

A SELECTABLE SLOPPY HEAP

Adrian Dumitrescu*

August 9, 2017

Abstract

We study the selection problem, namely that of computing the i th order statistic of n given elements. Here we offer a data structure called *selectable sloppy heap* handling a dynamic version in which upon request: (i) a new element is inserted or (ii) an element of a prescribed quantile group is deleted from the data structure. Each operation is executed in (ideal!) constant time—and is thus independent of n (the number of elements stored in the data structure)—provided that the number of quantile groups is fixed. This is the first result of this kind accommodating both insertion and deletion in constant time. As such, our data structure outperforms the soft heap data structure of Chazelle (which only offers constant amortized complexity for a fixed error rate $0 < \varepsilon \leq 1/2$) in applications such as dynamic percentile maintenance. The design demonstrates how slowing down a certain computation can speed up the data structure.

Keywords: online algorithm, approximate selection, i th order statistic, mediocre element, tournaments, quantiles, dynamic set, amortization.

1 Introduction

The following problem has been devised by Fredman about 25 years ago, for inclusion in homework assignments for an algorithms course. This paper both generalizes and strengthens that result.

A “very sloppy heap” (abbreviated *vsh*) is a data structure for performing the following operations on a set S : (i) *insert* and (ii) *delete-small*. The latter operation deletes (and returns) an element x which is among the $\lceil n/2 \rceil$ smallest elements in the set, where n is the current size of the set. Explain how to implement a *vsh* in constant amortized time per operation.

Together with sorting, selection is one of the most widely used procedure in computer algorithms. Given a sequence A of n numbers and an integer (selection) parameter $1 \leq i \leq n$, the selection problem asks to find the i th smallest element in A . Sorting trivially solves the selection problem; however, a higher level of sophistication is required by a linear time algorithm. A now classic approach for selection [6, 18, 25, 36, 39] from the 1970s is to use an element in A as a pivot to partition A into two smaller subsequences and recurse on one of them with a (possibly different) selection parameter i .

The time complexity of this kind of algorithms is sensitive to the pivots used. If a good pivot is used, many elements in A can be discarded, while if a bad pivot is used, the size of the problem may be only reduced by a constant in the worst case, leading to a quadratic worst-case running time. But carefully choosing a good pivot can be time consuming. Choosing the pivots randomly (and

*Department of Computer Science, University of Wisconsin–Milwaukee, USA. Email: dumitres@uwm.edu

thus without much effort) yields a well-known randomized selection algorithm with expected linear running time; see e.g., [11, Ch. 9.2], [31, Ch. 13.5], or [34, Ch. 3.4]; however its worst case running time is quadratic in n .

The first deterministic linear time selection algorithm `SELECT` is due to Blum et al. [6]; it is recursive in nature. By using the median of medians of small disjoint groups of the input array (of constant size at least 5) good pivots that reduce the size of the problem by a constant fraction and thereby lead to $O(n)$ time overall, can be chosen at low cost in each recursive invocation. More recently, suitable variants of `SELECT` with groups of 3 and 4 also running in $O(n)$ time have been also put forward [10, 40]. The selection problem, and computing the median in particular are in close relation with the problem of finding the quantiles of a set, that we describe next.

Quantiles. The k th *quantiles* of an n -element set are the $k - 1$ order statistics that divide the sorted set in k equal-sized groups (to within 1); see, e.g., [11, p. 223]. It is known that the k th quantiles of a set can be computed by a recursive algorithm running in $O(n \log k)$ time. Such an algorithm can be modified, if needed, so that the k groups can be also output, say, each as a linked list, within the same overall time. For $2 \leq i \leq k - 1$, the i th group of elements (bounded by the $(i - 1)$ th and the i th quantile) is referred to as the i th *quantile group*; the first quantile group consists of the elements less or equal to the first quantile, and the k th quantile group consists of the elements greater or equal to the $(k - 1)$ th quantile.

Our main result is the following; for the optimality of the dependence in k , see the first remark in Section 4.

Theorem 1. *For any fixed integer k , a data structure for dynamic sets exists accommodating each of the following operations in constant time (that depends on k): (i) `INSERT` a new element and (ii) `DELETE` (and return) an element of the i th quantile group of the current set, where $1 \leq i \leq k$. The time per operation is $O(\log k)$, which is optimal in a comparison-based model even in an amortized setting.*

Background and related problems. Since the selection problem is of primary importance, the interest in selection algorithms has remained high ever since; see for instance [2, 4, 5, 12, 14, 15, 16, 18, 20, 21, 22, 23, 27, 29, 35, 37, 39]. In particular, determining the comparison complexity for computing various order statistics including the median has lead to many exciting questions, some of which are still unanswered today; in this respect, Yao’s hypothesis on selection [38, Ch. 4] remains an inspiring endeavor [15, 35, 36]. A comprehensive review of early developments in selection is provided by Knuth [32]. Computational results on the exact comparison complexity of finding the i th smallest out of n for small i, n , have been obtained in [21, 24]. We also refer the reader to the dedicated book chapters on selection in [1, 3, 11, 13, 31] and the more recent articles [10, 26, 30].

The selection problem is also of interest in an online or dynamic setting, where elements are inserted or deleted. A balanced binary search tree on n distinct elements can be augmented with a *size* attribute for each node, thereby allowing the retrieval of an element of a given rank in $O(\log n)$ time [11, Ch. 14]. We say that the i th smallest element has *rank* i , where $i = 1, \dots, n$, with ties broken arbitrarily. Further, determining the rank of an element in the data structure can also be done in $O(\log n)$ time. Consequently, a dynamic order statistic operation (inserting a new element or deleting an element of a given rank) can be accomplished within the same time bound.

