

The effective entropy of next/previous larger/smaller value queries

Dekel Tsur*

Abstract

We study the problem of storing the minimum number of bits required to answer next/previous larger/smaller value queries on an array A of n numbers, without storing A . We show that these queries can be answered by storing at most $3.701n$ bits. Our result improves the result of Jo and Satti [TCS 2016] that gives an upper bound of $4.088n$ bits for this problem.

Keywords data structures, encoding model.

1 Introduction

A recent research area in data structures is designing data structures in the encoding model [1–4, 6–12]. In this model, the goal is to design a data structure for answering queries on some object A , without storing A . The space complexity of the data structure should be close to the minimum space required in order to answer the queries, without storing A . The minimum space required to answer the queries is called the *effective entropy* of the problem.

Let A be an array of n numbers. Consider the following four queries on A .

- Previous smaller value (PSV(i)): Given i , return $\max(\{j: j < i, A[j] < A[i]\} \cup \{0\})$.
- Previous larger value (PLV(i)): Given i , return $\max(\{j: j < i, A[j] > A[i]\} \cup \{0\})$.
- Next smaller value (NSV(i)): Given i , return $\min(\{j: j > i, A[j] < A[i]\} \cup \{n+1\})$.
- Next larger value (NLV(i)): Given i , return $\min(\{j: j > i, A[j] > A[i]\} \cup \{n+1\})$.

The effective entropy of answering one type of queries from the four types above is $2n - \Theta(\log n)$ bits. If the problem is to support more than one type of queries, the effective entropy becomes larger. Fischer [4] showed that the effective entropy of answering both PSV and NSV queries is $\log(3 + 2\sqrt{2}) \cdot n - \Theta(\log n) < 2.544n$ bits. Gawrychowski and Nicholson [7] showed that the effective entropy of answering both PSV and PLV queries is at most $3n$ bits and at least $3n - \Theta(\log n)$ bits. For each of the two problems above, it is possible to build a data structure that answers queries in constant time and with space that is equal to the effective entropy plus $o(n)$ bits [4, 7].

Jo and Satti [12] studied the problem of answering all four queries. They showed that the effective entropy is at most $4n + o(n)$ bits on arrays with no consecutive equal elements, and at most $\log_2 17 \cdot n + o(n) < 4.088n$ bits on general arrays. They also showed that it is possible to build data structures that answer queries in constant time with space

*Department of Computer Science, Ben-Gurion University of the Negev. Email: dekelts@cs.bgu.ac.il

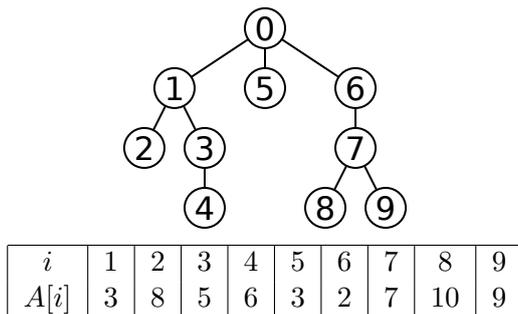


Figure 1: An example of a 2d-min heap of an array A .

complexity $4n + o(n)$ bits on arrays with no consecutive equal elements, and $4.585n$ bits for general arrays. In this paper we improve the results of Jo and Satti. We show that the effective entropy for answering all four queries is at most $(2 + \log 3)n + o(n) < 3.585n$ bits on array with no consecutive equal elements, and at most $\log 13 \cdot n + o(n) < 3.701n$ bits on general arrays.

