dj = \alpha_t^{-1} (\log_2(\beta) + 1/\ln 2)$$

(splitting the integral at $j = \log_2 \beta$). The sublemma follows from $1/\ln 2 \leq 2$. \square

Next we lower-bound the expected increase in the \log^* of the congestion in this step. We use $\sqrt{2}$ as the base of the iterated \log .⁵ Then $\log^*(2^{\alpha_t/2}) = 1 + \log^* \alpha_t$, so, conditioned on $X \in S_t$,

$$E[\log^* \alpha'] \geq \Pr[\alpha' \geq \alpha_t] \log^* \alpha_t + \Pr[\alpha' \geq 2^{\alpha_t/2}].$$

Bounding the two probabilities above via Sublemma 1 with $\beta = \alpha_t$ and $\beta = 2^{\alpha_t/2}$, the right-hand side above is

$$\begin{aligned} &\geq [1 - \alpha_t^{-1}(2 + \log_2 \alpha_t)] \log^* \alpha_t + [1 - \alpha_t^{-1}(2 + \alpha_t/2)] \\ &= \log^*(\alpha_t) + 1/2 - [2 + (2 + \log_2 \alpha_t) \log^* \alpha_t] / \alpha_t \\ &\geq \log^*(\alpha_t) + 1/2 - 3/10 = \log^*(\alpha_t) + 1/5, \end{aligned}$$

where the last inequality follows from $\alpha_t \geq \kappa$ (t is congested) and by setting $\kappa \geq 64$. It follows that $E[\log^* \alpha' - \log^* \alpha_t | X \in S_t] \geq 1/5$. That is, in each step, the expected increase in the iterated logarithm of the congestion is at least $1/5$.

Let random variable $L = |W(X)|$ be the length of the random walk. Let random variable α'_i be the congestion of the i th node on the walk. By the previous section, for each i , given that $i < L$, $E[\log^* \alpha'_{i+1} - \log^* \alpha'_i | \alpha'_i] \geq 1/5$. It follows by Wald's equation (see [14, p. 370] and [54, Lemma 4.1]) that $E[\log^* \alpha'_L - \log^* \alpha'_1] \geq E[L]/5$. Since $\alpha'_L = 2^\Delta$ and $\log^* \alpha'_1 \geq 0$, we have $E[\log^* \alpha'_L - \log^* \alpha'_1] \leq \log^* 2^\Delta$. It follows that $E[L] \leq 5 \log^* 2^\Delta$. Recall that the base of the iterated logarithm is $\sqrt{2}$; so, $\log^* 2^\Delta = 2 + \log^* \Delta$, yielding $E[L] \leq 10 + 5 \log^* \Delta$. That is, the expected length of the random walk is $O(\log^* \Delta)$. By the discussion at the start of the proof, this implies the lemma. \square

To recap, for each component S_t built by the algorithm, the (over-counted) build cost is charged item by item to those components in the optimal solution C^* that currently contain the item. In this way, the algorithm's total over-counted build cost $\sum_t \text{wt}(S_t)$ is charged to components in C^* . By Lemmas 2.5–2.8, each component S^* in the optimal solution C^* is charged $O(1)$ times its contribution $t_2 - t_1$ to the query cost of C^* plus (in expectation) $O(\log^* m)$ times its contribution $\text{wt}(S^*)$ to the build cost of C^* . It follows that the expected build cost incurred by the algorithm is $O(\log^* m)$ times the cost of C^* .

By Lemma 2.2, the total *query* cost incurred by the algorithm is at most twice the algorithm's over-counted build cost plus the cost of C^* . It follows that the total (build and query) cost incurred by the algorithm is $O(\log^*(m))$ times the cost of C^* . That is, the competitive ratio is $O(\log^* m)$, proving Part (i) of Theorem 2.1.⁶

2.2 Part (ii): the competitive ratio is $\Omega(\log^* m)$

LEMMA 2.9. *The competitive ratio of the Adaptive-Binary algorithm (Figure 3) is $\Omega(\log^* m)$.*

PROOF. We will show a ratio of $\Omega(\log^* m)$ on a particular class of inputs, one for each integer $D \geq 0$. (Figure 5 describes the input I for $D = 2$ and the resulting merge tree, of depth $D + 1$.)




Fig. 5. The “merge tree” for an execution of the Adaptive-Binary algorithm (Figure 3). The input sequence starts with $m = 132$ inserts I_1, I_2, \dots, I_{132} — one for each leaf, of weight equal to leaf’s label. It continues with $2^{16} - 132$ empty inserts ($I_t = \emptyset$). At each time $t = 2^9, 2^{10}, 2^{11}, \dots, 2^{17}$ (during the empty inserts) the algorithm merges all components of weight t to form a single new component, their parent. In this way, the algorithm builds a component for each node, with weight equal to the node’s label. At time $t = 2^{17}$ the final component is built—the root, of weight 2^{18} , containing all items. The algorithm merges each item four times, so pays build cost 4×2^{18} .




Fig. 6. The top three levels of T_∞ . Each node i has $2^{i-p(i)}$ children, where $p(i)$ is the parent of i (exc. $p(1) = 0$). The merge tree T_2^N (Figure 5) consists of these three levels, with each node i given weight $2^{N-p(i)}$, so the nodes with weight 2^{N-i} are the $2^{i-p(i)}$ children of node i , and their total weight equals the weight of node i . Note that the label of a node in the merge tree is its weight, and the merge tree of Figure 5 is T_2^{18} .

1. make node 1 the root
2. for $i \leftarrow 1, 2, 3, \dots$ do:
 - 2.1. let $p(i)$ be the parent of i , except $p(1) = 0$
 - 2.2. give node i the $2^{i-p(i)}$ children $\{c(i-1) + j : 1 \leq j \leq 2^{i-p(i)}\}$, where $c(i-1)$ is the max child of $i-1$, except $c(0) = 1$

Fig. 7. An algorithm defining the tree T_∞ , with nodes $\{1, 2, 3, \dots\}$.

The desired merge tree. For reference, define an infinite rooted tree T_∞ with node set $\{1, 2, 3, \dots\}$ by the iterative process shown in Figure 7. Each iteration i defines the children of node i . Node i has $2^{i-p(i)}$ children, where $p(i)$ is the parent of i (exc. $p(1) = 0$); these children are allocated greedily from the “next available” nodes, so that each node $i \geq 2$ is given exactly one parent. The depth of i is non-decreasing with i .⁷

Figure 6 shows the top three levels of T_∞ . Let n_d be the number of nodes of depth d or less in T_∞ . Each such node i satisfies $i \leq n_d$ (as depth is non-decreasing with i), so, inspecting Line 2.2,

⁵Defined by $\log_{\sqrt{2}}^* \alpha_t = 0$ if $\alpha_t \leq 8$, else $1 + \log_{\sqrt{2}}^*(\log_{\sqrt{2}} \alpha_t)$. Note that $\log_{\sqrt{2}}^* \alpha_t = \Theta(\log_{\epsilon}^* \alpha_t)$.

⁶Curiously, the algorithm’s cost is in fact $O(1)$ times the query cost of C^* plus $O(\log^* m)$ times its build cost.

⁷This follows by induction: Line 2.2 ensures that $p(i') \leq p(i)$ for $i' < i$, so inductively $\text{depth}(i') = 1 + \text{depth}(p(i')) \leq 1 + \text{depth}(p(i)) = \text{depth}(i)$.

node i has at most $2^i \leq 2^{n_d}$ children. Each node of depth $d + 1$ or less is either the root or a child of a node of depth d or less, so $n_{d+1} \leq 1 + n_d 2^{n_d} \leq 2^{2^{n_d}}$. Taking the \log^* of both sides gives $\log^* n_{d+1} \leq 2 + \log^* n_d$. Inductively, $\log^* n_d \leq 2d$ for each d .

Now fix an integer $D \geq 0$. Define the desired merge tree, T_D^N , to be the subtree of T_∞ induced by the nodes of depth at most $D + 1$. Let m be the number of leaves in T_D^N . By the previous paragraph (and $m \leq n_{D+1}$), every leaf in T_D^N has depth $\Omega(\log^* m)$.

Assign weights to the nodes in T_D^N as follows. Fix $N = 2n_D$. Give each node i weight $2^{N-p(i)}$, where $p(i)$ is the parent of i (except $p(1) = 0$). Each weight is a power of two, and the nodes of any given weight 2^{N-i} are exactly the $2^{i-p(i)}$ children of node i . The weight of each parent i equals the total weight of its children.

The input. Define the input I as follows. For each time $t \in \{1, 2, \dots, m\}$, insert a set I_t containing just one item whose weight equals the weight of the t th leaf of T_D^N . Then, at each time $t \in \{m + 1, m + 2, \dots, 2^{N-1}\}$, insert an empty set $I_t = \emptyset$.

In the following, we place a lower bound on the cost of the algorithm on input I . For this, we establish a matching between the algorithm's cover and the leaves of T_D^N , which guides our bound on the build cost of the algorithm.