Priority queues (*heaps* in particular) are ubiquitous data structures; a typical operation set might include `Insert`, `Delete-Min`, `Decrease-Key` (and perhaps also `Find-Min`, `Increase-Key`, or other operations). See for instance [7] and the references therein for some new variants and recent developments.

A variant of a priority queue that allows dynamic maintenance of percentiles is the *soft heap* data structure due to Chazelle [8]; in addition to `Create` and `Meld`, the data structure accommodates `Insert`, `Delete` and `Findmin` (`Delete x` removes item x). Consider a sequence of $m \geq n$ operations after its creation, that includes n `Insert` operations. For any $0 < \varepsilon \leq 1/2$, a soft heap with error rate ε supports each operation in constant amortized time, except for `Insert` which takes $O(\log 1/\varepsilon)$ amortized time. The data structure ensures that at most εn of its items are *corrupted* (their keys have been artificially raised). `Findmin` returns the minimum current key (which might be corrupted). In contrast, our data structure uses $O(\log 1/\varepsilon)$ time per operation, where $\varepsilon \sim 1/k$, and thereby outperforms the soft heap with respect to the worst-case time per update (insert or delete) operation.

Definition. Let k be a fixed positive integer. A “selectable sloppy heap” (abbreviated *ssh*) is a data structure for performing the following operations on a set S with n elements:

- (a) `INSERT x` : a new element x is inserted.
- (b) `DELETE i` : this deletes (and returns) some element x which belongs to the i th quantile group of the current set, where $1 \leq i \leq k$; if $n < k$, the deleted element is not subject to any requirement.

Outline of the paper. To explain the main challenges and introduce the main ideas, we first sketch several preliminary implementations of the data structure, meeting suboptimal (in k) benchmarks: (i) $O(k \log k)$ amortized time per operation for the first variant in Section 2; (ii) $O(k \log k)$ worst-case time per operation for the second variant in Section 2; (iii) $O(\log k)$ amortized time per operation for the third variant in Section 2. We then further refine these methods in Section 3 to obtain an optimal implementation of the data structure running in $O(\log k)$ worst-case time per operation. In particular, for constant k , the second variant in Section 2 and the (main) variant in Section 3 run in $O(1)$ worst-case time per operation. We conclude in Section 4 with applications and technical remarks.

Comments and notations. Duplicate elements are easily handled by the design. Each operation request is associated with a *discrete time* step, i.e., the j th operation occurs during the j th time step, where $j = 1, 2, \dots$. Without affecting the results, the floor and ceiling functions are omitted in the descriptions of the algorithms and their analyses.

Let A be a set; we write $x \leq A$ if $x \leq a$ for every $a \in A$. The size of a bucket \mathbf{b} , i.e., the number of elements in \mathbf{b} , is denoted by $s(\mathbf{b})$. If buckets \mathbf{b}' and \mathbf{b}'' merge into a new bucket \mathbf{b} , written as $\mathbf{b}' \cup \mathbf{b}'' \rightarrow \mathbf{b}$, we clearly have $s(\mathbf{b}') + s(\mathbf{b}'') = s(\mathbf{b})$. Similarly, if a bucket \mathbf{b} splits into two buckets $\mathbf{b}', \mathbf{b}''$, written as $\mathbf{b} \rightarrow \mathbf{b}' \cup \mathbf{b}''$, then we have $s(\mathbf{b}) = s(\mathbf{b}') + s(\mathbf{b}'')$ as well.

2 Some ideas and preliminary solutions

A first variant with $O(k \log k)$ amortized time per operation. Let n denote the current number of elements in the data structure, with the elements being stored in a linear list (or an array). By default (if n is large enough) the algorithm proceeds in *phases*.

If $n < 3k$, and no phase is under way, proceed by brute force: for `INSERT`, add the new element to the list; for `DELETE i` , compute the elements in the i th quantile group and delete (return) one of them arbitrarily. Since k is constant, each operation takes $O(3k)$, i.e., runs in constant time.

If $n \geq 3k$, and no phase is under way, start a new phase: reorganize the data structure by computing the $3k$ th quantiles of the set and the corresponding groups; store each group in a list

(bucket) of size $n/(3k)$. The next $n/(3k)$ operations make up the current phase. Process each of these $n/(3k)$ operations by using exclusively elements stored in these $3k$ buckets. For INSERT, add the new element to an initially empty overflow bucket (an empty overflow bucket is created at each reorganization). For DELETE i , remove any element from the bucket $(i-1)\frac{n}{k} + \frac{n}{3k}$, i.e., from the middle third of the i th quantile group (out of the total k).

The reorganization takes $O(n \log k)$ time and is executed about every $n/(3k)$ steps, where each operation counts as one step. The resulting amortized cost per operation is

$$O\left(\frac{n \log k}{n/k}\right) = O(k \log k),$$

namely $O(1)$ for constant k . The above idea is next refined so as to obtain this as a worst-case time bound per operation.

