# How to answer a small batch of RMQs or LCA queries in practice

Mai Alzamel[1], Panagiotis Charalampopoulos[1], Costas S. Iliopoulos[1], and
Solon P. Pissis[1]

Department of Informatics, King's College London, UK
[mai.alzamel,panagiotis.charalampopoulos,
costas.iliopoulos,solon.pissis]@kcl.ac.uk

**Abstract.** In the Range Minimum Query (RMQ) problem, we are given
an array $A$ of $n$ numbers and we are asked to answer queries of the fol-
lowing type: for indices $i$ and $j$ between 0 and $n-1$, query $\mathrm{RMQ}_A(i,j)$
returns the index of a minimum element in the subarray $A[i \mathinner{.\,.} j]$. Answer-
ing a small batch of RMQs is a core computational task in many real-
world applications, in particular due to the connection with the Lowest
Common Ancestor (LCA) problem. With *small batch*, we mean that the
number $q$ of queries is $o(n)$ and we have them all at hand. It is therefore
not relevant to build an $\Omega(n)$-sized data structure or spend $\Omega(n)$ time
to build a more succinct one. It is well-known, among practitioners and
elsewhere, that these data structures for online querying carry high con-
stants in their pre-processing and querying time. We would thus like to
answer this batch efficiently in practice. With *efficiently in practice*, we
mean that we (ultimately) want to spend $n + \mathcal{O}(q)$ time and $\mathcal{O}(q)$ space.
We write $n$ to stress that the number of operations per entry of $A$ should
be a very small constant. Here we show how existing algorithms can be
easily modified to satisfy these conditions. The presented experimental
results highlight the practicality of this new scheme. The most significant
improvement obtained is for answering a small batch of LCA queries. A
library implementation of the presented algorithms is made available.

## 1   Introduction

In the Range Minimum Query (RMQ) problem, we are given an array $A$ of $n$
numbers and we are asked to answer queries of the following type: for indices $i$
and $j$ between 0 and $n-1$, query $\mathrm{RMQ}_A(i,j)$ returns the index of a minimum
element in the subarray $A[i \mathinner{.\,.} j]$.

The RMQ problem and the linearly equivalent Lowest Common Ancestor
(LCA) problem [4] are very well-studied and several optimal algorithms exist
to solve them. It was first shown by Harel and Tarjan [14] that a tree can be
pre-processed in $\mathcal{O}(n)$ time so that LCA queries can be answered in $\mathcal{O}(1)$ time
per query. A major breakthrough in practicable constant-time LCA-computation
was made by Berkman and Vishkin [6]. Farach and Bender [4] further simplified
this algorithm by showing that the RMQ problem is linearly equivalent to the

LCA problem (shown also in [10]). The constants due to the reduction, however, remained quite large, making these algorithms impractical in most realistic cases. To this end, Fischer and Heun [9] presented yet another optimal, but also direct, algorithm for the RMQ problem. The same authors (but also others [15]) showed that due to large constants in the pre-processing and querying time implementations of this algorithm are often slower than implementations of the naive ones. Continuous efforts for engineering these solutions are being made [8].

In this article we try to address this problem, in particular when one wants to answer a relatively small batch of RMQs efficiently. This version of the problem is a core computational task in many real-world applications such as in *object inheritance* during static compilation of code [5] or in several *string matching* problems (see Section 5 for some). With *small batch*, we mean that the number $q$ of the queries is $o(n)$ *and* we have them all at hand. It is therefore not relevant to build an $\Omega(n)$-sized data structure or spend $\Omega(n)$ time to build a more succinct one. It is well-known, among practitioners and elsewhere, that these data structures carry high constants in both their pre-processing and querying time. (Note that when $q = \Omega(n)$ one can use these data structures for this computation.) We would thus like to answer this batch efficiently in practice. With *efficiently in practice*, we mean that we (ultimately) want to spend $n + \mathcal{O}(q)$ time and $\mathcal{O}(q)$ space. We write $n$ to stress that the number of operations per entry of $A$ should be a very small constant; e.g. scan the array once or twice. In what follows, we show how existing algorithms can be easily modified to satisfy these conditions. Experimental results presented here highlight the practicality of this scheme. The most significant improvement obtained is for answering a small batch of LCA queries. The RMQ Batch problem can be defined as follows.

---

RMQ Batch
**Input:** An array $A$ of size $n$ of numbers and a list $Q$ of $q$ pairs of indices $(i, j)$, $0 \leq i \leq j \leq n - 1$
**Output:** $\mathrm{RMQ}_A(i, j)$ for each $(i, j) \in Q$

---

The LCA Queries Batch problem can be defined as follows.

---

LCA Queries Batch
**Input:** A rooted tree $T$ with $n$ labelled nodes $0, 1, \ldots, n - 1$ and a list $Q$ of $q$ pairs of nodes $(u, v)$
**Output:** $\mathrm{LCA}_T(u, v)$ for each $(u, v) \in Q$

---

**Our computational model.** We assume the word-RAM model with word size $w = \Omega(\log n)$. For the RMQ Batch problem, we assume that we are given a rewritable array $A$ of size $n$, each entry of which may be increased by $n$ and still fit in a computer word. For the LCA Queries Batch problem, we assume that we are given (an $\mathcal{O}(n)$-sized representation of) a rewritable tree $T$ allowing constant-time access to (at least) the nodes of $T$ that are in some query in $Q$ (see the representation in [12], for instance). All presented algorithms are deterministic.