No merges until last non-empty insertion. The algorithm does no merges before time $\min_{i=1}^m \text{wt}(I_i)$, which is the minimum leaf weight in T_D^N . This is because if a merge occurs at a time t then there must be more than one component of weight at most t at that time (by step 2.3 of the algorithm, see Figure 3). The lightest leaves are the children of node n_D , of weight 2^{N-n_D} . Since the total leaf weight is the weight of the root, 2^N , it follows that $m 2^{N-n_D} \leq 2^N$, that is, $m \leq 2^{n_D} = 2^{N-n_D}$ (using $N = 2n_D$). So, *the algorithm does no merges until time $t(n_D) = 2^{N-n_D}$ (after all non-empty insertions).*

The algorithm's merge tree matches T_D^N . By the previous two paragraphs, just before time $t(n_D) = 2^{N-n_D}$ the algorithm's cover matches the leaves of T_D^N , meaning that the cover's components correspond to the leaves, with each component weighing the same as its corresponding node. The leaves are $\{j : p(j) \leq n_d < j\}$. So the following invariant holds initially, for $i = n_D$:

For each $i \in \{n_D, n_D - 1, \dots, 2, 1\}$, just before time $t(i) = 2^{N-i}$, the algorithm's cover $C_{t(i)}$ matches the nodes in Q_i , defined as

$$Q_i \doteq \{j : 2^{N-j} < t(i) \leq 2^{N-p(j)}\} = \{j : p(j) \leq i < j\}.$$

Informally, these are the nodes j that have not yet been merged by time $t(i)$, because their weight $2^{N-p(j)}$ is at least $t(i)$, but whose children (the nodes of weight 2^{N-j}) if any, have already been merged.

We establish the invariant for all i using a backward induction. Assume inductively that the invariant holds for a given i . We show it holds for $i - 1$. At time $t(i)$, the algorithm merges the components of weight at most $\mu(t(i)) = t(i) = 2^{N-i}$ in its cover. By the invariant, these are the components of weight $t(i) = 2^{N-i}$, corresponding to the children of node i (which are all in Q_i). They leave the cover and are replaced by their union, whose weight equals $2^{N-p(i)}$. Likewise, by the definition (and $p(j) < j$)

$$Q_{i-1} = \{i\} \cup Q_i \setminus \{j : p(j) = i\},$$

so the resulting cover matches Q_{i-1} , with the new component corresponding to node i . The minimum-weight nodes in Q_{i-1} are then $\{j : p(j) = i - 1\}$, the children of node $i - 1$. These have weight $2^{N-(i-1)} = t(i - 1)$, so the algorithm keeps this cover until just before time $t(i - 1)$, so that the invariant is maintained for $i - 1$.

Inductively, we arrive at the validity of the invariant for $i = 1$: just before time $t(1) = 2^{N-1} = n$, the algorithm's cover contains the components corresponding to $\{j : p(j) = 1 < j\}$, with weight $2^{N-p(j)} = 2^{N-1} = n$. At time n they are merged form the final component of weight 2^N , corresponding to the root node 1. So the algorithm's merge tree matches T_D^N .

Competitive ratio. Each leaf in the merge tree has depth $\Omega(\log^* m)$, so every item is merged $\Omega(\log^* m)$ times, and the algorithm's build cost is $\Omega(\text{wt}(1) \log^* m) = \Omega(n \log^* m)$ (using $\text{wt}(1) = 2n$).

But the optimal cost is $\Theta(n)$. (Consider the solution that merges all input sets into one component at time m , just after all non-empty insertions. Its query cost is $\sum_{t=1}^{m-1} t + \sum_{t=m}^n 1 = O(m^2 + n)$. Its merge cost is $2 \text{wt}(1) = O(n)$. Recalling that $m \leq 2^{nD} = 2^{N/2} = O(\sqrt{n})$, the optimal cost is $O(n)$.)

So the competitive ratio is $\Omega(\log^* m)$. \square

Note that in Lemma 2.9, $n \approx m^2$, so $\log^* m = \Omega(\log^* n)$. The upper bound in Section 2.1 and the lower bound in Lemma 2.9 prove Theorem 2.1.

3 K -Component Dynamization and variants (Theorems 3.1–3.9)

3.1 Lower bound on optimal competitive ratio

THEOREM 3.1. *For k -Component Dynamization (and consequently for its generalizations) no deterministic online algorithm has ratio competitive ratio less than k .*

To develop intuition before we give the detailed proof for the general case, here is a sketch of how the proof goes for $k = 2$. The adversary begins by inserting one item of weight 1 and one item of infinitesimal weight $\varepsilon > 0$, followed by a sequence of $n - 2$ weight-zero items just until the algorithm's cover has just one component. (This must happen, or the competitive ratio is unbounded — OPT pays only at time 1, while the algorithm continues to pay at least ε each time step.) By calculation the algorithm pays at least $2 + (n - 1)\varepsilon$, while OPT pays $\min(2 + \varepsilon, 1 + (n - 1)\varepsilon)$, giving a ratio of $1.5 - O(\varepsilon)$.

This lower bound does not reach 2 (in contrast to the standard “rent-or-buy” lower bound) because the algorithm and OPT both pay a “setup cost” of 1 at time 1. However, at the end of sequence, the algorithm and OPT are left with a component of weight ~ 1 in place. The adversary can now continue, doing a second phase without the setup cost, by inserting an item of weight $\sqrt{\varepsilon}$, then zeros just until the algorithm's cover has just one component (again this must happen or the ratio is unbounded). Let m be the length of this second phase. By calculation, for *this* phase, the algorithm pays at least $(m - 1)\sqrt{\varepsilon} + 1$ while OPT pays at most $\min(1 + \sqrt{\varepsilon} + \varepsilon, (m - 1)(\sqrt{\varepsilon} + \varepsilon))$, giving a ratio of $2 - O(\sqrt{\varepsilon})$ for just the phase.

The ratio of the whole sequence (both phases together) is now $1.75 - O(\sqrt{\varepsilon})$. By doing additional phases (using infinitesimal $\varepsilon^{1/i}$ in the i th phase), the adversary can drive the ratio arbitrarily close to 2. This is the idea for $k = 2$. Next we give the detailed proof for the general case (arbitrary $k \geq 2$).

PROOF OF THEOREM 3.1. Fix an arbitrarily small $\varepsilon > 0$. Define $k+1$ sequences of items (weights) as follows. Sequence $\sigma(k+1)$ has just one item, $\sigma_1(k+1) = \varepsilon$. For $j \in \{k, k-1, \dots, 1\}$, in decreasing order, define sequence $\sigma(j)$ to have $n_j = \lceil k/\sigma_1(j+1) \rceil$ items, with the i th item being $\sigma_i(j) = \varepsilon^{n_k + n_{k-1} + \dots + n_j - i + 2}$. Each sequence $\sigma(j)$ is strictly increasing, and all items in $\sigma(j)$ are smaller than all items in $\sigma(j+1)$. Every two items differ by a factor of at least $1/\varepsilon$, so the cost to build any component will be at most $1/(1 - \varepsilon)$ times the largest item in the component.

Adversarial input sequence I. Fix any deterministic online algorithm A. Define the input sequence I to interleave the $k+1$ sequences in $\{\sigma(j) : 1 \leq j \leq k+1\}$ as follows. Start by inserting the only item from sequence $\sigma(k+1)$: take $I_1 = \{\sigma_1(k+1)\} = \{\varepsilon\}$. For each time $t \geq 1$, after A responds

to the insertion at time t , determine the next insertion $I_{t+1} = \{x\}$ as follows. For each sequence $\sigma(j)$, call the most recent (and largest) item inserted so far from $\sigma(j)$, if any, the *representative* of the sequence. Define index $\ell(t)$ so that the largest representative in any new component at time t is the representative of $\sigma(\ell(t))$. (The item inserted at time t is necessarily a representative and in at least one new component, so $\ell(t)$ is well-defined.) At time $t + 1$ choose the inserted item x to be the next unused item from sequence $\sigma(\ell(t) - 1)$. Define the *parent* of x , denoted $p(x)$, to be the representative of $\sigma(\ell(t))$ at time t . (Note: A 's build cost at time t was at least $p(x) \gg x$.) Stop when the cumulative cost paid by A reaches k . This defines the input sequence I .

The input I is well-defined. Next we verify that I is well-defined, that is, that (a) $\ell(t) \neq 1$ for all t (so x 's specified sequence $\sigma(\ell(t) - 1)$ exists) and (b) each sequence $\sigma(j)$ is chosen at most n_j times. First we verify (a). Choosing x as described above forces the algorithm to maintain the following invariants at each time t :

- (i) Each of the sequences in $\{\sigma(j) : \ell(t) \leq j \leq k + 1\}$ has a representative, and
- (ii) no two of these $k - \ell(t) + 2$ representatives are in any one component.

Indeed, the invariants hold at time $t = 1$ when $\ell(t) = k + 1$. Assume they hold at some time t . At time $t + 1$ the newly inserted element x is the new representative of $\sigma(\ell(t) - 1)$ and is in some new component, so $\ell(t + 1) \geq \ell(t) - 1$. These facts imply that Invariant (i) is maintained. By the definition of $\ell(t + 1)$, the components built at time $t + 1$ contain the representative from $\sigma(\ell(t + 1))$ but no representative from any $\sigma(j)$ with $j > \ell(t + 1)$. This and $\ell(t + 1) \geq \ell(t) - 1$ imply that Invariant (ii) is also maintained.

By inspection, Invariants (i) and (ii) imply that A has at least $k - \ell(t) + 2$ components at time t . But A has at most k components, so $\ell(t) \geq 2$.

