

Improved Time and Space Bounds for Dynamic Range Mode

Hicham El-Zein

Cheriton School of Computer Science, University of Waterloo, Canada
helzein@uwaterloo.ca

Meng He

Faculty of Computer Science, Dalhousie University, Canada
mhe@cs.dal.ca

J. Ian Munro

Cheriton School of Computer Science, University of Waterloo, Canada
imunro@uwaterloo.ca

Bryce Sandlund

Cheriton School of Computer Science, University of Waterloo, Canada
bcsandlund@uwaterloo.ca

Abstract

Given an array A of n elements, we wish to support queries for the most frequent and least frequent element in a subrange $[l, r]$ of A . We also wish to support updates that change a particular element at index i or insert/ delete an element at index i . For the range mode problem, our data structure supports all operations in $O(n^{2/3})$ deterministic time using only $O(n)$ space. This improves two results by Chan et al. [3]: a linear space data structure supporting update and query operations in $\tilde{O}(n^{3/4})$ time and an $O(n^{4/3})$ space data structure supporting update and query operations in $\tilde{O}(n^{2/3})$ time. For the range least frequent problem, we address two variations. In the first, we are allowed to answer with an element of A that may not appear in the query range, and in the second, the returned element must be present in the query range. For the first variation, we develop a data structure that supports queries in $\tilde{O}(n^{2/3})$ time, updates in $O(n^{2/3})$ time, and occupies $O(n)$ space. For the second variation, we develop a Monte Carlo data structure that supports queries in $O(n^{2/3})$ time, updates in $\tilde{O}(n^{2/3})$ time, and occupies $\tilde{O}(n)$ space, but requires that updates are made independently of the results of previous queries. The Monte Carlo data structure is also capable of answering k -frequency queries; that is, the problem of finding an element of given frequency in the specified query range. Previously, no dynamic data structures were known for least frequent element or k -frequency queries.

2012 ACM Subject Classification Theory of computation → Data structures design and analysis

Keywords and phrases dynamic data structures, range query, range mode, range least frequent, range k -frequency

Digital Object Identifier 10.4230/LIPIcs.ESA.2018.25

Related Version A full version of the paper is available at <https://arxiv.org/abs/1807.03827>.

1 Introduction

The mode of a sample is a fundamental data statistic along with median and mean. Given an ordered sequence, the range mode of interval $[l, r]$ is the mode of the subsequence from index l to r . Building a data structure to efficiently compute range modes allows data analysis to



© Hicham El-Zein, Meng He, J. Ian Munro, and Bryce Sandlund;
licensed under Creative Commons License CC-BY

26th Annual European Symposium on Algorithms (ESA 2018).

Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 25; pp. 25:1–25:13

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

■ **Table 1** Known static and dynamic range query upper bounds.

Query Type	Query Time	Update Time	Space	Citation
Sum	$O(1)$	-	$O(n)$	trivial
	$O(\lg n)$	$O(\lg n)$	$O(n)$	[15]
Min/ Max	$O(1)$	-	$O(n)$	[9]
	$O(\lg n / \lg \lg n)$	$O(\lg^{1/2+\varepsilon} n)$	$O(n)$	[5]
Median	$O(\lg n / \lg \lg n)$	-	$O(n)$	[2]
	$O((\lg n / \lg \lg n)^2)$	$O((\lg n / \lg \lg n)^2)$	$O(n)$	[13]
Majority	$O(1)$	-	$O(n)$	[7]
	$O(\lg n / \lg \lg n)$	$O(\lg n)$	$O(n)$	[7, 10]
Mode	$O(n^\varepsilon \lg \lg n)$	-	$O(n^{2-2\varepsilon})$	[14]
	$O(1)$	-	$O(n^2 \lg \lg n / \lg n)$	[14]
	$O(\sqrt{n / \lg n})$	-	$O(n)$	[3]
	$O(n^{3/4} \lg n / \lg \lg n)$	$O(n^{3/4} \lg \lg n)$	$O(n)$	[3]
	$O(n^{2/3} \lg n / \lg \lg n)$	$O(n^{2/3} \lg n / \lg \lg n)$	$O(n^{4/3})$	[3]
	$O(n^{2/3})$	$O(n^{2/3})$	$O(n)$	new
Least Frequent	$O(\sqrt{n})$	-	$O(n)$	[4]
	$O(n^{2/3} \lg n \lg \lg n)$	$O(n^{2/3})$	$O(n)$	new
k -Frequency	$O(n^{2/3})$	$O(n^{2/3} \lg n)$	$O(n)$	new

be conducted over any window of one-dimensional data. Techniques to answer such queries are relevant to the design of database systems.

The range least frequent problem can be seen as a low-frequency variant of range mode. Instead of searching for the most frequent element in an interval of the sequence, we query for an element that occurs the fewest number of times. We may either restrict our attention to elements that occur at least once in the query range or allow the answer to be an element that occurs zero times in the interval but is present elsewhere in the sequence.

A third range frequency query we consider in this paper is the problem of identifying an element with given frequency k in the specified query interval. This problem has been called the range k -frequency problem.