## 2   Contracting the Input Array

Consider any two adjacent array entries $A[i]$ and $A[i+1]$. Observe that if no query in $Q$ starts or ends at $i$ or at $i+1$ then, if $A[i] \neq A[i+1]$, $\max(A[i], A[i+1])$ will never be the answer to any of the queries in $Q$. Hence, the idea is that we want to contract array $A$, so that each block that does not contain the left or right endpoint of any query gets replaced by one element: its minimum. A similar idea, based on sorting the list $Q$, has been considered in the *External Memory* model [1] (see also [2]). In this section, we present a solution for our computational model, which avoids using $\Omega(n)$ space or time, but also avoids using $\Omega(\text{sort}(Q))$ time.

There are some technical details in order to update the queries for $A$ into queries for the new array using only $\mathcal{O}(q)$ time and extra space. We first scan the array $A$ once and find $\mu = \max_i A[i]$. We also create two auxiliary arrays $Z_0[0 \mathbin{.\,.} 2q-1]$ and $Z_1[0 \mathbin{.\,.} 2q-1]$. For each query $(i,j) \in Q$ we mark positions $i$ (and $j$) in the array $A$ as follows. If $A[i] \leq \mu$, then $i$ has not been marked before. Let this be the $k$-th position, $k > 0$, that gets marked (we just store a counter for that). We store $A[i]$ in $Z_0[\mu + k \mod 2q]$ and replace the value that is stored in $A[i]$ by $\mu + k$. We also start a linked list at $Z_1[\mu + k \mod 2q]$, where we insert a pointer to query $(i,j)$, so that we can update it later. If $A[i] > \mu$, then the position has already been marked; we just add a pointer to the respective query in the linked list starting at $Z_1[A[i] \mod 2q]$.

We then scan array $A$ again and create a new array $A_Q$ as follows: for each marked position $j$ (i.e. $A[j] > \mu$), we copy the original value (i.e. $Z_0[A[j] \mod 2q]$) in $A_Q$, while each maximal block in $A$ that does not contain a marked position is replaced by a single entry—its minimum. When we insert the original entry of a marked position $j$ of $A$ (i.e. $Z_0[A[j] \mod 2q]$) in $A_Q$ at position $p$, we go through the linked list that is stored in $Z_1[A[j] \mod 2q]$, where we have stored pointers to all the queries of the form $(i,j)$ or $(j,k)$, and replace $j$ by $p$ in each of them. Thus, after we have scanned $A$, for each query $(i,j) \in Q$ on $A$, we will have stored the respective pair $(i', j')$ on $A_Q$. Note that we need to scan array $A$ only *once* if we know $\mu$ a priori (e.g. in LCP array [7]), or *twice* otherwise.

*Example 1.* Assume we are given array $A$ and $Q = \{(4, 18), (0, 6), (6, 10)\}$.

| $A$ | 17 | 22 | 38 | 4 | 5 | 8 | 2 | 8 | 9 | 21 | 0 | 12 | 8 | 7 | 13 | 3 | 6 | 14 | 1 | 36 | 0 | 4 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|

Then $A_Q$ is as follows.

| $A_Q$ | 17 | 4 | 5 | 8 | 2 | 8 | 0 | 3 | 1 | 0 |
|---|---|---|---|---|---|---|---|---|---|---|

While creating $A_Q$, we also store in an auxiliary array the function $f : \{0, 1, \ldots, |A_Q| - 1\} \to \{0, 1, \ldots, n-1\}$ between positions of $A_Q$ and the respective original positions in $A$.

Now notice that $A_Q$ and the auxiliary arrays are all of size $\mathcal{O}(q)$ since in the worst case we mark $2q$ distinct elements of $A$ and contract $2q + 1$ blocks that do not contain a marked position. (We can actually throw away everything before the first marked position and everything after the last marked position and get $4q - 1$ instead.) The whole procedure takes $n + \mathcal{O}(q)$ time and $\mathcal{O}(q)$ space. Note that if $\mathrm{RMQ}_{A_Q}(i', j') = \ell$ then $\mathrm{RMQ}_A(i, j) = f(\ell)$.

We can finally retrieve the original input array if required by replacing $A[f(j)]$ by $A_Q[j]$ for every $j$ in the domain of $f$ in $\mathcal{O}(q)$ time.

## 3 Small RMQ Batch

### 3.1 An $n + \mathcal{O}(q \log q)$-time and $\mathcal{O}(q)$-space Algorithm

The algorithm presented in this section is a modification of the *Sparse Table* algorithm by Bender and Farach-Colton [4] applied on array $A_Q$; we denote it by ST-RMQ. The modification is based on the fact that **(i)** we do not want to consume $\Omega(q \log q)$ extra space to answer the $q$ queries; and **(ii)** we do not want to necessarily do all the pre-processing work of the algorithm in [4], which is designed to answer any of the $\Theta(q^2)$ possible queries online. We denote this modified algorithm by ST-RMQ$_{\mathsf{CON}}$ and formalise it below.