Next we verify (b), that I takes at most n_j items from each sequence $\sigma(j)$. This holds for $\sigma(k + 1)$ just because, by definition, after time 1, I cannot insert an item from $\sigma(k + 1)$. Consider any $\sigma(j)$ with $j \leq k$. For each item $\sigma_i(j)$ in $\sigma(j)$, when I inserted $\sigma_i(j)$, algorithm A paid at least $p(\sigma_i(j)) \geq \sigma_1(j + 1)$ at the previous time step. So, before all n_j items from $\sigma(j)$ are inserted, A must pay at least $n_j \sigma_1(j + 1) \geq k$ (by the definition of n_j), and the input stops. It follows that I is well-defined.

Upper-bound on optimum cost. Next we upper-bound the optimum cost for I . For each $j \in \{1, \dots, k\}$, define $C(j)$ to be the solution for I that partitions the items inserted so far into the following k components: one component containing items from $\sigma(j)$ and $\sigma(j + 1)$, and, for each $h \in \{1, \dots, k + 1\} \setminus \{j, j + 1\}$, one containing items from $\sigma(h)$.

To bound $\text{cost}(C(j))$, i.e., the total cost of new components in $C(j)$, first consider the new components such that the largest item in the new component is the just-inserted item, say, x . The cost of such a component is at most $x/(1 - \varepsilon)$. Each item x is inserted at most once, so the total cost of all such components is at most $1/(1 - \varepsilon)$ times the sum of all defined items, and therefore at most $\sum_{i=1}^{\infty} \varepsilon^i / (1 - \varepsilon) = \varepsilon / (1 - \varepsilon)^2$. For every other new component, the just-inserted item x must be from sequence $\sigma(j + 1)$, so the largest item in the component is the parent $p(x)$ (in $\sigma(j)$) and the build cost is at most $p(x)/(1 - \varepsilon)$. Defining $m_j \leq n_j$ to be the number of items inserted from $\sigma(j)$, the total cost of building all such components is at most $\sum_{i=1}^{m_j} p(\sigma_i(j)) / (1 - \varepsilon)$. So $\text{cost}(C(j))$ is at most $\varepsilon / (1 - \varepsilon)^2 + \sum_{i=1}^{m_j} p(\sigma_i(j)) / (1 - \varepsilon)$.

The cost of OPT is at most $\min_j \text{cost}(C(j))$. The minimum is at most the average, so

$$(1 - \varepsilon)^2 \text{cost}(\text{OPT}) \leq \min_{j=1, \dots, k} \varepsilon + \sum_{i=1}^{m_j} p(\sigma_i(j)) \leq \varepsilon + \frac{1}{k} \sum_{j=1}^k \sum_{i=1}^{m_j} p(\sigma_i(j)).$$

<pre> algorithm Greedy-Dual(I_1, I_2, \dots, I_n) 1. maintain a cover (collection of components), initially empty 2. for each time $t = 1, 2, \dots, n$ such that $I_t \neq \emptyset$: 2.1. if there are k current components: 2.1.1. increase all components' credits continuously until some component S has $\text{credit}[S] \geq \text{wt}_t(S)$ 2.1.2. let S_0 be the oldest component such that $\text{credit}[S_0] \geq \text{wt}_t(S_0)$ 2.1.3. merge I_t, S_0 and all components newer than S_0 into one new component S' 2.1.4. initialize $\text{credit}[S']$ to 0 2.2. else: 2.2.1. create a new component from I_t, with zero credit </pre>	<i>— for k-Component Dynamization with decreasing weights</i>
---	--

Fig. 8. A newest-first k -competitive algorithm for k -Component Dynamization with decreasing weights (Theorem 3.2).

Lower bound on algorithm cost. The right-hand side of the above inequality is at most $(\varepsilon/k + 1/k) \text{cost}(A)$, because $\text{cost}(A) \geq k$ (by the stopping condition) and $\sum_{j=1}^k \sum_{i=1}^{m_j} p(\sigma_i(j)) \leq \text{cost}(A)$. (Indeed, for each $j \in \{1, \dots, k\}$ and $i \in \{1, \dots, m_j\}$, the item $\sigma_i(j)$ was inserted at some time $t \geq 2$, and A paid at least $p(\sigma_i(j))$ at the previous time $t - 1$.) So the competitive ratio is at least $(1 - \varepsilon)^2 / (\varepsilon/k + 1/k) \geq (1 - 3\varepsilon)k$. This holds for all $\varepsilon > 0$, so the ratio is at least k . \square

3.2 Upper bound for k -Component Dynamization with decreasing weights

Recall that, in k -Component Dynamization with decreasing weights, each item $x \in I_t$ has weights $\text{wt}_t(x) \geq \text{wt}_{t+1}(x) \geq \dots \geq \text{wt}_n(x)$. The cost of building a component $S \subseteq U_t$ at time t is redefined as $\text{wt}_t(S) = \sum_{x \in S} \text{wt}_t(x)$. This variant is technically useful, as a stepping stone to the LSM variant.

THEOREM 3.2. *For k -Component Dynamization with decreasing weights (and plain k -Component Dynamization) the deterministic online algorithm in Figure 8 is k -competitive.*

PROOF. Consider any execution of the algorithm on any input I_1, I_2, \dots, I_n . Let δ_t be such that each component's credit increases by δ_t at time t . (If Block 2.2 is executed, $\delta_t = 0$.) To prove the theorem we show the following lemmas.

LEMMA 3.3. *The cost incurred by the algorithm is at most $k \sum_{t=1}^n (\text{wt}_t(I_t) + \delta_t)$.*

LEMMA 3.4. *The cost incurred by the optimal solution is at least $\sum_{t=1}^n (\text{wt}_t(I_t) + \delta_t)$.*

PROOF OF LEMMA 3.3. As the algorithm executes, keep the components ordered by age, oldest first. Assign each component a *rank* equal to its rank in this ordering. Say that the rank of any *item* is the rank of its current component, or $k + 1$ if the item is not yet in any component. At each time t , when a new component is created in Line 2.1.3, the ranks of the items in S_0 stay the same, but the ranks of all other items decrease by at least 1. Divide the cost of the new component into two parts: the contribution from the items that decrease in rank, and the remaining cost.

Throughout the execution of the algorithm, each item's rank can decrease at most k times, so the total contribution from items as their ranks decrease is at most $k \sum_{t=1}^n \text{wt}_t(I_t)$ (using here that the weights are non-increasing with time). To complete the proof of the lemma, observe that the remaining cost is the sum, over times t when Line 2.1.3 is executed, of the weight $\text{wt}_t(S_0)$ of the component S_0 at time t . This sum is at most the total credit created, because, when a component S_0

is destroyed in Line 2.1.3, at least the same amount of credit (on S_0) is also destroyed. But the total credit created is $k \sum_{t=1}^n \delta_t$, because when Line 2.1.1 executes it increases the total component credit by $k\delta_t$. \square

PROOF OF LEMMA 3.4. Let C^* be an optimal solution. Let C denote the algorithm's solution. At each time t , when the algorithm executes Line 2.1.1, it increases the credit of each of its k components in C_{t-1} by δ_t . So *the total credit the algorithm gives is* $k \sum_t \delta_t$.

For each component $S \in C_{t-1}$, think of the credit given to S as being distributed over the component's items $x \in S$ in proportion to their weights, $\text{wt}_t(x)$: at time t , each item $x \in S$ receives credit $\delta_t \text{wt}_t(x)/\text{wt}_t(S)$. Have each x , in turn, charge this amount to one component in OPT's current cover C_t^* that contains x . In this way, *the entire credit* $k \sum_{t=1}^n \delta_t$ *is charged to components in* C^* .

Recall that $[t, t']$ denotes $\{t, t+1, \dots, t'\}$.

SUBLEMMA 2. *Let x be any item. Let $[t, t']$ be any time interval throughout which x remains in the same component in C . The cumulative credit given to x during $[t, t']$ is at most $\text{wt}_t(x)$.*

PROOF. Let S be the component in C that contains x throughout $[t, t']$. Assume that $\delta_{t'} > 0$ (otherwise reduce t' by one). Let $\text{credit}_{t'}[S]$ denote $\text{credit}[S]$ at the end of iteration t' . Weights are non-increasing with time, so the credit that x receives during $[t, t']$ is

$$\sum_{i=t}^{t'} \frac{\text{wt}_i(x)}{\text{wt}_i(S)} \delta_i \leq \frac{\text{wt}_t(x)}{\text{wt}_{t'}(S)} \sum_{i=t}^{t'} \delta_i \leq \frac{\text{wt}_t(x)}{\text{wt}_{t'}(S)} \text{credit}_{t'}[S].$$

The right-hand side is at most $\text{wt}_t(x)$. (Indeed, in iteration t' Line 2.1.1 increased all components' credits by $\delta_{t'} > 0$, while maintaining the invariant that $\text{credit}[S] \leq \text{wt}_{t'}(S)$, so $\text{credit}_{t'}[S] \leq \text{wt}_{t'}(S)$.) \square

Next we bound how much charge OPT's components (in C^*) receive. For any time t , let $\mathcal{N}_t^* = C_t^* \setminus C_{t-1}^*$ contain the components that OPT creates at time t , and let $N_t^* = \bigcup_{S \in \mathcal{N}_t^*} S$ contain the items in these components. Call the charges received by components in \mathcal{N}_t^* from components created by the algorithm before time t *forward charges*. Call the remaining charges (from components created by the algorithm at time t or after) *backward charges*.