Both range mode query and range least frequent query have theoretical connections to matrix multiplication. In particular, the ability to answer n range mode queries on an array of size $O(n)$ in faster than $O(n^{\omega/2})$ time, where ω is the constant in the exponent of the running time of matrix multiplication, would imply a faster algorithm for boolean matrix multiplication [3, 4]. Upon closer examination, this lower bound also applies to the range k -frequency problem.

The range mode, range least frequent, and range k -frequency problems are part of a set of questions one can ask on a subrange of a sequence. Other queries in this area include:

- Sum: Return the sum of elements in the query range.
- Min/ Max: Return the minimum/ maximum in the query range.
- Median: Determine the median element in the query range.
- Majority: Return the element that occurs more than $1/2$ the time, if such an element exists.

Note that the mean of a range can be reduced to the range sum problem. Table 1 gives an overview of known upper bounds on related static and dynamic range query data structures.

1.1 Our Results

We improve the results of Chan et al. [3] by giving a dynamic range mode data structure that takes $O(n)$ space and supports updates and queries in $O(n^{2/3})$ time. This improves the query/ update time of their linear space data structure by a polynomial factor and additionally improves the query/ update time of their $O(n^{4/3})$ space data structure by a $O(\log n / \log \log n)$ factor. We also include in our update procedures the ability to insert or delete elements in the middle of the array, an operation not addressed in previous dynamic range mode data structures.

Our improvements are based on the observation that knowing how many of a type of element occur in an interval can be as valuable as knowing the elements themselves. Specifically, instead of storing the frequency counts of elements per span, we store the number of elements with a particular frequency count per span. This information can be dynamically maintained, and uses $O(\log n)$ bits per frequency per span, rather than $O(\log n)$ bits per unique element per span.

Our technique is general enough to also apply to the range least frequent problem in a dynamic setting. To our knowledge, this is the first data structure to do so. In the version of the problem where we allow an answer to not occur in the specified query interval, our data structure supports queries in $O(n^{2/3} \log n \log \log n)$ time, updates in $O(n^{2/3})$ time, and occupies $O(n)$ space. In the version where the least frequent element must be present in the query interval, we develop a Monte Carlo data structure that supports queries in $O(n^{2/3})$ time, updates in $O(n^{2/3} \log n)$ time, and occupies $O(n \log^2 n)$ space. This data structure is correct with high probability for any polynomial sequence of updates and queries, with the restriction that updates are made independently of the results of previous queries. Notably, if the set of queries and updates is fixed in advance or given to the algorithm all at once, this property holds.

Furthermore, our Monte Carlo data structure is powerful enough to apply to the dynamic range k -frequency problem, also supporting queries in $O(n^{2/3})$ time, updates in $O(n^{2/3} \log n)$ time, and occupying $O(n \log^2 n)$ space. This data structure can be augmented to count the number of elements below, above, or at a given frequency, supporting both queries and updates in $O(n^{2/3})$ time and using $O(n)$ space, without the need for an independence assumption.

We organize our results as follows. In Section 2, we review previous work on static and dynamic range mode and least frequent element queries. In Section 3, we briefly give some notation that will be used for the rest of the paper. In Section 4, we give the basic setup of the $O(n)$ space data structure. Section 5 describes how to answer a range mode query in $O(n^{2/3})$ time. Section 6 explains how to support updates to our base data structures in $O(n^{2/3})$ time. In Section 7, we discuss how to answer range least frequent queries in $\tilde{O}(n^{2/3})$ time. Section 8 describes how to find an element of a given frequency in a specified range in $O(n^{2/3})$ time. Here we also mention additional frequency operations our data structure can support. Due to space limitations, we will not describe how the data structure can be made to support insertion and deletion of elements in the conference version of this paper.

2 Previous Work

2.1 Static Range Mode Query

The static range mode query problem was first studied by Krizanc et al. [14]. Their focus is primarily on subquadratic solutions with fast queries, achieving $O(n^{2-2\varepsilon})$ space and $O(n^\varepsilon \log n)$ query time, with $0 < \varepsilon \leq 1/2$, and $O(n^2 \log \log n / \log n)$ space and $O(1)$ query

time. If we set $\varepsilon = 1/2$ with the first approach, this gives a linear space static range mode data structure with query time $O(\sqrt{n} \log n)$. By substituting an $O(\log \log n)$ data structure for predecessor search, such as van Emde Boas trees, the query time can immediately be improved to $O(\sqrt{n} \log \log n)$.

Chan et al. [3] focus on linear space solutions to static range mode. They achieve a clever array-based solution with $O(n)$ space and $O(\sqrt{n})$ query time. By using bit-packing tricks and more advanced data structures, they reduce the query time to $O(\sqrt{n/\log n})$.

As with many range query data structures, the range mode problem has also been studied in an approximate setting [12, 1].

2.2 Static Range Least Frequent Query

The range least frequent problem was first studied by Chan et al. [4]. They again focus on linear-space solutions, this time achieving $O(n)$ space and $O(\sqrt{n})$ time query. In their paper, they focus on the version of range least frequent element where the element must occur in the query range.

2.3 Dynamic Range Mode

Chan et al. [3] also study the dynamic range mode problem. They give a solution tradeoff that at linear space, achieves $O(n^{3/4} \log n / \log \log n)$ worst-case time range mode query and $O(n^{3/4} \log \log n)$ amortized expected time update. At minimal update/ query time, this tradeoff gives $O(n^{4/3})$ space and $O(n^{2/3} \log n / \log \log n)$ worst-case time range mode query and amortized expected time update.

