

MIT Open Access Articles

Near-Optimal Quantum Algorithms for String Problems

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Akmal, Shyan and Jin, Ce. 2023. "Near-Optimal Quantum Algorithms for String Problems."

As Published: <https://doi.org/10.1007/s00453-022-01092-x>

Publisher: Springer US

Persistent URL: <https://hdl.handle.net/1721.1/147771>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution





Near-Optimal Quantum Algorithms for String Problems

Shyan Akmal¹ · Ce Jin¹

Received: 26 April 2022 / Accepted: 29 December 2022
© The Author(s) 2023

Abstract

We study quantum algorithms for several fundamental string problems, including *Longest Common Substring*, *Lexicographically Minimal String Rotation*, and *Longest Square Substring*. These problems have been widely studied in the stringology literature since the 1970s, and are known to be solvable by near-linear time classical algorithms. In this work, we give quantum algorithms for these problems with *near-optimal* query complexities and time complexities. Specifically, we show that: *Longest Common Substring* can be solved by a quantum algorithm in $\tilde{O}(n^{2/3})$ time, improving upon the recent $\tilde{O}(n^{5/6})$ -time algorithm by Le Gall and Seddighin (in: Proceedings of the 13th innovations in theoretical computer science conference (ITCS 2022), pp 97:1–97:23, 2022. <https://doi.org/10.4230/LIPIcs.ITCS.2022.97>). Our algorithm uses the MNRS quantum walk framework, together with a careful combination of string synchronizing sets (Kempa and Kociumaka, in: Proceedings of the 51st annual ACM SIGACT symposium on theory of computing (STOC 2019), ACM, pp 756–767, 2019. <https://doi.org/10.1145/3313276.3316368>) and generalized difference covers. *Lexicographically Minimal String Rotation* can be solved by a quantum algorithm in $n^{1/2+o(1)}$ time, improving upon the recent $\tilde{O}(n^{3/4})$ -time algorithm by Wang and Ying (in: Quantum algorithm for lexicographically minimal string rotation. CoRR, 2020. [arXiv:2012.09376](https://arxiv.org/abs/2012.09376)). We design our algorithm by first giving a new classical divide-and-conquer algorithm in near-linear time based on exclusion rules, and then speeding it up quadratically using nested Grover search and quantum minimum finding. *Longest Square Substring* can be solved by a quantum algorithm in $\tilde{O}(\sqrt{n})$ time. Our algorithm is an adaptation of the algorithm by Le Gall and Seddighin (2022) for the Longest Palindromic Substring problem, but uses additional techniques to overcome the difficulty that binary search no longer applies. Our techniques naturally extend to

Shyan Akmal supported by NSF CCF-1909429 and a Siebel Scholarship.
Ce Jin supported by an Akamai Presidential Fellowship and NSF CCF-2129139.

✉ Ce Jin
cejin@mit.edu
Shyan Akmal
naysh@mit.edu

¹ MIT, Cambridge, USA

other related string problems, such as Longest Repeated Substring, Longest Lyndon Substring, and Minimal Suffix.

Keywords String processing · Quantum walks · Longest common substring · String synchronizing sets

1 Introduction

The study of string processing algorithms is an important area of research in theoretical computer science, with applications in numerous fields including bioinformatics, data mining, plagiarism detection, etc. Many fundamental problems in this area have been known to have linear-time algorithms since over 40 years ago. Examples include *Exact String Matching* [1, 2], *Longest Common Substring* [3–5], and *(Lexicographically) Minimal String Rotation* [6–8]. These problems have also been studied extensively in the context of data structures, parallel algorithms, and low-space algorithms.

More recently, there has been growing interest in developing efficient *quantum algorithms* for these basic string problems. Given quantum query access to the input strings (defined in Sect. 2.3), it is sometimes possible to solve such problems in *sub-linear* query complexity and time complexity. The earliest such result was given by Ramesh and Vinay [9], who combined Vishkin’s deterministic sampling technique [10] with Grover search [11] to obtain a quantum algorithm for the Exact String Matching problem with near-optimal $\tilde{O}(\sqrt{n})$ time complexity.¹ More recently, Le Gall and Seddighin [12] obtained sublinear-time quantum algorithms for various string problems, among them an $\tilde{O}(n^{5/6})$ -time algorithm for Longest Common Substring (LCS) and an $\tilde{O}(\sqrt{n})$ -time algorithm for Longest Palindromic Substring (LPS). In developing these algorithms, they applied the quantum Exact String Matching algorithm [9] and Ambainis’ Element Distinctness algorithm [13] as subroutines, and used periodicity arguments to reduce the number of candidate solutions to be checked. Another recent work by Wang and Ying [14] showed that Minimal String Rotation can be solved in $\tilde{O}(n^{3/4})$ quantum time. Their algorithm was also based on quantum search primitives (including Grover search and quantum minimum finding [15]) and techniques borrowed from parallel string algorithms [10, 16, 17].

On the lower bound side, it has been shown that Longest Common Substring requires $\tilde{\Omega}(n^{2/3})$ quantum query complexity (by a reduction [12] from the Element Distinctness problem [18–20]), and that Exact String Matching, Minimal String Rotation, and Longest Palindromic Substring all require $\Omega(\sqrt{n})$ quantum query complexity (by reductions [12, 14] from the unstructured search problem [21]). Le Gall and Seddighin [12] observed that although the classical algorithms for LCS and LPS are almost the same (both based on suffix trees [3]), the latter problem (with time complexity $\tilde{\Theta}(\sqrt{n})$) is strictly easier than the former (with an $\tilde{\Omega}(n^{2/3})$ lower bound) in the quantum query model.

¹ Throughout this paper, $\tilde{O}(\cdot)$ hides poly $\log n$ factors where n denotes the input length, and $\tilde{\Omega}(\cdot)$, $\tilde{\Theta}(\cdot)$ are defined analogously. In particular, $\tilde{O}(1)$ means $O(\text{poly } \log n)$.

Problem	Time UB	Reference	Query LB
Longest Common Substring	$\tilde{O}(n^{5/6})$	[12]	$\tilde{\Omega}(n^{2/3})$
Longest Repeated Substring	$\tilde{O}(n^{2/3})$	This work (Theorem 3.1)	$\tilde{\Omega}(n^{2/3})$
	$\tilde{O}(n^{2/3})$	This work (Remark 3.8)	$\tilde{\Omega}(n^{2/3})$
Minimal String Rotation	$\tilde{O}(n^{3/4})$	[14]	$\Omega(\sqrt{n})$
Minimal Suffix	$n^{1/2+o(1)}$	This work (Theorem 4.6)	$\Omega(\sqrt{n})$
Maximal Suffix	$n^{1/2+o(1)}$	This work (Theorem 4.6)	$\Omega(\sqrt{n})$
Longest Lyndon Substring	$n^{1/2+o(1)}$	This work (Theorem 4.9)	$\Omega(\sqrt{n})$
Longest Palindromic Substring	$\tilde{O}(\sqrt{n})$	[12]	$\Omega(\sqrt{n})$
Longest Square Substring	$\tilde{O}(\sqrt{n})$	This work (Theorem 5.1)	$\Omega(\sqrt{n})$

Fig. 1 Near-optimal quantum algorithms for string problems (see definitions in Sect. 2.2). Problems are grouped based on similarity. All problems listed here have near-linear time classical algorithms

Despite these results, our knowledge about the quantum computational complexities of basic string problems is far from complete. For the LCS problem and the Minimal String Rotation problem mentioned above, there are $n^{\Omega(1)}$ gaps between current upper bounds and lower bounds. Better upper bounds are only known in special cases: Le Gall and Seddighin [12] gave an $\tilde{O}(n^{2/3})$ -time algorithm for $(1 - \varepsilon)$ -approximating LCS in non-repetitive strings, matching the query lower bound in this setting. Wang and Ying [14] gave an $\tilde{O}(\sqrt{n})$ -time algorithm for Minimum String Rotation in randomly generated strings, and showed a matching average-case query lower bound. However, these algorithms do not immediately extend to the general cases. Moreover, there remain many other string problems which have near-linear time classical algorithms with no known quantum speed-up.

1.1 Our Results

In this work, we develop new quantum query algorithms for many fundamental string problems. *All our algorithms are near-optimal and time-efficient*: they have time complexities that match the corresponding query complexity lower bounds up to $n^{o(1)}$ factors. In particular, we close the gaps for Longest Common Substring and Minimal String Rotation left open in previous work [12, 14]. We summarize our contributions (together with some earlier results) in Fig. 1. See Sect. 2.2 for formal definitions of the studied problems.

1.2 Technical Overview

We give high-level overviews of our quantum algorithms for Longest Common Substring (LCS), Minimal String Rotation, and Longest Square Substring.

1.2.1 Longest Common Substring

We consider the decision version of LCS with threshold length d : given two length- n input strings s, t , decide whether they have a common substring of length d .

Le Gall and Seddighin [12, Section 3.1.1] observed a simple reduction from this decision problem to the (bipartite version of) Element Distinctness problem, which asks whether the two input lists A, B contain a pair of identical items $A_i = B_j$. Ambainis [13] gave a comparison-based algorithm for this problem in $\tilde{O}(n^{2/3} \cdot T)$ time, where T denotes the time complexity of comparing two items. In the LCS problem of threshold length d , each item is a length- d substring of s or t (specified by the starting position), and the lexicographical order between two length- d substrings can be compared in $T = \tilde{O}(\sqrt{d})$ using binary search and Grover search (see Lemma 2.5). Hence, this problem can be solved in $\tilde{O}(n^{2/3} \cdot \sqrt{d})$ time.

The anchoring technique The inefficiency of the algorithm described above comes from the fact that there are $n - d + 1 = \Omega(n)$ positions to be considered in each input string. This seems rather unnecessary for larger d , since intuitively there is a lot of redundancy from the large overlap between these length- d substrings.

This is the idea behind the so-called *anchoring* technique, which has been widely applied in designing classical algorithms for various versions of the LCS problem [22–28]. In this technique, we carefully pick subsets $C_1, C_2 \subseteq [n]$ of *anchors*, such that in a YES input instance there must exist an *anchored common substring*, i.e., a common string with occurrences $s[i \dots i + d) = t[j \dots j + d)$ and a shift $0 \leq h < d$ such that $i + h \in C_1$ and $j + h \in C_2$. Then, the task reduces to the *Two String Families LCP problem* [23], where we want to find a pair of anchors $i' \in C_1, j' \in C_2$ that can be extended in both directions to get a length- d common substring, or equivalently, the longest common prefix of $s[i' \dots]$, $t[j' \dots]$ and the longest common suffix of $s[\dots i' - 1]$, $t[\dots j' - 1]$ have total length at least d . Intuitively, finding a smaller set of anchors would make our algorithm have better running time.

Small and explicit anchor sets One can construct such anchor sets based on *difference covers* [29, 30], with size $|C_1|, |C_2| \leq n/\sqrt{d}$. The construction is very simple and explicit (see Sect. 3.3.1), and is oblivious to the content of the input strings (in fact, it just consists of several arithmetic progressions of fixed lengths). In comparison, there exist much smaller constructions if the anchors are allowed to depend on the input strings: for example, in their time-space tradeoff algorithm for LCS, Ben-Nun, Golan, Kociumaka, and Kraus [26] used *partitioning sets* [31] to construct an anchor set of size $O(n/d)$. However, this latter anchor set takes too long time to construct to be used in our sublinear-time quantum algorithm.

Our key idea is to combine the oblivious approach and non-oblivious approach, and design anchor sets with a balance between the *size* and the *construction time*: the number of anchors is $m = O(n/d^{3/4})$, and, given any index $i \in [m]$, the i^{th} anchor can be reported in $T = \tilde{O}(\sqrt{d})$ quantum time. Our construction (Sect. 3.3) is based on an *approximate version* of difference covers, combined with the *string synchronizing sets* recently introduced by Kempa and Kociumaka [32] (adapted to the sublinear setting using tools from pseudorandomness). Roughly speaking, allowing errors in the difference cover makes the size much smaller, while also introducing slight misalignments between the anchors, which are then to be fixed by the string synchronizing sets.

Anchoring via quantum walks Now we explain how to use small and explicit anchor sets to obtain better quantum LCS algorithms with time complexity $\tilde{O}(m^{2/3} \cdot (\sqrt{d} + T)) = \tilde{O}(n^{2/3})$, where $m = \tilde{O}(n/d^{3/4})$ is the number of anchors, and $T = \tilde{O}(\sqrt{d})$ is the time complexity of computing the i^{th} anchor. Our algorithm uses the *MNRS quantum walk* framework [33] (see Sect. 2.5) on Johnson graphs. Informally speaking, to apply this framework, we need to solve the following dynamic problem: maintain a *subset* of r anchors which undergoes insertions and deletions (called *update steps*), and in each query (called a *checking step*) we need to solve the Two String Families LCP problem on this subset, i.e., answer whether the current subset contains a pair of anchors that can extend to a length- d common substring. If each update step takes time U , and each checking step takes time C , then the MNRS quantum walk algorithm has overall running time $\tilde{O}(r U + \frac{m}{r} (\sqrt{r} U + C))$. We will achieve $U = \tilde{O}(\sqrt{d} + T)$ and $C = \tilde{O}(\sqrt{rd})$, and obtain the claimed time complexity by setting $r = m^{2/3}$.

To solve this dynamic problem, we maintain the lexicographical ordering of the length- d substrings specified by the current subset of anchors, as well as the corresponding LCP array which contains the length of the longest common prefix between every two lexicographically adjacent substrings. Note that the maintained information uniquely defines the *compact trie* of these substrings. This information can be updated easily after each insertion (or deletion) operation: we first compute the inserted anchor in T time, and then use binary search with Grover search to find its lexicographical rank and the LCP values with its neighbors, in $\tilde{O}(\sqrt{d})$ quantum time.

The maintained information will be useful for the checking step. In fact, if we only care about query complexity, then we are already done, since the maintained information already uniquely determines the answer of the Two String Families LCP problem, and no additional queries to the input strings are needed. The main challenge is to implement this checking step time-efficiently. Unfortunately, the classical near-linear-time algorithm [23] for solving the Two String Families LCP problem is too slow compared to our goal of $C = \tilde{O}(\sqrt{rd})$, and it is not clear how to obtain a quantum speedup over this classical algorithm. Hence, we should try to dynamically maintain the solution using data structures, instead of solving it from scratch every time. In fact, such a data structure with $\text{poly log}(n)$ time per operation was already given by Charalampopoulos, Gawrychowski, and Pokorski [27], and was used to obtain a classical data structure for maintaining Longest Common Substring under character substitutions. However, this data structure cannot be applied to the quantum walk algorithm, since it violates two requirements that are crucial for the correctness of quantum walk algorithms: (1) It should have worst-case time complexity (instead of being amortized), and (2) it should be *history-independent* (see the discussion in Sect. 3.2.1 for more details). Instead, we will design a different data structure that satisfies these two requirements, and can solve the Two String Families LCP problem on the maintained subset in $\tilde{O}(\sqrt{rd})$ quantum time. This time complexity is worse than the $\text{poly log}(n)$ time achieved by the classical data structure of [27], but suffices for our application.

A technical hurdle: limitations of 2D range query data structures Our solution for the Two String Families LCP problem is straightforward, but a key component in the algorithm relies on *dynamic 2-dimensional orthogonal range queries*. This is a

well-studied problem in the data structure literature, and many poly log (n) -time data structures are known (see [34–36] and the references therein). However, for our results, the 2-dimensional (2D) range query data structure in question has to satisfy not only the two requirements mentioned above, but also a third requirement of being *comparison-based*. In particular, we are not allowed to treat the coordinates of the 2D points as poly (n) -bounded integers, because the coordinates actually correspond to substrings of the input string, and should be compared by lexicographical order. Unfortunately, no data structures satisfying all three requirements are known.

To bypass this difficulty, our novel idea is to use a sampling procedure that lets us estimate the rank of a coordinate of the inserted 2D point among all the possible coordinates, which effectively allows us to convert the non-integer coordinates into integer coordinates. By a version of the Balls-and-Bins hashing argument, the inaccuracy incurred by the sampling can be controlled for *most* of the vertices on the Johnson graph which the quantum walk operates on. This then lets us apply 2D range query data structures over integer coordinates (see Sect. 3.2.3 for the details of this argument), which can be implemented with worst-case time complexity and history-independence as required. Combining this method with the tools and ideas mentioned before lets us get a time-efficient implementation of the quantum walk algorithm for computing the LCS.

We believe this sampling idea will find further applications in improving the time efficiency of quantum walk algorithms (for example, it can simplify the implementation of Ambainis’ $\tilde{O}(n^{2/3})$ -time Element Distinctness algorithm, as noted in Sect. 6).

1.2.2 Minimal String Rotation

In the Minimal String Rotation problem, we are given a string s of length n and are tasked with finding the cyclic rotation of s which is lexicographically the smallest. We sketch the main ideas of our improved quantum algorithm for Minimal String Rotation by comparing it to the previous best solution for this problem.

The simplest version of Wang and Ying’s algorithm [14, Theorem 5.2] works by identifying a small prefix of the minimal rotation using Grover search, and then applying pattern matching with this small prefix to find the starting position of the minimum rotation. More concretely, let B be some size parameter. By quantum minimum finding over all prefixes of length B among the rotations of s , we can find the length- B prefix P of the minimal rotation in asymptotically $\sqrt{B} \cdot \sqrt{n}$ time. Next, split the string s into $\Theta(n/B)$ blocks of size $\Theta(B)$ each. Within each block, we find the *leftmost* occurrence of P via quantum Exact String Matching [9]. It turns out that one of these positions is guaranteed to be a starting position of the minimal rotation (this property is called an “exclusion rule” or “Ricochet Property” in the literature). By minimum finding over these $O(n/B)$ candidate starting positions (and comparisons of length- n strings via Grover search), we can find the true minimum rotation in asymptotically $\sqrt{n/B} \cdot \sqrt{n}$ time. So overall the algorithm takes asymptotically

$$\sqrt{Bn} + (n/\sqrt{B})$$

time, which is minimized at $B = \sqrt{n}$ and yields a runtime of $\tilde{O}(n^{3/4})$.

This algorithm is inefficient in its first step, where it uses quantum minimum finding to obtain the minimum length- B prefix P . The length- B prefixes we are searching over all come from rotations of the same string s . Due to this common structure, we should be able to find their minimum more efficiently than just using the generic algorithm for minimum finding. At a high level, we improve this step by finding P using *recursion* instead. Intuitively, this is possible because the Minimal Rotation problem is already about finding the minimum “prefix” (just of length n) among rotations of s . This then yields a recursive algorithm running in $n^{1/2+o(1)}$ quantum time.

In the presentation of this algorithm in Sect. 4, we use a chain of reductions and actually solve a more general problem to get this recursion to work. The argument also relies on a new “exclusion rule,” adapted from previous work, to prove that we only need to consider a constant number of candidate starting positions of the minimum rotation within each small block of the input string.

1.2.3 Longest Square Substring

A *square string* is a string of even length with the property that its first half is identical to its second half. In other words, a string is square if it can be viewed as some string repeated twice in a row.

We show how to find the longest square substring in an input string of length n using a quantum algorithm which runs in $\tilde{O}(\sqrt{n})$ time. Our algorithm mostly follows the ideas used in [12] to solve the Longest Palindromic Substring problem, but makes some modifications due to the differing structures of square substrings and palindromic substrings (for example, [12] exploits the fact that if a string contains a large palindromic substring it has smaller palindromic substrings centered at the same position; in contrast, it is possible for a string to have a large square substring but not contain any smaller square substrings, so we cannot leverage this sort of property).

At a high level, our algorithm starts by guessing the size of the longest square substring within a $(1 + \varepsilon)$ factor for some small constant $\varepsilon > 0$. We then guess a large substring P contained in the first half of an optimal solution, and then use the quantum algorithm for Exact String Matching to find a copy of this P in the second half of the corresponding solution. If we find a unique copy of P , we can use a Grover search to extend outwards from our copies of P and recover a longest square substring. Otherwise, if we find multiple copies, it implies our substring is periodic, so we can use a Grover search to find a maximal periodic substring containing a large square substring, and then employ some additional combinatorial arguments to recover the solution.

1.3 Related Work

Quantum algorithms on string problems Wang and Ying [14] improved the logarithmic factors of the quantum Exact String Matching algorithm by Ramesh and Vinay [9] (and filled in several gaps in their original proof), and showed that the same technique can be used to find the smallest period of a periodic string [14, Appendix D].