Consider first the *backward* charges to components in \mathcal{N}_t^* . These charges come from components in C_{t-1} , via items x in $N_t^* \cap U_{t-1}$, from time t until the algorithm destroys the component in C_{t-1} that contains x . By Sublemma 2, the total charge via a given x from time t until its component is destroyed is at most $\text{wt}_t(x)$, so the cumulative charge to components in \mathcal{N}_t^* from older components is at most $\text{wt}_t(N_t^* \cap U_{t-1}) = \text{wt}_t(N_t^*) - \text{wt}_t(I_t)$ (using that $N_t^* \setminus U_{t-1} = I_t$). Using that OPT pays at least $\text{wt}_t(N_t^*)$ at time t , and summing over t , *the sum of all backward charges is at most* $\text{cost}(\text{OPT}) - \sum_t \text{wt}_t(I_t)$.

Next consider the *forward* charges, from components created at time t or later, to any component S^* in \mathcal{N}_t^* . Component S^* receives no forward charges at time t , because components created by the algorithm at time t receive no credit at time t . Consider the forward charges S^* receives at any time $t' \geq t+1$. At most one component (in $C_{t'-1}$) can contain items in N_t^* , namely, the component in $C_{t'-1}$ that contains I_t . (Indeed, the algorithm merges components "newest first", so any other component in $C_{t'-1}$ created after time t only contains items inserted after time t , none of which are in N_t^* .) At time t' , the credit given to that component is $\delta_{t'}$, so the components created by the algorithm at time t' charge a total of at most $\delta_{t'}$ to S^* . Let $m(t, t') = |\mathcal{N}_t^* \cap C_{t'}^*|$ be the number of components S^* that OPT created at time t that remain at time t' . Summing over $t' \geq t+1$ and $S^* \in \mathcal{N}_t^*$, the forward charges to components in \mathcal{N}_t^* total at most $\sum_{t'=t+1}^n m(t, t') \delta_{t'}$. Summing over

t , the sum of all forward charges is at most

$$\sum_{t=1}^n \sum_{t'=t+1}^n m(t, t') \delta_{t'} = \sum_{t'=2}^n \delta_{t'} \sum_{t=1}^{t'-1} m(t, t') \leq \sum_{t'=1}^n \delta_{t'} (k-1)$$

(using that $\sum_{t=1}^{t'-1} m(t, t') \leq k-1$ for all t , because OPT has at most k components at time t' , at least one of which is created at time t').

Recall that the entire credit $k \sum_{t=1}^n \delta_t$ is charged to components in C^* . Summing the bounds from the two previous paragraphs on the (forward and backward) charges, this implies that

$$k \sum_{t=1}^n \delta_t \leq \text{cost}(\text{OPT}) - \sum_{t=1}^n \text{wt}_t(I_t) + (k-1) \sum_{t=1}^n \delta_t.$$

This proves the lemma, as it is equivalent to the desired bound $\text{cost}(\text{OPT}) \geq \sum_{t=1}^n \text{wt}_t(I_t) + \delta_t$. \square

This proves Theorem 3.2. \square

3.3 Bootstrapping newest-first algorithms

THEOREM 3.5. *Any newest-first online algorithm A for k -Component (or Min-Sum) dynamization with decreasing weights can be converted into an equally competitive online algorithm A' for the LSM variant.*

PROOF. Fix an instance (I, wt) of LSM k -Component (or Min-Sum) Dynamization. For any solution C to this instance, let $\text{wt}(C)$ denote its build cost using build-cost function wt . For any set S of items and any item $x \in S$, let $\text{nr}(x, S)$ be 0 if x is redundant in S (that is, there exists a newer item in S with the same key) and 1 otherwise. Then $\text{wt}_t(S) = \sum_{x \in S} \text{nr}(x, S) \text{wt}_t(x)$, where $\text{wt}_t(x)$ is $\text{wt}(x)$ unless x is expired, in which case $\text{wt}_t(x)$ is the tombstone weight of x . The tombstone weight of x must be at most $\text{wt}(x)$, so $\text{wt}_t(x)$ is non-increasing with t .

For any time t and item $x \in U_t$, define $\text{wt}'_t(x) = \text{nr}(x, U_t) \text{wt}_t(x)$. For any item x , $\text{wt}'_t(x)$ is non-increasing with t , so (I, wt') is an instance of k -Component Dynamization with decreasing weights. For any solution C for this instance, let $\text{wt}'(C)$ denote its build cost using build-cost function wt' .

LEMMA 3.6. *For any time t and set $S \subseteq U_t$, we have $\text{wt}'_t(S) \leq \text{wt}_t(S)$.*

PROOF. Redundant items in S are redundant in U_t , so

$$\text{wt}'_t(S) = \sum_{x \in S} \text{wt}'_t(x) = \sum_{x \in S} \text{nr}(x, U_t) \text{wt}_t(x) \leq \sum_{x \in S} \text{nr}(x, S) \text{wt}_t(x) = \text{wt}_t(S). \quad (3)$$

\square

LEMMA 3.7. *Let C be any newest-first solution for (I, wt') and (I, wt) . Then $\text{wt}'(C) = \text{wt}(C)$.*

PROOF. Consider any time t with $I_t \neq \emptyset$. Let S be C 's new component at time t (so $C_t \setminus C_{t-1} = \{S\}$). Consider any item $x \in S$. Because C is newest-first, S includes all items inserted with or after x . So x is redundant in U_t iff x is redundant in S , that is, $\text{nr}(x, U_t) = \text{nr}(x, S)$, so $\text{wt}'_t(S) = \text{wt}_t(S)$ (because Bound (3) above holds with equality). Summing over all t gives $\text{wt}'(C) = \text{wt}(C)$. \square

Given an instance (I, wt) of LSM k -Component Dynamization, the algorithm A' simulates A on the instance (I, wt') defined above. Using Lemma 3.7, that A is c -competitive, and $\text{wt}'(\text{OPT}(I, \text{wt}')) \leq \text{wt}(\text{OPT}(I, \text{wt}))$ (by Lemma 3.6), we get

$$\text{wt}(A'(I, \text{wt})) = \text{wt}'(A(I, \text{wt}')) \leq c \text{wt}'(\text{OPT}(I, \text{wt}')) \leq c \text{wt}(\text{OPT}(I, \text{wt})).$$

So A' is c -competitive. \square

<p><u>algorithm $B_1(I_1, I_2, \dots, I_n)$</u> — for $k = 1$</p> <p>1. for $t = 1, 2, \dots, n$: use cover $C_t = \{U_t\}$ where $U_t = \bigcup_{i=1}^t I_i$ — one component holding all items</p>
<p><u>algorithm $B_k(I_1, I_2, \dots, I_n)$</u> — for $k \geq 2$</p> <p>1. initialize $t' = 1$ — variable t' holds the start time of the current phase</p> <p>2. for $t = 1, 2, \dots, n$:</p> <p>2.1. let $C' = B_{k-1}(I_{t'}, I_{t'+1}, \dots, I_t)$ — the solution generated by B_{k-1} for the current phase so far</p> <p>2.2. if the total cost of C' exceeds $(k-1) \text{wt}_t(U_t)$: take $C_t = \{U_t\}$ and let $t' = t + 1$ — end the phase</p> <p>2.3. else: use cover $C_t = \{U_{t'}\} \cup C'_t$, where C'_t is the last cover in C' — C'_t has at most $k-1$ components</p>

Fig. 9. Recursive algorithm for general k -Component Dynamization (Theorem 3.9).

Combined with the observation that the Greedy-Dual algorithm (Figure 8) is newest-first, Theorems 3.2 and 3.5 yield a k -competitive algorithm for LSM k -Component Dynamization:

COROLLARY 3.8. *There is a deterministic online algorithm for LSM k -Component Dynamization with competitive ratio k .*

3.4 Upper bound for general variant

THEOREM 3.9. *For general k -Component Dynamization, the deterministic online algorithm B_k in Figure 9 is k -competitive.*

PROOF. The proof is by induction on k . For $k = 1$, Algorithm B_1 is 1-competitive (optimal) because there is only one solution for any instance. Consider any $k \geq 2$, and assume inductively that B_{k-1} is $(k-1)$ -competitive. Fix any input (I, w) with $I = (I_1, \dots, I_n)$. Let OPT_k denote the optimal (offline) algorithm, and let $C^* = \text{OPT}_k(I_1, \dots, I_n)$ be an optimal solution for I .

Let $\mathcal{N}_t^* = C_t^* \setminus C_{t-1}^*$ denote OPT 's new components at time t . For $a, b \in [n]$, let $\Delta_a^b(\text{OPT}_k)$ denote the cost incurred by OPT_k during time interval $[a, b]$, that is, $\sum_{i=a}^b \sum_{S \in \mathcal{N}_i^*} w_i(S)$. (Recall $[a, b]$ denotes $\{a, a+1, \dots, b\}$.) Likewise, let $\Delta_a^b(B_k)$ denote the cost incurred by B_k during $[a, b]$. Let $\mathcal{I}_a^b = (I_a, I_{a+1}, \dots, I_b)$ denote the subproblem formed by the insertions during $[a, b]$, with build-costs inherited from w .

Recall that B_k partitions the input sequence into *phases*, each of which (except possibly the last) ends with B_k doing a *full merge* (i.e., at a time t with $|C_t| = 1$). Assume without loss of generality that B_k ends the last phase with a full merge. (Otherwise, append a final empty insertion at time $n+1$ and define $w_{n+1}(U_{n+1}) = 0$. This does not increase the optimal cost, and causes the algorithm to do a full merge at time $n+1$ unless its total cost in the phase is zero.) Consider any phase. Now fix a and b to be the first and last time steps during the phase. To prove the theorem, we show $\Delta_a^b(B_k) \leq k \Delta_a^b(\text{OPT}_k)$. The theorem follows by summing over the phases.