2.4 Lower Bounds

Both [3] and [4] give conditional lower bounds for range mode and range least frequent problems, respectively. Chan et al. [3] reduces multiplication of two $\sqrt{n} \times \sqrt{n}$ boolean matrices to n range mode queries in an array of size $O(n)$. This indicates that with current knowledge, preprocessing an $O(n)$ -sized range mode query data structure and answering n range mode queries cannot be done in better than $O(n^{\omega/2})$ time, where ω is the constant in the exponent of the running time of matrix multiplication. With $\omega = 2.3727$ [17], this implies with current knowledge that a range mode data structure must either have preprocessing time at least $\Omega(n^{1.18635})$ or query time at least $\Omega(n^{0.181635})$. Since we may choose to update an array rather than initializing it, this lower bound also indicates that a dynamic range mode data structure must have update/ query time at least $\Omega(n^{\omega/2-1})$.

In [3], Chan et al. also give another conditional lower bound for dynamic range mode. They reduce the multiphase problem of Pătraşcu [16] to dynamic range mode. A reduction from 3-SUM (given n integers, find three that sum to zero) is given by Pătraşcu [16] to the multiphase problem. Based on the conjecture that the 3-SUM problem cannot be solved in $O(n^{2-\varepsilon})$ time for any positive constant ε , this chain of reductions implies a dynamic range mode data structure must have polynomial time query or update.

The reduction of $\sqrt{n} \times \sqrt{n}$ boolean matrix multiplication to n $O(n)$ -sized range mode queries can be adopted to achieve the same conditional lower bound for the range least frequent problem [4]. Upon examination, the conditional lower bound also applies to k -frequency queries.

An unconditional lower bound also exists in the cell probe model for the range mode and k -frequency problem. Any range mode/ k -frequency data structure that uses S memory cells of w -bit words needs $\Omega(\frac{\log n}{\log(Sw/n)})$ time to answer a query [12].

3 Preliminaries

Before we discuss the technical details of our results, it will be helpful to develop some notation.

As in [3, 4], we will denote the subarray of A from index i to index j as $A[i : j]$ and use array notation $A[i]$ to denote the element of A at index i . We will assume zero-based indexing throughout this paper.

Furthermore, for range frequency data structures, the actual type that array A stores is irrelevant; we only care about how many times each element occurs. It will be useful to think of the identity of an element as a color. Therefore, we may say the color c occurs with frequency f in range $A[l : r]$. This is to distinguish that we are not referring to a particular index but rather the identity of multiple indices in the range.

For simplicity, we will assume the number of elements n is a perfect cube; however, the results discussed easily generalize for arbitrary n . All log's in this paper are assumed to be base 2.

In all our proofs, we analyze space cost in words, that is, $O(\log n)$ collections of bits.

4 Data Structure Setup

The idea of our data structure will be to break up array A into $O(n^{1/3})$ evenly-spaced endpoints, so that there are $O(n^{2/3})$ elements between each endpoint. We will use capital letters L and R when referring to particular endpoints. We will also occasionally refer to the elements between two consecutive endpoints as a segment; therefore, there are $O(n^{1/3})$ segments in A .

Each color that occurs in A will be split into the following two disjoint categories:

- **Frequent Colors:** Any color that appears more than $n^{1/3}$ times in A .
- **Infrequent Colors:** Any color that appears at most $n^{1/3}$ times in A .

Note that there can be at most $n/n^{1/3} = n^{2/3}$ frequent colors in A at any point in time.

Our data structure will need to use dynamic arrays as auxiliary data structures. These dynamic arrays will be a modification of the simple two-level version of the data structure described by Goodrich and Koss [11]. We have the following lemma regarding the performance of these dynamic arrays:

► **Lemma 1.** *There is a dynamic array data structure D that occupies $O(n)$ space and supports:*

1. $D[i]$: Retrieve/ Set the element at rank i , in $O(1)$ time.
2. $Insert(i, x)$: Insert the element x at rank i , in $O(\sqrt{n})$ time.
3. $Delete(i, x)$: Delete the element x at rank i , in $O(\sqrt{n})$ time.
4. $Rank(ptr)$: Determine the rank of the element pointed to by ptr , in $O(1)$ time.

The modified data structure will be discussed in the full version of this paper. If desired, the data structure can be replaced by a balanced binary search tree, at the cost of additional logarithmic factors in the query times of our data structure.

We can now describe the base set of auxiliary data structures used throughout the paper:

1. Arrays $B_{L,R}$, for all pairs of endpoints L, R , indexed from 0 to $n^{1/3}$, so that

$$B_{L,R}[i] := \text{The number of infrequent colors with frequency } i \text{ in } A[L : R].$$

25:6 Improved Time and Space Bounds for Dynamic Range Mode

2. Dynamic arrays D_c , for every color c , so that

$D_c[i] :=$ The index in A of the i th occurrence of color c .

3. An array E parallel to A so that

$E[i] :=$ A pointer to the location in memory of index i in dynamic array $D_{A[i]}$.

4. A binary search tree F of endpoints. At each endpoint R , we store

$F[R] :=$ A binary search tree on frequent colors, giving their frequency in $A[0 : R]$.