Another important string problem is computing the *edit distance* between two strings (the minimum number of deletions, insertions, and substitutions needed to turn one string into the other). The best known classical algorithm has $O(n^2/\log^2 n)$ time complexity [37], which is near-optimal under the Strong Exponential Time Hypothesis [38]. It is open whether quantum algorithms can compute edit distance in truly subquadratic time. For the approximate version of the edit distance problem, the breakthrough work of Boroujeni et al. [39] gave a truly subquadratic time quantum algorithm for computing a constant factor approximation. The quantum subroutines of this algorithm were subsequently replaced with classical randomized algorithms in [40] to get a truly subquadratic classical algorithm that approximates the edit distance to a constant factor.

Le Gall and Seddighin [12] also considered the $(1 + \varepsilon)$ -approximate Ulam distance problem (i.e., edit distance on non-repetitive strings), and showed a quantum algorithm with near-optimal $\tilde{O}(\sqrt{n})$ time complexity. Their algorithm was based on the classical algorithm by Naumovitz, Saks, and Seshadhri [41].

Montarano [42] gave quantum algorithms for the d -dimensional pattern matching problem with random inputs. Ambainis et al. [43] gave quantum algorithms for deciding Dyck languages. There are also some results [44, 45] on string problems with non-standard quantum queries to the input.

Quantum walks and time-efficient quantum algorithms Quantum walks [13, 33, 46] are a useful method to obtain query-efficient quantum algorithms for many important problems, such as Element Distinctness [13] and Triangle Finding [47–49]. Ambainis showed that the query-efficient algorithm for element distinctness [13] can also be implemented in a time-efficient manner with only a poly $\log(n)$ blowup, by applying history-independent data structures in the quantum walk. Since then, this “quantum walk plus data structure” strategy has been used in many quantum algorithms to obtain improved time complexity. For example, Belovs, Childs, Jeffery, Kothari, and Magniez [50] used nested quantum walk with Ambainis’ data structure to obtain time-efficient algorithms for the 3-distinctness problem. Bernstein, Jeffery, Lange, and Meurer [51] designed a simpler data structure called quantum radix tree [52], and applied it in their quantum walk algorithms for the Subset Sum problem on random input. Aaronson, Chia, Lin, Wang, and Zhang [53] gave a quantum walk algorithm for the Closest-Pair problem in $O(1)$ -dimensional space with near-optimal time complexity $\tilde{O}(n^{2/3})$. The previous $\tilde{O}(n^{2/3})$ -time algorithm for approximating LCS in non-repetitive strings [12] also applied quantum walks.

On the other hand, query-efficient quantum algorithms do not always have time-efficient implementations. This motivated the study of *quantum fine-grained complexity*. Aaronson et al. [53] formulated the QSETH conjecture, which is a quantum analogue of the classical Strong Exponential Time Hypothesis, and showed that Orthogonal Vectors and Closest-Pair in poly $\log(n)$ -dimensional space require $n^{1-o(1)}$ quantum time under QSETH. In contrast, these two problems have simple quantum walk algorithms with only $O(n^{2/3})$ query complexity. Buhrman, Patro, and Speelman [54] formulated another version of QSETH, which implies a conditional $\Omega(n^{1.5})$ -time lower bound for quantum algorithms solving the edit distance problem. Recently, Buhrman, Loff, Patro, and Speelman [55] proposed the quantum 3SUM hypothesis,

and used it to show that the quadratic quantum speedups obtained by Ambainis and Larka [56] for many computational geometry problems are conditionally optimal. Notably, in their fine-grained reductions, they employed a quantum walk with data structures to bypass the linear-time preprocessing stage that a naive approach would require.

Classical string algorithms We refer readers to several excellent textbooks [57–59] on string algorithms.

Weiner [3] introduced the suffix tree and gave a linear-time algorithm for computing the LCS of two strings over a constant-sized alphabet. For polynomially-bounded integer alphabets, Farach’s construction of suffix trees [4] implies an linear-time algorithm for LCS. Babenko and Starikovskaya [5] gave an algorithm for LCS based on suffix arrays. Recently, Charalampopoulos, Kociumaka, Pissis, and Radoszewski [28] gave faster word-RAM algorithms for LCS on compactly represented input strings over a small alphabet. The LCS problem has also been studied in the settings of time-space tradeoffs [22, 26, 60], approximate matching [5, 23, 61–66], and dynamic data structures [24, 25, 27].

Booth [6] and Shiloach [7] gave the first linear time algorithms for the Minimal String Rotation problem. Later, Duval [8] gave a constant-space linear-time algorithm for computing the *Lyndon factorization* of a string, which can be used to compute the minimal rotation, maximal suffix, and minimal suffix. Duval’s algorithm can also compute the minimal suffix and maximal suffix for every prefix of the input string. Apostolico and Crochemore [67] gave a linear-time algorithm for computing the minimal rotation of every prefix of the input string. Parallel algorithms for Minimal String Rotation were given by Iliopoulos and Smyth [17]. There are data structures [68–70] that, given a substring specified by its position and length in the input string, can efficiently answer its minimal suffix, maximal suffix, and minimal rotation. The Longest Lyndon Substring problem can be solved in linear time [8] by simply outputting the longest segment in the Lyndon factorization. There are data structures [24, 71] for dynamically maintaining the longest Lyndon substring under character substitutions.

There are $O(n \log n)$ -time algorithms for finding all the square substrings of the input string [16, 72, 73]. There are data structures [74] for dynamically maintaining the longest square substring under character substitutions.

The construction of difference cover [29, 30] has been previously used in many string algorithms, e.g., [28, 29, 75, 76]. The string synchronizing set recently introduced by Kempa and Kociumaka [32] has been applied in [28, 32, 77, 78]. The local-consistency idea behind the construction of string synchronizing set had also appeared in previous work [31, 79, 80].

1.4 Subsequent Works

Some of our results were improved after the conference version of this paper was published. Jin and Nogler [81] gave a quantum algorithm that decides whether the Longest Common Substring of two input strings has length at least d in $\tilde{O}((n/d)^{2/3} \cdot d^{1/2+o(1)})$ quantum query and time complexity. Childs, Kothari, Kovacs-Deak, Sundaram, and

Wang [82] gave a quantum algorithm for a *decision version* of the Lexicographical Minimal String Rotation problem in $\tilde{O}(n^{1/2})$ quantum query complexity.

1.5 Paper Organization

In Sect. 2, we provide useful definitions and review some quantum primitives which will be used in our algorithms. In Sect. 3, we present our algorithm for *Longest Common Substring*. In Sect. 4, we present our algorithm for *Minimal String Rotation* and several related problems. In Sect. 5, we present our algorithm for *Longest Square Substring*. Finally, we mention several open problems in Sect. 6.

2 Preliminaries

2.1 Notations and Basic Properties of Strings

We define sets $\mathbb{N} = \{0, 1, 2, 3, \dots\}$ and $\mathbb{N}^+ = \{1, 2, 3, \dots\}$. For every positive integer n we introduce the set $[n] = \{1, 2, \dots, n\}$. Given two integers $i \leq j$, we let $[i \dots j] = \{i, i + 1, \dots, j\}$ denote the set of integers in the closed interval $[i, j]$. We define $[i \dots j)$, $(i \dots j]$, and $(i \dots j)$ analogously. For an integer x and an integer set A , let $x + A$ denote $\{x + a : a \in A\}$, and let xA denote $\{x \cdot a : a \in A\}$.

As is standard in the literature, we consider strings over a *polynomially-bounded integer alphabet* $\Sigma = [1 \dots |\Sigma|]$ of size $|\Sigma| \leq n^{O(1)}$. A string $s \in \Sigma^*$ of length $|s| = n$ is a sequence of characters $s = s[1]s[2] \dots s[n]$ from the alphabet Σ (we use 1-based indexing). The *concatenation* of two strings $s, t \in \Sigma^*$ is denoted by st . The *reversed string* of s is denoted by $s^R = s[n]s[n - 1] \dots s[1]$.

Given a string s of length $|s| = n$, a *substring* of s is any string of the form $s[i \dots j] = s[i]s[i + 1] \dots s[j]$ for some indices $1 \leq i \leq j \leq n$. For $i > j$, we define $s[i \dots j]$ to be the empty string ε . When i, j may be out of bounds, we use the convention $s[i \dots j] = s[\max\{1, i\} \dots \min\{n, j\}]$ to simplify notation. We sometimes use $s[i \dots j) = s[i]s[i + 1] \dots s[j - 1]$ and $s(i \dots j) = s[i + 1] \dots s[j - 1]s[j]$ to denote substrings. A substring $s[1 \dots j]$ is called a *prefix* of s , and a substring $s[i \dots n]$ is called a *suffix* of s . For two strings s, t , let $\text{lcp}(s, t) = \max\{j : j \leq \min\{|s|, |t|\}, s[1 \dots j] = t[1 \dots j]\}$ denote the length of their *longest common prefix*.

We say string s is *lexicographically smaller* than string t (denoted $s < t$) if either s is a proper prefix of t (i.e., $|s| < |t|$ and $s = t[1 \dots |s|]$), or $\ell = \text{lcp}(s, t) < \min\{|s|, |t|\}$ and $s[\ell + 1] < t[\ell + 1]$. The notations $>, \leq, \geq$ are defined analogously. The following easy-to-prove and well-known fact has been widely used in string data structures and algorithms.

Lemma 2.1 (e.g. [83, Lemma 1]) *Given strings $s_1 \leq s_2 \leq \dots \leq s_m$, we have $\text{lcp}(s_1, s_m) = \min_{1 \leq i < m} \{\text{lcp}(s_i, s_{i+1})\}$.*

For a positive integer $p \leq |s|$, we say p is a *period* of s if $s[i] = s[i + p]$ holds for all $1 \leq i \leq |s| - p$. One can compute all the periods of s by a classical deterministic algorithm in linear time [1]. We refer to the minimal period of s as *the period* of s , and denote it by $\text{per}(s)$. If $\text{per}(s) \leq |s|/2$, we say that s is *periodic*. If $\text{per}(s)$ does

not divide $|s|$, we say that s is *primitive*. We will need the following classical results regarding periods of strings for some of our algorithms.

Lemma 2.2 (Weak Periodicity Lemma, [84]) *If a string s has periods p and q such that $p + q \leq |s|$, then $\gcd(p, q)$ is also a period of s .*

Lemma 2.3 (e.g., [79, 85]) *Let s, t be two strings with $|s|/2 \leq |t| \leq |s|$, and let $s[i_1 \dots i_1 + |t|] = s[i_2 \dots i_2 + |t|] = \dots = s[i_m \dots i_m + |t|] = t$ be all the occurrences of t in s (where $i_k < i_{k+1}$). Then, i_1, i_2, \dots, i_m form an arithmetic progression. Moreover, if $m \geq 3$, then $\text{per}(t) = i_2 - i_1$.*

We say string s is a (cyclic) rotation of string t , if $|s| = |t| = n$ and there exists an index $1 \leq i \leq n$ such that $s = t[i \dots n]t[1 \dots i - 1]$. If string s is primitive and is lexicographically minimal among its cyclic rotations, we call s a *Lyndon word*. Equivalently, s is a Lyndon word if and only if $s \leq t$ for all proper suffixes t of s . For a periodic string s with minimal period $\text{per}(s) = p$, the *Lyndon root* of s is defined as the lexicographically minimal rotation of $s[1 \dots p]$, which can be computed by a classical deterministic algorithm in linear time (e.g., [6–8]).

2.2 Problem Definitions

We give formal definitions of the string problems considered in this paper.

Longest Common Substring

Input: Two strings s, t

Task: Output the maximum length ℓ such that $s[i \dots i + \ell] = t[j \dots j + \ell]$ for some $i \in [|s| - \ell + 1], j \in [|t| - \ell + 1]$.

Longest Repeated Substring

Input: A string s

Task: Output the maximum length ℓ such that $s[i \dots i + \ell] = s[j \dots j + \ell]$ for some $i, j \in [|s| - \ell + 1], i \neq j$.

Longest Square Substring

Input: A string s

Task: Output the maximum shift Δ such that $s[i \dots i + \Delta] = s[i + \Delta \dots i + 2\Delta]$ for some $i \in [1 \dots |s| - 2\Delta + 1]$.

Longest Lyndon Substring

Input: A string s

Task: Output the maximum length ℓ such that $s[i \dots i + \ell]$ is a Lyndon word for some $i \in [|s| - \ell + 1]$.

Exact String Matching

Input: Two strings s, t with $|s| \geq |t|$

Task: Output the minimum position i such that $s[i \dots i + |t|] = t$.

Minimal String Rotation

Input: A string s

Task: Output a position $i \in [1 \dots |s|]$ such that $s[i \dots |s|]s[1 \dots i - 1] \preceq s[j \dots |s|]s[1 \dots j - 1]$ holds for all $j \in [1 \dots |s|]$. If there are multiple solutions, output the smallest such i .

Maximal Suffix

Input: A string s

Task: Output the position $i \in [1 \dots |s|]$ such that $s[i \dots |s|] \succ s[j \dots |s|]$ holds for all $j \in [|s|] \setminus \{i\}$.

Minimal Suffix

Input: A string s

Task: Output the position $i \in [1 \dots |s|]$ such that $s[i \dots |s|] \prec s[j \dots |s|]$ holds for all $j \in [|s|] \setminus \{i\}$.

In the first four problems, we only require the algorithm to output the maximum length. The locations of the witness substrings can be found by a binary search.

2.3 Computational Model

We assume the input strings can be accessed in a quantum query model [86, 87], which is standard in the literature of quantum algorithms. More precisely, letting s be an input string of length n , we have access to an oracle O_s that, for any index $i \in [n]$ and any $b \in \Sigma$, performs the unitary mapping $O_s: |i, b\rangle \mapsto |i, b \oplus s[i]\rangle$, where \oplus denotes the XOR operation on the binary encodings of characters. The oracles can be queried in superposition, and each query has unit cost. Besides the input queries, the algorithm can also apply intermediate unitary operators that are independent of the input oracles. Finally, the query algorithm should return the correct answer with success probability at least $2/3$ (which can be boosted to high probability² by a majority vote over $O(\log n)$ repetitions). The *query complexity* of an algorithm is the number of queries it makes to the input oracles.

In this paper, we are also interested in the *time complexity* of the quantum algorithms, which counts not only the queries to the input oracles, but also the elementary gates [88] for implementing the unitary operators that are independent of the input. In order to implement the query algorithms in a time-efficient manner, we also need the *quantum random access gate*, defined as

$$|i, b, z_1, \dots, z_m\rangle \mapsto |i, z_i, z_1, \dots, z_{i-1}, b, z_{i+1}, \dots, z_m\rangle,$$

to access at unit cost the i^{th} element from the quantum working memory $|z_1, \dots, z_m\rangle$. Assuming quantum random access, a classical time- T algorithm that uses random access memory can be converted into a quantum subroutine in time $O(T)$, which can be

² We say an algorithm succeeds with *high probability* (w.h.p.), if the success probability can be made at least $1 - 1/n^c$ for any desired constant $c > 1$.

invoked by quantum search primitives such as Grover search. Quantum random access has become a standard assumption in designing time-efficient quantum algorithms (for example, all the time-efficient quantum walk algorithms mentioned in Sect. 1.3 relied on this assumption).

2.4 Basic Quantum Primitives

Grover search (Amplitude amplification) [11, 89]. Let $f: [n] \rightarrow \{0, 1\}$ be a function, where $f(i)$ for each $i \in [n]$ can be evaluated in time T . There is a quantum algorithm that, with high probability, finds an $x \in f^{-1}(1)$ or report that $f^{-1}(1)$ is empty, in $\tilde{O}(\sqrt{n} \cdot T)$ time. Moreover, if it is guaranteed that either $|f^{-1}(1)| \geq M$ or $|f^{-1}(1)| = 0$ holds, then the algorithm runs in $\tilde{O}(\sqrt{n}/M \cdot T)$ time.

Quantum minimum finding [15]. Let x_1, \dots, x_n be n items with a total order, where each pair of x_i and x_j can be compared in time T . There is a quantum algorithm that, with high probability, finds the minimum item among x_1, \dots, x_n in $\tilde{O}(\sqrt{n} \cdot T)$ time.

Remark 2.4 If the algorithm for evaluating $f(i)$ (or for comparing x_i, x_j) has some small probability of outputting the wrong answer, we can first boost it to high success probability, and then the Grover search (or Quantum minimum finding) still works, since quantum computational errors only accumulate linearly. It is possible to improve the log-factors in the query complexity of quantum search when the input has errors [90], but in this paper we do not seek to optimize the log-factors.

Lemma 2.5 (Computing LCP) *Given two strings s, t of lengths $|s|, |t| \leq n$, there is a quantum algorithm that computes $\text{lcp}(s, t)$ and decides whether $s \preceq t$, in $\tilde{O}(\sqrt{n})$ time.*

Proof Note that we can use Grover search to decide whether two strings are identical in $\tilde{O}(\sqrt{n})$ time. Then we can compute $\text{lcp}(s, t)$ by a simple binary search over the length of the prefix. After that we can easily compare their lexicographical order by comparing the next position. \square

Given a string s and positions g and h such that $s[g] = s[h]$, we often use Lemma 2.5 to “extend” these common characters to larger identical strings to some bound d while keeping them equivalent (i.e. find the largest positive integer $j \leq d$ such that $s[g \dots g+j] = s[h \dots h+j]$). We will often refer to this process (somewhat informally) as “extending strings via Grover search.”

As a final useful subroutine, we appeal to the result of Ramesh and Vinay [9], who combined Grover search with the deterministic sampling technique of Vishkin [10], and obtained a quantum algorithm for Exact String Matching.

Theorem 2.6 (Quantum Exact String Matching [9]) *We can solve the Exact String Matching problem with a quantum algorithm on input strings s, t of length at most n using $\tilde{O}(\sqrt{n})$ query complexity and time complexity.*

2.5 Quantum Walks

We use the quantum walk framework developed by Magniez, Nayak, Roland, and Santha [33], and apply it on Johnson graphs,

The Johnson graph $J(m, r)$ has $\binom{m}{r}$ vertices, each being an subset of $[m]$ with size r , where two vertices in the graph $A, B \in \binom{[m]}{r}$ are connected by an edge if and only if $|A \cap B| = r - 1$, or equivalently there exist $a \in A, b \in [m] \setminus A$ such that $B = (A \setminus \{a\}) \cup \{b\}$. Depending on the application, we usually identify a special subset of the vertices $V_{\text{marked}} \subseteq \binom{[m]}{r}$ as being *marked*. The quantum walk is analogous to a random walk on the Johnson graph attempting to find a marked vertex, but provides quantum speed-up compared to the classical random walk. The vertices in the Johnson graph are also called the states of the walk.

In the quantum walk algorithm, each vertex $K \in \binom{[m]}{r}$ is associated with a data structure $D(K)$. The setup cost S is the cost to set up the data structure $D(K)$ for any $K \in \binom{[m]}{r}$, where the cost could be measured in query complexity or time complexity. The checking cost C is the cost to check whether K is a marked vertex, given the data structure $D(K)$. The update cost U is the cost of updating the data structure from $D(K)$ to $D(K')$, where $K' = (K \setminus \{a\}) \cup \{b\}$ is an adjacent vertex specified by $a \in K, b \in [m] \setminus K$. The MNRS quantum walk algorithm can be summarized as follows.

Theorem 2.7 (MNRS Quantum Walk [33]) *Suppose $|V_{\text{marked}}|/\binom{m}{r} \geq \varepsilon$ whenever V_{marked} is non-empty. Then there is a quantum algorithm that with high probability determines if V_{marked} is empty or finds a marked vertex, with cost of order $S + \frac{1}{\sqrt{\varepsilon}}(\sqrt{r} \cdot U + C)$.*

Readers unfamiliar with the quantum walk approach are referred to [91, Section 8.3.2] for a quick application of this theorem to solve the Element Distinctness problem using $O(n^{2/3})$ quantum queries. This algorithm can be implemented in $\tilde{O}(n^{2/3})$ time by carefully designing the data structures to support time-efficient insertion, deletion, and searching [13, Section 6.2]. We elaborate on the issue of time efficiency when we apply quantum walks in our algorithm in Sect. 3.2.

3 Longest Common Substring

In this section, we prove the following theorem.

Theorem 3.1 *The Longest Common Substring (LCS) problem can be solved by a quantum algorithm with $\tilde{O}(n^{2/3})$ query complexity and time complexity.*

In Sect. 3.1, we will give an outline of our quantum walk algorithm based on the notion of *good anchor sets*, and show that this algorithm achieves good *query complexity*. Then in Sect. 3.2, we describe how to use data structures to implement the quantum walk algorithm in a time-efficient manner. Finally, in Sect. 3.3, we present the construction of good anchor sets used by the algorithm. We have organized the arguments so that Sects. 3.2 and 3.3 are independent of one another, and can be read separately.