A second variant with $O(k \log k)$ worst-case time per operation. It suffices to consider the case of large n , namely $n \geq 3k$. The algorithm proceeds in phases: each phase starts with a reorganization, namely computing the $(3k)$ th quantiles and the corresponding quantile groups; see [11, Ch.15]. The time taken is $O(n \log(3k)) = O(n \log k)$. There are $n/(3k)$ INSERT and DELETE operations, i.e., $n/(3k)$ steps following each reorganization until the results of the next reorganization become available: assume that the data structure holds n items when the current reorganization starts (n is redefined at the start of each reorganization) and there are old buckets to use until the reorganization is finalized. That is, after that many steps new buckets (however, with an old content) become available. Use these buckets for the next $n/(3k)$ operations, and so on.

To transform the constant amortized time guarantee into a constant worst-case guarantee for each operation, spread the execution of reorganization over the next $n/(3k)$ operations, i.e., over the entire phase. The resulting time per operation is bounded from above by

$$O\left(\frac{n \log k}{n/(3k)}\right) = O(k \log k).$$

Since DELETE is serviced from the existent buckets and INSERT is serviced using the overflow bucket, each operation involves inserting or deleting one element from a list; thus for constant k , each operation takes $O(1)$ overall time to process.

New buckets become available about every other $n/(3k)$ operations; this is a rather imprecise estimate because the current number of elements, n , changes. Since a reorganization is spread out over multiple steps, the result becomes available with some delay, and moreover, its content is (partially) obsolete. To verify that the data structure operates as intended, one needs to check that the rank of a deleted element belongs to the required interval (quantile group); we omit the calculation details. The key observation is that any one operation can affect the rank of any element by at most 1: to be precise, only INSERT or DELETE of a smaller element can increase the rank of an element by 1 or decrease the rank of an element by 1, respectively.

A third variant with $O(\log k)$ amortized time per operation. We briefly describe an implementation achieving $O(\log k)$ amortized time per operation that is tailored on the idea of a B-tree; this variant is due to Fredman [19]. Use a balanced binary search tree with $\Theta(\log k)$ levels storing $O(k)$ splitting keys at the leafs. Each leaf comprises a bucket of $\Theta(n/k)$ items, with order maintained between buckets, but *not within* buckets. When an insertion causes a bucket to become too large, it gets split into two buckets by performing a median selection operation. Small buckets (due to deletions) get merged. A bucket of size m , once split, won't get split sooner than another $\Omega(m)$ operations.

When the number of elements doubles (or halves), a tree reorganization is triggered that partitions the present items into $6k$ uniform sized buckets, so that the new common bucket size m is $n/(6k)$ items. These event-triggered reorganizations ensure that buckets do not become too small unless they are target of deletions; similarly, buckets do not become too large unless they are target of insertions. Since the reorganization cost is $O(n \log k)$, and $\Omega(n)$ operation requests take place between successive reorganizations, this scheme yields $O(\log k)$ amortized time per operation, namely $O(1)$ for constant k .

3 A variant with optimal $O(\log k)$ worst-case time per operation

A brief examination of the approach in the 3rd variant reveals two bottlenecks in achieving $O(\log k)$ worst-case time per operation: the median computation that comes with splitting large buckets and the tree reorganization that occurs when n doubles (or halves). We briefly indicate below how ideas from the 2nd and 3rd early variants are refined to obtain $O(\log k)$ worst-case time per operation; in particular, $O(1)$ time for constant k . It is shown in the process how execution in *parallel* of several sequential procedures can lead to a *speed up* of the data structure.

A balanced BST for $\Theta(k)$ keys is used, subdividing the data into $O(k)$ buckets. A *size* attribute is associated with each node reflecting the number of elements stored in the subtree rooted at the respective node. Modifications in the data structure at each operation are reflected in appropriate updates of the size attributes at the $O(\log k)$ nodes of the $O(1)$ search paths involved, in $O(\log k)$ time overhead per operation.

If n is the current number of elements, each bucket holds at most $n/(3k)$ elements, and so each of the k quantile groups contains at least one of the buckets entirely. By choosing one such bucket at the bottom of the search path in the BST for executing a DELETE operation from the i th quantile group guarantees its correctness.

The n elements are kept in $\Theta(k)$ buckets of maximum size $O(n/k)$ that form the $\Theta(k)$ leaves of a balanced binary search tree \mathcal{A} . As in the third variant (in Section 2), each leaf holds a bucket, with order maintained between buckets, but *not within* buckets. Buckets that become too large are split in order to enforce a suitable upper limit on the bucket size. Median finding procedures along with other preparatory and follow-up procedures accompanying bucket splits are scheduled as background computation, as in the second variant (in Section 2).

Our data structure merges small buckets in order to keep the number of buckets under control and renounces the periodic tree reorganizations by introducing new elements of design: a round robin process and a priority queue jointly control the maximum bucket size and the number of buckets in the BST. These mechanisms are introduced to prevent buckets becoming too small or too large as an effect of changes in the total number of elements, n , and *not* necessarily as an effect of operations directed to them.

Outline and features. Let $N := 12k$. For illustrating the main ideas, assume now that $n \geq N$. The buckets are linked in a doubly-linked linear list \mathcal{B} , in key order; adding two links between the last and the first bucket yields a circular list \mathcal{C} , referred to as the *round robin* list. We note that \mathcal{B} and \mathcal{C} are two views of the same data structure.

Each operation request translates to locating a suitable bucket for implementing the request. The circular list is traversed in a round robin fashion, so that the current round robin bucket in the list is also examined during the current operation request. The round robin process ensures that (i) the buckets do not exceed their maximum capacity, and (ii) certain “long-term” preparatory bucket-splitting procedures are run in the background over a succession of non-consecutive discrete time steps allocated to the same bucket.