The proof is via a series of lemmas. Recall that U_t denotes $\bigcup_{i=1}^t I_i$.

LEMMA 3.10. *For any integer $j \in [a, b]$, $\text{cost}(B_{k-1}(\mathcal{I}_a^j)) \leq (k-1) \text{cost}(\text{OPT}_{k-1}(\mathcal{I}_a^j))$.*

PROOF. The instance (I, w) obeys Restrictions **(R1)**–**(R3)**. So, by inspection of those restrictions, \mathcal{I}_a^j also obeys them. That is, \mathcal{I}_a^j is a valid instance of general $(k-1)$ -component Dynamization. So, by the inductive assumption, B_{k-1} is $(k-1)$ -competitive for \mathcal{I}_a^j . \square

For $j \in [a, b]$, say that OPT rebuilds by time j if $U_{a-1} \subseteq \bigcup_{i=a}^j \bigcup_{S \in \mathcal{N}_i} S$. That is, every element inserted before time a is in some new component during $[a, j]$. (Equivalently, $\bigcup_{i=a}^j \bigcup_{S \in \mathcal{N}_i} S = U_j$.)

LEMMA 3.11. *Suppose OPT rebuilds by time j . Then $\Delta_a^j(\text{OPT}_k) \geq w_j(U_j)$.*

PROOF.

$$\begin{aligned}
w_j(U_j) &= w_j\left(\bigcup_{i=a}^j \bigcup_{S \in \mathcal{N}_i} S\right) && \text{(OPT rebuilds by time } j\text{)} \\
&\leq \sum_{i=a}^j \sum_{S \in \mathcal{N}_i} w_j(S) && \text{(by sub-additivity (R1))} \\
&\leq \sum_{i=a}^j \sum_{S \in \mathcal{N}_i} w_i(S) && \text{(by temporal monotonicity (R3))} \\
&= \Delta_a^j(\text{OPT}_k). && \text{(by definition)}
\end{aligned}$$

□

LEMMA 3.12. *Suppose OPT does not rebuild by time $j \in [a, b]$. Then $\text{cost}(\text{OPT}_{k-1}(\mathcal{I}_a^j)) \leq \Delta_a^j(\text{OPT}_k)$.*

PROOF. Because OPT does not rebuild by time j , some element x in U_{a-1} is not in any new component during $[a, j]$. Let S be the component in C_j^* containing x . Since S^* is not new during $[a, j]$, it must be that S^* is in C_i^* for every $i \in [a-1, j]$, and $S^* \subseteq U_{a-1}$.

For the subproblem \mathcal{I}_a^j , let C' be the solution defined by $C'_i = \{S \setminus U_{a-1} : S \in C_i^*\} \setminus \{\emptyset\}$ for $i \in [a, j]$. Because each C_i^* has at most k components, one of which is S^* , and $S^* \subseteq U_{a-1}$, it follows that each C_i has at most $k-1$ components. So $\text{cost}(\text{OPT}_{k-1}(\mathcal{I}_a^j)) \leq \text{cost}(C')$.

If a given component $S \setminus U_{a-1}$ is new in C' at time $i \in [a, j]$, then the corresponding component S is new in C^* at time i . Further, by suffix monotonicity (R2), the cost $w_i(S \setminus U_{a-1})$ paid by C' for $S \setminus U_{a-1}$ is at most the cost $w_i(S)$ paid by C^* for S . (Inspecting the definition of (R2), we require that $S \neq U_i$, which holds because OPT has not rebuilt by time j .) So $\text{cost}(C') \leq \Delta_a^j(\text{OPT}_k)$. □

LEMMA 3.13. $\text{cost}(B_{k-1}(\mathcal{I}_a^{b-1})) \leq (k-1)\Delta_a^{b-1}(\text{OPT}_k)$

PROOF. If $a = b$ then $\text{cost}(B_{k-1}(\mathcal{I}_a^{b-1})) = 0$, so assume $a < b$. If OPT rebuilds by time $b-1$, then $\text{cost}(B_{k-1}(\mathcal{I}_a^{b-1})) \leq (k-1)w_{b-1}(U_{b-1})$ (B_k does not end the phase at time $b-1$)
 $\leq (k-1)\Delta_a^{b-1}(\text{OPT}_k)$ (Lemma 3.11 with $j = b-1$).

Otherwise OPT does not rebuild by time $b-1$, so

$\text{cost}(B_{k-1}(\mathcal{I}_a^{b-1})) \leq (k-1)\text{cost}(\text{OPT}_{k-1}(\mathcal{I}_a^{b-1}))$ (Lemma 3.10 with $j = b-1$)
 $\leq (k-1)\Delta_a^{b-1}(\text{OPT}_k)$ (Lemma 3.12 with $j = b-1$).

□

LEMMA 3.14. $w_b(U_b) \leq \Delta_a^b(\text{OPT}_k)$

PROOF. If OPT rebuilds by time b , then

$$w_b(U_b) \leq \Delta_a^b(\text{OPT}_k) \quad \text{(Lemma 3.11 with } j = b\text{)}$$

Otherwise OPT does not rebuild by time b , so

$$\begin{aligned}
w_b(U_b) &< \text{cost}(B_{k-1}(\mathcal{I}_a^b))/(k-1) && \text{(because } B_k \text{ ends the phase at time } b\text{)} \\
&\leq \text{cost}(\text{OPT}_{k-1}(\mathcal{I}_a^b)) && \text{(Lemma 3.10 with } j = b\text{)} \\
&\leq \Delta_a^b(\text{OPT}_k) && \text{(Lemma 3.12 with } j = b\text{)}.
\end{aligned}$$




Fig. 10. Replacing intervals V and V' by V'_1 and V'_2 (proof of Theorem 4.1).

□

The lemmas imply that the algorithm is k -competitive for the phase:

$$\begin{aligned}
 \Delta_a^b(\mathbf{B}_k) &= \text{cost}(\mathbf{B}_{k-1}(\mathcal{I}_a^{b-1})) + w_b(U_b) && \text{(by definition of } \mathbf{B}_k) \\
 &\leq \text{cost}(\mathbf{B}_{k-1}(\mathcal{I}_a^{b-1})) + \Delta_a^b(\text{OPT}_k) && \text{(Lemma 3.14)} \\
 &\leq (k-1)\Delta_a^{b-1}(\text{OPT}_k) + \Delta_a^b(\text{OPT}_k) && \text{(Lemma 3.13)} \\
 &\leq k(\Delta_a^b \text{OPT}_k) && \text{(as } \Delta_a^{b-1}(\text{OPT}_k) \leq \Delta_a^b(\text{OPT}_k))
 \end{aligned}$$

Theorem 3.9 follows by summing over the phases. □

4 Properties of optimal offline solutions

THEOREM 4.1. *Every instance of k -Component or Min-Sum Dynamization has an optimal solution that is newest-first and lightest-first.*

PROOF. Fix an instance $I = (I_1, \dots, I_n)$. Recall that $[t, t']$ denotes $\{t, t+1, \dots, t'\}$. For any component S that is new at some time t of a given solution C , we say that S uses (time) interval $[t, t']$, where $t' = \max\{j \in [t, n] : (\forall i \in [t, j]) S \in C_i\}$ is the time that (this occurrence of) S is destroyed. We refer to $[t, t']$ as the interval of (this occurrence of) S . For the proof we think of any solution C as being constructed in two steps: (i) choose the set T of time intervals that the components of C will use, then (ii) given T , for each interval $[t, t'] \in T$, choose a set S of items for $[t, t']$, then form a component S in C with interval $[t, t']$ (that is, add S to C_i for $i \in [t, t']$). We shall see that the second step (ii) decomposes by item: an optimal solution can be found by greedily choosing the intervals for each item $x \in U_n$ independently. The resulting solution has the desired properties. Here are the details.

Fix an optimal solution C^* for the given instance, breaking ties by choosing C^* to minimize the total query cost $\sum_{[t, t'] \in T^*} t' - t + 1$ where T^* is the set of intervals of components in C^* . Assume without loss of generality that, for each $t \in [1, n]$, if $I_t = \emptyset$, then $C_t^* = C_{t-1}^*$ (interpreting C_0^* as \emptyset). (If not, replace C_t^* by C_{t-1}^* .) For each item $x \in U_n$, let $\alpha^*(x)$ denote the set of intervals in T^* of components that contain x . The build cost of C^* equals $\sum_{x \in U_n} \text{wt}(x) |\alpha^*(x)|$. For each time t and item $x \in I_t$, the intervals $\alpha^*(x)$ of x cover $[t, n]$, meaning that the union of the intervals in $\alpha^*(x)$ is $[t, n]$.

Next construct the desired solution C' from T^* . For each time t and item $x \in I_t$, let $\alpha(x) = \{V_1, \dots, V_\ell\}$ be a sequence of intervals chosen greedily from T^* as follows. Interval V_1 is the latest-ending interval starting at time t . For $i \geq 2$, interval V_i is the latest-ending interval starting at time $t'_{i-1} + 1$ or earlier, where t'_{i-1} is the end-time of V_{i-1} . The final interval has end-time $t'_\ell = n$. By a standard argument, this greedy algorithm chooses from T^* a minimum-size interval cover of $[t, n]$, so $|\alpha(x)| \leq |\alpha^*(x)|$.