3.1 Anchoring Via Quantum Walks

As mentioned in Sect. 1.2.1, our algorithm for LCS is based on the anchoring technique which previously appeared in classical LCS algorithms. Here, we will implement this technique using the MNRS quantum walk framework (Sect. 2.5).

Notations and input assumptions To simplify the presentation, we concatenate the two input strings s, t into $S := s\$t$, where $\$$ is a delimiter symbol that does not appear in the input strings, and let $n = |S| = |s| + 1 + |t|$. So $s[i] = S[i]$ for all $i \in [1 \dots |s|]$, and $t[j] = S[|s| + 1 + j]$ for all $j \in [1 \dots |t|]$.

We will only solve the *decision version* of LCS: given a length threshold d , determine whether s and t have a common substring of length d . The algorithm for computing the length of the longest common substring then follows from a binary search over the threshold d . We assume $d \geq 100$ to avoid corner cases in later analysis; for smaller d , the problem can be solved in $\tilde{O}(n^{2/3}d^{1/2}) = \tilde{O}(n^{2/3})$ time by reducing to the (bipartite version of) element distinctness problem [12, Section 3.1.1] and applying Ambainis' algorithm [13] (see Sect. 1.2.1).

Anchoring We begin by introducing the notion of *good anchor sets*.

Definition 3.2 (*Good anchor sets*) For input strings s, t and threshold length d , we call $C \subseteq [1 \dots n]$ a *good anchor set* if the following holds: if the longest common substring of s and t has length at least d , then there exist positions $i \in [1 \dots |s| - d + 1]$, $j \in [1 \dots |t| - d + 1]$ and a shift $h \in [0 \dots d]$, such that $s[i \dots i + d] = t[j \dots j + d]$, and $i + h, |s| + 1 + j + h \in C$.

In this definition, the anchor set C is allowed to depend on s and t . If $C = \{C(1), C(2), \dots, C(m)\}$ and there is a (quantum) algorithm that, given any index $1 \leq j \leq m$, computes the element $C(j)$ in $T(n, d)$ time, then we say C is $T(n, d)$ -(quantum)-time constructible. The elements $C(1), C(2), \dots, C(m)$ are allowed to contain duplicates (i.e., C could be a multiset), and are not necessarily sorted in any particular order.

The set $[1 \dots n]$ is trivially a good anchor set, but there are constructions of much smaller size. As a concrete example, one can directly construct good anchor sets using *difference covers*.

Definition 3.3 (*Difference cover* [29, 30]) A set $D \subseteq \mathbb{N}^+$ is called a d -cover, if for every $i, j \in \mathbb{N}^+$, there exists an integer $h(i, j) \in [0 \dots d]$ such that $i + h(i, j), j + h(i, j) \in D$.

The following construction of d -cover has optimal size (up to a constant factor).

Lemma 3.4 (Construction of d -cover [29, 30]) *For every positive integer $d \geq 1$, there is a d -cover D such that $D \cap [n]$ contains $O(n/\sqrt{d})$ elements. Moreover, given integer $i \geq 1$, one can compute the i^{th} smallest element of $D \cap [n]$ in $\tilde{O}(1)$ time.*

Here we omit the proof of Lemma 3.4, as a more general version (Lemma 3.19) will be proved later in Sect. 3.3.1. Using difference covers, we immediately have the following simple construction of good anchor sets.

Corollary 3.5 (A simple good anchor set) *There is a $\tilde{O}(1)$ -time constructible good anchor set C of size $m = O(n/\sqrt{d})$.*

Proof Let D be the d -cover from Lemma 3.4. Then, for input strings s, t and threshold length d , it immediately follows from definition that $C := (D \cap [|s|]) \cup (|s| + 1 + (D \cap [|t|]))$ is a good anchor set. \square

Note that the construction in Corollary 3.5 is deterministic, and oblivious to the content of the input strings s and t . The following lemma (which will be proved in Sect. 3.3) states that we can achieve a smaller size by a probabilistic non-oblivious construction that takes longer time to compute.

Lemma 3.6 (A smaller good anchor set) *There is an $\tilde{O}(\sqrt{d})$ -quantum-time constructible anchor set C of size $m = O(n/d^{3/4})$. This set C depends on the input strings s, t and $O(\log n)$ many random coins, and is a good anchor set with at least $2/3$ probability over the random coins.*

Let $C = \{C(1), \dots, C(m)\} \subseteq [n]$ be a good anchor set of size $|C| = m$. For every anchor $C(k)$ indexed by $k \in [m]$, we associate it with a pair of strings $(P(k), Q(k))$, where

$$\begin{aligned} P(k) &:= S[C(k) \dots C(k) + d], \\ Q(k) &:= (S(C(k) - d \dots C(k) - 1))^R \end{aligned}$$

are substrings (or reversed substrings) of S obtained by extending from the anchor $C(k)$ to the right or reversely to the left. The length of $P(k)$ is at most³ d , and the length of $Q(k)$ is at most $d - 1$. We say the string pair $(P(k), Q(k))$ is *red* if $C(k) \in [1 \dots |s|]$, or *blue* if $C(k) \in [|s| + 1 \dots n]$. We also say $k \in [m]$ is a *red index* or a *blue index*, depending on the color of the string pair $(P(k), Q(k))$. Then, from the definition of good anchor sets, we immediately have the following simple observation.

Proposition 3.7 (Witness Pair) *The longest common substring of s and t has length at least d , if and only if there exist a red string pair $(P(k), Q(k))$ and a blue string pair $(P(k'), Q(k'))$ where $k, k' \in [m]$, such that $\text{lcp}(P(k), P(k')) + \text{lcp}(Q(k), Q(k')) \geq d$. In such case, (k, k') is called a witness pair.*

Proof Suppose s and t have LCS of length at least d . Then the property of the good anchor set C implies the existence of a shift $h \in [0 \dots d]$ and a length- d common substring $s[i \dots i + d] = t[j \dots j + d]$ such that $i + h = C(k)$, $|s| + 1 + j + h = C(k')$ for some $k, k' \in [m]$. Then, we must have $\text{lcp}(P(k), P(k')) \geq d - h$ and $\text{lcp}(Q(k), Q(k')) \geq h$, implying that (k, k') is a witness pair.

Conversely, the existence of a witness pair immediately implies a common substring of length at least d . \square

Remark 3.8 The algorithm we are going to describe can be easily adapted to the Longest Repeated Substring problem: we only have one input string $S[1 \dots n]$, and we drop the red-blue constraint in the definition of witness pairs in Proposition 3.7.

³ Recall that we use the convention $S[x \dots y] := S[\max\{1, x\} \dots \min\{y + d, n + 1\}]$ for a length- n string S .

Now we shall describe our quantum walk algorithm that solves the decision version of LCS by searching for a witness pair.

Definition of the Johnson graph Recall that $C = \{C(1), \dots, C(m)\}$ is a good anchor set of size $|C| = m$. We perform a quantum walk on the Johnson graph with vertex set $\binom{[m]}{r}$, where r is a parameter to be determined later. A vertex $K = \{k_1, k_2, \dots, k_r\} \subseteq [m]$ in the Johnson graph is called a *marked vertex*, if and only if $\{k_1, k_2, \dots, k_r\}$ contains a witness pair (Proposition 3.7). If s and t have a common substring of length d , then at least $\binom{m-2}{r-2} / \binom{m}{r} = \Omega(r^2/m^2)$ fraction of the vertices are marked. Otherwise, there are no marked vertices.

Associated data In the quantum walk algorithm, each state $K = \{k_1, \dots, k_r\} \subseteq [m]$ is associated with the following data.

- The indices k_1, \dots, k_r themselves.
- The corresponding anchors $C(k_1), \dots, C(k_r) \in [n]$.
- An array (k_1^P, \dots, k_r^P) , which is a permutation of k_1, \dots, k_r , such that $P(k_i^P) \leq P(k_{i+1}^P)$ for all $1 \leq i < r$.
- The LCP array h_1^P, \dots, h_{r-1}^P , where $h_i^P = \text{lcp}(P(k_i^P), P(k_{i+1}^P))$
- An array (k_1^Q, \dots, k_r^Q) , which is a permutation of k_1, \dots, k_r , such that $Q(k_i^Q) \leq Q(k_{i+1}^Q)$ for all $1 \leq i < r$.
- The LCP array h_1^Q, \dots, h_{r-1}^Q , where $h_i^Q = \text{lcp}(Q(k_i^Q), Q(k_{i+1}^Q))$.

Note that we stored the *lexicographical orderings* of the strings $P(k_1), \dots, P(k_r)$ and $Q(k_1), \dots, Q(k_r)$ (for identical substrings, we break ties by comparing the indices themselves), as well as the *LCP arrays* which include the length of the longest common prefix of every pair of lexicographically adjacent substrings. By Lemma 2.1, these arrays together uniquely determine the values of $\text{lcp}(P(k_i), P(k_j))$ and $\text{lcp}(Q(k_i), Q(k_j))$, for every pair of $i, j \in [r]$.⁴

In the checking step of the quantum walk algorithm, we decide whether the state is marked, by searching for a witness pair (Proposition 3.7) in $\{k_1, \dots, k_r\}$. Note that the contents of the involved strings $\{P(k_i)\}_{i \in [r]}$, $\{Q(k_i)\}_{i \in [r]}$ are no longer needed in order to solve this task, as long as we already know their lexicographical orderings and the LCP arrays. This task is termed as the *Two String Families LCP* problem in the literature [23], formalized as below.

Two String Families LCP

Input: r red/blue pairs of strings $(P_1, Q_1), (P_2, Q_2), \dots, (P_r, Q_r)$ of lengths $|P_i|, |Q_i| \leq d$, which are represented by the lexicographical orderings of P_1, \dots, P_r and of Q_1, \dots, Q_r , and their LCP arrays

Task: Decide if there exist a red pair (P, Q) and a blue pair (P', Q') , such that $\text{lcp}(P, P') + \text{lcp}(Q, Q') \geq d$.

We will show how to solve this task time-efficiently in Sect. 3.2. For now, we only consider the query complexity of the algorithm, and we have the following simple

⁴ To better understand this fact, observe that they uniquely determine the *compact tries* of $P(k_1), \dots, P(k_r)$ and of $Q(k_1), \dots, Q(k_r)$, where the LCP of two strings equals the depth of the lowest common ancestor of the corresponding nodes in the compact trie.

observation, due to the fact that our associated information already uniquely determines the LCP value of every pair.

Proposition 3.9 (Query complexity of checking step is zero) *Using the associated data defined above, we can determine whether $\{k_1, \dots, k_r\} \subseteq [m]$ is a marked state, without making any additional queries to the input.*

Now, we consider the cost of maintaining the associated data when the subset $\{k_1, \dots, k_r\}$ undergoes insertion and deletion during the quantum walk algorithm.

Proposition 3.10 (Update cost) *Assume the anchor set C is T -time constructible. Then, each update step of the quantum walk algorithm has query complexity $U = \tilde{O}(\sqrt{d} + T)$.*

Proof Let us consider how to update the associated data when a new index k is being inserted into the subset $\{k_1, \dots, k_r\}$. The deletion process is simply the reverse operation of insertion.

The insertion procedure can be summarized by the pseudocode in Algorithm 1. First, we compute and store $C(k)$ in time T . Then we use a binary search to find the correct place to insert k into the lexicographical orderings (k_1^P, \dots, k_r^P) (and (k_1^Q, \dots, k_r^Q)). Since the involved substrings have length d , each lexicographical comparison required by this binary search can be implemented in $\tilde{O}(\sqrt{d})$ time by Lemma 2.5. After inserting k into the list, we update the LCP array by computing its LCP values $h_{\text{pre}}, h_{\text{suc}}$ with two neighboring substrings, and removing (by “uncomputing”) the LCP value h_{old} between their neighbors which were adjacent at first, in $\tilde{O}(\sqrt{d})$ time (Lemma 2.5). \square

Algorithm 1: The insertion procedure

- 1 Given an index $k \in [m]$
 - 2 Compute $C(k)$
 - 3 Store the data $(k, C(k))$
 - 4 Compute the rank i of $P(k)$ among $P(k_1^P), \dots, P(k_r^P)$
 - 5 Compute $h_{\text{pre}} = \text{lcp}(P(k_{i-1}^P), P(k))$
 - 6 Compute $h_{\text{suc}} = \text{lcp}(P(k_i^P), P(k))$
 - 7 Compute $h_{\text{old}} = \text{lcp}(P(k_{i-1}^P), P(k_i^P))$
 - 8 Update $(k_1^P, \dots, k_r^P) \leftarrow (k_1^P, \dots, k_{i-1}^P, k, k_i^P, \dots, k_r^P)$
 - 9 Update $(h_1^P, \dots, h_{r-1}^P) \leftarrow (h_1^P, \dots, h_{i-2}^P, h_{\text{pre}}, h_{\text{suc}}, h_i^P, \dots, h_{r-1}^P)$
 - 10 Update (k_1^Q, \dots, k_r^Q) and $(h_1^Q, \dots, h_{r-1}^Q)$ similarly as in Lines 4–9
-

Proposition 3.11 (Setup cost) *The setup step of the quantum walk has query complexity $S = \tilde{O}(r \cdot (\sqrt{d} + T))$.*

Proof We can set up the initial state for the quantum walk by simply performing r insertions successively using Proposition 3.10. \square

Remark 3.12 Observe that, in the insertion procedure in Algorithm 1, Lines 2 and 4–7 can be implemented also in *time complexity* $\tilde{O}(\sqrt{d} + T)$. The time-consuming steps in Algorithm 1 are those that actually modify the data. For example, in Lines 8 and 9, the insertion causes some elements in the array to shift to the right, and would take $O(r)$ time if implemented naively. Later in Sect. 3.2 we will describe appropriate data structures to implement these steps time-efficiently.

Finally, by Theorem 2.7, the query complexity of our quantum walk algorithm (omitting poly log(n) factors) is

$$\begin{aligned} & S + \sqrt{\frac{m^2}{r^2}} \cdot (C + \sqrt{r} \cdot U) \\ &= r \cdot (\sqrt{d} + T) + \frac{m}{r} \cdot (0 + \sqrt{r} \cdot (\sqrt{d} + T)) \\ &= m^{2/3} \cdot (\sqrt{d} + T), \end{aligned} \quad (1)$$

by choosing $r = m^{2/3}$. The construction of good anchor sets from Corollary 3.5 has $m = O(n/\sqrt{d})$, $T = \tilde{O}(1)$, achieving query complexity $\tilde{O}(n^{2/3} \cdot d^{1/6})$. The improved construction from Lemma 3.6 has $m = O(n/d^{3/4})$, $T = \tilde{O}(\sqrt{d})$, achieving query complexity $\tilde{O}(n^{2/3})$.

3.2 Time-Efficient Implementation

In this section, we will show how to implement the $\tilde{O}(n^{2/3})$ -query quantum walk algorithm from Sect. 3.1 in *time complexity* $\tilde{O}(n^{2/3})$.

3.2.1 Overview

Recall that our algorithm described in Sect. 3.1 for input strings s, t and threshold length d performs a quantum walk on the Johnson graph $\binom{[m]}{r}$. In this section, we have to measure the quantum walk costs S, C, U in terms of the *time complexity* instead of query complexity. Inspecting Equation 1, we observe that the quantum walk algorithm can achieve $\tilde{O}(n^{2/3})$ time complexity, as long as we can implement the setup, checking and update steps with time complexities $S = \tilde{O}(r(\sqrt{d} + T))$, $C = \tilde{O}(\sqrt{rd})$, and $U = \tilde{O}(\sqrt{d} + T)$.

As mentioned in Sect. 3.1, there are two parts in the described quantum walk algorithm that are time-assuming:

- Maintaining the arrays of associated data under insertions and deletions ((Remark 3.12).
- Solving the Two String Families LCP problem in the checking step.

Now we give an overview of how we address these two problems.

Dynamic arrays under insertions and deletions A natural solution to speed up the insertions and deletions is to maintain the arrays of using appropriate data structures,

which support the required operations in $\tilde{O}(1)$ time. This “quantum walk plus data structures” framework was first used in Ambainis’ element distinctness algorithm [13], and have been applied to many time-efficient quantum walk algorithms (see the discussion in Sect. 1.3). However, as noticed by Ambainis [13, Section 6.2], such data structures have to satisfy the following requirements in order to be applicable in quantum walk algorithms.

1. The data structure needs to be *history-independent*, that is, the representation of the data structure in memory should only depend on the set of elements stored (and the random coins used) by the data structure, *not* on the sequence of operations leading to this set of elements.
2. The data structure should guarantee *worst-case* time complexity (with high probability over the random coins) per operation.

The first requirement guarantees that each vertex of the Johnson graph corresponds to a unique quantum state, which is necessary since having multiple possible states would destroy the interference during the quantum walk algorithm. This requirement rules out many types of self-balancing binary search trees⁵ such as AVL Tree and Red-Black Tree.

The second requirement rules out data structures with amortized or expected running time, which may take very long time in some of the operations. The reason is that, during the quantum algorithm, each operation is actually applied to a superposition of many instances of the data structure, so we would like the time complexity of an operation to have a fixed upper bound that is independent of the particular instance being operated on.

Ambainis [13] designed a data structure satisfying both requirements based on hash tables and skip lists, which maintains a sorted list of items, and supports insertions, deletions, and searching in $\tilde{O}(1)$ time with high probability. Buhrman, Loff, Patro, and Speelman [55] modified this data structure to also support indexing queries, which ask for the k^{th} item in the current list (see Lemma 3.14 below). Using this data structure to maintain the arrays in our quantum walk algorithm, we can implement the update steps and the setup steps time-efficiently.

Dynamic Two String Families LCP The checking step of our quantum walk algorithm (Proposition 3.9) requires solving an *Two String Families LCP* instance with r string pairs of lengths bounded by d . We will not try to solve this problem from scratch for each instance, since it is not clear how to solve it significantly faster than the $\tilde{O}(r)$ -time classical algorithm [23, Lemma 3] even using quantum algorithms. Instead, we *dynamically maintain* the solution using some data structure, which efficiently handles each update step during the quantum walk where we insert one string pair (P, Q) into (and remove one from) the current Two String Families LCP instance. As mentioned in Sect. 1.2.1, the classical data structure for this task given by Charalampopoulos, Gawrychowski, and Pokorski [27] is not applicable here, since it violates both requirements mentioned above: it maintains a heavy-light decomposition of the compact tries of the input strings, and rebuilds them from time to time to ensure amortized poly log(n)

⁵ One exception is Treap.

time complexity. It is not clear how to implement this strategy in a history-independent way and with worst-case time complexity per operation.

Instead, we will design a different data structure that satisfies the history-independence and worst-case update time requirements, and can solve the Two String Families LCP problem on the maintained instance in $\tilde{O}(\sqrt{rd})$ quantum time. This time complexity is much worse than the $\text{poly} \log(n)$ time achieved by the classical data structure of [27], but is sufficient for our purpose. As mentioned in Sect. 1.2.1, one challenge is the lack of a *comparison-based* data structure for 2D range query that also satisfies the two requirements above. We remark that there exist comparison-based data structures with history-independence but only with expected time complexity (e.g., [92]). There also exist folklore data structures for *integer coordinates* that have history-independence and worst-case time complexity (e.g., Lemma 3.15). For the easier problem of 1-dimensional range query, there exist folklore data structures (e.g., Lemma 3.14) that satisfy all three requirements. To get around this issue, we will use a sampling procedure and a version of the Balls-and-Bins argument, which can effectively convert the involved non-integer coordinates into integer coordinates. Then, we are able to apply 2D range query data structures over integer coordinates. Details will be given in Sect. 3.2.3.

3.2.2 Basic Data Structures

In this section, we will review several existing constructions of classical history-independent data structures.

Let D be a classical data structure using $\tilde{O}(1)$ many random coins r that maintains a dynamically changing data set S . We say D is *history-independent* if for each possible S and r , the data structure has a unique representation $D(S, r)$ in the memory. Furthermore, we say D has *worst-case update time $O(T)$ with high probability*, if for every possible S and update operation $S \rightarrow S'$, with high probability over r , the time complexity to update from $D(S, r)$ to $D(S', r)$ is $O(T)$. Similarly we can define worst-case query time with high probability.

Since our quantum walk algorithm is over the Johnson graph $\binom{[m]}{r}$, for consistency we will use r to denote the size of the data structure instances in the following statements.

Hash tables We use hash tables to implement efficient lookup operations without using too much memory.

Lemma 3.13 (Hash tables) *There is a history-independent data structure of size $\tilde{O}(r)$ that maintains a set of at most r key-value pairs $\{(\text{key}_1, \text{value}_1), (\text{key}_2, \text{value}_2), \dots, (\text{key}_r, \text{value}_r)\}$ where key_i 's are distinct integers from $[m]$, and supports the following operations in worst-case $\tilde{O}(1)$ time with high probability:*

- *Lookup* Given a key $\in [m]$, find the value corresponding to key (or report that key is not present in the set).
- *Insertion* Insert a key-value pair into the set.
- *Deletion* Delete a key-value pair from the set.