In regards to $B_{L,R}$, we will sometimes refer to the set of elements between endpoints L and R as the span of L and R .

We now analyze the space complexity and construction time.

► **Lemma 2.** *The base data structures take $O(n)$ space.*

Proof. The arrays $B_{L,R}$ have size $O(n^{1/3})$ and there are $O(n^{2/3})$ of them, so in total these take $O(n)$ space. Every index of A is present in exactly one of the D_c arrays, and each dynamic array takes linear space. Therefore, in total all D_c arrays take $O(n)$ space. Array E has exactly the same size as A and thus takes $O(n)$ space. The binary search trees in each node of F have size equal to the number of frequent colors, which is at most $n^{2/3}$. Since there are $O(n^{1/3})$ endpoints and thus nodes of F , this structure takes $O(n)$ space. ◀

► **Lemma 3.** *The base data structures can be initialized in $O(n^{4/3})$ time.*

Proof. We can count the number of occurrences of each color in $O(n \log n)$ time and determine for each color whether it is frequent or infrequent. Let A' be the array A without any frequent colors. We can scan A' $n^{1/3}$ times, starting from each endpoint, to build the arrays $B_{L,R}$. This will be done as follows. For each scan, we maintain an array T so that $T[c]$ denotes the number of occurrences of color c found so far. We also maintain the array $B_{L,*}$ which is the array B with endpoint L and a variable right endpoint, that is maintained as elements are scanned. When $A[i] = c$, we check the number of occurrences of c to update $B_{L,*}$ to the correct state. In total, each element scanned results in $O(1)$ operations, until we reach a right endpoint. When we reach a right endpoint, we write $B_{L,*}$ to array $B_{L,R}$, where R denotes the right endpoint just encountered. In this way, for all endpoints L, R , we spend $O(n^{1/3})$ time to create the array. Thus the element scan dominates the time complexity, requiring $O(n^{4/3})$ time to create all $B_{L,R}$ arrays.

The dynamic arrays D_c can be built in linear time overall by walking through A and appending indices to the ends of D_c arrays. At this same time E can be built. The BSTs for all endpoints in F can also be built in linear time overall, since there are at most $n^{2/3}$ frequent colors per list, there are $n^{1/3}$ lists in total, and counting each frequent element in each interval can be done at the same time in one scan through A . ◀

Throughout the next few sections, we will use the following Lemma, as in [3]:

► **Lemma 4.** *Let $A[i] = c$. Then, given frequency f and right endpoint j , our base data structures may answer the following questions in constant time:*

Does color c occur at least f times in $A[i : j]$? (1)

Does color c occur at most f times in $A[i : j]$? (2)

Does color c occur exactly f times in $A[i : j]$? (3)

Proof. We call $\text{rank}(E[i])$ on D_c to get the rank r of i in D_c . Since array D_c stores the indices of every occurrence of color c in A , the values of $D_c[r+f]$ and $D_c[r+f-1]$ determine the answers to the above questions. ◀

A similar strategy can be used to answer the above questions given an index j and left endpoint i .

5 Range Mode Query

The range mode query will make use of the following lemma, originating from [14] and also used by [3]:

► **Lemma 5** (Krizanc et al. [14]). *Let A_1 and A_2 be any multisets. If c is a mode of $A_1 \cup A_2$ and $c \notin A_1$, then c is a mode of A_2 .*

The query algorithm can be summarized as follows:

Algorithm 1 Range Mode Query in $A[l:r]$.

Let L and R be the first and last endpoints in $[l, r]$, respectively.

1. Check the frequency of every frequent color in $A[L:R]$ via BSTs $F[R]$ and $F[L]$. Let f be the highest frequency found so far.
 2. Ask question (1) for all colors in $A[l:L-1] \cup A[R+1:r]$ with frequency f and right endpoint r / left endpoint l . If (1) is answered in the affirmative for color c , linearly scan D_c to count the number of occurrences of color c in $[l, r]$, update f , and continue.
 3. Find the largest nonzero index of $B_{L,R}$. If this is larger than f , update f and do the following:
 - a. Find the next endpoint R' to the left of R .
 - b. Check $B_{L,R'}[f]$. If $B_{L,R'}[f] < B_{L,R}[f]$, search $A[R'+1:R]$ for a color that occurs f times in $A[L:R]$, via question (3) with left endpoint L .
 - c. Otherwise, repeat from step (a) with $R \leftarrow R'$.
 4. Return f and the corresponding color found from either step 1, 2, or 3(b).
-

► **Theorem 6.** *Algorithm 1 finds the current range mode of $A[l:r]$ in $O(n^{2/3})$ time.*

Proof. Endpoints L and R can be found from $[l, r]$ by appropriate floors and ceilings in constant time.

In step 1, we can iterate through $F[R]$ and $F[L]$ in $O(n^{2/3})$ time, determining frequency counts for all frequent elements.