Each bucket split entails a merge-test for the pair of adjacent buckets with the minimum sum of sizes. The process of merging adjacent buckets in \mathcal{B} is controlled by a priority queue in the form a *binary min-heap* \mathcal{H} . If $|\mathcal{B}| = t$, i.e., there are t buckets, \mathcal{H} holds the $t - 1$ sums of sizes $s(\mathbf{b}) + s(\mathbf{b}^+)$, for buckets $\mathbf{b} \in \mathcal{B}$; here \mathbf{b}^+ denotes the bucket that follows \mathbf{b} in \mathcal{B} . A merge is made provided the minimum value at the top of \mathcal{H} is below some threshold; and \mathcal{A} , \mathcal{B} and \mathcal{H} are updated. Merging adjacent buckets ensures that the total number of buckets remains under control, regardless of which buckets are accessed by operation requests.

Elements of the design. We highlight two: (i) use of the priority queue \mathcal{H} to keep the number of buckets under control, and (ii) running the procedures involved at different rates as needed to ensure that certain task deadlines and precedence constraints among them are met. The data structure maintains the following two invariants:

- $\mathcal{I}1$ Each bucket contains between 1 and $n/(3k)$ elements; there is no limit from below imposed on the bucket size.
- $\mathcal{I}2$ The number of buckets is between $3k$ and $N = 12k$, as implied by the maximum bucket size, and a later argument based on the rules of operation (on merging adjacent buckets); see Action 2 and Lemma 1.

Recall that all $\Theta(k)$ buckets are linked in a circular round robin list: in key order, and with the last bucket linked to the first one. A pointer to the current *round robin bucket* (initialized arbitrarily, say, to the first bucket) is maintained. Each *operation request* advances this position in the list by one slot. Observe that every bucket becomes the round robin bucket about every $\Theta(k)$ discrete time steps. Further, each operation request leads via the search path in \mathcal{A} to one of the buckets, referred to as the *operation bucket*. Executing one operation is done by a sequence of actions performed in the operation bucket and the round robin bucket; these are referred to as the *current* buckets. All these actions are therefore associated with the same discrete time step. Every bucket update (this includes creation or deletion of buckets) entails a corresponding update of the binary heap \mathcal{H} in the (at most) two heap elements $s(\mathbf{b}) + s(\mathbf{b}^+)$ that depend on it.

A bucket is declared *large* if its current size exceeds $9/10$ of the maximum allowed, i.e., if the current size exceeds $9n/(30k)$. All other buckets are declared *regular*. Each operation may cause an update in the *status* of the round robin bucket and the operation bucket (among large and regular).

Action 1. Execute the requested operation: either add a new element to the respective bucket, or delete an arbitrary element from it. If the operation bucket becomes empty as a result of the current DELETE operation, it is deleted, and the BST is correspondingly updated in $O(\log k)$ time. Status updates in the current two buckets are made, if needed. If a median finding procedure or a follow up procedure is under way in the current operation bucket or the current round robin bucket (see details below), the next segment consisting of 500 (or less) *elementary operations* of this procedure is executed.

Action 2 (merge-test). Let $\sigma := s(\mathbf{b}) + s(\mathbf{b}^+)$ be the minimum value at the top of the heap \mathcal{H} . If $\sigma \leq 5n/(30k)$, merge the two buckets into one: $\mathbf{b} \cup \mathbf{b}^+ \rightarrow \mathbf{b}$, and update the tree \mathcal{A} , the bucket list \mathcal{B} and the binary heap \mathcal{H} to reflect this change: (i) delete the two buckets that are merged and insert the new one that results into \mathcal{A} and \mathcal{B} ; (ii) extract the minimum element from \mathcal{H} and insert the two new sum of sizes of two adjacent buckets formed by the new bucket into \mathcal{H} (if they exist).

Handling of \mathcal{A} , \mathcal{B} and \mathcal{H} take $O(\log k)$ time, $O(1)$ time, and $O(\log k)$ time, respectively. In particular, this action ensures that the size of any new bucket obtained by merge is at most $5n/(30k)$.

It is worth noting that a bucket for which a partitioning procedure is under way, as described below, cannot be part of a pair of buckets to merge (i.e., passing the merge-test).

Action 3 (finalizing a split). Similar to the merge operations in Action 2, finalizing the split can be completed in $O(\log k)$ time: it essentially involves updating \mathcal{A} , \mathcal{B} and \mathcal{H} . It will be shown subsequently that the size of any *new* bucket resulting from a split is in the range $[4n/(30k), 6n/(30k)]$.

Besides Actions 1 – 3, there are actions associated with procedures running in the background, whose execution is spread out over a succession of non-consecutive discrete time steps allocated to the same bucket. The procedures are in preparation of splitting large buckets.

Splitting a large bucket. For simplicity of exposition, we assume that $n = \Omega(k)$, for a sufficiently large constant factor. If the current bucket \mathbf{b} is large, i.e., $s(\mathbf{b}) \geq 9n/(30k)$, and no procedure is active in the current bucket, let $n_0 := n$ (the number of elements existent when the procedure is initiated); and place $8.8n_0/(30k)$ elements into (a main part) P_1 and the remaining $s(\mathbf{b}) - |P_1|$ elements into (a secondary part) P_2 . Observe that $0.2n_0/(30k) \leq |P_2| \leq 0.4n_0/(30k)$. Any insertions and deletions from the current bucket until the split is finalized are performed using P_2 . A *balanced partition* of the current bucket will be obtained in at most $n_0/10$ time steps.