2 Preliminaries

2.1 Encoding PSV queries

The *2d-min heap* [5] of an array A of size n , denoted $\text{Min}(A)$, is an ordinal tree with nodes $0, \dots, n$. For every $i > 0$, the parent of node i is $\text{PSV}(i)$. The children of a node i are ordered in increasing order of their names. See Figure 1 for an example. Note that the preorder of the nodes is $0, \dots, n$. Therefore, the tree $\text{Min}(A)$ can be encoded by just storing its topology. Since $\text{Min}(A)$ is an ordinal tree with $n + 1$ nodes, it follows that the effective entropy of PSV queries is at most $\lceil \log C_{n+1} \rceil = 2n - \Theta(\log n)$, where $C_n = \frac{1}{n+1} \binom{2n}{n}$ is the number of ordinal trees with n nodes. We also define the 2d-max heap, denoted $\text{Max}(A)$, to be an ordinal tree in which the parent of node $i > 0$ is $\text{PLV}(i)$.

2.2 Encoding PSV/NSV queries

In order to encode both PSV and NSV queries, Fischer [4] defined the *colored 2d-min heap* of an array A , denoted $\text{cMin}(A)$, to be the tree $\text{Min}(A)$ with the following coloring of its nodes. If node i has right siblings and $A[i] \neq A[j]$, where j is the immediately right sibling of i , then i is colored red. Otherwise, i is colored blue. See Figure 2 for an example. Fischer showed that if $\text{cMin}(A)$ is known, both PSV and NSV queries on A can be answered without storing A . Since $\text{cMin}(A)$ is a Schröder tree, it follows that the effective entropy of PSV/NSV queries is at most $\log(3 + 2\sqrt{2}) \cdot n - \Theta(\log n)$ bits [13].

2.3 Encoding PSV/PLV queries

To answer both PSV and PLV queries, we can store both $\text{Min}(A)$ and $\text{Max}(A)$. Gawrychowski and Nicholson [7] showed that for an array A with no consecutive equal elements, $\text{Min}(A)$ and $\text{Max}(A)$ can be encoded using $3n - 1$ bits. The proof of this result is based on the following claim.

Claim 1. *If A has no consecutive equal elements, then for every $0 < i < n$, i is a leaf in $\text{Min}(A)$ if and only if i is an internal node of $\text{Max}(A)$*

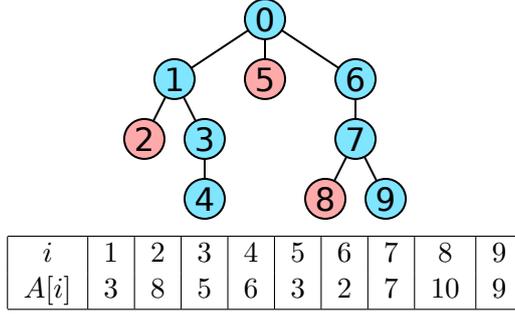


Figure 2: An example of a colored 2d-min heap of an array A .

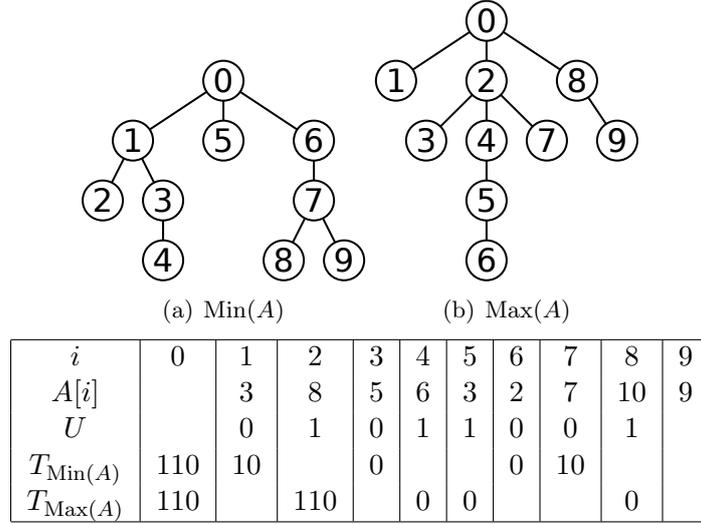


Figure 3: An example of the trees $\text{Min}(A)$ and $\text{Max}(A)$ of an array A and the encoding of these trees using $3n - 1$ bits.