In step 2, answering question (1) takes $O(1)$ time per element via Lemma 4. Note that we use left endpoint l for elements in range $A[R+1:r]$ and right endpoint r for elements in range $A[l:L-1]$. Although we do not check (1) for the full range $[l, r]$, we will check the first/ last occurrence of any color in $A[l:L-1] \cup A[R+1:r]$, thereby effectively checking for all of $[l, r]$. When (1) is answered in the affirmative at index i , we linearly scan D_c , starting at $D_c[i+f+1]$ ($D_c[i-f-1]$ if $i \in [R+1:r]$) to determine a new highest frequency. Let us determine the cost of these linear scans. Let c be the most frequent color found from steps 1 and 2. If c is an infrequent color, it cannot occur more than $n^{1/3}$ times, so the total cost of the linear scans is no more than $O(n^{1/3})$. If it is a frequent color, its frequency cannot have increased by more than $O(n^{2/3})$, since its frequency was checked in step 1 and only $O(n^{2/3})$ elements exist in $A[l:L-1] \cup A[R+1:r]$. Thus the total cost of the linear scans is no more than $O(n^{2/3})$. In either case, step 2 takes $O(n^{2/3})$ time.

For step 3, finding the largest nonzero index of $B_{L,R}$ takes $O(n^{1/3})$ time. If this is larger than the frequencies found in steps 1 or 2, we execute steps 3(a) - 3(c). We can only repeat the steps at most $O(n^{1/3})$ times. The condition $B_{L,R'}[f] < B_{L,R}[f]$ will happen for one of these iterations, since eventually $R' = L$ in which case there are no elements accounted for in the interval. When $B_{L,R'}[f] < B_{L,R}[f]$, this implies a color with frequency f in range $[L, R]$ appears somewhere in $A[R' + 1 : R]$. There are only $O(n^{2/3})$ elements in $A[R' + 1 : R]$. Checking if color c occurs exactly f times in $A[L : R]$ can be done by asking question (3).

For the correctness of the value returned in step 4, note that the mode of $A[l : r]$ is either an element in $A[l : L - 1] \cup A[R + 1 : r]$ or the mode of $A[L : R]$, by Lemma 5. The frequency of all colors in $A[l : L - 1] \cup A[R + 1 : r]$ is checked. Further, the frequency of infrequent and frequent colors for interval $A[L : R]$ is also checked in steps 1 and 3, respectively. Therefore the color and frequency returned in step 4 must be the mode of $A[l : r]$. Putting it together, we see Algorithm 1 is correct and takes $O(n^{2/3})$ time. ◀

6 Update Operation

The update operation will require us to keep the base data structures up to date. Given update $A[i] \leftarrow c$, there are two similar procedures that must occur: adjusting the data structures for the removal of current color $A[i]$ and adjusting the data structures for the addition of color c at index i .

The following algorithm can be used for both procedures. We note that if either the color removed is the last occurrence of its type or the color added is a new color, the list D_c will have to also be constructed/ deleted and $B_{L,R}$ should be modified to only reflect colors present in A . For simplicity we omit these details from Algorithm 2.

Algorithm 2 Update Base Data Structures for Addition/ Removal of Color c at Index i .

1. If color c is infrequent prior to this operation, count how many times c occurs in each span via D_c , decrementing the corresponding index of each B array.
 2. Adjust D_c by adding/ removing i . If this is an add operation, set $E[i]$ to the memory location of i in D_c .
 3. If color c remains or becomes infrequent after step 2, again count how many times c occurs in each span via D_c , incrementing the corresponding index of each B array.
 4. If color c became infrequent in step 2, delete its entry in all nodes of F . If color c became frequent after step 2, add its frequencies to F . If color c remained frequent after step 2, increment the frequencies of all prefixes including index i in F .
-

► **Theorem 7.** *Algorithm 2 updates the base data structures for addition/ removal of color c at index i in $O(n^{2/3})$ time.*

Proof. If c is an infrequent color prior to the update, step 1 removes its contribution to all B arrays. Counting the frequency of color c in each interval can be done via $O(n^{1/3})$ searches through D_c , each taking $O(n^{1/3})$ time. We first start at the beginning, finding its frequency for all right endpoints with left endpoint fixed, then move the left endpoint to the next endpoint and repeat, etc. Step 1 in total takes $O(n^{2/3})$ time.

Adding or removing i from D_c takes $O(\sqrt{n})$ time, as given in Lemma 1. Adjusting $E[i]$ can be done during this operation, so in total step 2 takes $O(\sqrt{n})$ time. The analysis of step 3 is identical to step 1. In step 4, adding, deleting, or modifying the entry of color c in all/

some nodes of F takes $O(n^{1/3} \log n)$ time, since each node stores the list of frequent colors and their frequency counts in a BST. If color c just became frequent, it only occurs $O(n^{1/3})$ times in A , and thus counting its frequency from the beginning to each endpoint can be done in $O(n^{1/3})$ time. In total, step 4 takes at most $O(n^{1/3} \log n)$ time.

After the completion of Algorithm 2, B arrays, dynamic array D_c , array E , and binary search tree F has been updated to reflect the current state of array A . Since all steps execute in no more than $O(n^{2/3})$ time, Algorithm 2 is correct and runs in $O(n^{2/3})$ time. ◀

7 Range Least Frequent Query, Allowing Zero

To answer range least frequent queries, we require the use of one more auxiliary data structure: an array $C_{L,R}$ for all pairs of endpoints L, R , indexed from 1 to $n^{1/3}$. At index $C_{L,R}[i]$ we store a list of all colors that occur i times in A such that the smallest span enclosing all occurrences is $[L, R]$.