ST-RMQ$_{\mathrm{CON}}(A, Q)$

```
 1   A_Q ← CONTRACT(A, Q)
 2   Store function f; store (i', j') for every (i, j) ∈ Q
 3   for each (i, j) ∈ Q do
 4        if i = j then
 5             REPORT((i, i), i)
 6        else Add (i, j) in bucket B_⌊log(j'−i')⌋
 7   t ← max{r|B_r ≠ ∅} + 1
 8   for m = 0 to |A_Q| − 1 do
 9        D[m] ← (A_Q[m], m)
10   for k = 0 to t − 1 do
11        for each (i, j) ∈ B_k do
12             (a, p) ← min(D[i'], D[j' − 2^k + 1])
13             REPORT((i, j), f(p))
14        for m = 0 to |A_Q| − 1 do
15             if m + 2^k ≤ |A_Q| − 1 then
16                  D[m] ← min(D[m], D[m + 2^k])
```

The idea is to first put each $(i, j) \in Q$ with $i \neq j$ in a bucket $B_k$ based on the $k$ for which $2^k \leq j' - i' < 2^{k+1}$—we can have at most $\lceil \log(|A_Q| - 1) \rceil$ such buckets. In this process, if we find queries of the form $(i, i) \in Q$, we answer them on the spot. We can do this in $\mathcal{O}(q)$ time.

We then create an array $D$ of size $|A_Q|$ where we will store 2-tuples $(a, p)$. In Step $k$, $D[m]$ will store the minimum value across $A_Q[m .. m + 2^k - 1]$, as well as the position $p$, $m \leq p < m + 2^k$ where it occurs. We initialise it as

$D[m] = (A_Q[m], m)$ and we will then update it by utilising the *doubling technique*. At Step 0 we answer all (trivial) queries that are stored in $B_0$; they are of the form $(i, i+1)$ and the answer can be found by looking at $\min(D[i'], D[i'+1])$—note that we compare elements of $D$ lexicographically. When we are done with $B_0$ we have to update $D$ by setting $D[m] = \min(D[m], D[m+2^0])$ for all $m < |A_Q| - 1$.

Generally, in Step $k$, we answer the queries of $B_k$ as follows. For query $(i, j)$, we find the answer by obtaining $\min(D[i'], D[j'-2^k+1]) = (a, p)$. We then return $f(p)$. The point is that $\{i', \ldots, i'+2^k-1\} \cup \{j'-2^k+1, \ldots, j'\} = \{i', \ldots, j'\}$. When we are done with $B_k$ we set $D[m] = \min(D[m], D[m+2^k])$ if $m+2^k \leq |A_Q| - 1$.

We do this until we have gone through all $t$ non-empty buckets (i.e. $t = \max\{r | B_r \neq \emptyset\} + 1$). Updating $D$ takes $\mathcal{O}(q)$ time in each step, and we need in total $\mathcal{O}(q)$ time for the queries. We thus need $\mathcal{O}(qt)$ time for this part of the algorithm. Since $t = \max\{\lfloor \log(j'-i') \rfloor | (f(i'), f(j')) \in Q\} = \mathcal{O}(\log q)$, this time is $\mathcal{O}(q \log q)$. The overall time complexity of the algorithm is thus $n + \mathcal{O}(q \log q)$. Notably, the space required is only $\mathcal{O}(q)$ as we overwrite $D$ in each step.

### 3.2 $n + \mathcal{O}(q)$-time and $\mathcal{O}(q)$-space Algorithms

*Offline-based algorithm.* Given an array $A$ of $n$ numbers its *Cartesian tree* is defined as follows. The root of the Cartesian tree is $A[i] = \min\{A[0], \ldots, A[n-1]\}$, its left subtree is computed recursively on $A[0], \ldots, A[i-1]$ and its right subtree on $A[i+1], \ldots, A[n-1]$. An LCA instance can be obtained from an RMQ instance on an array $A$ by letting $T$ be the Cartesian tree of $A$ that can be constructed in $\mathcal{O}(n)$ time [10]. It is easy to see that $\mathrm{RMQ}_A(i, j)$ in $A$ translates to $\mathrm{LCA}_T(A[i], A[j])$ in $T$. The first step of this algorithm is to create array $A_Q$ in $n + \mathcal{O}(q)$ time similarly to algorithm ST-RMQ$_{\mathsf{CON}}$. The second step is to construct the Cartesian tree $T_Q$ of $A_Q$ in $\mathcal{O}(q)$ time and extra space. Finally, we apply the offline algorithm by Gabow and Tarjan [11] to answer $q$ $\mathrm{LCA}_{T_Q}$ queries in $\mathcal{O}(q)$ time and extra space. This takes overall $n + \mathcal{O}(q)$ time and $\mathcal{O}(q)$ extra space. We denote this algorithm by OFF-RMQ$_{\mathsf{CON}}$. We denote by OFF-RMQ the same algorithm applied on array $A$.