Proof (Sketch) The construction is similar to [13, Section 6.2]. The hash table has r buckets, each with the capacity for storing $O(\log m)$ many key-value pairs. A pair (key, value) is stored in the $(h(\text{key}))^{\text{th}}$ bucket, and the pairs inside each bucket are sorted in increasing order of keys. If some buckets overflow, we can collect all the leftover pairs into a separate buffer of size r and store them in sorted order. This ensures that any set of r key-value pairs has a unique representation in the memory. And, each basic operation can be implemented in $\text{poly} \log(m)$ time, unless there is an overflow. Using an $O(\log m)$ -wise independent hash function $h: [m] \rightarrow [r]$, for any possible r -subset of keys, with high probability none of the buckets overflow.⁶ \square

Dynamic arrays We will need a dynamic array that supports indexing, insertion, deletion, and some other operations.

The *skip list* [93] is a probabilistic data structure which is usually used as an alternative to balanced trees, and satisfies the history-independence property. Ambainis' quantum Element Distinctness algorithm [13] used the skip list to maintain a *sorted* array, supporting insertions, deletions, and binary search. In order to apply the skip list in the quantum walk, a crucial adaptation in Ambainis' construction is to show that the random choices made by the skip list can be simulated using $O(\log n)$ -wise independent functions [13, Section 6.2], which only take $\text{poly} \log(n)$ random coins to sample. In the recent quantum fine-grained reduction result by Buhrman, Loff, Patro, and Speelman [55, Section 3.2], they used a more powerful version of skip lists that supports *efficient indexing*. We will use this version of skip lists with some slight extension.

Lemma 3.14 (Dynamic arrays) *There is a history-independent data structure of size $\tilde{O}(r)$ that maintains an array of items $(\text{key}_1, \text{value}_1), (\text{key}_2, \text{value}_2), \dots, (\text{key}_r, \text{value}_r)$ with distinct keys (note that neither the keys nor the values are necessarily sorted in increasing order), and supports the following operations with worst-case $\tilde{O}(1)$ time complexity and high success probability:*

- *Indexing* Given an index $1 \leq i \leq r$, return the i^{th} item $(\text{key}_i, \text{value}_i)$
- *Insertion* Given an index $1 \leq i \leq r + 1$ and a new item, insert it into the array between the $(i - 1)^{\text{st}}$ item and the i^{th} item (shifting later items to the right).
- *Deletion* Given an index $1 \leq i \leq r$, delete the i^{th} item from the array (shifting later items to the left).
- *Location* Given a key, return its position i in the array (i.e., $\text{key}_i = \text{key}$).
- *Range-minimum query* Given $1 \leq a \leq b \leq r$, return $\min_{a \leq i \leq b} \{\text{value}_i\}$.

Proof (Sketch) We will use (a slightly modified version of) the data structure described in [55, Section 3.2], which extends the construction of [13, Section 6.2] to support *insertion, deletion, and indexing*. Their construction is a (bidirectional) skip list of the items, where a pointer (a “skip”) from an item (key, value) to another item

⁶ We remark that Ambainis only used a fixed hash function $h(i) = \lfloor r \cdot i/m \rfloor$, which ensures the buckets do not overflow with high probability over a random r -subset $K \subseteq [m]$ of keys. Ambainis showed that this property is already sufficient for the correctness of the quantum walk algorithm. Here we choose to state a different version that achieves high success probability for every fixed r -subset of keys, merely for keeping consistency with later presentation.

$(key', value')$ is stored in a hash table as a key-value pair (key, key') . To support efficient indexing, for each pointer they also store the *distance* of this skip, which is used during an indexing query to keep track of the current position after following the pointers (similar ideas were also used in, e.g., [94, Section 3.4]). After every insertion or deletion, the affected distance values are updated recursively, by decomposing a level- i skip into $O(\log n)$ many level- $(i - 1)$ skips.

A difference between their setting and ours is that they always keep the array sorted in increasing order of value's, and the position of an inserted item is decided by its relative order among the values in the array, instead of by a given position $1 \leq i \leq r+1$. Nevertheless, it is straightforward to adapt their construction to our setting, by using the distance values of the skips to keep track of the current position, instead of by comparing the values of items.

Note that using the distance values we can also efficiently implement the *Location* operation in a reversed way compared to *Indexing*, by following the pointers backwards and moving up levels.

To implement the *range-minimum query* operations, we maintain the range-minimum value of each skip in the skip list, in a similar way to maintaining the distance values of the skips. They can also be updated recursively after each update. Then, to answer a query, we can travel from the a^{th} item to the b^{th} by following the pointers (this is slightly trickier if $a \neq 1$, where we may first move up levels and then move down). \square

We also need a 2D range sum data structure for points with integer coordinates.

Lemma 3.15 (2D range sum) *Let integer $N \leq n^{O(1)}$. There is a history-independent data structure of size $\tilde{O}(r)$ that maintains a multi-set of at most r points $\{(x_1, y_1), \dots, (x_r, y_r)\}$ with integer coordinates $x_i \in [N]$, $y_i \in [N]$, and supports the following operations with worst-case $\tilde{O}(1)$ time complexity and high success probability:*

- *Insertion* Add a new point (x, y) into the multiset (duplicates are allowed).
- *Deletion* Delete the point (x, y) from the multiset (if it appears more than once, only delete one copy of them).
- *Range sum* Given $1 \leq x_1 \leq x_2 \leq N$, $1 \leq y_1 \leq y_2 \leq N$, return the number of points (x, y) in the multiset that are in the rectangle $[x_1 \dots x_2] \times [y_1 \dots y_2]$.

Proof (Sketch) Without loss of generality, assume N is a power of two. We use a simple folklore construction that resembles a 2D segment tree (sometimes called 2D range tree or 2D radix tree). Define a class $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \dots \cup \mathcal{C}_{\log N}$ of sub-segments of the segment $[1 \dots N]$ as follows:

$$\begin{aligned} \mathcal{C}_1 &= \{[1 \dots N]\}, \\ \mathcal{C}_2 &= \{[1 \dots N/2], [N/2 + 1 \dots N]\}, \\ \mathcal{C}_3 &= \{[1 \dots N/4], [N/4 + 1 \dots 2N/4], [2N/4 + 1 \dots 3N/4], [3N/4 + 1 \dots N]\}, \\ &\dots \\ \mathcal{C}_{\log N} &= \{[1 \dots 1], [2 \dots 2], \dots, [N \dots N]\}. \end{aligned}$$

Then it is not hard to see that every segment $[a \dots b] \subseteq [1 \dots N]$ can be represented as the disjoint union of at most $2 \log N$ segments in \mathcal{C} . Consequently, the query rectangle $[x_1 \dots x_2] \times [y_1 \dots y_2]$ can always be represented as the disjoint union of $O(\log^2 N)$ rectangles of the form $\mathcal{I} \times \mathcal{J}$ where $\mathcal{I}, \mathcal{J} \in \mathcal{C}$.

Hence, for every $\mathcal{I}, \mathcal{J} \in \mathcal{C}$ with *non-zero* range sum $s(\mathcal{I} \times \mathcal{J})$, we store this range sum into a hash table, indexed by the canonical encoding of $(\mathcal{I}, \mathcal{J})$. Then we can efficiently answer all the range-sum queries by decomposing the rectangles and summing up their stored range sums.

When a point (x, y) is updated, we only need to update the range sums of $\log^2 N$ many rectangles that are affected, since each $a \in [1 \dots N]$ is only included by $\log N$ intervals in \mathcal{C} . We may also need to insert a new rectangle into the hash table, or remove a rectangle once its range sum becomes zero. \square

Data structures in quantum walk Ambainis [13] showed that a history-independent classical data structure D with worst-case time complexity T (with high probability over the random coins r) can be applied to the quantum walk framework by creating a uniform superposition over all possible r , i.e., the data structure storing data S corresponds to the quantum state $\sum_r |D(S, r)\rangle |r\rangle$. During the quantum walk algorithm, each data structure operation is aborted after running for T time steps. By doing this, some components in the quantum state may correspond to malfunctioning data structures, but Ambainis showed that this will not significantly affect the behavior of the quantum walk algorithm. We do not repeat the error analysis here, but instead refer interested readers to the proof of [13, Lemma 5 and 6] (see also [55, Lemma 1 and 2]).

3.2.3 Applying the Data Structures

Now we will use the data structures described in Sect. 3.2.2 to implement our quantum walk algorithm from Sect. 3.1 time-efficiently.

Recall that C is the T -quantum-time-constructible good anchor set of size $|C| = m$ (Definition 3.2). The states of our quantum walk algorithms are r -subsets $K = \{k_1, k_2, \dots, k_r\} \subseteq [m]$, where each index $k \in K$ is associated with an anchor $C(k) \in [n]$, which specifies the color (red or blue) of k and the pair $(P(k), Q(k))$ of strings of lengths at most d . We need to maintain the lexicographical orderings (k_1^P, \dots, k_r^P) and LCP arrays $(h_1^P, \dots, h_{r-1}^P)$, so that $P(k_1^P) \leq P(k_2^P) \leq \dots \leq P(k_r^P)$ and $h_i^P = \text{lcp}(P(k_i^P), P(k_{i+1}^P))$, and similarly maintain (k_1^Q, \dots, k_r^Q) , $(h_1^Q, \dots, h_{r-1}^Q)$ for the strings $\{Q(k)\}_{k \in K}$.

For $k \in K$, we use $\text{pos}^P(k)$ to denote the position i such that $k_i^P = k$, i.e., the lexicographical rank of $P(k)$ among all $P(k_1), \dots, P(k_r)$. Similarly, let $\text{pos}^Q(k)$ denote the position i such that $k_i^Q = k$.

We can immediately see that all the steps in the *update step* (Algorithm 1) of our quantum walk can be implemented time-efficiently. In particular, we use a hash table (Lemma 3.13) to store the anchor $C(k)$ corresponding to each $k \in K$, and use Lemma 3.14 to maintain the lexicographical orderings and LCP arrays under insertions and deletions. Each update operation on these data structures takes $\tilde{O}(1)$ time. Additionally, these data structures allow us to efficiently compute some useful information, as summarized below.

Proposition 3.16 *Given indices $k, k' \in K$, the following information can be computed in $\tilde{O}(1)$ time.*

1. The anchor $C(k)$, the color of k , and the lengths $|P(k)|, |Q(k)| \leq d$.
2. $\text{pos}^P(k)$ and $\text{pos}^Q(k)$.
3. $\text{lcp}(P(k), P(k'))$ and $\text{lcp}(Q(k), Q(k'))$.

Proof For 1, rather than use T time to compute $C(k)$ (Definition 3.2), we instead look up the value of $C(k)$ from the hash table. Then, $C(k) \in [n]$ determines the color of k and the string lengths.

For 2, we use the location operation of the dynamic array data structure (Lemma 3.14).

For 3, we first compute $i = \text{pos}^P(k), i' = \text{pos}^P(k')$, and assume $i < i'$ without loss of generality. Then, by Lemma 2.1, we can compute $\text{lcp}(P(k), P(k')) = \text{lcp}(P(k_i^P), P(k_{i'}^P)) = \min\{h_i^P, h_{i+1}^P, \dots, h_{i'-1}^P\}$ using a range-minimum query (Lemma 3.14). \square

The remaining task is to efficiently implement the checking step, where we need to solve the Two String Families LCP problem. The goal is to find a red index $k^{\text{red}} \in K$ and a blue index $k^{\text{blue}} \in K$, such that $\text{lcp}(P(k^{\text{red}}), P(k^{\text{blue}})) + \text{lcp}(Q(k^{\text{red}}), Q(k^{\text{blue}})) \geq d$. Now we give an outline of the algorithm for solving this task.

Algorithm 2: Solving the Two String Families LCP problem in the checking step

```

1 Grover-Search over red indices  $k^{\text{red}} \in K$ , and integers  $d' \in [0..d]$ 
2   Find  $\ell^P, r^P$  such that  $\text{lcp}(P(k_i^P), P(k^{\text{red}})) \geq d'$  if and only if  $\ell^P \leq i \leq r^P$ .
3   Find  $\ell^Q, r^Q$  such that  $\text{lcp}(Q(k_i^Q), Q(k^{\text{red}})) \geq d - d'$  if and only if  $\ell^Q \leq i \leq r^Q$ .
4   if exists a blue index  $k^{\text{blue}} \in K$  such that  $\text{pos}^P(k^{\text{blue}}) \in [\ell^P..r^P], \text{pos}^Q(k^{\text{blue}}) \in [\ell^Q..r^Q]$ 
      then return True
5 return False

```

In the Algorithm 2, we use Grover search to find a red index $k^{\text{red}} \in K$ and an integer $d' \in [0..d]$, such that there exists a blue index $k^{\text{blue}} \in K$ with $\text{lcp}(P(k^{\text{red}}), P(k^{\text{blue}})) \geq d'$ and $\text{lcp}(Q(k^{\text{red}}), Q(k^{\text{blue}})) \geq d - d'$. The number of Grover iterations is $\tilde{O}(\sqrt{|K| \cdot d}) = \tilde{O}(\sqrt{rd})$, and we will implement each iteration in poly log(n) time. By Lemma 2.1, all the strings $P(k)$ that satisfy $\text{lcp}(P(k), P(k^{\text{red}})) \geq d'$ form a *contiguous segment* in the lexicographical ordering $P(k_1^P) \preceq \dots \preceq P(k_r^P)$. In Line 2, we find the left and right boundaries ℓ^P, r^P of this segment, using a binary search with Proposition 3.16 (3). Line 3 is similar to Line 2. Then, Line 4 checks the existence of such a blue string pair. It is clear that this procedure correctly solves the Two String Families LCP problem. The only remaining problem is how to implement Line 4 efficiently.

Note that Line 4 can be viewed as a 2D orthogonal range query, where each 2D point is a blue string pair $(P(k), Q(k))$, with coordinates being strings to be compared in lexicographical order. We cannot simply replace the coordinates by their ranks

$\text{pos}^P(k)$ and $\text{pos}^Q(k)$ among the r substrings in the current state, since their ranks will change over time. It is also unrealistic to replace the coordinates by their ranks among all the possible substrings $\{P(k)\}_{k \in [m]}$, since m could be much larger than the desired overall time complexity $n^{2/3}$. These issues seem to require our 2D range query data structure to be comparison-based, which is also difficult to achieve as mentioned before.

Instead, we will use a sampling technique, which effectively converts the non-integer coordinates into integer coordinates. At the very beginning of the algorithm (before running the quantum walk), we uniformly sample r distinct indices $x_1, x_2, \dots, x_r \in [m]$, and sort them so that $P(x_1) \leq P(x_2) \leq \dots \leq P(x_r)$ (breaking ties by the indices), in $\tilde{O}(r(\sqrt{d} + T))$ total time (this complexity is absorbed by the time complexity of the setup step $S = O(r(\sqrt{d} + T))$). Then, during the quantum walk algorithm, when we insert an index $k \in [m]$ into K , we assign it an *integer label* $\rho^P(k)$ defined as the unique $i \in [0..r]$ satisfying $P(x_i) \leq s' < P(x_{i+1})$, which can be computed in $\tilde{O}(\sqrt{d})$ time by a binary search on the sorted sequence $P(x_1) \leq \dots \leq P(x_r)$. We also sample $y_1, \dots, y_r \in [m]$ and sort them so that $Q(y_1) \leq Q(y_2) \leq \dots \leq Q(y_r)$, and similarly define the integer labels $\rho^Q(k)$. Intuitively, the (scaled) label $\rho^P(k) \cdot (m/r)$ estimates the rank of $P(k)$ among all the strings $\{P(k')\}_{k' \in [m]}$.

The following lemma formalizes this intuition. It states that in a typical r -subset $K = \{k_1, k_2, \dots, k_r\} \subseteq [m]$, not too many indices can receive the same label.

Lemma 3.17 *For any $c > 1$, there is a $c' > 1$, such that the following statement holds:*

For positive integers $r \leq m$, let $A, B \subseteq [m]$ be two independently uniformly random r -subsets. Let $A = \{a_1, a_2, \dots, a_r\}$ where $a_1 < a_2 < \dots < a_r$, and denote

$$A_0 := [1..a_1), A_1 := [a_1..a_2), \dots, A_{r-1} := [a_{r-1}..a_r), A_r := [a_r..m].$$

Then,

$$\Pr_{A,B} [|A_i \cap B| \geq c' \log m \text{ for some } 0 \leq i \leq r] \leq \frac{1}{m^c}.$$

Proof Let $k = c' \log m$ for some $c' > 1$ to be determined later, and we can assume $k \leq r$. Observe that, $|A_i \cap B| \geq k$ holds for some i only if there exist $b, b' \in [m]$, such that $|[b..b'] \cap B| \geq k$ and $[b+1..b'] \cap A = \emptyset$.

Let $b, b' \in [m]$, $b \leq b'$. For $b' - b \geq (c+2)(m \ln m)/r$, we have

$$\begin{aligned} \Pr_A [[b+1..b'] \cap A = \emptyset] &= \frac{\binom{m-(b'-b)}{r}}{\binom{m}{r}} \\ &\leq \left(1 - \frac{b'-b}{m}\right)^r \\ &\leq 1/m^{c+2}. \end{aligned}$$

For $b' - b < (c + 2)(m \ln m)/r$, we have

$$\begin{aligned} \Pr_B \left[|[b \dots b'] \cap B| \geq k \right] &\leq \frac{\binom{b'-b+1}{k} \cdot \binom{m-k}{r-k}}{\binom{m}{r}} \\ &= \binom{b'-b+1}{k} \cdot \frac{\binom{r}{k}}{\binom{m}{k}} \\ &\leq \left(\frac{e(b'-b+1)}{k} \right)^k \cdot \left(\frac{r}{m} \right)^k \\ &< \left(\frac{e(c+3) \ln m}{k} \right)^k \\ &< 1/m^{c+3}, \end{aligned}$$

where we set $k = c' \log m = 3(c + 3) \log m$.

The proof then follows from a union bound over all pairs of $b, b' \in [m]$. \square

Then, we can use the 2D point $(\rho^P(k), \rho^Q(k))$ with integer coordinates to represent the string pair $(P(k), Q(k))$, and use the data structure from Lemma 3.15 to handle the 2D range sum queries. To correctly handle the points near the boundary of a query, we need to check them one by one, and Lemma 3.17 implies that in average case this brute force step is not expensive.

The pseudocode in Algorithm 3 describes the additional steps to be performed during each insertion step of the quantum walk (the deletion step is simply the reversed operation of the insertion step).

Algorithm 3: Extra steps in the insertion procedure (in addition to the steps in Algorithm 1)

- 1 Given an index $k \in [m]$
 - 2 Compute the integer labels $\rho^P(k)$ and $\rho^Q(k)$ using binary search, and store them in hash table
 - 3 **if** k is blue **then**
 - 4 Insert the 2D point $(\rho^P(k), \rho^Q(k))$ into the 2D range sum data structure
-

The pseudocode in Algorithm 4 describes how to implement Line 4 in Algorithm 2 for solving the Two String Families LCP problem. Line 4 correctly handles all the “internal” blue pairs $(P(k^{\text{blue}}), Q(k^{\text{blue}}))$, which must satisfy $\text{pos}^P(k^{\text{blue}}) \in [\ell^P, r^P]$ and $\text{pos}^Q(k^{\text{blue}}) \in [\ell^Q, r^Q]$ by the definition of our integer labels $\rho^P(\cdot), \rho^Q(\cdot)$ and Lines 2 and 3. In Line 4 we handle the remaining possible blue pairs, which must have $\rho^P(k^{\text{blue}}) \in \{\tilde{\ell}^P, \tilde{r}^P\}$ or $\rho^Q(k^{\text{blue}}) \in \{\tilde{\ell}^Q, \tilde{r}^Q\}$, and can be found by binary searches on the lexicographical orderings (to be able to do this, we need to maintain the lexicographical orderings of $P(k_1), \dots, P(k_r)$ and the sampled strings $P(x_1), \dots, P(x_r)$ combined).