Since each infrequent color is represented exactly once in the $C_{L,R}$ lists and there are $O(n^{2/3} \cdot n^{1/3}) = O(n)$ total indices present, this data structure takes linear space. It can also be initialized in linear time, and we can modify the update procedure of Algorithm 2 to update $C_{L,R}$ at the same time as $B_{L,R}$ for no additional time cost.

It is additionally worth noting that since the Range Least Frequent Query will require $\tilde{O}(n^{2/3})$ time, the use of dynamic arrays for data structures D_c is not necessary for this section. Instead, we can use binary search trees, augmented to support lookup by index in $O(\log n)$ time, or Dietz' data structure [6]. For the best time complexity, we will use augmented binary search trees to count occurrences of colors in a specific range and an augmented dynamic linear-space van Emde Boas tree to count occurrences of colors between endpoints. The van Emde Boas tree stores a single node for any endpoint R that a color appears in, which keeps the number of occurrences of that color from the beginning to endpoint R . The van Emde Boas tree can be updated in $O(n^{1/3} \log \log n)$ time upon insertion or removal and via predecessor/ successor queries, can support counting the number of occurrences of any color between any two endpoints in $O(\log \log n)$ time.

Algorithm 3 Range Least Frequent Query in $A[l:r]$, Allowing Zero.

Let L and R be the first and last endpoints in $[l, r]$, respectively.

1. Check the frequency of every frequent color in $A[l:r]$ via D_c . Let f be the lowest frequency found so far.
 2. Find the set of all infrequent colors that occur in $A[l:L-1] \cup A[R+1:r]$; call it U . Count the frequencies of all colors of U in range $A[l:r]$ and update f with the lowest frequency found so far.
 3. Compute $B'_{L,R}$, the array $B_{L,R}$ updated to erase the contribution of all colors in U .
 4. If the smallest positive-valued index of $B'_{L,R}$ is less than f , update f and check if any list $C_{L',R'}[f]$, $[L', R'] \subseteq [L, R]$, is non-empty. If so, return a color from the appropriate list. Otherwise, binary search from R to the last endpoint of A in the following way:
 - a. Let R' be a value in the middle of the search range. If $B'_{L,R'}[f] < B'_{L,R}[f]$, let R' be the new upper bound; otherwise, continue the search in the half of the range after R' .
 - b. When the search range is two consecutive endpoints, we search the range for a color that occurs f times in $A[L:R]$ and return its identity.
 - c. If the condition in a. is never satisfied, we must repeat a binary search on the other side, with an initial search range of the beginning of A to L .
-

► **Theorem 8.** *Algorithm 3 finds the least frequent element of $A[l : r]$, allowing zero, in $O(n^{2/3} \log n \log \log n)$ time.*

Proof. We can find a list of frequent colors in any node of the F BST. Using an augmented binary search tree for D_c , step 1 can then be done in $O(n^{2/3} \log n)$ time. In step 2, since there are $O(n^{2/3})$ elements in $A[l : L-1] \cup A[R+1 : r]$, we can find color set U and complete step 2 similarly to step 1 in $O(n^{2/3} \log n)$ time. Step 3 requires counting the frequency of each color of U in range $A[L : R]$ and decrementing the corresponding index to make $B'_{L,R}$; thus it can also be done in $O(n^{2/3} \log \log n)$ time using the augmented van Emde Boas tree.

In step 4 we are looking for a least frequent element of $A[L : R]$ that does not occur in $A[l : L-1] \cup A[R+1 : r]$. All colors of $A[l : L-1] \cup A[R+1 : r]$ are represented in set U , found in step 3. We effectively erase the contribution of colors of U via computing $B'_{L,R}$; therefore, we can proceed as if colors of U do not exist in A . We will refer to A' as array A without any colors of U .

If the least frequent color in $A'[L : R]$ does not exist elsewhere in A , then its identity will be stored in a list $C_{L',R'}[f]$, $[L', R'] \subseteq [L, R]$. Note colors of U need not be special-cased for this lookup, since by appearing in $A[l : L-1] \cup A[R+1 : r]$, they will not be present in any of the searched lists. Otherwise, we know the least frequent color in $A'[L : R]$ must occur somewhere else in A' .

Now, amongst all colors in A' , we know frequency f is minimal in range $A'[L : R]$. Therefore if we increase the range to $A'[L : R']$, the frequency of colors can only increase. For this property to hold, we must allow $f = 0$.

In each iteration, the smallest positive-valued index of $B'_{L,R'}$ will be f or greater and $B'_{L,R'}[f]$ will be no more than $B'_{L,R}[f]$. If it is less, we know one of the colors that occurred f times in $A'[L : R]$ now occurs more than f times in $A'[L : R']$. Therefore we may find it in the half of the search range before R' . If it is the same, we know none of the colors that occurred f times in $A'[L : R]$ appear in the half of the search range before R' . Either way we decrease the search range by a factor of 2. When the range represents two consecutive endpoints, we can search it for a color that occurs f times in $A[l : r]$ in $O(n^{2/3} \log n)$ time.