Obtain C' as follows: for each interval $[i, j] \in T^*$, add a component in C' with time interval $[i, j]$ containing the items x such that $[i, j] \in \alpha(x)$. This is a valid solution because, for each time t and

$x \in I_t$, $\alpha(x)$ covers $[t, n]$. Its build cost is at most the build cost of C^* , because $\sum_{x \in U_n} \text{wt}(x) |\alpha(x)| \leq \sum_{x \in U_n} \text{wt}(x) |\alpha^*(x)|$. At each time t , its query cost is at most the query cost of C^* , because it uses the same set T^* of intervals. So C' is an optimal solution.

C' is *newest-first*. The following properties hold:

- (1) α uses (assigns at least one item to) each interval $V \in T^*$. Otherwise removing V from T^* (and using the same α) would give a solution with the same build cost but lower query cost, contradicting the definition of C^* .
- (2) For all $t \in [1, n]$, the number of intervals in T^* starting at time t is 1 if $I_t \neq \emptyset$ and 0 otherwise. Among intervals in T^* that start at t , only one — the latest ending — can be used in any $\alpha(x)$. So by Property 1 above, T^* has at most interval starting at t . If $I_t \neq \emptyset$, C^* must have a new component at time t , so there is such an interval. If $I_t = \emptyset$ there is not (by the initial choice of C^* it has no new component at time t).
- (3) For every two consecutive intervals V_i, V_{i+1} in any $\alpha(x)$, V_{i+1} is the interval in T^* that starts just after V_i ends. Fix any such V_i, V_{i+1} . For every other item y with $V_i \in \alpha(y)$, the interval following V_i in $\alpha(y)$ must also (by the greedy choice) be V_{i+1} . That is, every item assigned to V_i is also assigned to V_{i+1} . If V_{i+1} were to overlap V_i , replacing V_i by the interval $V_i \setminus V_{i+1}$ (within T^* and every $\alpha(x)$) would give a valid solution with the same build cost but smaller total query cost, contradicting the choice of C^* . So V_{i+1} starts just after V_i ends. By Property 2 above, V_{i+1} is the only interval starting then.
- (4) For every pair of intervals V and V' in T^* , either $V \cap V' = \emptyset$, or one contains the other. Assume otherwise for contradiction, that is, two intervals *cross*: $V \cap V' \neq \emptyset$ and neither contains the other. Let $[a, a']$ and $[b, b']$ be a rightmost crossing pair in T^* , that is, such that $a < b < a' < b'$ and no crossing pair lies in $[a+1, n]$. By Property 1 above, $[a, a']$ is in some $\alpha(x)$. Also $a' < n$. Let $[a'+1, c]$ be the interval added greedily to $\alpha(x)$ following $[a, a']$. (It starts at time $a'+1$ by Property 3 above.) The start-time of $[b, b']$ is in $[a, a'+1]$ (as $a < b < a'$), so by the greedy choice (for $[a, a']$) $[b, b']$ ends no later than $[a'+1, c]$. Further, by the tie-breaking in the greedy choice, $c > b'$. So $[a'+1, c]$ crosses $[b, b']$, contradicting that no crossing pair lies in $[a+1, n]$.

By inspection of the definition of newest-first, Properties 2 and 4 imply that C' is newest-first.

C' is *lightest-first*. To finish we show that C' is lightest-first. For any time $t \in [1, n]$, consider any intervals $V, V' \in T^*$ where V ends at time t while V' includes t but does not end then. To prove that C' is lightest-first, we show $\text{wt}(V) < \text{wt}(V')$.

The intervals of C' are nested (Property 4 above), so $V \subset V'$ and the items assigned to $V = V_1$ are subsequently assigned (by Property 3 above) to intervals V_2, \dots, V_ℓ within V' as shown in Figure 10, with V_ℓ and V' ending at the same time. Since V' does not end when V does, $\ell \geq 2$. Consider modifying the solution C' as follows: remove intervals V and V' from T^* , and replace them by intervals V'_1 and V'_2 obtained by splitting V' so that V'_2 starts when V started. (See the right side of Figure 10.)

Reassign all of V' 's items to V'_1 and V'_2 . Reassign all of V 's items to V'_2 and unassign those items from each interval V_i . This gives another valid solution. It has lower query cost (as V is gone), so by the choice of C^* (including the tie-breaking) the new solution must have strictly larger build cost. That is, the change in the build cost, $\text{wt}(V)(1 - \ell) + \text{wt}(V')$, must be positive, implying that $\text{wt}(V') > \text{wt}(V)(\ell - 1) \geq \text{wt}(V)$ (using $\ell \geq 2$). Hence $\text{wt}(V') > \text{wt}(V)$. \square

5 Concluding remarks

This paper brings competitive analysis to bear on data-structure dynamization for non-uniform inputs, via two new online covering problems—Min-Sum Dynamization and k -Component Dynamization—for which it gives deterministic online algorithms with competitive ratios $\Theta(\log^* m)$ and k , respectively. The algorithms extend to handle lazy updates and deletions as they occur in industrial LSM systems.

The paper also shows the existence of optimal offline solutions that are newest-first and lightest-first. As mentioned in the introduction, one consequence is that Bentley and Saxe’s transforms give optimal solutions (up to lower-order terms) for uniform inputs. Another is that, for Min-Sum and k -Component Dynamization, optimal solutions can be computed in time $O(n^3)$ and $O(kn^3)$, respectively, because optimal newest-first solutions can be computed in these time bounds via natural dynamic programs.

5.1 Open problems

Here are a few of many interesting problems that remain open. For k -Component Dynamization:

- Is there an online algorithm with competitive ratio $O(\min(k, \log^* m))$?
- Is there an algorithm with competitive ratio $O(k/(k - h + 1))$ versus OPT_h (the optimal solution with maximum query cost $h \leq k$)?
- Is there a randomized algorithm with competitive ratio $o(k)$?
- A *memoryless* randomized algorithm with competitive $O(k)$?

For Min-Sum Dynamization:

- Is there an $O(1)$ -competitive algorithm?
- Some LSM architectures only support newest-first algorithms. Is there a newest-first algorithm with competitive ratio $O(\log^* m)$?
- What are the best ratios for the LSM and general variants?

For both problems:

- For instances I that occur in practice, the ratio $\max_{t,t'} \text{wt}(I_t)/\text{wt}(I_{t'})$ (for t' such that $\text{wt}(I_{t'}) > 0$) is often bounded. Does restricting to such instances allow smaller competitive ratios?
- For the decreasing-weights and LSM variants, is there always an optimal newest-first solution?

5.2 Variations on the model

Tombstones deleted during major compactions. Times when the cover C_t has just one component (containing all inserted items) are called *full merges* or *major compactions*. At these times, LSM systems delete all tombstone items (even non-redundant tombstones). Our definition of LSM k -Component Dynamization does not capture this, but our definition of General k -Component Dynamization does, so the algorithm B_k in Figure 9 is k -competitive in this case.

Monolithic builds. Our model underestimates query costs because it assumes that new components can be built in response to each query, before responding to the query. In reality, builds take time. Can this be modelled cleanly, perhaps via a problem that constrains the build cost at each time t (and $\text{wt}(I_t)$) to be at most 1, with the objective of minimizing the total query cost?

Splitting the key space. To avoid monolithic builds, when the data size reaches some threshold (e.g., when the available RAM can hold 1% of the stored data) some LSM systems “split”: they divide the workload into two parts—the keys above and below some threshold—then restart, handling

each part on separate servers. This requires a mechanism for routing insertions and queries by key to the appropriate server. Can this (including a routing layer supporting multiple splits) be cleanly modeled?

Other LSM systems (LevelDB and its derivatives) instead use many small (disk-block size) components, storing in the (cached) indices each component’s key interval (its minimum and maximum key). A query for a given key accesses only the components whose intervals contain the key. This suggests a natural modification of our model: redefine the query cost at time t to be the maximum number of such components for any key.

Bloom filters. Most practical LSM systems are configurable to use a Bloom filter for each component, so as to avoid (with some probability) accessing component that do not hold the queried key. However, Bloom filters are only cost-effective when they are small enough to be cached. They require about a byte per key, so are effective only for the smallest components (with a total number of keys no more than the bytes available in RAM). Used effectively, they can save a few disk accesses per query (see [26]). They do not speed up range queries (that is, efficient searches for all keys in a given interval, which LSM systems support but hash-based external-memory dictionaries do not).

External-memory. More generally, to what extent can we apply competitive analysis to the standard I/O (external-memory) model? Given an input sequence (rather than being constrained to maintain a cover) the algorithm would be free to use the cache and disk as it pleases, subject only to the constraints of the I/O model, with the objective of minimizing the number of disk I/O’s, divided by the minimum possible number of disk I/O’s for that particular input. This setting may be too general to work with. Is there a clean compromise?

The results below do not address this per se, but they do analyze external-memory algorithms using metrics other than standard worst-case analysis, with a somewhat similar flavor:

- [8] Studies competitive algorithms for allocating cache space to competing processes.
- [10] Analyzes external-memory algorithms while available RAM varies with time, seeking an algorithm such that, no matter how RAM availability varies, the worst-case performance is as good as that of any other algorithm.
- [17] Presents external-memory sorting algorithms that have per-input guarantees — they use fewer I/O’s for inputs that are “close” to sorted.
- [23, 37] Present external-memory dictionaries with a kind of *static-optimality* property: for any sequence of queries, they incur cost bounded in terms of the minimum achievable by any static tree of a certain kind. (This is analogous to the static optimality of splay trees [39, 49].)