*Online-based algorithm.* The first step of this algorithm is to create array $A_Q$ in $n + \mathcal{O}(q)$ time similarly to algorithm ST-RMQ$_{\mathsf{CON}}$. We can then apply the algorithm by Fischer and Heun [9] on array $A_Q$ to obtain overall an $n + \mathcal{O}(q)$-time and $\mathcal{O}(q)$-space algorithm. We denote this algorithm by ON-RMQ$_{\mathsf{CON}}$. We denote by ON-RMQ the same algorithm applied on array $A$.

Note that in the case when $q = \Omega(n)$, i.e. the batch is not so small, we can choose to apply algorithm OFF-RMQ or algorithm ON-RMQ on array $A$ directly thus obtaining an algorithm that always works in $n + \mathcal{O}(q)$ time and $\mathcal{O}(\min\{n, q\})$ extra space. We therefore obtain the following result asymptotically.

**Theorem 1.** *The RMQ Batch problem can be solved in $n + \mathcal{O}(q)$ time and $\mathcal{O}(\min\{n, q\})$ extra space.*

## 4 Small LCA Queries Batch

In the LCA problem, we are given a rooted tree $T$ having $n$ labelled nodes and we are asked to answer queries of the following type: for nodes $u$ and $v$, query $\text{LCA}_T(u, v)$ returns the node furthest from the root that is an ancestor of both $u$ and $v$. There exists a time-optimal algorithm by Gabow and Tarjan [11] to answer a batch $Q$ of $q$ LCA queries in $\mathcal{O}(n + q)$ time and $\mathcal{O}(n)$ extra space. We denote this algorithm by OFF-LCA. In this section, we present a simple but *non-trivial* algorithm for improving this, for $q = o(n)$, to $n + \mathcal{O}(q)$ time and $\mathcal{O}(q)$ extra space.

It is well-known (see [4] for the details) that an RMQ instance $A$ can be obtained from an LCA instance on a tree $T$ by writing down the depths of the nodes visited during an *Euler tour* of $T$. That is, $A$ is obtained by listing all node-visitations in a depth-first search (DFS) traversal of $T$ starting from the root. The LCA of two nodes translates to an RMQ (where we compare nodes based on their level) between the first occurrences of these nodes in $A$.

We proceed largely as in Section 2. For each query $(u, v) \in Q$, we mark nodes $u$ (and $v$) in $T$ as follows. If $u < n$ then $u$ has not been marked before. Let this be the $k$-th node, $k > 0$, that gets marked (we just store a counter for that). We also create two arrays $Z_0[0 \mathinner{.\,.} 2q - 1]$ and $Z_1[0 \mathinner{.\,.} 2q - 1]$. We store $u$ in $Z_0[n - 1 + k \mod 2q]$ and replace $u$ by $n - 1 + k$. We also start a linked list at $Z_1[n - 1 + k \mod 2q]$, where we insert a pointer to query $(u, v)$, so that we can update it later. If $u > n - 1$, the node has already been marked, and we just add a pointer to the respective query in the linked list starting at $Z_1[u \mod 2q]$.

We then do a single DFS traversal on $T$ and create two new arrays $E_Q$ and $L_Q$ as follows. When a marked node $v$ (i.e. $v > n - 1$) is visited for the *first time*, we write down in $E_Q$ its original value (i.e. $Z_0[v \mod 2q]$), while for each maximal sequence of visited nodes that are not marked we write down a single entry—the one with the *minimum tree level*. At the same time, we store in $L_Q[v]$ the level of the node added in $E_Q[v]$. While creating $E_Q$, we also store in an auxiliary array the function $f : \{0, 1, \ldots, |E_Q| - 1\} \to \{0, 1, \ldots, n - 1\}$ between positions of $E_Q$ and the respective node labels in $T$.

When we insert the original entry of a marked node $u$ of $T$ (i.e. $Z_0[u \mod 2q]$) in $E_Q$ at position $p$, we go through the linked list that is stored in $Z_1[u \mod 2q]$, where we have stored pointers to all the queries of the form $(u, v)$ or $(w, u)$, and replace $u$ by $p$ in each of these queries. Thus, after we have finished the traversal on $T$, for each LCA query $(u, v) \in Q$ on $T$, we will have stored the respective RMQ pair $(u', v')$ on $L_Q$; where $u'$ (resp. $v'$) corresponds to the *first occurrence* of node $u$ (resp. $v$) in the traversal. Thus we traverse $T$ only *once*.