However, if R' is the end of the array and $B'_{L,R'}[f] = B'_{L,R}[f]$, then none of the colors that appear f times in $A'[L : R]$ appear to the right of R in A' . In this case, we can repeat the same binary search on the other side, decreasing a left endpoint L' and checking the same condition. Since the least frequent color in $A'[L : R]$ must occur elsewhere in A , as it was not present in any of the lists $C_{L',R'}[f]$, the search on this side must identify a color that appears f times in $A'[L : R]$.

The time complexity of step 4 can be analyzed as follows. Checking the lists $C_{L',R'}$ takes $O(n^{2/3})$ time, since there can be $O(n^{2/3})$ endpoints $[L', R'] \subseteq [L, R]$. The binary search is on endpoints, of which there are $O(n^{1/3})$. Thus, there are $O(\log(n^{1/3})) = O(\log n)$ iterations of the binary search, and in each iteration we must compute $B'_{L,R'}$ or $B'_{L',R}$. This computation takes $O(n^{2/3} \log \log n)$ time as in step 3. Therefore the binary search process takes $O(n^{2/3} \log n \log \log n)$ time. In total, step 4 takes $O(n^{2/3} \log n \log \log n)$ time.

For correctness, the least frequent element in $A[l : r]$ is either a frequent or infrequent color. If it is a frequent color, it is identified in step 1. If it is an infrequent color, we have two cases. Either the color occurs in $A[l : L-1] \cup A[R+1 : r]$, and thus set U , or it does not occur in set U . Step 2 accounts for all infrequent colors in set U . Steps 3 and 4 account for the last case. By the above, these steps find an infrequent color in $A[L : R]$ that does not occur in $A[l : L-1] \cup A[R+1 : r]$ in $O(n^{2/3} \log n \log \log n)$ time. Thus, Algorithm 3 is correct and finds the least frequent element of $A[l : r]$, allowing zero, in $O(n^{2/3} \log n \log \log n)$ time. ◀

8 Range k -Frequency Query

The previous two sections make use of a monotonicity property to find a color of given frequency in a range: for range mode, we know the frequency of the most frequent element can only decrease if the query range is decreased; furthermore, for range least frequent, we know the frequency of the least frequent element can only increase if the query range is increased. For this monotonicity condition to hold for least frequent elements, we must allow answering with an element of frequency zero. To force our answer to be an element that occurs in the query range, we must use an additional data structure that allows retrieval of colors by frequency in the $B_{L,R}$ arrays. To achieve $\tilde{O}(n)$ space, we cannot afford to store a list of colors at each index $B_{L,R}[i]$, and storing a single color at each index will run into issues during updates. Instead, we will use the following data structure which is similar to randomized data structures in the dynamic streaming literature [8]:

► **Lemma 9.** *There is a Monte Carlo data structure that occupies expected $O(\log^2 n)$ space and supports:*

1. *Insert(x): Insert element x into the collection, in $O(\log n)$ expected time.*
2. *Delete(x): Delete element x from the collection, in $O(\log n)$ expected time.*
3. *Retrieve(): Return an element in the collection, in $O(1)$ expected time.*

The data structure requires the Delete(x) operation is executed independently of the results of Retrieve().

The low-space data structure will be described in the full version of this paper. With it, we can answer the general problem of finding an element of given frequency in a query range. The additional auxiliary data structures needed will be as follows:

1. An array $G_{L,R}$ parallel to $B_{L,R}$. At index $G_{L,R}[i]$, we store the number of infrequent colors with frequency i in $A[L : R]$, excluding colors that appear in segments immediately left of L or right of R .
2. An array $H_{L,R}$, parallel to $G_{L,R}$, so that at index $H_{L,R}[i]$, we store a collection of colors counted in $G_{L,R}[i]$ in the data structure of Lemma 9.

The array $G_{L,R}$ is similar to the item (ii) stored in table D of [4]. During preprocessing, G arrays can be built similarly to B ; however, when we fix left endpoint, we check all colors that occur in the segment to the left. We avoid counting such colors. Similarly, before we finalize the count for $G_{L,R}$, we move the right endpoint out as if to count the next range, keeping track of colors encountered. We subtract the frequency of such colors in $G_{L,R}$. Whenever we add to/ subtract from $G_{L,R}$, we can insert or delete the color from $H_{L,R}$.

As explained above, our space cost now becomes $\tilde{O}(n)$. Furthermore, our updates to G and H can be done alongside the update to B ; however, since each insertion takes $O(\log n)$ time, the update procedure now takes $O(n^{2/3} \log n)$ time. With this we have:

Algorithm 4 Range k -Frequency Query in $A[l:r]$.

Let L and R be the first and last endpoints in $[l, r]$, respectively.