Consider the following $n_0/10$ operation requests. By Lemma 1 below, the number of buckets is at most $12k$ at any time, and so at least $(n_0/10)/(12k) = n_0/(120k)$ time steps are allocated to \mathbf{b} in the round robin process by the $n_0/10$ time mark. Let t_1 , where $n_0/(120k) \leq t_1 < n_0/10$, mark the first occurrence when a total of $n_0/(120k)$ discrete time steps have been allocated to \mathbf{b} (as operation steps or round robin steps). As shown next, this number of steps suffices for finalizing the balanced partition and the split of \mathbf{b} . The process calls two procedures, labeled (A) and (B):

(A) Start a *median finding* procedure, i.e., for finding the two quantile groups Q_1, Q_2 of P_1 : an element $m \in Q_1 \cup Q_2 = P_1$ is found so that $Q_1 \leq m \leq Q_2$ and $||Q_1| - |Q_2|| \leq 1$. At the point when this procedure is launched, the computational steps are scheduled so that every time the bucket gets accessed (either as the operation bucket or as the round robin bucket), 500 computational steps for selecting the median element of P_1 take place. Assume for concreteness that median finding on an input set S takes at most $10|S|$ elementary operations; when applied to P_1 , we have $|P_1| \leq 8.8n_0/(30k)$, and thus $88n_0/(30k)$ elementary operations suffice. At the rate of 500 per discrete time step, it follows that $88n_0/(15000k) \leq n_0/(160k)$ time steps suffice for finding the median of P_1 .

(B) After the median has been determined, a *follow up* procedure is initiated in the same bucket. It aims at reducing and finally eliminating the leftover of P_2 , so that in another $n_0/(480k)$ discrete steps, a balanced partition of the current bucket is obtained. The follow up procedure starts with the two quantile groups of the same size, $4.4n_0/(30k)$, as created by the median finding procedure, by comparing each element of P_2 against m and properly placing it in one of the two groups. This procedure runs at the rate of 10 items per discrete time step accessing the current bucket (either as the operation bucket or as the round robin bucket), until finally the partitioning process is completed with all elements in the current bucket properly placed against the pivot m . Note that $|P_2| \leq 0.4n_0/(30k) + n_0/(480) \leq 0.5n_0/(30k)$ at any time; at the rate of 10 items per discrete time step, it follows that $0.5n_0/(300k) \leq n_0/(480k)$ time steps suffice for completing the follow up procedure.

The two procedures terminate within a total of at most $n_0/(160k) + n_0/(480k) = n_0/(120k)$ discrete time steps, as required. The parameters are chosen so that the split takes place before the large bucket \mathbf{b} becomes illegal.

Recall that $0.9n_0 \leq n \leq 1.1n_0$; thus, in terms of the current number of elements, n , the split of a large¹ bucket \mathbf{b} produces two smaller buckets \mathbf{b}' , \mathbf{b}'' , where

$$\frac{4n}{30k} \leq \frac{4.4n_0}{30k} \leq s(\mathbf{b}'), s(\mathbf{b}'') \leq \frac{4.4n_0}{30k} + \frac{0.4n_0}{30k} + \frac{n_0}{480k} \leq \frac{5n_0}{30k} \leq \frac{6n}{30k}. \quad (1)$$

Overlapping partitioning phases in the same bucket are precluded by the fact that no fewer than $3n/(30k)$ data structure operations accessing a given fresh bucket can trigger the next launching of a partitioning phase for that bucket.

Remark. Let \mathbf{b} be a new bucket produced in the current operation. If \mathbf{b} is generated by a split operation, then $\frac{4n}{30k} \leq s(\mathbf{b}) \leq \frac{6n}{30k}$ by (1). If \mathbf{b} is generated by a merge operation, then $s(\mathbf{b}) \leq \frac{5n}{30k}$ by the merge-test.

Analysis of merging buckets and maintaining the two invariants. To prove that the two invariants $\mathcal{I}1$ and $\mathcal{I}2$ are maintained, we need the following key fact.

Lemma 1. *The number of buckets is at most $N = 12k$ at any time.*

Proof. Let t denote the number of buckets after the current operation is executed, and $j = 1, 2, \dots$ denote the discrete time steps. Let B_1, \dots, B_t be the buckets after the current step in key order. Write $a_i = s(B_i)$, for $i = 1, \dots, t$. We proceed by induction on j and show that, if the number of buckets in \mathcal{A} (and \mathcal{B}) is at most N after each of the preceding N time steps, it remains at most N after the current time step. Observe that the number of buckets can only increase by one after a bucket split, and there can be at most two splits associated with a discrete time step. The induction basis is $j \leq N$, and then indeed, we have $t \leq j \leq N$, as required.

For the induction step, assume that the number of buckets is at most N after the previous operation. If no bucket splits occur during the execution of the current operation, the number of buckets remains unchanged, and is thus still at most N after the current operation. If bucket splits occur, it suffices to show that the number of buckets is at most N after each split. Consider a split operation, and let $\sigma := s(\mathbf{b}) + s(\mathbf{b}^+)$ be the minimum value at the top of the heap \mathcal{H} after the split. There are two cases:

Case 1. $\sigma \leq 5n/(30k)$, and thus the merge is executed. Consequently, the number of buckets after the split is still at most N , as required.