Proof. Since A has no consecutive equal elements, we have that either $A[i] < A[i + 1]$ or $A[i] > A[i + 1]$. In the former case, $\text{PSV}(i + 1) = i$. Therefore, $i + 1$ is a child of i in $\text{Min}(A)$, and thus i is an internal node in $\text{Min}(A)$. Additionally, $\text{PLV}(i + 1) < i$, so $i + 1$ is not a child of i in $\text{Max}(A)$. Since the names of the nodes of $\text{Max}(A)$ are according to their ranks in the preorder, the descendants of an internal node j are $j + 1, \dots, j'$ for some j' . Since $i + 1$ is not a descendant of i in $\text{Max}(A)$, it follows that i is a leaf in $\text{Max}(A)$.

The proof for the case when $A[i] > A[i + 1]$ is analogous and thus omitted. \blacksquare

The encoding $\text{Min}(A)$ and $\text{Max}(A)$ consists of three binary strings U , $T_{\text{Min}(A)}$, and $T_{\text{Max}(A)}$. The string U is a string of length $n - 1$ in which $U[i] = 1$ if and only if i is a leaf in $\text{Min}(A)$. The strings $T_{\text{Min}(A)}$ and $T_{\text{Max}(A)}$ are defined as follows. Start with empty strings $T_{\text{Min}(A)}$ and $T_{\text{Max}(A)}$. Then, for $i = 0, 1, \dots, n - 1$, if $i = 0$ or $U[i] = 0$ (namely, i is an internal node in $\text{Min}(A)$), append the string $1^{d_i-1}0$ to $T_{\text{Min}(A)}$, where d_i is the number of children of node i in $\text{Min}(A)$. Additionally, if $i = 0$ or $U[i] = 1$, append the string $1^{d'_i-1}0$ to $T_{\text{Max}(A)}$, where d'_i is the number of children of node i in $\text{Max}(A)$. See Figure 3 for an example.

By Claim 1, the trees $\text{Min}(A)$ and $\text{Max}(A)$ can be reconstructed from the strings $U, T_{\text{Min}(A)}, T_{\text{Max}(A)}$. It is easy to show that $|T_{\text{Min}(A)}| + |T_{\text{Max}(A)}| = 2n$. Therefore, the total size of this encoding is $3n - 1$ bits.

3 Arrays with no consecutive equal elements

Theorem 2. *The effective entropy of PSV/PLV/NSV/NLV queries on arrays with no consecutive equal elements is at most $(2 + \log 3)n + o(n) < 3.585n$.*

Proof. To answer PSV/PLV/NSV/NLV queries, it suffices to encode the trees $\text{cMin}(A)$ and $\text{cMax}(A)$. We will show that this can be done using $(2 + \log 3)n + o(n)$ bits. To reduce the size of the encoding, we use the following claim.

Claim 3. *If i is a leaf in $\text{cMin}(A)$ (resp., $\text{cMax}(A)$) and i has right siblings then i is red in $\text{cMin}(A)$ (resp., $\text{cMax}(A)$).*

Proof. Since i is a leaf, the immediate right sibling of i is $i + 1$. Due to the assumption that A has no consecutive equal elements we have that $A[i] \neq A[i + 1]$. Therefore, i is red. ■

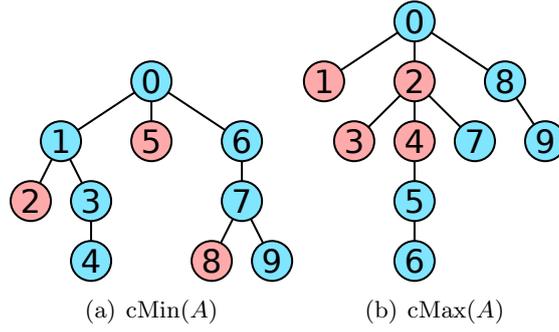
We say that an index $0 < i < n$ is *good* if i does not have right siblings in $\text{cMin}(A)$ and in $\text{cMax}(A)$. We say that $0 < i < n$ is *bad* if i has right siblings in $\text{cMin}(A)$ and in $\text{cMax}(A)$. If $0 < i < n$ is not good or bad we say that i is *neutral*. For an index $0 < i < n$, the *relevant tree of i* is the tree from $\text{cMin}(A), \text{cMax}(A)$ in which i is an internal node (note that by Claim 1 there is exactly one tree in which i is an internal node).