1. For color c at index i in the segment left of L , check via D_c to see if i is the first index of color c to appear in range $[l, r]$, or, if outside $[l, r]$, the next occurrence of color c lies in range $[l, r]$. If so, ask question (3) for the occurrence of color c in $[l, r]$ with frequency k and right endpoint r . If answered in the affirmative, return color c . Do the same, symmetrically, for colors in the segment right of R .
 2. For each frequent color not addressed in step 1, check its frequency in $A[L:R]$ via BSTs $F[R]$ and $F[L]$. If any occur with frequency k , return the color.
 3. If no color is found from step 1, check if $G_{L,R}[k] > 0$. If so, return $H_{L,R}[k].Retrieve()$. If not, return that no color has frequency k in range $A[l:r]$.
-

► **Theorem 10.** *Algorithm 4 returns an element of frequency k in $A[l:r]$ or indicates no such element exists, assuming update operations have been executed independently of results of query operations, in $O(n^{2/3})$ time.*

Proof. In step 1, we look at two full segments of $O(n^{2/3})$ total elements. For color c at index i , if $i \in [l, r]$, we must determine if index i is the first occurrence of color c in $[l, r]$. Let $m = D_c.rank(E[i])$. Index i is the first occurrence of color c in $[l, r]$ if $D_c[m-1]$ is outside $[l, r]$. Similarly, if $i \notin [l, r]$, we can again define $m = D_c.rank(E[i])$, then check if $D_c[m+1]$ is in $[l, r]$. In any case, for any color that occurs in segments immediately left of L or right of R , one of the indices will be the first outside $[l, r]$ or the first within $[l, r]$. Thus the frequency of the color will be checked in $A[l:r]$. Since we do a constant number of constant time operations for $O(n^{2/3})$ elements, step 1 takes $O(n^{2/3})$ time.

As in Algorithm 1, step 2 takes $O(n^{2/3})$ time. Step 3 takes $O(1)$ time. In any case, in step 1 we check all elements in segments immediately left of L or right of R to see if they occur k times in $A[l:r]$. Furthermore, in steps 2 and 3, we check every frequent and infrequent color that occurs in $A[L:R]$ but not in segments immediately left of L or right of R , via array G . Thus we have checked every color if it occurs k times in $A[l:r]$. Since no step takes more than $O(n^{2/3})$ time, this proves Theorem 10. ◀

Algorithm 4 can be easily modified to return the least frequent element present in the query range with the same time complexity and independence assumption. It can also be modified to count the number of elements above, below, or at a given frequency, as well as only determine the frequency of the least frequent element. Since these queries do not ask for a color, arrays $H_{L,R}$ are not needed. This reduces the space cost to $O(n)$, update cost to $O(n^{2/3})$, and requires no independence assumption.

References

- 1 Prosenjit Bose, Evangelos Kranakis, Pat Morin, and Yihui Tang. Approximate range mode and range median queries. In *Annual Symposium on Theoretical Aspects of Computer Science*, 2005.
- 2 Gerth Stølting Brodal, Beat Gfeller, Allan Grønlund Jørgensen, and Peter Sanders. Towards optimal range medians. *Theoretical Computer Science*, 412(24):2588–2601, 2011.
- 3 Timothy M. Chan, Stephane Durocher, Kasper Green Larsen, Jason Morrison, and Bryan T. Wilkinson. Linear-space data structures for range mode query in arrays. *Theory of Computing Systems*, 55:719–741, 2014.

- 4 Timothy M. Chan, Stephane Durocher, Matthew Skala, and Bryan T. Wilkinson. Linear-space data structures for range minority query in arrays. *Algorithmica*, 72:901–913, 2015.
- 5 Timothy M Chan and Konstantinos Tsakalidis. Dynamic orthogonal range searching on the ram, revisited. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 77. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- 6 Paul F. Dietz. Optimal algorithms for list indexing and subset rank. In *Workshop on Data Structures and Algorithms*, pages 39–46, 1989.
- 7 Amr Elmasry, Meng He, J Ian Munro, and Patrick K Nicholson. Dynamic range majority data structures. In *International Symposium on Algorithms and Computation*, pages 150–159. Springer, 2011.
- 8 Gereon Frahling, Piotr Indyk, and Christian Sohler. Sampling in dynamic data streams and applications. *International Journal of Computational Geometry & Applications*, 18(1/2):3–28, 2008.
- 9 Harold N. Gabow, Jon Louis Bentley, and Robert E. Tarjan. Scaling and related techniques for geometry problems. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 135–143, 1984.
- 10 Travis Gagie, Meng He, and Gonzalo Navarro. Compressed dynamic range majority data structures. In *Data Compression Conference (DCC), 2017*, pages 260–269. IEEE, 2017.
- 11 Michael T. Goodrich and John G. Koss II. Tiered vectors: Efficient dynamic arrays for rank-based sequences. In *Workshop on Data Structures and Algorithms*, 1999.
- 12 Mark Greve, Allan Grønlund Jørgensen, Kasper Dalgaard Larsen, and Jakob Truelsen. Cell probe lower bounds and approximations for range mode. In *International Colloquium on Automata, Languages, and Programming*, 2010.
- 13 Meng He, J. Ian Munro, and Patrick K. Nicholson. Dynamic range selection in linear space. In *International Symposium on Algorithms and Computation*, 2011.
- 14 Danny Krizanc, Pat Morin, and Michiel Smid. Range mode and range median queries on lists and trees. In *International Symposium on Algorithms and Computation.*, pages 517–526, 2003.
- 15 Mihai Pătraşcu and Emanuele Viola. Cell-probe lower bounds for succinct partial sums. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 117–122. Society for Industrial and Applied Mathematics, 2010.
- 16 Mihai Pătraşcu. Towards polynomial lower bounds for dynamic problems. In *Proceedings of the forty-second annual ACM symposium on Theory of computing*, pages 603–610, 2010.
- 17 Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *Proceedings of the forty-fourth annual symposium on theory of computing*, 2012.