Now notice that $E_Q$ and the auxiliary arrays are all of size $\mathcal{O}(q)$ since in the worst case we mark $2q$ distinct nodes of $T$ and contract $2q + 1$ sequences of visited nodes that do not contain a marked node. (We can actually throw away everything before the first marked node and everything after the last marked node and get $4q - 1$ instead.) The whole procedure takes $n + \mathcal{O}(q)$ time and $\mathcal{O}(q)$ space. We are now in a position to apply algorithm ON-RMQ on $L_Q$ to obtain the final bound. To answer the queries, note that if $\text{RMQ}_{L_Q}(u', v') = \ell$ then

$\text{LCA}_T(u,v) = E_Q[\ell]$. We denote this algorithm by ON-LCA$_{\text{CON}}$. Alternatively, we can apply algorithm ST-RMQ on $L_Q$ to solve this problem in $n + \mathcal{O}(q \log q)$ and $\mathcal{O}(q)$ extra space; we denote this algorithm by ST-LCA$_{\text{CON}}$.

We can finally retrieve the original input tree if required by replacing node $f(v)$ by $E_Q[v]$ for every $v$ in the domain of $f$ in $\mathcal{O}(q)$ time.

Note that in the case when $q = \Omega(n)$, i.e. the batch is not so small, we can choose to apply algorithm OFF-LCA on tree $T$ directly, thus obtaining an algorithm that always works in $n + \mathcal{O}(q)$ time and $\mathcal{O}(\min\{n,q\})$ extra space. We therefore obtain the following result asymptotically.

**Theorem 2.** *The LCA Queries Batch problem can be solved in $n + \mathcal{O}(q)$ time and $\mathcal{O}(\min\{n,q\})$ extra space.*

## 5 Applications

We consider the well-known application of answering $q$ LCA queries on the suffix tree of a string. The *suffix tree* $T(S)$ of a non-empty string $S$ of length $n$ is a compact trie representing all suffixes of $S$ (see [7], for details). The nodes of the trie which become nodes of the suffix tree are called *explicit* nodes, while the other nodes are called *implicit*. Each edge of the suffix tree can be viewed as an upward maximal path of implicit nodes starting with an explicit node. Moreover, each node belongs to a unique path of that kind. Then, each node of the trie can be represented in the suffix tree by the edge it belongs to and an index within the corresponding path. The *path-label* of a node $v$ is the concatenation of the edge labels along the path from the root to $v$. The nodes whose path-label corresponds to a suffix of $S$ are called *terminal*. Given two terminal nodes $u$ and $v$ in $T(S)$, representing suffixes $S[i \mathinner{.\,.} n-1]$ and $S[j \mathinner{.\,.} n-1]$, the *string depth* of node $\text{LCA}_{T(S)}(u,v)$ corresponds to the *length* of their longest common prefix, also known as their longest common extension (LCE) [15].

In many textbook solutions for classical string matching problems (e.g. maximal palindromic factors, approximate string matching with $k$-mismatches, approximate string matching with $k$-differences, online string search with the suffix array, etc.) we have that $q = \Omega(n)$ and/or the queries have to be answered *online*. In other algorithms, however, $q$ can be *much smaller* on average (in practice) and the queries can be answered *offline*. We describe here a few such solutions. The common idea, as in many fast average-case algorithms, is to minimise the number of queries by *filtering out* queries that can never lead to a valid solution.

*Text indexing.* Suppose we are given the suffix tree $T(S)$ of a text $S$ of length $n$ and we are asked to create the suffix links for the internal nodes. This may be necessary if the construction algorithm does not compute suffix links (e.g. construction via suffix array) but they are needed for an application of interest. The *suffix link* of a node $v$ with path-label $\alpha y$ is a pointer to the node path-labelled $y$, where $\alpha \in \Sigma$ is a single letter and $y$ is a string. The suffix link of $v$ exists if $v$ is a non-root internal node of $T$. The suffix links can be computed as follows.

The first step is to mark each internal node $v$ of the suffix tree with a pair of leaves $(i, j)$ such that leaves labelled $i$ and $j$ are in subtrees rooted at different children of $v$. This can be done by a DFS traversal of the tree. (Note that if an internal node $v$ has only one child then it must be terminal; assume that it represents the suffix $S[t \mathinner{.\,.} n-1]$. We thus create a suffix link to the node representing $S[t+1 \mathinner{.\,.} n-1]$.) Given an internal node $v$ marked with $(i, j)$, note that $v = \mathrm{LCA}_{T(S)}(i, j)$, and let $\alpha y$ be its path-label. To create the suffix link from $v$, node $u$ with path-label $y$ can be obtained by the query $\mathrm{LCA}_{T(S)}(i+1, j+1)$. We can create a batch of LCA queries consisting of all such pairs. Note that in randomly generated texts, the number of internal nodes of $T(S)$ is $\mathcal{O}(n/h)$ on average, where $h$ is the alphabet's entropy [20]; thus the standard $\Theta(n)$-time and $\Theta(n)$-space solution to this problem, building the LCA data structure over $T(S)$ [4], is not satisfactory.