In the following, we encode the color red by 0 and the color blue by 1. The encoding of $\text{cMin}(A), \text{cMax}(A)$ consists of the following strings.

- The strings $T_{\text{Min}(A)}$ and $T_{\text{Max}(A)}$ (these strings were defined in Section 2.3).
- A binary string $U_{\text{good/bad}}$ that is obtained by concatenating the characters $U[i]$ for every i which is either good or bad (the string U was defined in Section 2.3).
- A binary string V_{bad} obtained by concatenating the color of i in the relevant tree of i for every bad i .
- A ternary string V_{neutral} obtained by concatenating a character c_i for every neutral index i . If i does not have right siblings in its relevant tree, $c_i = 2$. Otherwise, c_i is the color of i in the relevant tree of i .

See Figure 4 for an example.

We now show that given the string $T_{\text{Min}(A)}, T_{\text{Max}(A)}, U_{\text{good/bad}}, V_{\text{bad}},$ and V_{neutral} we can reconstruct the trees $\text{cMin}(A)$ and $\text{cMax}(A)$. We initialize two trees T_1, T_2 to contain node 0. At the end of the following algorithm, $T_1 = \text{cMin}(A)$ and $T_2 = \text{cMax}(A)$. By reading the prefixes of $T_{\text{Min}(A)}$ and $T_{\text{Max}(A)}$ until the first zero in each string, we know the degrees of node 0 in $\text{Min}(A)$ and in $\text{Max}(A)$. We now go over $i = 1, \dots, n$ and add node i to the trees T_1 and T_2 . This is done as follows. From the previous iterations of the algorithm, we know the number of children of the nodes $0, \dots, i - 1$ in $\text{Min}(A)$ and in $\text{Max}(A)$. We make node i the rightmost child of node j in T_1 where j is the minimum integer such that the number of children of j in T_1 is less than the number of children of j in $\text{Min}(A)$. We also add node i to the tree T_2 similarly. We then find which tree is the relevant tree of i . After we know the relevant tree of i , we read unread characters from the string $T_{\text{Min}(A)}$ or $T_{\text{Max}(A)}$ that corresponds the relevant tree of i until reaching the first zero. This gives us the number of children of i in its relevant tree. Moreover, by Claim 1, i does not have children in the non-relevant tree. Finally, we find the color of i in $\text{Min}(A)$ and $\text{Max}(A)$.



i	0	1	2	3	4	5	6	7	8	9
$A[i]$		3	8	5	6	3	2	7	10	9
$T_{\text{Min}(A)}$	110	10		0			0	10		
$T_{\text{Max}(A)}$	110		110		0	0			0	
$U_{\text{good/bad}}$		0	1				0	0		
V_{bad}		1	0							
V_{neutral}				2	0	2			2	

Figure 4: An example of the trees $\text{cMin}(A)$ and $\text{cMax}(A)$ of an array A and the encoding of these trees. The good indices are 6,7 and the bad indices are 1,2.

Since we know the number of children of the parent of i in $\text{Min}(A)$, we know whether node i has right siblings in $\text{Min}(A)$ (i has right siblings if and only if the number of children of the parent of i in T_1 is less than the number of children of the parent of i in $\text{Min}(A)$). Similarly, we know whether i has a right sibling in $\text{Max}(A)$. Therefore, we know whether i is good, bad, or neutral.

If i is good, reading the next unread character from $U_{\text{good/bad}}$ gives us the relevant tree of i . Since i has no right siblings in $\text{cMin}(A)$ and in $\text{cMax}(A)$, the color of i is blue in both trees.

If i is bad, reading the next unread character from $U_{\text{good/bad}}$ gives us the relevant tree of i . We also read the next unread character from V_{bad} to know the color of i in the relevant tree. The color of i in the non-relevant tree is red by Claim 3.