Case 2. $\sigma > 5n/(30k)$, and thus no merge is executed. Since \mathcal{H} is a min-heap, we have

$$a_i + a_{i+1} > \frac{5n}{30k}, \text{ for } i = 1, \dots, t-1. \quad (2)$$

Adding these $t-1$ inequalities yields

$$2n = 2 \sum_{i=1}^t a_i > (t-1) \frac{5n}{30k},$$

or $t \leq 12k = N$, as claimed, and concluding the induction step. \square

Summary of the analysis. As seen from the preceding paragraphs and the fact that the number of buckets is $t = O(k)$, the total time per operation is $O(\log t) + O(\log k) + O(1) = O(\log k)$, as required.

¹This qualification refers to the time when the partitioning procedure was initiated.

4 Applications and technical remarks

1. An argument due to Fredman [19] shows that executing a sequence of n operations (from among INSERT and DELETE i , where $1 \leq i \leq k$) requires $\Omega(n \log k)$ time in the worst case, regardless of the implementation of the data structure. The argument relies on the information theory lower bound [32, Ch. 5.3.1]. (For completeness, it is included in the Appendix.)

2. In addition to the two operations provided INSERT and DELETE i , the following operation can be also accommodated at no increase in cost: READ i : this returns some element x which belongs to the i th quantile group of the current set, where $1 \leq i \leq k$; the element remains part of the data structure.

3. The new data structure finds applications in settings with large dynamic sets where insertions and deletions need to be handled fast, and there is no need to be very precise with the ranks of the elements handled. One such application is *approximate sorting*. An array of n elements is said to be L -sorted if for every $i \in [n]$, we have $|\text{rank}(a_i) - i| \leq L$; see also [17, 33]. As shown by the lower bound argument, the selectable sloppy heap can be used to L -sort n elements in $O(n \log(n/L))$ time.

4. As mentioned in the introduction, the selection problem, and computing the median in particular, are in close relation with the problem of finding the quantiles of a set. A typical use of the median relies on its property of being both larger or equal and smaller or equal than a constant fraction of the n elements. Following Yao [38], an element (think of a player, in the context of tournaments) is said to be (i, j) -mediocre if it is neither among the top i nor among the bottom j of a totally ordered set S . As such, an element from the middle third quantile group of an n -element set is $(\lfloor n/3 \rfloor, \lfloor n/3 \rfloor)$ -mediocre; or simply *mediocre*, for short. Repeatedly returning a mediocre element in a dynamic setting (with an initially empty set) can thus be accomplished via the new data structure, by setting $k = 3$ and then repeatedly executing DELETE 2 or READ 2, as needed, at a minimal (constant) query cost. Similarly, putting $k = 100$ sets up the data structure for dynamic percentile maintenance in $O(1)$ time per operation; to the best of our knowledge, this outperforms previously available data structures with respect to this application; e.g., the soft heap data structure of Chazelle [8] only offers constant amortized time per operation.

5. Here we continue our earlier discussion in Section 1 on using BSTs in the context of dynamic selection taking into account the new data structure. Traditionally, search trees do not allow duplicate keys; it is however an easy matter to augment each node of such a tree with a *multiplicity* attribute associated with each distinct key. In particular, the multiplicity attribute is taken into account when computing the size attribute of a node. Now having a balanced search tree augmented as described, insertion, deletion and search all take $O(\log n)$ time.

Our construction method can be used for any parameter k , with $2 \leq k \leq n$. The data structure obtained in this way can be viewed as an *approximate search tree*. In particular, one can search for a given key and the approximate rank of a search key can be determined. To be precise, the following hold: (1) For any $2 \leq k \leq n$, an approximate search tree on n items (duplicate keys allowed) can be constructed in $O(n \log k)$ time. (2) Search for a given key takes $O(n/k + \log k)$ time per operation; its approximate rank, i.e., the quantile group to which it belongs (out of k), can be reported in $O(\log k)$ time (regardless of the presence of the key!). (3) Insertion can be accommodated in $O(\log k)$ time per operation. (4) Deletion of an unspecified key from a given quantile group (out of k) takes $O(\log k)$ time; while deletion of a given key takes $O(n/k + \log k)$ time.

It is worth noting that the approximate search tree we just described appears competitive when k is small and the search function is infrequently used. Then insertion and deletion of an unspecified key from a given quantile group (out of k) takes $O(\log k)$ time (the deleted element is revealed

after the operation); e.g., if $k = O(\log n)$, these operations take $O(\log \log n)$ time. On the other hand, the approximate search tree we just described is no real competitor for the *exact* solution previously discussed; indeed, no choice of k in our data structure would allow an improvement in the performance of all the basic three operations.

6. An alternative solution (to that outlined in Section 3) achieving $O(\log k)$ worst-case time per operation was conceived by Fredman [19] (after the author of the current paper has communicated him the main result). (i) His solution avoids the need to merge small buckets (in order to keep the number of buckets under control) by maintaining two running copies of the data structure and performing periodic tree reorganizations that create uniform-sized buckets. Buckets that become large are split using a mechanism similar to that devised here—in Section 3. While a complete description of this solution is not publicly available, the fact is that the decision of avoid merging small buckets comes at a high price. The main reasons are the need to maintain multiple running copies of the data structure (with two copies under permanent construction, etc), handling subtle consistency issues generally hard to satisfy, and the likely reduced speed caused by the above items. (ii) Another cumbersome approach claimed to be a simplification can be found in his subsequent arXiv posting. As such, neither alternative solution ((i) or (ii) above) matches in elegance and simplicity the one given here.