*Finding frequent gapped factors in texts.* We are given a text $S$ of length $n$, and positive integers $\ell_1$, $\ell_2$, $d$, and $k > 1$. The problem is to find all couples $(u, v)$, such that string $uwv$, for *any* string $w$ (known as *gap* or *spacer*), $|w| = d$, occurs in $S$ at least $k$ times, $|u| = \ell_1$, $|v| = \ell_2$ [16, 19]. The first step is to build $T(S)$. We then locate all subtrees rooted at an explicit node with string depth at least $\ell_1$ and whose parent has string depth less than $\ell_1$, corresponding to factors $u$ repeated in $S$. From these subtrees, we only consider the ones with at least $k$ terminal nodes. Note that if $k$ is large enough, we may have only a few such subtrees. For each subtree with $k' \geq k$ terminal nodes, representing suffixes $S[i_1 \mathinner{.\,.} n-1], S[i_2 \mathinner{.\,.} n-1], \ldots, S[i_{k'} \mathinner{.\,.} n-1]$, we create a batch of LCA queries between all pairs $(i_j + \ell_1 + d, i_{j'} + \ell_1 + d)$ and report occurrences when LCA queries extend pairwise matches to length at least $\ell_2$ for a set of at least $k$ such suffixes. (This algorithm can be easily generalised for any number of gaps.)

*Pattern matching on weighted sequences.* A *weighted sequence* specifies the probability of occurrence of each letter of the alphabet for every position. A weighted sequence thus represents many different strings, each with the probability of occurrence equal to the product of probabilities of its letters at subsequent positions of the weighted sequence. The problem is to find all occurrences of a (standard) pattern $P$ of length $m$ with probability at least $1/z$ in a weighted sequence $S$ of length $n$ [17]. The first step is to construct the heavy string of $S$, denoted by $H(S)$, by assigning to $H(S)[i]$ the most probable letter of $S[i]$ (resolving ties arbitrarily). The second step is to build $T(P\$H(S))$, $\$ \notin \Sigma$. We can then compute the first mismatch between $P$ and every substring of $H(S)$. Note that the number of positions in $S$ where two or more letters occur with probability at least $1/z$ can be small, and so we consider only these positions to cause a legitimate mismatch between $P$ and a factor of $H(S)$. We then use $\mathcal{O}(\log z)$ batches of LCA queries per such starting position to extend a match to length at least $m$. This is because $P$ cannot match a weighted sequence $S$ with probability $1/z$ if more than $\lfloor \log z \rfloor$ mismatches occur between $P$ and $H(S)$ [17].

*Pattern matching with don't care letters.* We are given a pattern $P$ of length $m$, with $m - k$ letters from alphabet $\Sigma$ and $k$ occurrences of a don't care letter (matching itself and any letter from $\Sigma$), and a text $S$ of length $n$. The problem is to find all occurrences of $P$ in $S$ [18]. The first step is to build $T(P'\$S)$, $\$ \notin \Sigma$, where $P'$ is the string obtained from $P$ by replacing don't care letters with a letter $\# \notin \Sigma$. We then locate the subtree rooted at the highest explicit node corresponding to the longest factor $f$ of $P'$ without $\#$'s. We also locate, in the same subtree, all $V$ terminal nodes corresponding to starting positions of $f$ in $S$. Note that if $f$ is long enough, we may have only a few such nodes. Since we know where the don't care letters occur in $P$, we can create a batch of $kV$ LCA queries. An occurrence is then reported when LCA queries extend a match to length at least $m$. (This algorithm can be easily generalised for any number of patterns.)

*Circular string matching.* We are given a pattern $P$ of length $m$ and a text $S$ of length $n$. The problem is to find all occurrences of $P$ or any of its cyclic shifts in $S$ [3]. The first step is to build $T(PP\$P^RP^R\#S\%S^R)$, where $\$, \#, \% \notin \Sigma$, and $X^R$ denotes the reverse image of string $X$. We then conceptually split $P$ in two fragments of lengths $\lceil m/2 \rceil$ and $\lfloor m/2 \rfloor$. Any cyclic shift of $P$ contains as a factor at least one of the two fragments. We thus locate the two subtrees rooted at the highest explicit nodes corresponding to the fragments. We also locate in the same subtrees all $V$ terminal nodes corresponding to starting positions of the fragments in $S$. Note that if $m$ is long enough, we may have only a few such nodes. We create a batch of at most $2V$ LCA queries in order to extend to the left and to the right and report occurrences when LCA queries extend a match to length at least $m$. (This algorithm can be easily generalised for any number of patterns.)

## 6 Experimental Results

We have implemented algorithms ST-RMQ$_{CON}$, OFF-RMQ$_{CON}$, and ON-RMQ$_{CON}$ in the C++ programming language. We have also implemented the same algorithms applied on the original array $A$, denoted by ST-RMQ, OFF-RMQ, and ON-RMQ, respectively; as well as the brute-force algorithm for answering RMQs in the two corresponding flavours, denoted by BF-RMQ$_{CON}$ and BF-RMQ. For the implementation of ON-RMQ$_{CON}$ and ON-RMQ we used the sdsl-lite library [13]. If an algorithm requires $f(n, q)$ time and $g(n, q)$ extra space, we say that the algorithm has complexity $< f(n, q), g(n, q) >$. Table 1 summarises the implemented algorithms. The following experiments were conducted on a Desktop PC using one core of Intel Core i5-4690 CPU at 3.50GHz and 16GB of RAM. All programs were compiled with g++ version 5.4.0 at optimisation level 3 (-O3).