Note that in Line 5 of Algorithm 4 we abort if we have checked more than $4c' \log m$ boundary points, so that Algorithm 4 has worst-case $\tilde{O}(1)$ overall running time. But

Algorithm 4: Implementation of Line 4 in Algorithm 2

```

1  Given indices  $\ell^P, r^P, \ell^Q, r^Q$ .
2  Let  $\tilde{\ell}^P := \rho^P(k_{\ell^P}^P), \tilde{r}^P := \rho^P(k_{r^P}^P)$ .
3  Let  $\tilde{\ell}^Q := \rho^Q(k_{\ell^Q}^Q), \tilde{r}^Q := \rho^Q(k_{r^Q}^Q)$ .
4  if the 2D range sum of  $[\tilde{\ell}^P + 1 \dots \tilde{r}^P - 1] \times [\tilde{\ell}^Q + 1 \dots \tilde{r}^Q - 1]$  is non-zero then return True for
    blue index  $k^{\text{blue}} \in K$  such that  $\rho^P(k^{\text{blue}}) \in \{\tilde{\ell}^P, \tilde{r}^P\}$  or  $\rho^Q(k^{\text{blue}}) \in \{\tilde{\ell}^Q, \tilde{r}^Q\}$  do
5  |   if  $\text{pos}^P(k^{\text{blue}}) \in [\ell^P \dots r^P], \text{pos}^Q(k^{\text{blue}}) \in [\ell^Q \dots r^Q]$  then return True if already looped
    |    $4c' \log m$  times then exit for loop
6  return False
    
```

this early stopping would also introduce (one-sided) error if there are too many boundary points which we have no time to check. However, a straightforward application of Lemma 3.17 implies that, with high success probability over the initial samples $P(x_1) \leq P(x_2) \leq \dots \leq P(x_r)$ and $Q(y_1) \leq Q(y_2) \leq \dots \leq Q(y_r)$, only $1/\text{poly}(m)$ fraction of the r -subsets $K = \{k_1, \dots, k_r\} \in [m]$ in the Johnson graph can have more than $c' \log m$ strings receiving the same label. On these problematic states $K = \{k_1, \dots, k_r\} \in [m]$, the checking procedure may erroneously recognize K as unmarked, while other states are handled correctly by Algorithm 4 since there is no early aborting. This decreases the fraction of marked states in the Johnson graph by only a $1/\text{poly}(m)$ fraction, which does not affect the overall time complexity of our quantum walk algorithm.

3.3 Improved Construction of Good Anchor Sets

In this section, we will prove Lemma 3.6 by constructing a good anchor set with smaller size. Our construction of good anchor sets is based on a careful combination of a generalized version of *difference covers* [29, 30] and the *string synchronizing sets* [32].

3.3.1 Approximate Difference Covers

We first need to generalize the notion of difference covers.

Definition 3.18 (*Approximate Difference Covers*) A set $D \subseteq \mathbb{N}^+$ is called a (d, τ) -cover, if for every $i, j \in \mathbb{N}^+$, there exists two integers $h_1(i, j), h_2(i, j) \in [0 \dots d]$ such that $i + h_1(i, j), j + h_2(i, j) \in D$, and $|h_1(i, j) - h_2(i, j)| \leq \tau - 1$.

The notion of d -cover (Definition 3.3) used in previous algorithms corresponds to the $\tau = 1$ case of our new definition. Our generalization to larger τ can be viewed as an approximate version of difference covers with additive error $\leq \tau - 1$. As we shall see, allowing additive error makes the size of the (d, τ) -cover much smaller compared to Definition 3.3.

We present a construction of approximate difference covers, by adapting previous constructions from $\tau = 1$ to general values of τ .

Lemma 3.19 (Construction of (d, τ) -cover) *For every positive integers $1 \leq \tau \leq d$, there is a (d, τ) -cover D such that $D \cap [n]$ contains $O(n/\sqrt{d\tau})$ elements. Moreover, given integer $i \geq 1$, one can compute the i^{th} smallest element of $D \cap [n]$ in $\tilde{O}(1)$ time.*

Proof Let $M := \lfloor \sqrt{d/\tau} \rfloor \geq 1$. Define

$$I := \{zM \cdot \tau \mid z \in \mathbb{N}^+\},$$

and

$$J := \{(xM^2 - y) \cdot \tau \mid x \in \mathbb{N}^+, y \in [M]\}.$$

We claim that $D := I \cup J$ is a (d, τ) -cover that satisfies the desired properties.

First, observe that $|I \cap [n]| = \lfloor n/(M\tau) \rfloor \leq O(n/\sqrt{d\tau})$, and $|J \cap [n]| \leq \lfloor n/\tau \rfloor \cdot (M/M^2) = O(n/\sqrt{d\tau})$. Hence $D = I \cup J$ satisfies the claimed size bound.

Next, we verify D is indeed a (d, τ) -cover. For any $i, j \in \mathbb{N}^+$, let $i' := \lceil i/\tau \rceil$, $j' := \lceil j/\tau \rceil$. Let $z \cdot \tau \in J$ be the smallest integer in J such that $z \geq j'$ and $z \equiv j' - i' \pmod{M}$. By the construction of J , we have $z \leq j' + M^2 - 1$. Hence, let $h_1(i, j) = (z - j' + i') \cdot \tau - i$ and $h_2(i, j) = z \cdot \tau - j$. Note that

$$\begin{aligned} h_2(i, j) &\leq (j' + M^2 - 1) \cdot \tau - j \leq M^2\tau - 1 \leq d - 1, \\ h_2(i, j) &\geq j' \cdot \tau - j \geq 0, \end{aligned}$$

where we used $j' \cdot \tau - j \in [0 \dots \tau - 1]$. Similarly we can show $0 \leq h_1(i, j) \leq d - 1$, and we have

$$|h_2(i, j) - h_1(i, j)| = |(j' \cdot \tau - j) - (i' \cdot \tau - i)| \leq \tau - 1.$$

Moreover, $j + h_2(i, j) \in J \subseteq D$, and

$$i + h_1(i, j) = (z - j' + i') \cdot \tau \equiv 0 \pmod{M\tau}$$

which implies $i + h_1(i, j) \in I \subseteq D$. □

3.3.2 String Synchronizing

In Corollary 3.5 we obtained a good anchor set using a $(d, 1)$ -cover. If we simply replace it by a (d, τ) -cover with larger τ , the size of the obtained anchor set would become smaller, but it would no longer be a good anchor set, due to the misalignment introduced by approximate difference covers. To deal with the misalignment, we will use the *string synchronizing sets* recently introduced by Kempa and Kociumaka [32]. Informally, a synchronizing set of string S is a small set of synchronizing positions, such that every two sufficiently long matching substrings of S with no short periods should contain a pair of consistent synchronizing positions.

Definition 3.20 (*String synchronizing sets* [32, Definition 3.1]) For a string $S[1..n]$ and a positive integer $1 \leq \tau \leq n/2$, we say $A \subseteq [1..n - 2\tau + 1]$ is a τ -synchronizing set of S if it satisfies the following properties:

- **Consistency** If $S[i..i + 2\tau] = S[j..j + 2\tau]$, then $i \in A$ if and only if $j \in A$.
- **Density** For $i \in [1..n - 3\tau + 2]$, $A \cap [i..i + \tau] = \emptyset$ if and only if $\text{per}(S[i..i + 3\tau - 2]) \leq \tau/3$.

Kempa and Kociumaka gave a linear-time classical randomized algorithm (as well as a derandomized version, which we will not use here) to construct a τ -synchronizing set A of optimal size⁷ $|A| = O(n/\tau)$. However, this classical algorithm for constructing A has to query each of the n input characters, and is not directly applicable to our sublinear quantum algorithm.

To apply Kempa and Kociumaka's construction algorithm to the quantum setting, we observe that this algorithm is *local*, in the sense that whether an index i should be included in A is completely decided by its short context $S[i..i + 2\tau]$ and the random coins. Moreover, by suitable adaptation of their construction, one can compute all the synchronizing positions in an $O(\tau)$ -length interval in $\tilde{O}(\tau)$ time. We summarize all the desired properties of the synchronizing set in the following lemma.

Lemma 3.21 (Adaptation of [32]) *For a string $S[1..n]$ and a positive integer $1 \leq \tau \leq n/2$, given a sequence r of $O(\log n)$ many random coins, there exists a set A with the following properties:*

- **Correctness** With high probability over r , A is a τ -synchronizing set of S .
- **Locality** For every $i \in [1..n - 2\tau + 1]$, whether $i \in A$ or not is completely determined by the random coins r and the substring $s[i..i + 2\tau]$.
Moreover, given $s[i..i + 4\tau]$ and r , one can compute all the elements in $A \cap [i..i + 2\tau]$ by a classical algorithm in $\tilde{O}(\tau)$ time.
- **Sparsity** For every $i \in [1..n - 2\tau + 1]$, $\mathbb{E}_r[|A \cap [i..i + 2\tau]|] \leq 80$

In the following, we first inspect the (slightly adapted) randomized construction of string synchronized sets by Kempa and Kociumaka [32], and then show that it satisfies the properties in Lemma 3.21.

Construction of string synchronizing sets Fix an input string $S[1..n]$ and a positive integer $\tau \leq n/2$. Define sets

$$\begin{aligned} Q &= \{i \in [1..n - \tau + 1] : \text{per}(S[i..i + \tau]) \leq \tau/3\}, \\ B &= \{i \in [1..n - \tau + 1] \setminus Q : \text{per}(S[i..i + \tau - 1]) \leq \tau/3 \\ &\quad \vee \text{per}(S[i + 1..i + \tau]) \leq \tau/3\}, \end{aligned}$$

where we define $B = \emptyset$ in the special case of $\tau = 1$.

Let $\mathcal{P} = \{s \in \Sigma^\tau : s \text{ is a substring of } S\}$ denote the set of all the length- τ substrings in S (without duplicates). Let $\pi : \mathcal{P} \rightarrow [N]$ be any injection, and define the identifier

⁷ In the case where S has no highly periodic substrings, every τ -length interval should contain at least one index from A .

function $\text{id}: [1 \dots n - \tau + 1] \rightarrow \mathbb{N}^+$ by

$$\text{id}(i) := \begin{cases} \pi(S[i \dots i + \tau]) & i \in B, \\ \pi(S[i \dots i + \tau]) + N & i \notin B. \end{cases}$$

In this way, we have $\text{id}(i) = \text{id}(j)$ if and only if $S[i \dots i + \tau] = S[j \dots j + \tau]$. Moreover, for $i \in B$, $j \notin B$, we always have $\text{id}(i) < \text{id}(j)$. Finally, define

$$A = \{i \in [1 \dots n - 2\tau + 1] : \min\{\text{id}(j) : j \in [i \dots i + \tau] \setminus Q\} \in \{\text{id}(i), \text{id}(i + \tau)\}\}.$$

Kempa and Kociumaka proved the following fact.

Lemma 3.22 ([32, Lemma 8.2]) *The set A is always a τ -synchronizing set of string S .*

We first quickly verify the locality property of this construction.

Proposition 3.23 *For every $i \in [1 \dots n - 2\tau + 1]$, whether $i \in A$ or not is completely determined by π and the substring $s[i \dots i + 2\tau]$.*

Proof This immediately follows from the definition of Q , B , id , and A . □

Now, suppose $\pi: \mathcal{P} \rightarrow [N]$ is randomly chosen so that the 0.1-approximate min-wise independence property is satisfied: for any $x \in \mathcal{P}$ and subset $X \subseteq \mathcal{P} \setminus \{x\}$,

$$\Pr_{\pi} [\pi(x) < \min\{\pi(x') : x' \in X\}] \in \frac{1}{|X| + 1} \cdot (1 \pm 0.1).$$

Then the following holds.

Lemma 3.24 ([32, Fact 8.9], adapted) *The expected size of A satisfies $\mathbf{E}_{\pi}[|A|] \leq 20n/\tau$.*

Remark 3.25 We remark that in the original construction of [32], π was chosen to be a uniformly random bijection $\mathcal{P} \rightarrow [|\mathcal{P}|]$, and this is *the only part that differs from our modified version*. The main issue with this ideal choice is that, in our quantum algorithm, we do not have enough time to sample and store π , which could have size $\Omega(n)$. Observe that in their proof of Lemma 3.24, the only required property of π is that π satisfies (perfect) min-wise independence. Hence, here we can relax it to have approximate min-wise independence, and their proof of Lemma 3.24 still applies (with a slightly worse constant factor).

Now we describe how to design such a mapping π that is efficiently computable. First, we hash the substrings into integers using the standard rolling hash method [2]. Recall that the alphabet Σ is identified with the integer set $[\Sigma]$.

Definition 3.26 [Rolling hash] Let $p > |\Sigma|$ be a prime, and pick $y \in \mathbb{F}_p$ uniformly at random. Then, the rolling hash function $\rho_{p,y}: \Sigma^\tau \rightarrow \mathbb{F}_p$ on length- τ strings is defined as

$$\rho_{p,x}(s[1.. \tau]) := \sum_{i=1}^{\tau} s[i] \cdot y^i \pmod{p}.$$

We have two the following two folklore facts about rolling hash.

- Rolling hashes of substrings can be easily computed in batch: on any given string s of length $|s| \geq \tau$, one can compute the hash values $\rho_{p,y}(s[i..i+\tau])$ for all $i \in [1..|s|-\tau+1]$, in $O(|s| \cdot \text{poly} \log p)$ total time.
- By choosing $p = \text{poly}(n)$, we can ensure that with high probability over the choice of y , the rolling hash function $\rho_{p,y}$ takes distinct values over all the strings in \mathcal{P} (by Schwartz-Zippel lemma).

After hashing the strings in \mathcal{P} to a small integer set $[\text{poly}(n)]$, we can apply known constructions of approximate min-wise independent hash families.

Lemma 3.27 (Approximate min-wise independent hash family, e.g., [95]) *Given parameter $n \geq 1$, one can choose $N \leq n^{O(1)}$, so that there is a hash family $\mathcal{H} = \{h: [N] \rightarrow [N]\}$ that satisfies the following properties:*

- *Injectivity* For any subset $X \subseteq [N]$ of size $|X| \leq n$, with high probability over the choice of $h \in \mathcal{H}$, h maps X to distinct elements.
- *Approximate min-wise independence* For any $x \in [N]$ and subset $X \subseteq [N] \setminus \{x\}$,

$$\Pr_{h \in \mathcal{H}} [h(x) < \min\{h(x') : x' \in X\}] \in \frac{1}{|X|+1} \cdot (1 \pm 0.1).$$

- *Explicitness* Each hash function $h \in \mathcal{H}$ can be specified using $O(\log n)$ bits, and can be evaluated at any point in $\text{poly} \log(n)$ time.

Finally, we choose parameters $p = \text{poly}(n)$, $N = \text{poly}(n)$, $p \leq N$, and define the pseudorandom mapping $\pi: \mathcal{P} \rightarrow [N]$ by $\pi(s) := h(\rho_{p,y}(s))$, where $\rho_{p,y}: \mathcal{P} \rightarrow \mathbb{F}_p$ is the rolling hash function (identifying \mathbb{F}_p with $[p] \subseteq [N]$), and $h: [N] \rightarrow [N]$ is the approximate min-wise independent hash function.

Now we are ready to prove that the string synchronizing set A determined by the random mapping π satisfies the properties stated in Lemma 3.21.

Proof (of Lemma 3.21) First note that the random coins r are used to sample $y \in \mathbb{F}_p$ and $h \in \mathcal{H}$, which only take $O(\log n)$ bits of seed.

Correctness By Lemma 3.22, A is correct as long as π is an injection, which holds with high probability by the injectivity properties of $\rho_{p,y}$ and h .

Locality The first part of the statement is already verified in Proposition 3.23. To show the moreover part, first note that the values of $\pi(S[j..j+\tau])$ over all $j \in [i..i+3\tau)$ can be computed in $\tilde{O}(\tau)$ time, by the property of rolling hash and the explicitness of h . By [32, Lemma 8.8], the sets $Q \cap [i..i+3\tau)$ and $B \cap [i..i+3\tau)$ can be computed

in $O(\tau)$ time. Hence, we can compute $\text{id}(j)$ for all $j \in [i \dots i + 3\tau]$, and then construct $A \cap [i \dots i + 2\tau]$ by computing the sliding-window minima, in $\tilde{O}(\tau)$ overall time.

Sparsity Let $S' = S[i \dots i + 4\tau]$, and construct a τ -synchronizing set A' of S' using the *same* random coins r . Then, from the locality property we clearly have $|A'| \geq |A \cap [i \dots i + 2\tau]|$. Hence, by Lemma 3.24, $\mathbb{E}_r[|A \cap [i \dots i + 2\tau]|] \leq \mathbb{E}_r[|A'|] \leq 20 \cdot 4\tau/\tau = 80$ \square

3.3.3 Putting it Together

Now we will construct the good anchor set for input strings s, t and threshold length d . Recall that $S = s\$t$ and $|S| = n$, and we have assumed $d \geq 100$ in order to avoid corner cases. Our plan is to use string synchronizing sets to fix the misalignment introduced by the approximate different covers. However, in highly-periodic parts where synchronizing fails, we have to rely on periodicity arguments and Grover search.

Construction 3.28 (Anchor set C) Let D be a $(\lfloor d/2 \rfloor, \tau)$ -cover for some parameter $\tau \leq d/100$ to be determined later, and let $D_S := (D \cap [s]) \cup (|s| + 1 + (D \cap [t]))$. Let A be the τ -synchronizing set of S determined by random coins r .

For every $i \in D_S \cap [n - 3\tau + 2]$, let $L_i \subseteq [1 \dots n]$ be defined by the following procedure.

- *Step 1* If $A \cap [i \dots i + 2\tau]$ has at most 1000 elements, then add all the elements from $A \cap [i \dots i + 2\tau]$ into L_i . Otherwise, add the smallest 1000 elements from $A \cap [i \dots i + 2\tau]$ into L_i .
- *Step 2* If $p := \text{per}(S[i + \tau \dots i + 3\tau - 2]) \leq \tau/3$, then we do the following:
 - Define two boundary indices

$$r := \max \{r : r \leq \min\{n, i + d\} \wedge \text{per}(S[i + \tau \dots r]) = p\},$$

$$\ell := \min \{\ell : \ell \geq \min\{1, i - d\} \wedge \text{per}(S[\ell \dots i + 3\tau - 2]) = p\}.$$

Let P be the Lyndon root of $S[i + \tau \dots i + 3\tau - 2]$ (see Sect. 2.1). Then $|P| = p$, and let $P = S[i^{(b)} \dots i^{(b)} + p] = S[i^{(e)} \dots i^{(e)} + p]$ be the first and last occurrences of P in $S[\ell \dots r]$. We add three elements $i^{(b)}, i^{(b)} + p$, and $i^{(e)}$ into L_i .

Finally, the anchor set C is defined as $\bigcup_{i \in D_S \cap [n - 3\tau + 2]} L_i$.

Before proving the correctness of the anchor set C in Construction 3.28, we first observe that C has small size and is efficiently constructible.

Lemma 3.29 The anchor set C has size $|C| \leq O(n/\sqrt{d\tau})$, and is $\tilde{O}(\tau + \sqrt{d})$ -quantum-time constructible.

Proof For any given $i \in D_S \cap [n - 3\tau + 2]$, L_i contains at most 1003 elements. Hence, $|C| \leq 1003 \cdot |D_S| \leq O(n/\sqrt{d\tau})$ by Lemma 3.19.

In Construction 3.28, Step 1 takes $\tilde{O}(\tau)$ classical time by the Locality property in Lemma 3.21. In Step 2, we can find the period p and the Lyndon root P in $\tilde{O}(\tau)$

classical time (see Sect. 2.1). Then, finding the right boundary r is equivalent to searching for the largest $r \in [i + 3\tau - 2 \dots \min\{n, i + d\}]$ such that p is a period of $S[i + \tau \dots r]$ (this is because we already know $\text{per}(S[i + \tau \dots i + 3\tau - 2]) = p$, and the period is monotonically non-decreasing in r). We do this with a binary search over r , where each check can be performed by a Grover search in $\tilde{O}(\sqrt{d})$ time, since the length of $S[i + \tau \dots r]$ is at most d . The left boundary ℓ can be found similarly. Finally, the positions $i^{(b)}$ and $i^{(e)}$ can be found in $\tilde{O}(p)$ time classically, since we must have $i^{(b)} - \ell, r - i^{(e)} \leq 2p$. Hence, our anchor set C is $\tilde{O}(\tau + \sqrt{d})$ -quantum-time constructible. \square

Now we show that, with constant probability, C is a good anchor set (Definition 3.2).

Lemma 3.30 *For input strings s, t and threshold length d , with at least 0.8 probability over the random coins r , the set C is a good anchor set.*

The proof of this lemma has a similar structure to the proof of [28, Lemma 17], but is additionally complicated by the fact that (1) we have to deal with the misalignment introduced by approximate difference covers, and (2) we only considered a subset of the synchronizing set when defining the anchors.