7. As mentioned earlier, the $O(\log 1/\varepsilon)$ amortized complexity of the soft heap is optimal, however this does not hold for its worst-case complexity, in regard to its **Insert**, **Delete** and **Findmin** operations. In contrast, the $O(\log 1/\varepsilon)$ worst-case complexity per operation of the selectable sloppy heap is optimal (where $\varepsilon \sim 1/k$).

8. Interestingly enough, most applications envisioned by Chazelle [8] for the soft heap (i.e., four items from his list of five) can be dealt with using the selectable sloppy heap; these include dynamic maintenance of percentiles, linear-time selection, and two versions of approximate sorting. One can observe that the complexity analysis for the soft heap, of the error rate in particular, is more complicated than that for the selectable sloppy heap; moreover, special effort is required to ensure that the space needed by the soft heap is linear in the number of items present. In at least one application, approximate sorting, the analysis of doing that with a soft heap is more involved than that for using our structure for this task. Needless to say, the soft heap only allows deletion of small items. As such, the new data structure presented here compares favorably to the soft heap with respect to insertion and deletion. Moreover, the constant amortized guarantee per operation is replaced by the stronger constant worst-case guarantee per operation in our case.

It is worth recalling that the soft heap was designed with a specific application in mind, minimum spanning trees. Given a connected graph $G = (V, E)$, where $|V| = n$, $|E| = m$, with weighted edges, Chazelle [9] showed that a minimum spanning tree (MST) can be computed by a deterministic algorithm in $O(m\alpha(m, n))$ time, where α is the (slowly growing) inverse of the Ackermann function. (A randomized linear-time algorithm was given by Karger et al. [28].) The question of whether the selectable sloppy heap can be used for MST computation is left open.

References

- [1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman, *Data Structures and Algorithms*, Addison–Wesley, Reading, Massachusetts, 1983.
- [2] M. Ajtai, J. Komlós, W. L. Steiger, and E. Szemerédi, Optimal parallel selection has complexity $O(\log \log n)$, *Journal of Computer and System Sciences* **38**(1) (1989), 125–133.

- [3] S. Baase, *Computer Algorithms: Introduction to Design and Analysis*, 2nd edition, Addison-Wesley, Reading, Massachusetts, 1988.
- [4] S. Battiato, D. Cantone, D. Catalano, G. Cincotti, and M. Hofri, An efficient algorithm for the approximate median selection problem, *Proceedings of the 4th Italian Conference on Algorithms and Complexity (CIAC 2000)*, LNCS vol. 1767, Springer, 2000, pp. 226–238.
- [5] S. W. Bent and J. W. John, Finding the median requires $2n$ comparisons, *Proceedings of the 17th Annual ACM Symposium on Theory of Computing (STOC 1985)*, ACM, 1985, pp. 213–216.
- [6] M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, and R. E. Tarjan, Time bounds for selection, *Journal of Computer and System Sciences* **7(4)** (1973), 448–461.
- [7] T. M. Chan, Quake heaps: a simple alternative to Fibonacci heaps, in *Space-Efficient Data Structures, Streams, and Algorithms – Papers in Honor of J. Ian Munro on the Occasion of His 66th Birthday* (A. Brodnik, A. López-Ortiz, V. Raman, and A. Viola, editors), LNCS vol. 8066, Springer, 2013, pp. 27–32.
- [8] B. Chazelle, The soft heap: an approximate priority queue with optimal error rate, *Journal of ACM* **47(6)** (2000), 1012–1027.
- [9] B. Chazelle, A minimum spanning tree algorithm with inverse-Ackermann type complexity, *Journal of ACM* **47(6)** (2000), 1028–1047.
- [10] K. Chen and A. Dumitrescu, Select with groups of 3 or 4, *Proceedings of the 29th International Symposium on Algorithms and Data Structures*, (WADS 2015), Victoria, Canada, August 2015, Vol. 9214 of LNCS, Springer, pp. 189–199. Also available at [arXiv.org/abs/1409.3600](https://arxiv.org/abs/1409.3600).
- [11] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd edition, MIT Press, Cambridge, 2009.
- [12] W. Cunto and J. I. Munro, Average case selection, *Journal of ACM* **36(2)** (1989), 270–279.
- [13] S. Dasgupta, C. Papadimitriou, and U. Vazirani, *Algorithms*, Mc Graw Hill, New York, 2008.
- [14] D. Dor, J. Håstad, S. Ulfberg, and U. Zwick, On lower bounds for selecting the median, *SIAM Journal on Discrete Mathematics* **14(3)** (2001), 299–311.
- [15] D. Dor and U. Zwick, Finding the α th largest element, *Combinatorica* **16(1)** (1996), 41–58.
- [16] D. Dor and U. Zwick, Selecting the median, *SIAM Journal on Computing* **28(5)** (1999), 1722–1758.
- [17] V. Estivill-Castro and D. Wood, Sorting, measures of disorder, and worst-case performance, in *New Results and New Trends in Computer Science* (H. Maurer, editor), LNCS vol. 555, Springer, 1991, pp. 124–131.
- [18] R. W. Floyd and R. L. Rivest, Expected time bounds for selection, *Communications of ACM* **18(3)** (1975), 165–172.
- [19] M. Fredman, personal communication, 2016.
- [20] F. Fussenegger and H. N. Gabow, A counting approach to lower bounds for selection problems, *Journal of ACM* **26(2)** (1979), 227–238.
- [21] W. Gasarch, W. Kelly, and W. Pugh, Finding the i th largest of n for small i, n , *SIGACT News* **27(2)** (1996), 88–96.