*Experiment I.* We generated random (uniform distribution) input arrays of $n = 1,000,000$ and $n = 100,000,000$ entries (integers), and random (uniform distribution) lists of queries of sizes varying from $\sqrt{n}$ to $128\sqrt{n}$, doubling each time. We compared the runtime of the implementations of the algorithms in

| Non-Contracted | | Contracted | |
|---|---|---|---|
| ST-RMQ | $< \mathcal{O}(n \log n + q), \mathcal{O}(n \log n) >$ | ST-RMQ$_{\mathsf{CON}}$ | $< n + \mathcal{O}(q \log q), \mathcal{O}(q) >$ |
| ON-RMQ | $< \mathcal{O}(n + q), \mathcal{O}(n) >$ | ON-RMQ$_{\mathsf{CON}}$ | $< n + \mathcal{O}(q), \mathcal{O}(q) >$ |
| OFF-RMQ | $< \mathcal{O}(n + q), \mathcal{O}(n) >$ | OFF-RMQ$_{\mathsf{CON}}$ | $< n + \mathcal{O}(q), \mathcal{O}(q) >$ |
| BF-RMQ | $< \mathcal{O}(qn), \mathcal{O}(1) >$ | BF-RMQ$_{\mathsf{CON}}$ | $< n + \mathcal{O}(q^2), \mathcal{O}(q) >$ |

Table 1: Time and space complexities of algorithms for answering RMQs offline.



(a) $n = 100,000,000$

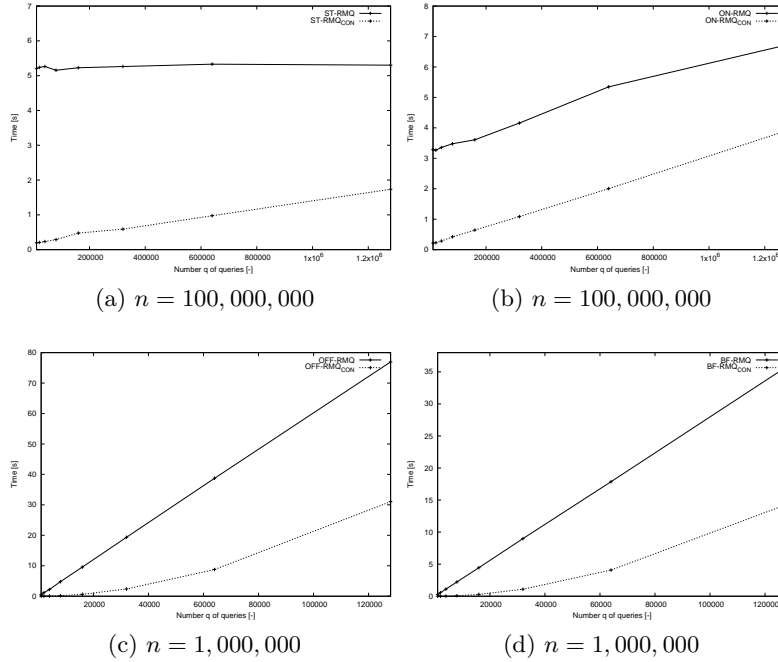(b) $n = 100,000,000$

(c) $n = 1,000,000$

(d) $n = 1,000,000$

Fig. 1: Impact of the proposed scheme on the RMQ algorithms of Table 1.

Table 1 on these inputs; in particular, for each algorithm, we compared the standard implementation against the one with the contracted array. We used the large array, $n = 100,000,000$, for ST-RMQ and ON-RMQ because they are significantly faster and the small one, $n = 1,000,000$, for OFF-RMQ and BF-RMQ. The results plotted in Figure 1 show that the proposed scheme of contracting the input array improves the performance for all implementations substantially.

*Experiment II.* We generated random input arrays of $n = 1,000,000,000$ entries, and random lists of queries of sizes varying from $\sqrt{n}$ to $128\sqrt{n}$, doubling each time. We then compared the runtime of ON-RMQ$_{\mathsf{CON}}$ and ST-RMQ$_{\mathsf{CON}}$ on these inputs. The results are plotted in Figure 2. We observe that ST-RMQ$_{\mathsf{CON}}$ becomes two times faster than ON-RMQ$_{\mathsf{CON}}$ as $q$ grows. Notably, it was not

Fig. 2: Elapsed-time comparison of ON-RMQ$_{\mathsf{CON}}$ and ST-RMQ$_{\mathsf{CON}}$ algorithms for $n = 1,000,000,000$.

possible to run this experiment with ON-RMQ, which implements a *succinct* data structure for answering RMQs, due to insufficient amount of main memory.

*Experiment III.* In addition, we have implemented algorithms ST-LCA$_{\mathsf{CON}}$ and OFF-LCA for answering LCA queries. We first generated a random input array of $n = 1,000,000$ entries and used this array to compute its Cartesian tree. Next we generated random lists of LCA queries of sizes varying from $\sqrt{n}$ to $128\sqrt{n}$, doubling each time. We then compared the runtime of OFF-LCA and ST-LCA$_{\mathsf{CON}}$ on these inputs. The results plotted in Figure 3 show that the implementation of ST-LCA$_{\mathsf{CON}}$ is more than two orders of magnitude faster than the implementation of OFF-LCA, highlighting the impact of the proposed scheme on LCA queries.