Finally, if i is neutral, we read the next unread character from V_{neutral} and let c be this character. Suppose without loss of generality that i does not have right siblings in $\text{cMin}(A)$ and it has right siblings in $\text{cMax}(A)$. The color of i in $\text{cMin}(A)$ is blue (since i does not have right siblings in $\text{cMin}(A)$). If $c = 2$ then the relevant tree of i is $\text{cMin}(A)$ and the color of i in $\text{cMax}(A)$ is red (by Claim 3). Otherwise, the relevant tree of i is $\text{cMax}(A)$ and the color of i in $\text{cMax}(A)$ is red if $c = 0$ and blue if $c = 1$.

We now analyze the size of the encoding. We need the following lemma.

Lemma 4. *The number of good indices is equal to the number of bad indices.*

Proof. For the purpose of the proof we also define the index n to be a good index. Thus, we now need to show that the number of good indices is equal to the number of bad indices plus one. We prove this claim using induction on n . The base of the induction, $n = 1$, is true since in this case there is one good index and no bad indices.

Now suppose that $n > 1$. Suppose without loss of generality that $A[n-1] > A[n]$. Then, node n is the only child of $n-1$ in $\text{cMax}(A)$. Moreover, node n is not a child of $n-1$ in $\text{cMin}(A)$. Let j be the parent of n in $\text{cMin}(A)$. Since the names of the nodes are according to their ranks in the preorder, we have that nodes $j+1, \dots, n-1$ are descendants

of j in $\text{cMin}(A)$, and therefore node $j + 1$ is a child of j in $\text{cMin}(A)$. Therefore, n has left siblings in $\text{cMin}(A)$. Let k be the immediate left sibling of n in $\text{cMin}(A)$.

Let A' be the array obtained by taking the first $n - 1$ elements of A . We have that the trees $\text{Min}(A')$ and $\text{Max}(A')$ are obtained by deleting node n from $\text{Min}(A)$ and $\text{Max}(A)$, respectively.

Since n is the single child of $n - 1$ in $\text{Max}(A)$, we conclude that a node i has right siblings in $\text{Max}(A')$ if and only if i has right siblings in $\text{Max}(A)$. Since n is the immediate right sibling of k in $\text{Min}(A)$, we have that a node $i \neq k$ has right siblings in $\text{Min}(A')$ if and only if i has right siblings in $\text{Min}(A)$. Moreover, k has right siblings in $\text{Min}(A)$ but not in $\text{Min}(A')$. It follows that

- n is a good index with respect to A , but not with respect to A' .
- Either k is a good index with respect to A' and neutral with respect to A (if k does not have right siblings in $\text{Max}(A')$), or k is a neutral index with respect to A' and bad with respect to A .
- For every $i \neq k, n$, i is good (resp., bad) with respect to A' if and only if i is good (resp., bad) with respect to A .

It follows that the difference between the number of good indices and bad indices with respect to A is equal to the difference of these numbers with respect to A' . By the induction hypothesis, the latter difference is 1. ■

Let g be the number of good indices. The combined size of $T_{\text{Min}(A)}$ and $T_{\text{Max}(A)}$ is $2n$ bits. The size of $U_{\text{good/bad}}$ is $2g$ bits and the size of V_{bad} is g bits. The string V_{neutral} has length $n - 1 - 2g$ so it can be encoded using $\lceil (n - 1 - 2g) \log 3 \rceil$ bits. We also need to store the lengths of the strings which requires $O(\log n)$ bits. The total size of the encoding is $2n + 3g + (n - 2g) \log 3 + O(\log n)$ bits. This expression is maximized when $g = 0$, and the theorem follows. ■

4 General arrays

Lemma 5. For a constant $c > 0$ and integers n, k , $c(n - k) + \log \binom{n}{k} \leq \log(2^c + 1) \cdot n$.