Here, we first give an informal overview of the proof. By the property of approximate difference covers, the length- d common substring of s and t should have a pair of slightly misaligned anchors within a shift of at most $\tau - 1$ from each other. If the context around these misaligned anchors are not highly-periodic (Case 1 in the proof below), then their $O(\tau)$ -neighborhood must contain a pair of synchronizing positions (by the density property of A), which are included in Step 1 of Construction 3.28, and form a pair of perfectly aligned anchors (by the consistency property of A). If the context around the misaligned anchors are highly-periodic (Case 2), we can extend the period to the left or to the right, and look at the first position where the period stops. If this stopping position is inside the common substring, then we have a pair of anchors (Cases 2(i), 2(ii)). Otherwise, the whole common substring is highly-periodic, and we can also obtain anchors by looking at the starting positions of its Lyndon roots (Case 2(iii)). These anchors for Case 2 are included in Step 2 of Construction 3.28.

Proof (of Lemma 3.30) Let $s[i_\star \dots i_\star + d] = t[j_\star \dots j_\star + d]$ be a length- d common substring of s and t . Our goal is to show the existence of positions $i \in [|s| - d + 1]$, $j \in [|t| - d + 1]$ and a shift $h \in [0 \dots d)$, such that $s[i \dots i + d] = t[j \dots j + d]$, and $i + h, |s| + 1 + j + h \in C$.

Recall that we assumed $d \geq 100\tau$. By the definition of D_S , there exist $h_1, h_2 \in [0 \dots d/2)$ such that $i_\star + h_1, |s| + 1 + j_\star + h_2 \in D_S$, and $|h_1 - h_2| \leq \tau - 1$. These h_1, h_2 form a pair of anchors that are slightly misaligned by a shift of at most $\tau - 1$. Then the plan is to obtain perfectly aligned anchors from h_1, h_2 , either by finding synchronizing positions in their $O(\tau)$ neighborhood, or by exploiting periodicity. Without loss of generality, we assume $h_1 \leq h_2$, and the case of $h_1 > h_2$ can be proved analogously by switching the roles of s and t . Now we consider two cases:

- *Case 1* $\text{per}(s[i_\star + h_1 + \tau \dots i_\star + h_1 + 4\tau - 2]) > \tau/3$.

In this aperiodic case, by the density condition of the τ -synchronizing set A , we know that $A \cap [i_\star + h_1 + \tau \dots i_\star + h_1 + 2\tau)$ is a non-empty set, and let a be an arbitrary

element of this set. Here a is a synchronizing position from s , and we let $b = a - i_\star + j_\star$ be the corresponding position in t . Since $s[i_\star \dots i_\star + d] = t[j_\star \dots j_\star + d]$ is a common substring, in particular we have $s[a \dots a + 2\tau] = t[b \dots b + 2\tau]$, or equivalently $S[a \dots a + 2\tau] = S[|s| + 1 + b \dots |s| + 1 + b + 2\tau]$, so $|s| + 1 + b \in A$ by the consistency condition of A . Hence, we have found a pair of perfectly aligned anchors, a and $|s| + 1 + b$, for the common substring. It remains to check that they are indeed included in our anchor set C defined in Construction 3.28.

Note that we have

$$\begin{aligned} b &= j_\star + h_2 + (a - i_\star - h_1) - (h_2 - h_1) \\ &\in [j_\star + h_2 + \tau - (h_2 - h_1) \dots j_\star + h_2 + 2\tau - (h_2 - h_1)] \\ &\subseteq [j_\star + h_2 + 1 \dots j_\star + h_2 + 2\tau], \end{aligned}$$

so $|s| + 1 + b \in A \cap [|s| + 1 + j_\star + h_2 + 1 \dots |s| + 1 + j_\star + h_2 + 2\tau]$.

From the sparsity property of A (Lemma 3.21), using Markov's inequality and a union bound, we can show that $|A \cap [i_\star + h_1 \dots i_\star + h_1 + 2\tau]| \leq 1000$ and $|A \cap [|s| + 1 + j_\star + h_2 \dots |s| + 1 + j_\star + h_2 + 2\tau]| \leq 1000$ hold simultaneously with probability at least $1 - 2 \cdot 80/1000 > 0.8$. In this case, in Step 1 of Construction 3.28 we must have $a \in L_{i_\star + h_1}$, and $|s| + 1 + b \in L_{|s| + 1 + j_\star + h_2}$. Then, setting $i = i_\star$, $j = j_\star$, $h = a - i_\star$ satisfies the requirement.

- *Case 2* $p = \text{per}(s[i_\star + h_1 + \tau \dots i_\star + h_1 + 4\tau - 2]) \leq \tau/3$.

In this highly-periodic case, we cannot rely on synchronizing sets. Instead, we have the following intuition (which is based on [28]): if the common substring contains the boundary of the highly-periodic region, then we can use these boundary positions (which are easy to locate) as our anchors (Case 2(i), 2(ii)). Otherwise, the common substring is highly-periodic, and it may repeatedly occur throughout the region that shares the same period. This allows us to ignore all but a constant number of these repeated occurrences of this common substring (Case 2(iii)).

By the assumption on p we have $\text{per}(s[i_\star + h_1 + \tau \dots i_\star + h_1 + 3\tau - 2]) = p$. From $s[i_\star \dots i_\star + d] = t[j_\star \dots j_\star + d]$ and $0 \leq h_2 - h_1 \leq \tau - 1$, we also have $\text{per}(t[j_\star + h_2 + \tau \dots j_\star + h_2 + 3\tau - 2]) = p$. Hence, for both $L_{i_\star + h_1}$ and $L_{|s| + 1 + j_\star + h_2}$, we triggered Step 2 in Construction 3.28. Then, we consider three subcases.

- *Case 2(i)* $\text{per}(s[i_\star + h_1 + \tau \dots i_\star + d]) \neq p$.

In this case, the period p of $s[i_\star + h_1 + \tau \dots i_\star + h_1 + 3\tau - 2]$ does not extend to its superstring $s[i_\star + h_1 + \tau \dots i_\star + d]$, so the right boundary

$$r_s := \max \{r : \text{per}(s[i_\star + h_1 + \tau \dots r]) = p\}$$

must satisfy $i_\star + h_1 + 4\tau - 2 \leq r_s < i_\star + d - 1$. Here we observe that r_s is the same as the right boundary r in Step 2 of Construction 3.28 for constructing $L_{i_\star + h_1}$.

Let $r_t := r_s - i_\star + j_\star$. Then $j_\star + h_2 + \tau < r_t$, and $t[j_\star + h_2 + \tau \dots r_t + 1] = s[i_\star + h_2 + \tau \dots r_s + 1]$. Then from the definition of r_s , we can observe that

$$r_t = \max \{r : \text{per}(t[j_\star + h_2 + \tau \dots r]) = p\},$$

and $|s|+1+r_t$ must be the same as the right boundary r in Step 2 of Construction 3.28 for constructing $L_{|s|+1+j_\star+h_2}$.

Let P denote the Lyndon root of $s[i_\star+h_1+\tau \dots r_s]$, and let $P = s[i^{(e)} \dots i^{(e)}+p] = t[j^{(e)} \dots j^{(e)}+p]$ be the last occurrences of P in $s[i_\star+h_1+\tau \dots r_s]$ and $t[j_\star+h_2+\tau \dots r_t]$. We must have $r_s-i^{(e)} = r_t-j^{(e)}$. Note that $i^{(e)} \in L_{i_\star+h_1}$ and $|s|+1+j^{(e)} \in L_{|s|+1+j_\star+h_2}$. So setting $i = i_\star$, $j = j_\star$, $h = i^{(e)} - i_\star$ satisfies the requirement.

- Case 2(ii) $\text{per}(s[i_\star+h_1+\tau \dots i_\star+d]) = p$, but $\text{per}(s[i_\star \dots i_\star+d]) \neq p$.
In this case, the period p fully extends to the right but not to the left. Using a similar argument as in Case 2(i), we can show that the left boundaries

$$\begin{aligned}\ell_s &:= \min\{\ell : \text{per}(s[\ell \dots i_\star+h_1+3\tau-2]) = p\}, \\ \ell_t &:= \min\{\ell : \text{per}(t[\ell \dots j_\star+h_2+3\tau-2]) = p\}\end{aligned}$$

must satisfy $\ell_s - i_\star = \ell_t - j_\star \geq 1$, and $\ell_s, |s|+1+\ell_t$ are the left boundaries in Step 2 of Construction 3.28 for constructing $L_{i_\star+h_1}, L_{|s|+1+j_\star+h_2}$ respectively. Then, letting $s[i^{(b)} \dots i^{(b)}+p] = t[j^{(b)} \dots j^{(b)}+p]$ be the first occurrences of the Lyndon root in $s[\ell_s \dots i_\star+h_1+3\tau-2]$ and $t[\ell_t \dots j_\star+h_2+3\tau-2]$, we can similarly see that setting $i = i_\star$, $j = j_\star$, $h = i^{(b)} - i_\star$ satisfies the requirement.

- Case 2(iii) $\text{per}(s[i_\star \dots i_\star+d]) = p$.
Let

$$\begin{aligned}\ell_s &:= \min\{\ell : \ell \geq \min\{1, i_\star+h_1-d\} \wedge \text{per}(s[\ell \dots i_\star+h_1+3\tau-2]) = p\}, \\ \ell_t &:= \min\{\ell : \ell \geq \min\{1, j_\star+h_2-d\} \wedge \text{per}(s[\ell \dots j_\star+h_2+3\tau-2]) = p\}.\end{aligned}$$

Then $\ell_s, |s|+1+\ell_t$ are the left boundaries in Step 2 of Construction 3.28 for constructing the sets $L_{i_\star+h_1}, L_{|s|+1+j_\star+h_2}$ respectively. We must have $\ell_s < i_\star$ and $\ell_t < j_\star$.

Let P be the Lyndon root of $s[i_\star \dots i_\star+d]$, and assume the first occurrence of P in $s[i_\star \dots i_\star+d]$ is $s[i' \dots i'+p]$. We also let $s[i^{(b)} \dots i^{(b)}+p] = t[j^{(b)} \dots j^{(b)}+p]$ be the first occurrences of P in $s[\ell_s \dots i_\star+d]$ and $t[\ell_t \dots j_\star+d]$. Then the second occurrence of P in $s[\ell_s \dots i_\star+d]$ is $s[i^{(b)}+p \dots i^{(b)}+2p]$. Observe that, if we find the first occurrence of the common substring $s[i_\star \dots i_\star+d]$ inside the entire periodic region $s[\ell_s \dots i_\star+d]$, this occurrence should align $s[i' \dots i'+p]$ with either the first occurrence of P or the second occurrence of P in this region. Formally, let $i'_\star = i^{(b)} - (i' - i_\star)$. Then, we have either $s[i_\star \dots i_\star+d] = s[i'_\star \dots i'_\star+d]$ or $s[i_\star \dots i_\star+d] = s[i'_\star+p \dots i'_\star+p+d]$. Similarly, letting $j'_\star = j^{(b)} - (j' - j_\star)$, we have $t[j_\star \dots j_\star+d] = t[j'_\star \dots j'_\star+d]$ or $t[j_\star \dots j_\star+d] = t[j'_\star+p \dots j'_\star+p+d]$. Note that $i^{(b)}, i^{(b)}+p \in L_{i_\star+h_1}$, and $|s|+1+j^{(b)}, |s|+1+j^{(b)}+p \in L_{|s|+1+j_\star+h_2}$. Hence, setting $i = i'_\star$ (or $i = i'_\star+p$), $j = j'_\star$ (or $j = j'_\star+p$), $h = i' - i_\star$ satisfies the requirement.

Hence, the desired i, j and h always exist. \square

Finally, Lemma 3.6 immediately follows from Construction 3.28, Lemma 3.29, and Lemma 3.30, by setting $\tau = \Theta(\sqrt{d})$.

4 Minimal String Rotation

4.1 Minimal Length- ℓ Substrings

Rather than work with the Minimal String Rotation problem directly, we present an algorithm for the following problem, which is more amenable to work with using our divide-and-conquer approach.

Minimal Length- ℓ Substrings

Input: A string $s[1..n]$ and an integer $n/2 \leq \ell \leq n$

Task: Output all elements in $\arg \min_{1 \leq i \leq n-\ell+1} s[i..i+\ell)$ represented as an arithmetic progression.

The elements in the output are guaranteed to be an arithmetic progression thanks to Lemma 2.3.

We will prove the following theorem.

Theorem 4.1 *Minimal Length- ℓ Substrings can be solved by a quantum algorithm with $n^{1/2+o(1)}$ query complexity and time complexity.*

For convenience, we also introduce the following problem.

Maximal String Rotation

Input: A string s

Task: Output a position $i \in [1..|s|]$ such that $s[j..|s|]s[1..j-1] \leq s[i..|s|]s[1..i-1]$ holds for all $j \in [1..|s|]$. If there are multiple solutions, output the smallest such i .

We now use a series of simple folklore reductions to show that the Minimal Length- ℓ Substrings problem generalizes the Minimal String Rotation problem.

Proposition 4.2 *The Minimal String Rotation problem reduces to the Maximal String Rotation problem.*

Proof Take an instance of the Minimal String Rotation problem, consisting of a string s over an alphabet Σ , which recall we identify with the set $[1..|\Sigma|]$. Consider the map $\varphi: \Sigma \rightarrow \Sigma$ defined by taking

$$\varphi(c) = |\Sigma| - c + 1$$

for each character $c \in \Sigma$. Let

$$t = \varphi(s[1]) \cdots \varphi(s[n])$$

be the result of applying this map to each character of s .

By construction, for any $c, c' \in \Sigma$ we have $\varphi(c) < \varphi(c')$ if and only if $c' < c$. Combining this observation together with the definition of lexicographic order, we deduce that for any indices $j, k \in [1..n]$ we have

$$t[j..n]t[1..j-1] \preceq t[k..n]t[1..k-1]$$

if and only if

$$s[k..n]s[1..k-1] \preceq s[j..n]s[1..j-1].$$

Thus the solution to the Maximal String Rotation problem on t recovers the solution to the Minimal String Rotation problem on s , which proves the desired result. \square

Proposition 4.3 *The Maximal String Rotation problem reduces to the Maximal Suffix problem.*

Proof Take an instance of the Maximal String Rotation problem, consisting of a string s of length n .

Let $t = ss$ be the string of length $2n$ formed by concatenating s with itself. Suppose i is the starting index of the maximal rotation of s . Then we claim that i is the starting index of the maximal suffix of t as well.

Indeed, take any position $j \in [1..2n]$ in string t with $j \neq i$.

If $j > n$, then we can write $j = n + \Delta$ for some positive integer $\Delta \leq n$. In this case we have

$$t[j..2n] < t[\Delta..2n]$$

because the string on the left hand side is a proper prefix of the string on the right hand side. Thus j cannot be the starting position of a maximal suffix for t .

Otherwise, $j \leq n$. Note that we can write

$$\begin{aligned} t[i..2n] &= s[i..n]s[1..i-1]s[i..n] \quad \text{and} \\ t[j..2n] &= s[j..n]s[1..j-1]s[j..n]. \end{aligned} \tag{2}$$

Since i is a solution to the Maximal String Rotation problem, we know that either

$$s[j..n]s[1..j-1] < s[i..n]s[1..i-1]$$

or $s[j..n]s[1..j-1] = s[i..n]s[1..i-1]$ and $i < j$.

In the first case, the decompositions from Eq. (2) immediately imply that

$$t[j..2n] < t[i..2n]$$

by considering the length n prefixes of the two strings. In the second case, since $s[j..n]s[1..j-1] = s[i..n]s[1..i-1]$ and $i < j$ the decompositions from Eq.

(2) imply that

$$t[j \dots 2n] \prec t[i \dots 2n]$$

because the string on the left hand side is a proper prefix of the string on the right hand side. Combining these results, we see that the solution to the Maximal Suffix problem on t is the index i which solves the Maximal String Rotation problem on s . \square

Proposition 4.4 *The Maximal Suffix problem reduces to the Minimal Suffix problem.*

Proof Take an instance of the Maximal Suffix problem, consisting of a string s of length n over an alphabet $\Sigma = [1 \dots |\Sigma|]$. Let $\sigma = |\Sigma| + 1$ denote a character lexicographically after all the characters in Σ . As in the proof of Proposition 4.2, consider the map $\varphi: \Sigma \rightarrow \Sigma$ defined by taking

$$\varphi(c) = |\Sigma| - c + 1$$

for each character $c \in \Sigma$. Now, build the string

$$t = \varphi(s[1])\varphi(s[2]) \cdots \varphi(s[n])\sigma$$

formed by applying φ to each character of s and then appending σ to the end.

Suppose that $s[i \dots n]$ is the maximal suffix of s . We claim that $t[i \dots n + 1]$ is the minimal suffix of t . Thus solving the Minimal Suffix problem on t recovers a solution to the Maximal Suffix problem on s .

To see this, note that take any index $1 \leq j \leq n$ with $j \neq i$. By assumption

$$s[j \dots n] \prec s[i \dots n]. \quad (3)$$

This can happen one of two ways.

First, it could be that $j > i$ and the string on the left hand side above is a proper prefix of the string on the right hand side. In this case we must have

$$t[i \dots n + 1] = \varphi(s[i]) \cdots \varphi(s[n])\sigma \prec \varphi(s[j]) \cdots \varphi(s[n])\sigma = t[j \dots n + 1]$$

because the string on the left hand side agrees with the string on the right hand side for the first j positions, but then at the $(n - j + 2)^{\text{th}}$ position, the string on the right hand side has the character σ , which is larger than the corresponding character $\varphi(s[n - (j - i)])$ from the string on the left hand side.

Otherwise, Eq. (3) holds because there exists some nonnegative integer Δ such that $s[j + \Delta] < s[i + \Delta]$ and $s[j + d] = s[i + d]$ for all nonnegative $d < \Delta$. By definition, $\varphi(c) \prec \varphi(c')$ if and only if characters $c' \prec c$ for all $c, c' \in \Sigma$. Thus in this case too we have

$$t[i \dots n + 1] = \varphi(s[i]) \cdots \varphi(s[n])\sigma \prec \varphi(s[j]) \cdots \varphi(s[n])\sigma = t[j \dots n + 1]$$

because the strings agree for the first Δ characters, but then at the $(\Delta + 1)^{\text{st}}$ position, the string on the right hand side has the character $\varphi(s[j + \Delta])$, which is larger than the corresponding character $\varphi(s[i + \Delta])$ from the string on the left hand side by our observation on φ . Finally, note that the suffix $t[n + 1] = \sigma$ is larger than every other suffix of t by construction, and is thus not a minimal suffix of t . Thus the minimal suffix of t corresponds to the maximal suffix of s , and the reduction is correct. \square

Proposition 4.5 *The Minimal Suffix problem reduces to the Minimal Length- ℓ Substrings problem.*

Proof Take an instance of the Minimal Suffix problem, consisting of a string s of length n . Consider the string of length $2n - 1$ of the form

$$t = s \underbrace{00 \cdots 0}_{n-1 \text{ times}}$$

formed by appending $n - 1$ copies of a character 0, smaller than every character from the alphabet Σ of s , to the end of s .

Let i be the starting index of the minimal suffix of s . We claim that i is also the unique index returned by solving the Minimal Length- n Substrings problem on t (note that n is at least half the length of t).

Indeed, take any index $j \in [1 \dots n]$ with $j \neq i$. By assumption we have

$$s[i \dots n] \prec s[j \dots n].$$

Because the string on the left hand side occurs strictly before the string on the right hand side in lexicographic order, appending any number 0s to the ends of the strings above cannot change their relative order. Thus

$$t[i \dots i + n] = s[i \dots n] \underbrace{00 \cdots 0}_{i-1 \text{ times}} \prec s[j \dots n] \underbrace{00 \cdots 0}_{j-1 \text{ times}} = t[j \dots j + n]$$

as well. Because this holds for all $j \neq i$ we get that i is the unique position output by solving the Minimal Length- n Substrings problem on t . This proves the reduction is correct. \square

By chaining the above reductions together, we obtain the following corollary of Theorem 4.1.

Theorem 4.6 *Minimal String Rotation, Maximal Suffix, and Minimal Suffix can be solved by a quantum algorithm with $n^{1/2+o(1)}$ query complexity and time complexity.*

Proof By combining the results of Propositions 4.2, 4.3, 4.4, and 4.5, we see that all the problems mentioned in the theorem statement reduce to the Minimal Length- ℓ Substrings problem. Each of the reductions only involves simple substitutions and insertions to the input strings.