5.3 Practical considerations

Heuristics for newest-first solutions. Some LSM systems *require* newest-first solutions. The Min-Sum Dynamization algorithm Adaptive-Binary (Figure 3) can produce solutions that are not newest-first. Here is one naive heuristic to make it newest-first: at time t , do the minimal newest-first merge that includes all the components that the algorithm would otherwise have selected to merge. This might result in only a small cost increase on some workloads.

Major compactions. For various reasons, it can be useful to force major compactions at specified times. An easy way to model this is to treat each interval between forced major compactions as a separate problem instance, starting each instance by inserting all items from the major compaction.

Estimating the build cost $wt_t(S)$. Our algorithms for the decreasing-weights, LSM, and general variants depend on the build costs $wt_t(S)$ of components S that are not yet built. The model assumes these become known at time t , but in practice they can be hard to compute. However, the algorithms

only depend on the build costs of components S that are unions of the current components. For the LSM variant, it may be possible to construct, along with each component S , a small signature that can be used to estimate the build costs of unions of such components (at later times t), using techniques for estimating intersections of large sets (e.g. [24, 46]). It would be desirable to show that dynamization algorithms are robust in this context—that their competitive ratios are approximately preserved if they use approximate build costs.

Exploiting slack in the Greedy-Dual algorithm. For paging, LEAST-RECENTLY-USED (LRU) is preferred in practice to FLUSH-WHEN-FULL (FWF), although their competitive ratios are equal. In practice, it can be useful to tune an algorithm while preserving its theoretical performance guarantee. In this spirit, consider the following variant of the Greedy-Dual algorithm in Figure 8. As the algorithm runs, maintain a “spare credit” ϕ . Initially $\phi = 0$. When the algorithm does a merge in Line 2.1.3, increase ϕ by the total credit of the components newer than S_0 , which the algorithm destroys. Then, at any time, optionally, reduce ϕ by some amount $\delta \leq \phi$, and increase the credit of any component in the cover by ϕ . The proof of Theorem 3.2, essentially unchanged, shows that the modified algorithm is still k -competitive. This kind of additional flexibility may be useful in tuning the algorithm. As an example, consider classifying the spare credit by the rank of the component that contributes it, and, when a new component S' of some rank r is created, transferring all spare credit associated with rank r to credit[S'] (after Line 2.1.4 initializes credit[S'] to 0). This natural BALANCE algorithm balances the work done for each of the k ranks.

Acknowledgments

Thanks to Carl Staelin for bringing the problem to our attention and for informative discussions about Bigtable.

References

- [1] Pankaj K. Agarwal, Lars Arge, Octavian Procopiuc, and Jeffrey Scott Vitter. 2001. A Framework for Index Bulk Loading and Dynamization. In *Automata, Languages and Programming (Lecture Notes in Computer Science)*, Fernando Orejas, Paul G. Spirakis, and Jan van Leeuwen (Eds.). Springer Berlin Heidelberg, 115–127.
- [2] Pankaj K. Agarwal, Sarel Har-Peled, and Kasturi R. Varadarajan. 2004. Approximating Extent Measures of Points. *J. ACM* 51, 4 (2004), 606–635.
- [3] Alok Aggarwal, Ashok K. Chandra, and Marc Snir. 1987. Hierarchical Memory with Block Transfer. In *28th Annual Symposium on Foundations of Computer Science*. IEEE, 204–216. <https://doi.org/10.1109/SFCS.1987.31>
- [4] Sattam Alsubaiee, Yasser Altowim, Hotham Altwaijry, Alexander Behm, Vinayak Borkar, Yingyi Bu, Michael Carey, Inci Cetindil, Madhusudan Cheelang, and Khurram Faraaz. 2014. AsterixDB: A Scalable, Open Source BDMS. *Proceedings of the VLDB Endowment* 7, 14 (2014), 1905–1916.
- [5] Lars Arge. 2002. External Memory Data Structures. In *Handbook of Massive Data Sets*, James Abello, Panos M. Pardalos, and Mauricio G. C. Resende (Eds.). Springer US, Boston, MA, 313–357. https://doi.org/10.1007/978-1-4615-0005-6_9
- [6] Lars Arge and Jan Vahrenhold. 2004. I/O-Efficient Dynamic Planar Point Location. *Computational Geometry* 29, 2 (Oct. 2004), 147–162. <https://doi.org/10.1016/j.comgeo.2003.04.001>
- [7] Amitabha Bagchi, Amitabh Chaudhary, David Eppstein, and Michael T. Goodrich. 2007. Deterministic Sampling and Range Counting in Geometric Data Streams. *ACM Transactions on Algorithms* 3, 2 (May 2007). <https://doi.org/10.1145/1240233.1240239>
- [8] Rakesh D. Barve, Edward F. Grove, and Jeffrey Scott Vitter. 2000. Application-Controlled Paging for a Shared Cache. *SIAM J. Comput.* 29, 4 (Jan. 2000), 1290–1303. <https://doi.org/10.1137/S0097539797324278>
- [9] Michael A. Bender, Rezaul A. Chowdhury, Rathish Das, Rob Johnson, William Kuszmaul, Andrea Lincoln, Quanquan C. Liu, Jayson Lynch, and Helen Xu. 2020. Closing the Gap Between Cache-Oblivious and Cache-Adaptive Analysis. In *Proceedings of the ACM Symposium on Parallelism in Algorithms and Architectures*. ACM, Virtual Event USA, 63–73. <https://doi.org/10.1145/3350755.3400274>
- [10] Michael A. Bender, Roozbeh Ebrahimi, Jeremy T. Fineman, Golnaz Ghasemiefteh, Rob Johnson, and Samuel McCauley. 2014. Cache-Adaptive Algorithms. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 958–971.

- [11] Michael A. Bender, Martin Farach-Colton, Jeremy T. Fineman, Yonatan R. Fogel, Bradley C. Kuszmaul, and Jelani Nelson. 2007. Cache-Oblivious Streaming b-Trees. In *Proceedings of the ACM Symposium on Parallel Algorithms and Architectures*. ACM, New York, NY, USA, 81–92. <https://doi.org/10.1145/1248377.1248393>
- [12] Jon Louis Bentley. 1979. Decomposable Searching Problems. *Inform. Process. Lett.* 8, 5 (June 1979), 244–251. [https://doi.org/10.1016/0020-0190\(79\)90117-0](https://doi.org/10.1016/0020-0190(79)90117-0)
- [13] Jon Louis Bentley and James B Saxe. 1980. Decomposable Searching Problems I. Static-to-Dynamic Transformation. *Journal of Algorithms* 1, 4 (Dec. 1980), 301–358. [https://doi.org/10.1016/0196-6774\(80\)90015-2](https://doi.org/10.1016/0196-6774(80)90015-2)
- [14] Allan Borodin and Ran El-Yaniv. 1998. *Online Computation and Competitive Analysis*. Cambridge University Press, USA.
- [15] Edward Bortnikov, Anastasia Braginsky, Eshcar Hillel, Idit Keidar, and Gali Sheffi. 2018. Accordion: Better Memory Organization for LSM Key-Value Stores. *Proceedings of the VLDB Endowment* 11, 12 (Aug. 2018), 1863–1875. <https://doi.org/10.14778/3229863.3229873>
- [16] Gerth Stolting Brodal and Rolf Fagerberg. 2003. Lower Bounds for External Memory Dictionaries. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 546–554.
- [17] Gerth Stø lting Brodal, Rolf Fagerberg, and Gabriel Moruz. 2005. Cache-Aware and Cache-Oblivious Adaptive Sorting. In *Automata, Languages and Programming (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 576–588. https://doi.org/10.1007/11523468_47
- [18] Hervé Brönnimann, Timothy M. Chan, and Eric Y. Chen. 2004. Towards In-Place Geometric Algorithms and Data Structures. In *Proceedings of the Symposium on Computational Geometry*. ACM, New York, NY, USA, 239–246. <https://doi.org/10.1145/997817.997854>
- [19] Niv Buchbinder, Shahar Chen, and Joseph Naor. 2014. Competitive analysis via regularization. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 436–444.
- [20] Niv Buchbinder and Joseph Naor. 2009. The Design of Competitive Online Algorithms via a Primal–Dual Approach. *Foundations and Trends® in Theoretical Computer Science* 3, 2–3 (2009), 93–263. <https://doi.org/10.1561/0400000024>
- [21] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. 2008. BigTable: A Distributed Storage System for Structured Data. *ACM Transactions on Computing Systems* 26, 2 (June 2008), 4:1–4:26. <https://doi.org/10.1145/1365815.1365816>
- [22] Y. Chiang and R. Tamassia. 1992. Dynamic Algorithms in Computational Geometry. *Proc. IEEE* 80, 9 (Sept. 1992), 1412–1434. <https://doi.org/10.1109/5.163409>
- [23] V. Ciriani, P. Ferragina, F. Luccio, and S. Muthukrishnan. 2002. Static Optimality Theorem for External Memory String Access. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings*. 219–227. <https://doi.org/10.1109/SFCS.2002.1181945>
- [24] Reuven Cohen, Liran Katzir, and Aviv Yehezkel. 2017. A Minimal Variance Estimator for the Cardinality of Big Data Set Intersection. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, Halifax NS Canada, 95–103. <https://doi.org/10.1145/3097983.3097999>
- [25] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, J. J. Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Yasushi Saito, Michal Szymaniak, Christopher Taylor, Ruth Wang, and Dale Woodford. 2013. Spanner: Google’s Globally Distributed Database. *ACM Transactions on Computing Systems* 31, 3 (Aug. 2013), 8:1–8:22. <https://doi.org/10.1145/2491245>
- [26] Niv Dayan, Manos Athanassoulis, and Stratos Idreos. 2017. Monkey: Optimal Navigable Key-Value Store. In *Proceedings of the ACM International Conference on Management of Data*. ACM, New York, NY, USA, 79–94. <https://doi.org/10.1145/3035918.3064054>
- [27] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Voshall, and Werner Vogels. 2007. Dynamo: Amazon’s Highly Available Key-Value Store. In *Proceedings of Twenty-First ACM SIGOPS Symposium on Operating Systems Principles*. ACM, New York, NY, USA, 205–220. <https://doi.org/10.1145/1294261.1294281>
- [28] Andy Dent. 2013. *Getting Started with LevelDB*. Packt Publishing Ltd.
- [29] Siying Dong, Mark Callaghan, Leonidas Galanis, Dhruva Borthakur, Tony Savor, and Michael Stumm. 2017. Optimizing Space Amplification in RocksDB. In *Proceedings of the Biennial Conference on Innovative Data Systems Research*. 3–12.
- [30] D.T. Lee and F.P. Preparata. 1984. Computational Geometry—A Survey. *IEEE Trans. Comput.* C-33, 12 (Dec. 1984), 1072–1101. <https://doi.org/10.1109/TC.1984.1676388>
- [31] Dan Feldman, Melanie Schmidt, and Christian Sohler. 2013. Turning Big Data into Tiny Data: Constant-Size Coresets for k-Means, PCA and Projective Clustering. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1434–1453.
- [32] Lars George. 2011. *HBase: The Definitive Guide: Random Access to Your Planet-Size Data*. O’Reilly Media, Inc.