Proof. Let $y = k/n$. We have that

$$c(n - k) + \log \binom{n}{k} = c(1 - y)n + \log \binom{n}{yn} \leq c(1 - y)n + n \log \left(\frac{1}{y^y(1 - y)^{1 - y}} \right).$$

Let $f(x) = c(1 - x) + \log \left(\frac{1}{x^x(1 - x)^{1 - x}} \right)$. The derivative of f is $-c + \log(1 - x) - \log x$. Therefore, f is maximized at $x^* = 1/(2^c + 1)$ and $f(x^*) = \log(2^c + 1)$. ■

Theorem 6. The effective entropy of PSV/PLV/NSV/NLV queries is at most $\log 13 \cdot n + o(n) < 3.701n$.

Proof. As in [12], we define a binary string C of length $n - 1$ in which $C[i] = 1$ if and only if $A[i] = A[i + 1]$. We also define an array A' that is obtained from A by deleting the elements $A[i]$ for every i such that $C[i] = 1$. Let k be the number of ones in C . In order to answer PSV/PLV/NSV/NLV queries on A , it suffices to store C and information for answering PSV/PLV/NSV/NLV queries on A' . By Theorem 2, the latter can be done using at most $(2 + \log 3)(n - k)$ bits. Moreover, C can be stored using $\lceil \log n \rceil + \lceil \log \binom{n}{k} \rceil$ bits. The theorem now follows from Lemma 5. ■

References

- [1] G. S. Brodal and P. Davoodi and S. S. Rao. On space efficient two dimensional range minimum data structures. *Algorithmica*, 63(4):815–830, 2012.
- [2] G. S. Brodal, A. Brodnik, and P. Davoodi. The encoding complexity of two dimensional range minimum data structures. In *Proc. 21st European Symposium on Algorithms (ESA)*, pages 229–240, 2013.
- [3] P. Davoodi, G. Navarro, R. Raman, and S. S. Rao. Encoding range minima and range top-2 queries. *Philosophical Transactions of the Royal Society*, 372(2016):20130131, 2014.
- [4] J. Fischer. Combined data structure for previous-and next-smaller-values. *Theoretical Computer Science*, 412(22):2451–2456, 2011.
- [5] J. Fischer and V. Heun. Space-efficient preprocessing schemes for range minimum queries on static arrays. *SIAM J. on Computing*, 4(2):465–492, 2011.
- [6] T. Gagie, G. Manzini, and R. Venturini. An encoding for order-preserving matching. In *Proc. 25th European Symposium on Algorithms (ESA)*, pages 38:1–38:15, 2017.
- [7] P. Gawrychowski and P. K. Nicholson. Optimal encodings for range top- k , selection, and min-max. In *Proc. 42nd International Colloquium on Automata, Languages and Programming (ICALP)*, pages 593–604, 2015.
- [8] M. Golin, J. Iacono, D. Krizanc, R. Raman, S. R. Satti, and S. Shende. Encoding 2d range maximum queries. *Theoretical Computer Science*, 609:316–327, 2016.
- [9] R. Grossi, J. Iacono, G. Navarro, R. Raman, and S. R. Satti. Asymptotically optimal encodings of range data structures for selection and top- k queries. *ACM Transactions on Algorithms*, 13(2):28:1–28:31, 2017.
- [10] V. Jayapaul, S. Jo, R. Raman, V. Raman, and S. R. Satti. Space efficient data structures for nearest larger neighbor. *J. of Discrete Algorithms*, 36:63–75, 2016.
- [11] S. Jo, R. Lingala, and S. R. Satti. Encoding two-dimensional range top- k queries. In *Proc. 27th Symposium on Combinatorial Pattern Matching (CPM)*, volume 54, 2016.
- [12] S. Jo and S. R. Satti. Simultaneous encodings for range and next/previous larger/smaller value queries. *Theoretical Computer Science*, 654:80–91, 2016.
- [13] D. Merlini, R. Sprugnoli, and M. C. Verri. Waiting patterns for a printer. *Discrete Applied Mathematics*, 144(3):359–373, 2004.