In particular, by inspecting the proofs of the propositions, we can verify that for an input string s and its image t under any of these reductions, any query to a character




Fig. 2 Proof of Lemma 4.8. Here, the answers in the Minimal Length- k Substrings problem on input string $s[a \dots a + 2k]$ is $J = \{j_1, j_2, j_3\}$, and $p = j_2 - j_1$. This period p extends to the right up to position r

of t can be simulated with $O(1)$ queries to the characters of s . Thus, we can get a $n^{1/2+o(1)}$ query and time quantum algorithm for each of the listed problems by using the algorithm of Theorem 4.1 and simulating the aforementioned reductions appropriately in the query model. \square

Remark 4.7 We remark that, from the $\Omega(\sqrt{n})$ quantum query lower bound for Minimal String Rotation [14], this chain of reductions also implies that Maximal Suffix and Minimal Suffix require $\Omega(\sqrt{n})$ quantum query complexity.

It remains to prove Theorem 4.1. To solve the Minimal Length- ℓ Substrings problem, it suffices to find any individual solution

$$i \in \operatorname{argmin}_{1 \leq i \leq n-\ell+1} s[i \dots i + \ell),$$

and then use the quantum Exact String Matching algorithm to find all the elements (represented as an arithmetic progression) in $\tilde{O}(\sqrt{n})$ time. Our approach will invoke the following “exclusion rule,” which simplifies the previous approach used in [14]. We remark that similar kinds of exclusion rules have been applied previously in parallel algorithms for Exact String Matching [10] and Minimal String Rotation [17] (under the name of “Ricochet Property” or “duel”), as well as the quantum algorithm by Wang and Ying [14, Lemma 5.1]. The advantage of our exclusion rule is that it naturally yields a *recursive* approach for solving the Minimal Length- ℓ Substrings problem.

Lemma 4.8 (Exclusion Rule) *In the Minimal Length- ℓ Substrings problem with input $s[1 \dots n]$ with $n/2 \leq \ell \leq n$, let*

$$I := \operatorname{argmin}_{1 \leq i \leq n-\ell+1} s[i \dots i + \ell)$$

denote the set of answers forming an arithmetic progression. For integers $a \geq 1, k \geq 1$ such that $a + k \leq n - \ell + 1$, let J denote the set of answers in the Minimal Length- k Substrings problem on the input string $s[a \dots a + 2k]$. Then if $\{\min J, \max J\} \cap I = \emptyset$, we must have $J \cap I = \emptyset$.

Proof First observe that

$$a + 2k - 1 \leq n - \ell + k \leq 2(n - \ell) \leq n,$$

so $s[a \dots a + 2k]$ is a length- $2k$ substring of s . Since the statement is trivial for $|J| \leq 2$, we assume J consists of $j_1 < j_2 < \dots < j_m$ where $m \geq 3$. Let $p = j_2 - j_1$. Then

$p = (j_m - j_1)/(m - 1) \leq k/2$. Then from

$$s[j_1 \dots j_1 + k] = s[j_2 \dots j_2 + k] = \dots = s[j_m \dots j_m + k]$$

we know that p must be a period of $s[j_1 \dots j_m + k]$.⁸ We consider the first position r where this period p stops, that is, $r := \min\{j_m + k \leq r \leq n : s[r] \neq s[r - p]\}$. If such r does not exist, let $r = n + 1$. See Fig. 2. With this setup, we now proceed to prove the contrapositive of the original claim.

Suppose $j_q \in I$ for some $1 \leq q \leq m$. We consider three cases.

- *Case 1* $r \geq j_m + \ell$.
In this case, we must have $s[j_1 \dots j_1 + \ell] = s[j_2 \dots j_2 + \ell] = \dots = s[j_m \dots j_m + \ell]$. Then, $j_q \in I$ implies $j_1 \in I$.
- *Case 2* $r < j_m + \ell$, and $s[r] < s[r - p]$.
For every $1 \leq t \leq m - 1$, by the definition of r , we must have $s[j_{t+1} \dots r] = s[j_t \dots r - p]$. Then from $s[r] < s[r - p]$ we have $s[j_t \dots j_t + \ell] \geq s[j_{t+1} \dots j_{t+1} + \ell]$. Hence, $j_q \in I$ implies $j_{q+1}, j_{q+2}, \dots, j_m \in I$.
- *Case 3* $r < j_m + \ell$, and $s[r] > s[r - p]$.
By an argument similar to Case 2, we can show $s[j_t \dots j_t + \ell] \leq s[j_{t+1} \dots j_{t+1} + \ell]$. Then, $j_q \in I$ implies $j_{q-1}, j_{q-2}, \dots, j_1 \in I$.

Thus $\{j_1, j_m\} \cap I \neq \emptyset$ in all of the cases, which proves the desired result. \square

4.2 Divide and Conquer Algorithm

To motivate our quantum algorithm, we first describe a classical algorithm for the Minimal $n/2$ -length Substring problem which runs in $O(n \log n)$ time (note that other classical algorithms can solve the problem faster in $O(n)$ time). Our quantum algorithm will use the same setup, but obtain a speed-up via Grover search. For the purpose of this overview, we assume n is a power of 2. The classical algorithm works as follows:

Suppose we are given an input string s of length n and target substring size $\ell = n/2$. Set $m = \ell/2 = n/4$. Then the half of the solution (i.e. the first m characters of a minimum length ℓ -substring) are contained entirely in either the block $s_1 = s[1 \dots n/2]$ or the block $s_2 = s[n/4 \dots 3n/4]$.

With that in mind, we recursively solve the problem on the strings s_1 and s_2 with target size m in both cases. Let u_1 and v_1 be the smallest and largest starting positions returned by the recursive call to s_1 respectively. Define u_2 and v_2 as the analogous positions returned by the recursive call to s_2 . Then by Lemma 4.8, the true starting position of the minimal ℓ -length substring of s is in $\{u_1, u_2, v_1, v_2\}$.

We identify the ℓ -length substrings starting at each of these positions, and find their lexicographic minimum in $O(n)$ time via linear-time string comparison. This lets us find at least one occurrence of the minimum substring of length ℓ . Then, to find all occurrences of this minimum substring, we use a linear time string matching algorithm (such as the classic Knuth-Morris-Pratt algorithm [1]) to find the first two occurrences of the minimum length ℓ substring in s . The difference between the starting positions

⁸ In fact, p is the minimum period of this substring.

then lets us determine the common difference of the arithmetic sequence of positions encoding all starting positions of the minimum substring.

If we let $T(n)$ denote the runtime of this algorithm, the recursion above yields a recurrence

$$T(n) = 2T(n/2) + O(n)$$

which solves to $T(n) = O(n \log n)$.

4.3 Quantum Speedup

Next, we show how to improve the runtime of this divide-and-conquer approach in the quantum setting. The key change is to break the string into b blocks, and apply quantum minimum finding over these blocks which only takes $\tilde{O}(\sqrt{b})$ recursive calls, instead of b recursive calls needed by the classical algorithm. We will set b large enough to get a quantum speedup.

Proof (of Theorem 4.1) Let b be some parameter to be set later. For convenience assume that b divides both ℓ and n (this assumption does not affect the validity of our arguments, and is only used to let us avoid working with floor and ceiling functions). Set $m = \ell/b$.

For each nonnegative integer $k \leq \lfloor n/m \rfloor - 2$ we define the substring

$$s_k = s(km \dots (k+2)m).$$

Also set $s_{\lfloor n/m \rfloor - 1} = s(n - 2m \dots n)$.

These s_k blocks each have length $2m$, and together cover every substring of length m in s . Let P be the minimum length- ℓ substring in s . By construction, the first $m = \ell/b$ characters of P is contained entirely in one of the s_k blocks.

For each block s_k , let P_k denote its minimum length- m substring and let u_k and v_k be the smallest and largest starting positions respectively of an occurrence of P_k in s_k . The lexicographically smallest prefix P_k will make up the first m characters of the minimum length- ℓ substring. Thus by Lemma 4.8, we know the minimum length- ℓ substring of s must start at position u_k or v_k for some index k .

We now use quantum minimum finding to find P . We search over the $\Theta(n/m) = \Theta(b)$ blocks above. To compare blocks s_i and s_j , we recursively solve the Minimal Length- m Substrings problem on s_i and s_j to find positions u_i, v_i and u_j, v_j . Then we look at the substrings of length ℓ starting at these four positions. By binary search and Grover search (Lemma 2.5), in $\tilde{O}(\sqrt{n})$ time we can determine which of these four substrings is lexicographically the smallest. If the smallest of these substrings came from s_i we say block s_i is smaller than block s_j , and vice versa.

After running the minimum finding algorithm, we will have found P . To return all occurrences of P , we can then use the quantum algorithm for Exact String Matching to find the two leftmost occurrences and the rightmost occurrence of P in s in $\tilde{O}(\sqrt{n})$ time. Together they determine the positions of all copies of P in s as an arithmetic sequence, which we can return to solve the original problem.

It remains to check the runtime of the algorithm. Let $T(n)$ denote the runtime of the algorithm with error probability at most $1/n$. Recall that our algorithm solves Minimum Finding over $\Theta(b)$ blocks, where each comparison involves a recursive call on strings of size $2m = \Theta(n/b)$ and a constant number of string comparisons of length n (via Lemma 2.5), and finally solves Exact String Matching for strings of size $\Theta(n)$. Hence we have the recurrence (assuming all logarithms are base 2)

$$T(n) \leq \tilde{O}(\sqrt{b}) \cdot (T(n/b) + \tilde{O}(\sqrt{n})) + \tilde{O}(\sqrt{n}) = c(\log b)^e \sqrt{b} (T(n/b) + \sqrt{n})$$

for some constants $c, e > 0$, where the polylogarithmic factors are inherited from the subroutines we use and the possibility of repeating our steps $O(\log n)$ times to drive down the error probability. Now set

$$b = 2^{d(\log n)^{2/3}}$$

for some constant d . We claim that for sufficiently large d , we recover a runtime of $T(n) = n^{1/2} \cdot 2^{d(\log n)^{2/3}}$.

We prove this by induction. The result holds when n is a small constant by taking d large enough. Now, suppose we want to prove the result for some arbitrary n , and that the claimed runtime bound holds on inputs of size less than n . Then using the recurrence above and the inductive hypothesis we have

$$\begin{aligned} T(n) &\leq c(\log b)^e \sqrt{b} (T(n/b) + \sqrt{n}) \\ &\leq c(\log b)^e \sqrt{n} (2^{d(\log(n/b))^{2/3}} + \sqrt{b}) \\ &\leq 2c(\log b)^e \sqrt{n} \cdot 2^{d(\log(n/b))^{2/3}}, \end{aligned}$$

where the last inequality follows from $d(\log(n/b))^{2/3} \geq d(\log(\sqrt{n}))^{2/3} > \frac{1}{2}d(\log n)^{2/3} = \log(\sqrt{b})$ for large enough n . Equivalently, this means that

$$\frac{T(n)}{n^{1/2}2^{d(\log n)^{2/3}}} \leq 2c \cdot 2^{e(\log \log b) - d((\log n)^{2/3} - (\log n - \log b))^{2/3}}. \quad (4)$$

Using the mean value theorem, we can bound

$$\begin{aligned} (\log n)^{2/3} - (\log n - \log b)^{2/3} &\geq (2/3)(\log b)(\log n)^{-1/3} \\ &= (2/3)d(\log n)^{1/3} \\ &\geq \omega(\log \log b), \end{aligned}$$

where the last inequality follows from $\log \log b = \log d + (2/3) \log \log n$. Thus, by taking d to be a large enough constant in terms of c and e , we can force the right hand side of Equation (4) to be less than 1, which proves that

$$T(n) \leq n^{1/2}2^{d(\log n)^{2/3}}.$$

This completes the induction, and proves that we can solve the Minimum Length- ℓ Substrings problem in the desired runtime as claimed. \square

4.4 Longest Lyndon Substring

The technique we use to solve the Minimal String Rotation problem can also be adapted to get a quantum speed-up for solving the Longest Lyndon Substring problem.

Theorem 4.9 *The Longest Lyndon Substring problem can be solved by a quantum algorithm with $n^{1/2+o(1)}$ query complexity and time complexity.*

A difficulty in solving Longest Lyndon Substring compared to other string problems such as LCS and Longest Palindromic Substring is that the lengths of Lyndon Substrings do not have the monotone property, and hence we cannot use binary search (the Longest Square Substring problem in Sect. 5 also has the same issue). To overcome this issue, we first present a simple reduction.

Theorem 4.10 *For any constant $0 < \varepsilon < 1$, suppose there is a $T(d)$ -time quantum algorithm (where $T(d) \geq \Omega(\sqrt{d})$) for solving the Longest Lyndon Substring problem on string s of length $|s| = (1+2\varepsilon)d$ with the promise that the longest Lyndon substring of s has length in the interval $[d, (1+\varepsilon)d]$. And suppose there is an $T(d)$ -time quantum algorithm for checking whether an $O(d)$ -length string is a Lyndon word.*

Then, there is an algorithm in time $\tilde{O}(T(n))$ for solving the Longest Lyndon Substring problem on length- n strings in general case.

Proof Let s be the input string of length n . For each nonnegative integer $i \leq \lceil (\log n)/(\log(1+\varepsilon)) \rceil - 1$, we look for a longest Lyndon substring of s whose length is in the interval $[(1+\varepsilon)^i, (1+\varepsilon)^{i+1}]$, and return the largest length (after certifying that it is indeed a Lyndon substring) found. This only blows up the total time complexity by an $O(\log n)$ factor.

For each i , we define the positions $j_k := 1 + k \cdot \varepsilon d/2$ for all $0 \leq k < 2n/(\varepsilon d)$, and consider the substrings

$$s[j_0 \dots j_0 + (1+2\varepsilon)d], s[j_1 \dots j_1 + (1+2\varepsilon)d], \dots$$

Note that, if the longest Lyndon substring of s has length in the interval $[d, (1+\varepsilon)d]$, then it must be entirely covered by some of these substrings. For each of these substrings, its longest Lyndon substring can be computed in $T(d)$ -time by the assumption. Then, we use the quantum maximum finding algorithm (see Sect. 2.4) to find the longest among these $2n/(\varepsilon d)$ answers, in $\tilde{O}(\sqrt{2n/(\varepsilon d)} \cdot T(d)) = \tilde{O}(\sqrt{n} \cdot T(d)/\sqrt{d}) \leq \tilde{O}(T(n))$ overall time, where we used the assumption of $T(d) \geq \Omega(\sqrt{d})$. \square

Now, we are going to describe an $d^{1/2+o(1)}$ -time quantum algorithm for solving the Longest Lyndon Substring problem on string s of length $|s| = (1+2\varepsilon)d$, with the promise that the longest Lyndon substring of s has length in the interval $[d, (1+\varepsilon)d]$. Combined with the reduction above, this proves Theorem 4.9, since a string is Lyndon

if and only if its minimal suffix is itself (see Sect. 2.1) and can be checked by our Minimal Suffix algorithm. We will set $\varepsilon = 0.1$.

We will make use of the following celebrated fact related to Lyndon substrings.

Definition 4.11 (*Lyndon Factorization* [8, 96]) Any string s can be written as a concatenation

$$s = s_1 s_2 \cdots s_k$$

where each string s_i is a Lyndon word, and $s_1 \geq s_2 \geq \cdots \geq s_k$. This decomposition is unique, and called the *Lyndon factorization* of s . The s_i are called *Lyndon factors* of s .

The following fact characterizes the longest Lyndon substring in a given string.

Proposition 4.12 (e.g., [71, Lemma 3]) *The longest Lyndon substring of a string s is necessarily a longest Lyndon factor of s .*

Then, given the promise about the input string s of length $(1 + 2\varepsilon)d$, we know s has Lyndon factorization $s_1 \cdots s^* \cdots s_k$, where $|s^*| \in [d, (1 + \varepsilon)d]$. The remaining task is to identify the position and length of the Lyndon factor s^* .

Lemma 4.13 (The position of s^*) *Suppose $s[i \dots i + |s^*|] = s^*$. Then, $s[i \dots |s|]$ must be the minimal suffix among all $i \in [1 \dots \varepsilon d + 1]$.*

Proof Note that $i \in [1 \dots \varepsilon d + 1]$ due to $|s| = (1 + 2\varepsilon)d$ and $|s^*| \in [d, (1 + \varepsilon)d]$. For any other starting position $j \in [1 \dots \varepsilon d + 1]$, we will prove that $s[j \dots |s|] \succ s[i \dots |s|]$. We consider two cases.

Case 1 $j > i$. In this case we must have $j \in (i \dots i + |s^*|)$ due to the length constraints. Then, we have $s[j \dots i + |s^*|] \succ s[i \dots i + |s^*|]$ due to the fact that s^* is a Lyndon word, which immediately implies $s[j \dots |s|] \succ s[i \dots |s|]$, since $|s[i \dots i + |s^*|]| > |s[j \dots i + |s^*|]|$.

Case 2 $j < i$. Then, suppose a Lyndon factor s_t prior to s^* occurs at $s_t = s[j' \dots j'']$ with $j' \leq j \leq j''$. Then, we have $s[j \dots j''] \geq s[j' \dots j''] = s_t \geq s^*$ by the property of Lyndon factorization. Then, from the length constraint $|s[j \dots j'']| < |s^*|$, we necessarily have $s[j \dots |s|] \succ s[i \dots |s|]$. \square

Then we can find the starting position of s^* , by looking for the minimal suffix of s whose starting position is in $[1 \dots \varepsilon d + 1]$. We observe that, this task can be reduced to the Minimum Length- ℓ Substrings problem using the same reduction as in Proposition 4.5 by appropriately adjusting the lengths, and we omit the proof here. Hence, we can find the starting position of s^* in $d^{1/2+o(1)}$ time.

We can now without loss of generality assume that s^* appears as the first Lyndon factor of the input string $s = s^* s_2 s_3 \cdots s_m$ of length $|s| \leq (1 + 2\varepsilon)d$. It remains to find the ending position of s^* . We need the following definition.

Definition 4.14 We say a string s is *pre-Lyndon*, if there is a Lyndon word t such that s is a prefix of t .

We have the following characterization of pre-Lyndon strings.

Proposition 4.15 (e.g., [71, Lemma 10]) *For any pre-Lyndon string w , there exists a unique Lyndon word x such that $w = x^k x'$ where $k \geq 1$, and $x' = x[1 \dots i]$ for some $i \in [0 \dots |x| - 1]$. Here x^k denotes concatenating x for k times.*

Note that we can check whether a string w is pre-Lyndon, in $|w|^{1/2+o(1)}$ time.

Lemma 4.16 *Given any string w of length d , we can check whether it is pre-Lyndon in $d^{1/2+o(1)}$ quantum time. Moreover, if w is pre-Lyndon, we can find its decomposition described in Proposition 4.15 also in $d^{1/2+o(1)}$ quantum time.*

Proof We assume w is indeed a pre-Lyndon string, and has decomposition $w = x^k x'$ as described in Proposition 4.15.

We first observe that, the minimal rotation of $w = x^k x'$ must equal $x' x^k$. This observation can be easily proved from the fact that the Lyndon word x is strictly smaller than all other rotations of x . In the case of $|x'| \geq 1$, we can compute the shift of the minimal rotation of w (which must be unique), and obtain the length $|x'|$. We can also detect the case of $|x'| = 0$, by finding that w itself equals the minimal rotation of w .

After finding $|x'|$, we are left with the part $w' = x^k$, and we know that $|x| = \text{per}(|w'|)$. We can then compute $\text{per}(|w'|)$ by finding the second earliest occurrence of w' in the string $w'w'$, using the quantum Exact String Matching algorithm [9]. (Alternatively, we can use the periodicity algorithm of Wang and Ying [14])

Finally, after obtaining x and x' , we certify that w is indeed a pre-Lyndon string, by checking that x is a Lyndon word, x' is a prefix of x , and $w = x^k x'$, in $\tilde{O}(\sqrt{|w|})$ time by Grover search. \square

Then, on the input string s with Lyndon factorization $s = s^* s_2 s_3 \dots s_m$, we apply Lemma 4.16 with binary search to find the maximum position $i \in [|s|]$ such that $s[1 \dots i]$ is pre-Lyndon. We must have $i \geq |s^*|$, by the definition of pre-Lyndon string and the fact that s^* is Lyndon. We also obtain the decomposition of $s[1 \dots i] = x^k x'$ described in Proposition 4.15, where x is a Lyndon word with proper prefix x' . Note that the longest Lyndon prefix of $x^k x'$ must be x , since any longer prefix of $x^k x'$ can be written as $x^j x''$ and obviously has a smaller suffix than itself. Then, from the fact that s^* is the longest Lyndon prefix of $s[1 \dots i]$, we know $x = s^*$. Hence, we have completely determined s^* .