- [33] Goetz Graefe. 2010. Modern B-Tree Techniques. *Foundations and Trends in Databases* 3, 4 (2010), 203–402. <https://doi.org/10.1561/19000000028>
- [34] Sarel Har-Peled and Soham Mazumdar. 2004. On Coresets for K-Means and k-Median Clustering. In *Proceedings of the ACM Symposium on Theory of Computing*. ACM, New York, NY, USA, 291–300. <https://doi.org/10.1145/1007352.1007400> See also <https://doi.org/10.48550/arXiv.1810.12826>.
- [35] Anna R. Karlin, Claire Kenyon, and Dana Randall. 2003. Dynamic TCP Acknowledgment and Other Stories about $e/(e-1)$. *Algorithmica* 36, 3 (July 2003), 209–224. <https://doi.org/10.1007/s00453-003-1013-x>
- [36] Jeremy Kepner, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Vijay Gadepally, Matthew Hubbell, Peter Michaleas, Julie Mullen, Andrew Prout, Albert Reuther, Antonio Rosa, and Charles Yee. 2014. Achieving 100,000,000 Database Inserts per Second Using Accumulo and D4M. In *2014 IEEE High Performance Extreme Computing Conference (HPEC)*. 1–6. <https://doi.org/10.1109/HPEC.2014.7040945>
- [37] Pang Ko and Srinivas Aluru. 2007. Optimal Self-Adjusting Trees for Dynamic String Data in Secondary Storage. In *String Processing and Information Retrieval*, Nivio Ziviani and Ricardo Baeza-Yates (Eds.). Vol. 4726. Springer Berlin Heidelberg, Berlin, Heidelberg, 184–194. https://doi.org/10.1007/978-3-540-75530-2_17
- [38] Avinash Lakshman and Prashant Malik. 2010. Cassandra: A Decentralized Structured Storage System. *SIGOPS Operating Systems Review* 44, 2 (April 2010), 35–40. <https://doi.org/10.1145/1773912.1773922>
- [39] Caleb Levy and Robert Tarjan. 2019. A New Path from Splay to Dynamic Optimality. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1311–1330.
- [40] Hyeontaek Lim, David G. Andersen, and Michael Kaminsky. 2016. Towards Accurate and Fast Evaluation of Multi-Stage Log-Structured Designs. In *Proceedings of the Usenix Conference on File and Storage Technologies*. USENIX Association, Berkeley, CA, USA, 149–166.
- [41] Chen Luo and Michael J. Carey. 2020. LSM-Based Storage Techniques: a survey. *VLDB J.* 29, 1 (2020), 393–418. <https://doi.org/10.1007/s00778-019-00555-y>
- [42] Claire Mathieu, Rajmohan Rajaraman, Neal E Young, and Arman Yousefi. 2021. Competitive data-structure dynamization. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2269–2287.
- [43] Kurt Mehlhorn. 1981. Lower Bounds on the Efficiency of Transforming Static Data Structures into Dynamic Structures. *Mathematical systems theory* 15, 1 (Dec. 1981), 1–16. <https://doi.org/10.1007/BF01786969>
- [44] Patrick O’Neil, Edward Cheng, Dieter Gawlick, and Elizabeth O’Neil. 1996. The Log-Structured Merge-Tree (LSM-Tree). *Acta Informatica* 33, 4 (June 1996), 351–385. <https://doi.org/10.1007/s002360050048>
- [45] Mark H. Overmars. 1987. *The Design of Dynamic Data Structures*. Number 156 in Lecture Notes in Computer Science. Springer, Berlin.
- [46] Rasmus Pagh, Morten Stöckel, and David P. Woodruff. 2014. Is Min-Wise Hashing Optimal for Summarizing Set Intersection?. In *Proceedings of the ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. ACM Press, Snowbird, Utah, USA, 109–120. <https://doi.org/10.1145/2594538.2594554>
- [47] Mendel Rosenblum and John K. Ousterhout. 1991. The Design and Implementation of a Log-Structured File System. In *Proceedings of the ACM Symposium on Operating Systems Principles*. ACM, New York, NY, USA, 1–15. <https://doi.org/10.1145/121132.121137>
- [48] Dennis G. Severance and Guy M. Lohman. 1976. Differential Files: Their Application to the Maintenance of Large Databases. *ACM Transaction on Database Systems* 1, 3 (Sept. 1976), 256–267. <https://doi.org/10.1145/320473.320484>
- [49] Daniel Dominic Sleator and Robert Endre Tarjan. 1983. Self-Adjusting Binary Trees. In *Proceedings of the ACM Symposium on Theory of Computing*. ACM, New York, NY, USA, 235–245. <https://doi.org/10.1145/800061.808752>
- [50] Carl Staelin. 2013. Personal communication.
- [51] Jan van Leeuwen and Mark H. Overmars. 1981. The Art of Dynamizing. In *Mathematical Foundations of Computer Science (Lecture Notes in Computer Science)*, Jozef Gruska and Michal Chytil (Eds.). Springer Berlin Heidelberg, 121–131.
- [52] Jeffrey Scott Vitter. 2008. *Algorithms and Data Structures for External Memory*. Number 2:4 in Foundations and Trends in Theoretical Computer Science. Now Publishers, Boston.
- [53] Ke Yi. 2012. Dynamic Indexability and the Optimality of B-Trees. *J. ACM* 59, 4 (Aug. 2012), 1–19. <https://doi.org/10.1145/2339123.2339129>
- [54] Neal E Young. 2000. K-medians, facility location, and the Chernoff-Wald bound. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. 86–95.

A Deferred proofs

LEMMA A.1. *For the min-sum dynamization problem, the competitive ratio of the naive adaptation of Bentley’s binary transform, which treats each insertion as a size-1 item and applies the transform, is $\Omega(\log n)$.*

PROOF. Consider an input sequence in which an item of weight n^2 is inserted in step 1, and an item of weight ε for infinitesimally small $\varepsilon > 0$ is inserted in each step t , for $2 \leq t \leq n$. The naive adaptation of Bentley's binary transform ignores the weights and treats each insertion as a size-1 item. Recall that the binary transform maintains at most one component of size 2^i for each integer i . Since the input sequence inserts one item each step, for each t that is a power of two, the binary transform has exactly one component of size t immediately after step t (for instance, see Figure 1). Thus, each step t that is a power of two incurs build cost at least n^2 (owing to the item of weight n^2). This yields a total build cost of $\Omega(n^2 \log n)$.

An alternative solution, such as the one computed by the algorithm of Figure 3, maintains at most two components, one consisting of the weight n^2 item and the other consisting of any remaining items. The build cost for step 1 is n^2 . For step i , $2 \leq i \leq n$, the build cost is $(i-1)\varepsilon$ since $i-1$ items of weight ε are merged into a component. This yields a total build cost of at most $n^2(1 + \varepsilon/2)$. Since there are at most two components, the query cost is at most $2n$. We thus have an $\Omega(\log n)$ bound on the competitive ratio of the naive adaptation of the binary transform. \square

LEMMA A.2. *The naive generalization of Bentley and Saxe's k -binomial transform to k -Component Dynamization has competitive ratio $\Omega(kn^{1/k})$ for any $k \geq 2$.*

PROOF. Recall that the naive algorithm treats each insertion I_t as one size-1 item, then applies the k -binomial transform. Consider inserting a single item of weight 1, then $n-1$ single items of weight 0. The naive algorithm merges its largest component $\Theta(d)$ times where $\binom{d}{k} \approx n$, so $d = \Theta(kn^{1/k})$. Each such merge costs 1. So the naive algorithm incurs total cost $\Omega(kn^{1/k})$.

The optimum keeps the weight-1 item in one component, then does all remaining merges into the other (size-zero) component, for total cost of 1. \square

Received 1 January 1960; revised 1 January 1960; accepted 1 January 1960