## 7  Final Remarks

In this article, we presented a new family of algorithms for answering a small batch of RMQs or LCA queries in practice. The main purpose was to show that if the number $q$ of queries is small with respect to $n$ and we have them all at hand existing algorithms for RMQs and LCA queries can be easily modified to perform in $n + \mathcal{O}(q)$ time and $\mathcal{O}(q)$ extra space. The presented experimental results indeed show that with this new scheme significant practical improvements can be obtained; in particular, for answering a small batch of LCA queries.

Specifically, algorithms ST-RMQ$_{\mathsf{CON}}$ and ST-LCA$_{\mathsf{CON}}$, our modifications to the *Sparse Table* algorithm whose main catch is $\Theta(n \log n)$ space [4], seem to be the best way to answer in practice a small batch of RMQs and LCA queries, respectively. A library implementation of ST-RMQ$_{\mathsf{CON}}$ is available at `https://github.com/solonas13/rmqo` under the GNU General Public License.

Fig. 3: Elapsed-time comparison of OFF-LCA and ST-LCA$_{CON}$ algorithms for $n = 1,000,000$.

# References

1. P. Afshani and N. Sitchinava. I/O-efficient range minima queries. In *SWAT 2014*, volume 8503 of *LNCS*, pages 1–12. Springer, 2014.

2. L. Arge, J. Fischer, P. Sanders, and N. Sitchinava. On (dynamic) range minimum queries in external memory. In *WADS 2013*, volume 8037 of *LNCS*, pages 37–48. Springer, 2013.

3. T. Athar, C. Barton, W. Bland, J. Gao, C. S. Iliopoulos, C. Liu, and S. P. Pissis. Fast circular dictionary-matching algorithm. *Mathematical Structures in Computer Science*, 27(2):143–156, 2017.

4. M. A. Bender and M. Farach-Colton. The LCA problem revisited. In *LATIN 2000*, volume 1776 of *LNCS*, pages 88–94. Springer-Verlag, 2000.

5. M. A. Bender, M. Farach-Colton, G. Pemmasani, S. Skiena, and P. Sumazin. Lowest common ancestors in trees and directed acyclic graphs. *Journal of Algorithms*, 57(2):75–94, 2005.

6. O. Berkman and U. Vishkin. Recursive star-tree parallel data structure. *SIAM J. Comput.*, 22(2):221–242, 1993.

7. M. Crochemore, C. Hancart, and T. Lecroq. *Algorithms on strings*. Cambridge University Press, 2007.

8. H. Ferrada and G. Navarro. Improved range minimum queries. *J. Discrete Algorithms*, 43:72–80, 2016.

9. J. Fischer and V. Heun. Theoretical and practical improvements on the rmq-problem, with applications to lca and lce. In *CPM 2006*, volume 4009 of *LNCS*, pages 36–48. Springer Berlin Heidelberg, 2006.

10. H. N. Gabow, J. L. Bentley, and R. E. Tarjan. Scaling and related techniques for geometry problems. In *STOC 1984*, pages 135–143. ACM, 1984.

11. H. N. Gabow and R. E. Tarjan. A linear-time algorithm for a special case of disjoint set union. *Journal of Computer and System Sciences*, 30(2):209–221, 1985.

12. R. F. Geary, N. Rahman, R. Raman, and V. Raman. A simple optimal representation for balanced parentheses. *Theor. Comput. Sci.*, 368(3):231–246, 2006.

13. S. Gog, T. Beller, A. Moffat, and M. Petri. From theory to practice: Plug and play with succinct data structures. In *SEA*, volume 8504 of *LNCS*, pages 326–337, 2014.

14. D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984.

15. L. Ilie, G. Navarro, and L. Tinta. The longest common extension problem revisited and applications to approximate string searching. *J. Discrete Algorithms*, 8(4):418–428, 2010.

16. C. Iliopoulos, J. Mchugh, P. Peterlongo, N. Pisanti, W. Rytter, and M.-F. Sagot. A first approach to finding common motifs with gaps. *International Journal of Foundations of Computer Science*, 16(6):1145–1155, 2005.

17. T. Kociumaka, S. P. Pissis, and J. Radoszewski. Pattern Matching and Consensus Problems on Weighted Sequences and Profiles. In *ISAAC 2016*, volume 64 of *LIPIcs*, pages 46:1–46:12. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016.

18. R. Y. Pinter. Efficient string matching with don't-care patterns. In *Combinatorial Algorithms on Words*, volume F12 of *NATO ASI Series*, pages 11–29. Springer Berlin Heidelberg, 1985.

19. S. P. Pissis. MoTeX-II: structured motif extraction from large-scale datasets. *BMC Bioinformatics*, 15:235, 2014.

20. M. Régnier and P. Jacquet. New results on the size of tries. *IEEE Trans. Information Theory*, 35(1):203–205, 1989.