- [22] A. Hadian and M. Sobel, Selecting the t -th largest using binary errorless comparisons, *Combinatorial Theory and Its Applications* **4** (1969), 585–599.
- [23] C. A. R. Hoare, Algorithm 63 (PARTITION) and algorithm 65 (FIND), *Communications of the ACM* **4(7)** (1961), 321–322.
- [24] M. Hofri, Optimal selection and sorting via dynamic programming, *ACM Journal of Experimental Algorithmics* **18(2)** (2013), Article 2.3, 14 pages.
- [25] L. Hyafil, Bounds for selection, *SIAM Journal on Computing* **5(1)** (1976), 109–114.
- [26] V. Jayapaul, J. I. Munro, V. Raman, and S. R. Satti, Sorting and selection with equality comparisons, *Proceedings of the 14th International Symposium on Algorithms and Data Structures*, LNCS vol. 9214, Springer, 2015, pp. 434–445.
- [27] J. W. John, A new lower bound for the set-partitioning problem, *SIAM Journal on Computing* **17(4)** (1988), pp. 640–647.
- [28] D. R. Karger, P. N. Klein, and R. E. Tarjan, A Randomized linear-time algorithm to find minimum spanning trees, *Journal of ACM* **42(2)** (1995), 321–328.
- [29] D. G. Kirkpatrick, A unified lower bound for selection and set partitioning problems, *Journal of ACM* **28(1)** (1981), 150–165.
- [30] D. G. Kirkpatrick, Closing a long-standing complexity gap for selection: $V_3(42) = 50$, in *Space-Efficient Data Structures, Streams, and Algorithms – Papers in Honor of J. Ian Munro on the Occasion of His 66th Birthday* (A. Brodnik, A. López-Ortiz, V. Raman, and A. Viola, editors), LNCS vol. 8066, Springer, 2013, pp. 61–76.
- [31] J. Kleinberg and É. Tardos, *Algorithm Design*, Pearson & Addison–Wesley, Boston, Massachusetts, 2006.
- [32] D. E. Knuth, *The Art of Computer Programming, Vol. 3: Sorting and Searching*, 2nd edition, Addison–Wesley, Reading, Massachusetts, 1998.
- [33] H. Manilla, Measures of presortedness and optimal sorting algorithms, *IEEE Trans. Computers* **34(4)** (1985), 318–325.
- [34] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*, Cambridge University Press, 2005.
- [35] M. Paterson, Progress in selection, *Proceedings of the 5th Scandinavian Workshop on Algorithm Theory (SWAT 1996)*, LNCS vol. 1097, Springer, 1996, pp. 368–379.
- [36] A. Schönhage, M. Paterson, and N. Pippenger, Finding the median, *Journal of Computer and System Sciences* **13(2)** (1976), 184–199.
- [37] A. Yao and F. Yao, On the average-case complexity of selecting the k th best, *SIAM Journal on Computing* **11(3)** (1982), 428–447.
- [38] F. Yao, On lower bounds for selection problems, Technical report MAC TR-121, Massachusetts Institute of Technology, Cambridge, 1974.
- [39] C. K. Yap, New upper bounds for selection, *Communications of the ACM* **19(9)** (1976), 501–508.
- [40] U. Zwick, personal communication, September 2014.

Appendix

In this section it is shown that executing a sequence of n operations (from among INSERT and DELETE i , where $1 \leq i \leq k$) may require $\Omega(n \log k)$ time in the worst case, regardless of the implementation of the data structure.

Lower bound. Suppose the delete operation always removes one of the L smallest items. Then to sort n items: (a) insert them, (b) do repeated deletion, and (c) sort the sequence of deleted items in the order created by the successive deletions by executing only $n \lceil \log L \rceil$ comparisons.

To accomplish (a) and (b) using the data structure, call INSERT n times followed by DELETE 1 n times (as we will set $L = \lceil n/k \rceil$). To accomplish (c), start from the right end (the last of the deleted items), and sort from right-to-left by successive insertions into a sorted list being formed, inserting each successive item from the left end. A given item to be inserted would end up in a position no further than L from the front of the sorted list under formation, since when it was deleted, it was one of the L smallest items. Therefore its correct position can be obtained by a $\lceil \log L \rceil$ -comparison binary search (as in binary insertion sort [32, Ch. 5.2.1]).

Since by the information theory lower bound [32, Ch. 5.3.1], (a)-(c) accomplish a task requiring $\lceil \log n! \rceil = n \log n - O(n)$ comparisons in the worst case, and (c) requires only $n \lceil \log L \rceil = n \log L + O(n)$ comparisons, it must be that (a) and (b) jointly require $n \log(n/L) - O(n)$ comparisons in the worst case. Setting $L = \lceil n/k \rceil$ yields that (a) and (b) jointly require $n \log k - O(n) = \Omega(n \log k)$ comparisons.

Finally, substitute n for $2n$. It follows that there exists a sequence of n operations that require $\Omega(n \log k)$ time, thus $\Omega(\log k)$ amortized time per operation.