Remark 4.17 We can show that the Longest Lyndon Substring problem requires $\Omega(\sqrt{n})$ quantum queries, by a simple reduction from the unstructured search problem. Suppose we are given a string $s \in \{0, 1\}^n$ and want to decide whether there exists $i \in [n]$ such that $s[i] = 1$. We create another string $s' := s 0^n 1$ by appending n zeros and a one after s . Then, if $s = 0^n$, then the longest Lyndon substring of s' will be s' itself. If there is at least a one in s , then s' cannot be a Lyndon word, since its suffix $0^n 1$ must be smaller than s' . Hence, by [21], this implies that Longest Lyndon Substring problem requires query complexity $\Omega(\sqrt{n})$.

5 Longest Square Substring

Recall that in the Longest Square Substring problem, we are given a string s of length n and tasked with finding the largest positive integer Δ such that there exists some index $1 \leq i \leq n - 2\Delta + 1$ with $s[i \dots i + \Delta) = s[i + \Delta \dots i + 2\Delta)$. In other words, we want to find the maximum size $\ell = 2\Delta$ such that s contains a Δ -periodic substring of length ℓ . We call Δ the *shift* and ℓ the *length* of the longest square substring. We refer to the substrings $s[i \dots i + \Delta)$ and $s[i + \Delta \dots i + 2\Delta)$ as the *first half* and *second half* of the solution respectively.

In this section, we present a quantum algorithm which solves this problem on strings of length n in $\tilde{O}(\sqrt{n})$ time. We follow this up with a brief argument indicating why this algorithm is optimal up to polylogarithmic factors in the query complexity.

Theorem 5.1 *The Longest Square Substring problem can be solved by a quantum algorithm with $\tilde{O}(\sqrt{n})$ query complexity and time complexity.*

Proof Let s be the input string of length n .

Set $\varepsilon = 1/10$. For each nonnegative integer $i \leq \lceil (\log n)/(\log(1+\varepsilon)) \rceil - 1$, we look for a longest square substring of s whose length is in the interval $[(1+\varepsilon)^i, (1+\varepsilon)^{i+1})$. We begin with i equal to its upper bound, and then keep decrementing i until we find a square substring in the relevant interval. The first time we find such a string we return it and halt. If we never find such a string we report that s has no square substring. We try out $O(\log n)$ values of i , so it suffices to solve each of these subproblems in $\tilde{O}(\sqrt{n})$ time (this is very similar to the argument used in Theorem 4.10).

If s has no square substring, our algorithm will never find a solution and will correctly detect that there is none. Hence in the remainder of this proof, suppose that s contains a square substring, and let $\ell = 2\Delta$ be the length of the longest square substring in s . Let i be the unique positive integer with $\ell \in [(1+\varepsilon)^i, (1+\varepsilon)^{i+1})$. We will eventually reach this value of i since we cannot have found any square substrings of larger size. Write $d = (1+\varepsilon)^i$ and, for convenience, assume that d is an integer multiple of 10 (this will not affect the correctness of the arguments below).

We start by sampling a uniform random position g in the string. We say g is *good* if there exists a square substring A of size ℓ in s with the property that g is among the first $d/10$ positions in A . Note that g is good with probability at least $\Omega(d/n)$.

Suppose g is a good position. Now, consider the substring

$$P = s(g \dots g + 2d/5]$$

of length $2d/5$ starting immediately after g . Since g is good, the end of P is at most the

$$d/10 + 2d/5 \leq d/2 \leq \ell/2$$

position character in A . Hence P is contained completely in the first half of A .

Now define $S = s(g + 2d/5 \dots g + (1+\varepsilon)d]$ to be an $O(d)$ length substring of s starting immediately after P . Since g is good and A has length at most $(1+\varepsilon)d$,




Fig. 3 An example of how we extract a square substring when there is only one copy of P . The black vertical line segments in the top three images bound the first half and second half of the true optimal square substring A . When we extend the patterns backwards we cover the entire prefix, we are guaranteed to reach the boundaries of A , but may end up extending further. In the final step we extend the patterns forward and find a square substring bounded by the ends of the resulting extensions, marked by vertical line segments in the bottom image. This square substring may be different from A , but is always at least as long as A

we know that the second half of A is contained completely in S . In particular, S must contain at least one copy of P .

By solving the Exact String Matching problem with S and P , we can find the leftmost and rightmost occurrences of P in S in $\tilde{O}(\sqrt{d})$ time. If these copies occur at the same position, S actually contains only a single copy of P , and otherwise S contains multiple copies of P . We consider these cases separately.

Case 1: Unique Copy Suppose S has only one copy of P . Let this copy begin after position h , so that $s(h \dots h + 2d/5) = P$. Thus we get that $\Delta = h - g$, since the shift of the longest square substring equals the distance between copies of P in the first and second half of A . Now, find the largest nonnegative integer $j \leq 2d/5$ such that $s(g - j \dots g) = s(h - j \dots h)$. All we are doing in this step is extending the copies of P backwards while keeping them identical (pictured in the second image of Fig. 3 as light blue rectangles). This takes $\tilde{O}(\sqrt{d})$ time via Grover search.

Next, in a similar fashion, we find the largest positive integer $k \leq \Delta - j$ such that $s(g, \dots, g + k) = s(h, \dots, h + k)$. In this step we are maximally extending the copies of P forwards while making sure they do not overlap with our previous extension.

Now, let j' be the positive integer such that $A = s(g - j' \dots g - j' + \ell)$ is the optimal solution whose first half and second half each contain the copies of P we are considering. Then because A is a square substring with shift $\Delta = h - g$, we have $s(g - j' \dots g) = s(h - j' \dots h)$, which implies that $j \geq j'$ by construction. But since A is square we also have $s(g \dots g + \Delta - j') = s(h \dots h + \Delta - j')$. Then the definition of k together with the observation that $j \geq j'$ forces $k = \Delta - j$ (this is pictured in the bottom two images of Fig. 3, where the left green substring and right blue substring are bordering each other). Combining these observations together, we get that $s(g - j \dots h + k)$ is a square substring of size $2\Delta = \ell$. Thus returning this substring produces the desired solution.

Case 2: Multiple Copies It remains to consider the case where S contains multiple copies of P . In this case, we use the quantum algorithm for Exact String Matching to find the rightmost and second rightmost copies of P in S in $\tilde{O}(\sqrt{d})$ time. Suppose that these copies start after positions h and h' respectively, so that

$$P = s(h \dots h + 2d/5] = s(h' \dots h' + 2d/5]$$

with $h' < h$. Then since $2|P| = 4d/5 > (3/5 + \varepsilon)d = |S|$, we know by Lemma 2.3 that P has minimum period $p = h - h'$ and every appearance of P in S starts some multiple of p away from h . Moreover, all the copies of P in S overlap each other and together form one larger p -periodic substring in s . Our next step will be to extend these periodic parts to maximal periodic substrings, which will help us locate a large square substring.

By Exact String Matching, we can find the leftmost copy $s(l \dots l + 2d/5]$ of P in S in $\tilde{O}(\sqrt{d})$, where the integer $l + 1$ is the starting position of this copy. By our earlier discussion, we know that the string $s(l \dots h + 2d/5]$ is p -periodic.

We now extend the original pattern P as well as the leftmost copy of P in S backwards while maintaining the property of being p -periodic.

Formally, we find the largest nonnegative integer $j_1 \leq 2d/5$ such that $s(g - j_1 \dots g + 2d/5]$ is p -periodic and the largest nonnegative integer $j_2 \leq l + 1 - g - 2d/5$ such that $s(l - j_2 \dots h + 2d/5]$ is p -periodic. Because we upper bound $j_1, j_2 \leq O(d)$, extending the strings in this way takes $\tilde{O}(\sqrt{d})$ time via Grover search. We now split into two further subcases, depending on how far back the strings are extended.

Case 2a: Single Periodic Substring

Suppose we get $j_2 = l + 1 - g - 2d/5$. This means that we were able to extend the leftmost copy of P in S so far back that it overlapped with our original pattern P contained in the first half of A . It follows that the substring $s(g \dots h + 2d/5]$ is p -periodic. In particular, we deduce that the substring starting from the original pattern P in the first half of A to its Δ -shifted copy in the second half of A is contained in this p -periodic part. Since A is a square substring, it follows that its prefix which ends at position $g + 2d/5$ of s is also p -periodic. Thus, position $g - j_1$ of s occurs before the first character of A . This reasoning is depicted in the second image of Fig. 4.

We now extend this entire p -periodic substring forward. Find the largest nonnegative integer $k \leq (1 + \varepsilon)d$ such that $s(g - j_1 \dots g + k]$ is p -periodic. Since $j_1, k \leq O(\log d)$ this takes $\tilde{O}(\sqrt{d})$ time via Grover search. As pictured in the bottom image of Fig. 4, since A is a square string and the end of its first half is p -periodic, the end of its second half is p -periodic as well. Thus position $g + k$ in s occurs after the final character of A .

We now have a p -periodic string $s(g - j_1 \dots g + k]$ which contains A , and thus has length at least ℓ . This means that A has period p as well. We claim the shift Δ associated with A is an integer multiple of p .

Indeed, by definition, A is Δ -periodic. Then because A has length $2\Delta \geq \Delta + p$, by Lemma 2.2 we know that A is $\gcd(p, \Delta)$ -periodic as well. Since P is a substring of A , P must have period $\gcd(p, \Delta)$ too. But p is the minimum period of P . Hence $p = \gcd(p, \Delta)$, so $\Delta = pm$ for some positive integer m as claimed.




Fig. 4 An example of how we extract a square substring when there are multiple copies of P and the light blue substring spanning the copies of P across both halves of A is p -periodic. In this case, because A is square, when we extend the initial pattern P backwards to a p -periodic string we cover a prefix of A . Similarly, when we extend copies of P forward to a p -periodic string we cover a suffix of A . Curly braces indicate parts of the strings guaranteed to be identical because A is square (Color figure online)



Fig. 5 If extending P backwards (pictured by the light blue substrings) to a p -periodic substring does not allow us to cover the prefix of A , then these p -periodic parts are identical because A is square. Thus the distance Δ' between the beginnings of the light blue substrings equals the true shift Δ (Color figure online)

Thus A has length $\ell = 2m \cdot p$. Our p -periodic string $s(g - j_1 \dots g + k)$ is guaranteed to be at least this long. Thus, we can simply return the substring $s(g - j_1 \dots g + 2mp)$. Because this substring is p -periodic and has length an even multiple of p , it is a square substring. Because it has length equal to A , it is a longest square substring as desired.

Case 2b: Disjoint Periodic Substrings

If we do not fall into **Case 2a**, then we must have $j_2 < l + 1 - g - 2d/5$, so that the p -periodic substrings $s(g - j_1 \dots g + 2d/5)$ and $s(l - j_2 \dots h + 2d/5)$ do not overlap. In this case, we construct two candidate solutions, and afterwards prove that one of them is guaranteed to be a longest square substring.

Define $\Delta' = (l - j_2) - (g - j_1)$. Via Grover search in $\tilde{O}(\sqrt{d})$ time, we find the largest nonnegative integers $b, b' \leq \Delta'$ such that $s(g - j_1 - b \dots g - j_1 + 1) = s(l - j_2 - b \dots h - j_1 + 1)$ and $s(g - j_1 \dots g - j_1 + b') = s(l - j_2 \dots l - j_2 + b')$. We then set string $B = s(g - j_1 - b \dots l - j_2 + b')$ to be our first candidate solution. Intuitively, this candidate corresponds to a guess that $\Delta = \Delta'$.

To construct the second candidate, we use a similar procedure, but first extend the strings forward. Using Grover search, we find the largest positive integers $k_1 \leq l - j_2 - g$ and $k_2 \leq (1 + \varepsilon)d$ such that $s(g \dots g + k_1)$ and $s(l \dots l + k_2)$ are each p -periodic, in $\tilde{O}(\sqrt{d})$ time. Set $\Delta'' = (l + k_2) - (g + k_1)$. Then, as before, we find the largest nonnegative integers $c, c' \leq \Delta''$ such that $s(g + k_1 - c \dots g + k_1) = s(l + k_2 - c \dots l + k_2)$ and $s(g + k_1 \dots g + k_1 + c') = s(l + k_2 \dots l + k_2 + c')$. The string $C = s(g + k_1 - c \dots l + k_2 + c')$ is then our second candidate. Intuitively, this corresponds to a guess that $\Delta = \Delta''$.

We can check if B and C are square in $\tilde{O}(\sqrt{d})$ time by Grover search. If neither of them are square we report that we find no square substring. Otherwise, we return



Fig. 6 If in **Case 2b**, extending P backwards (pictured by the light blue substring on the left) to a p -periodic substring covers a prefix of A , then when we extend P forward in the same way (pictured by the left green substring) the result cannot cross into the second half of A (if it did, we would have a single connected periodic substring and fall into **Case 2a**). Then the p -periodic parts at the ends of each half must be identical because A is square. Thus the distance Δ'' between the ends of the green substrings equals the true shift Δ (Color figure online)

the largest square substring among these two. It remains to prove that this procedure is correct. There are two cases to consider, based off how large j_1 is relative to the position of A .

First, suppose that position $g - j_1 + 1$ in s is a character in the first half of A . Then, as depicted in Fig. 5, since A is square, $l - j_2 + 1$ must also be in the second half of A , and in fact be exactly Δ characters to the right of $g - j_1 + 1$ (because if this position was earlier, it would mean we could have picked j_1 larger and still had a p -periodic string). Thus $\Delta = (l - j_2 + 1) - (g - j_1 + 1) = \Delta'$ is forced. Then when we construct the string B by searching backwards and forwards from positions $g - j_1 + 1$ and $l - j_2 + 1$ we will in fact find a square string of length A , and B will our desired longest square substring.

Otherwise, position $g - j_1 + 1$ in s is placed before every character of A . Then as depicted in Fig. 6, since A is square, position $l - j_2$ must be in the first half of A . Consequently, when we extend P forward to position $g + k_1$, this position is also in the first half of A (otherwise the p -periodic parts would overlap, and we would have been in **Case 2a** instead). As in the previous case, using the fact that A is a square again, we get that position $l + k_2$ must be exactly Δ characters to the right of $g + k_1$. So $\Delta = (l + k_2) - (g + k_1) = \Delta''$ is forced. Then when we construct the string C by searching backwards and forwards from positions $g + k_1$ and $l + k_2$ we find a square string of length A , so C will be our desired longest square substring.

This handles all of the cases. So far, we have a described an algorithm that, for any integer i , will find the longest square substring of s with size in $[d, (1 + \varepsilon)d)$ with probability at least $\Omega(d/n)$ (recall this is the probability that g is good), in time $\tilde{O}(\sqrt{d})$. By amplitude amplification and trying out the $O(\log n)$ choices of i in decreasing order, we recover an algorithm for the Longest Square Substring problem which runs in

$$\tilde{O}(\sqrt{d} \cdot \sqrt{n/d}) = \tilde{O}(\sqrt{n})$$

time, as desired. \square

We show that our algorithm is optimal by giving a quantum query lower bound of $\Omega(\sqrt{n})$ for finding the longest square substring. This proof is essentially already present in [12], where the authors give a lower bound for finding the longest *palindromic* substring, but we sketch the argument here for completeness.

Proposition 5.2 *Any quantum algorithm that computes the longest square substring of a string of length n requires $\Omega(\sqrt{n})$ queries.*

Proof Let S be the set of strings of length $2n$ over the alphabet $\{0, 1\}$ which contain at most one occurrence of the character 1. In [21] the authors prove that deciding with whether a given string $s \in S$ is the string consisting of all 0s requires $\Omega(\sqrt{n})$ queries in the quantum setting.

The longest square substring of the 0s string of length $2n$ is just the entire string, and has length $2n$. However, every other string in S has an odd number of 1s, and thus has longest square substring of size strictly less than $2n$. So solving the Longest Square Substring problem lets us decide if a string from S is the all 0s string, which means that any quantum algorithm solving this problem requires $\Omega(\sqrt{n})$ queries as well. \square

6 Open Problems

We conclude by mentioning several open questions related to our work.

- Our $\tilde{O}(n^{2/3})$ -time algorithm for LCS assumes that the input characters are integers in $[\text{poly}(n)]$. This assumption was used for constructing string synchronizing sets sublinearly (Sect. 3.3.2). However, the previous $\tilde{O}(n^{5/6})$ -time algorithm by Le Gall and Seddighin [12] can work with *general ordered alphabet*, where the only allowed query is to compare two symbols $S[i]$, $S[j]$ in the input strings (with three possible outcomes $S[i] > S[j]$, $S[i] = S[j]$, or $S[i] < S[j]$). Is $\tilde{O}(n^{2/3})$ query complexity (or even time complexity) achievable in this more restricted setting? Alternatively, can we show a better query lower bound?
- Our algorithm for the Minimal String Rotation problem (and other related problems in Sect. 4) has time complexity (and query complexity) $n^{1/2+o(1)}$. Can we reduce the $n^{o(1)}$ factor down to $\text{poly} \log(n)$? A subsequent work by Childs, Kothari, Kovacs-Deak, Sundaram, and Wang [82] showed such an improvement for the decision version of Minimal String Rotation, but the question remains open for the search version.
- In our time-efficient implementation of the LCS algorithm, we used a simple sampling technique to bypass certain restrictions on 2D range query data structures (Sect. 3.2.3). Can this idea have further applications in designing time-efficient quantum walk algorithms? As a simple example, we can use this idea to get an $\tilde{O}(n^{2/3})$ -time comparison-based algorithm for the element distinctness problem with *simpler implementation*. At the beginning, uniformly sample r items x_1, \dots, x_r from the input array, and sort them so that $x_1 \leq \dots \leq x_r$. Then, we create a hash table with $r + 1$ buckets each having $O(\log n)$ capacity, where the hash function $h(x)$ is defined as the index i such that $x_i \leq x < x_{i+1}$, which can be found by binary search. Then, each insertion, deletion, and search operation can be performed in $O(\log n)$ time, provided that the buckets do not overflow. The error caused by overflows can be analyzed using Ambainis' proof of [13, Lemma 6]. In comparison, Ambainis' implementation [13] additionally used a skip list, and Jeffery's (non-comparison-based) implementation used a quantum radix tree [52, Section 3.3.4].

Acknowledgements We thank Virginia Vassilevska Williams, Ryan Williams, and Yinzhan Xu for several helpful discussions. We additionally thank Virginia Vassilevska Williams for several useful comments on the writeup of this paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Knuth, D.E., Morris, J.H., Jr., Pratt, V.R.: Fast pattern matching in strings. *SIAM J. Comput.* **6**(2), 323–350 (1977). <https://doi.org/10.1137/0206024>
- Karp, R.M., Rabin, M.O.: Efficient randomized pattern-matching algorithms. *IBM J. Res. Dev.* **31**(2), 249–260 (1987). <https://doi.org/10.1147/rd.312.0249>
- Weiner, P.: Linear pattern matching algorithms. In: *Proceedings of the 14th Annual Symposium on Switching and Automata Theory*, pp. 1–11 (1973). <https://doi.org/10.1109/SWAT.1973.13>
- Farach, M.: Optimal suffix tree construction with large alphabets. In: *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS 1997)*, pp. 137–143 (1997). <https://doi.org/10.1109/SFCS.1997.646102>
- Babenko, M.A., Starikovskaya, T.: Computing longest common substrings via suffix arrays. In: *Proceedings of the 3rd International Computer Science Symposium in Russia (CSR 2008)*, Theory and Applications, pp. 64–75 (2008). https://doi.org/10.1007/978-3-540-79709-8_10
- Booth, K.S.: Lexicographically least circular substrings. *Inf. Process. Lett.* **10**(4/5), 240–242 (1980). [https://doi.org/10.1016/0020-0190\(80\)90149-0](https://doi.org/10.1016/0020-0190(80)90149-0)
- Shiloach, Y.: Fast canonization of circular strings. *J. Algorithms* **2**(2), 107–121 (1981). [https://doi.org/10.1016/0196-6774\(81\)90013-4](https://doi.org/10.1016/0196-6774(81)90013-4)
- Duval, J.-P.: Factorizing words over an ordered alphabet. *J. Algorithms* **4**(4), 363–381 (1983). [https://doi.org/10.1016/0196-6774\(83\)90017-2](https://doi.org/10.1016/0196-6774(83)90017-2)
- Ramesh, H., Vinay, V.: String matching in $\tilde{O}(\sqrt{n} + \sqrt{m})$ quantum time. *J. Discrete Algorithms* **1**(1), 103–110 (2003). [https://doi.org/10.1016/S1570-8667\(03\)00010-8](https://doi.org/10.1016/S1570-8667(03)00010-8)
- Vishkin, U.: Deterministic sampling: a new technique for fast pattern matching. *SIAM J. Comput.* **20**(1), 22–40 (1991). <https://doi.org/10.1137/0220002>
- Grover, L.K.: A fast quantum mechanical algorithm for database search. In: *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing (STOC 1996)*, pp. 212–219 (1996). <https://doi.org/10.1145/237814.237866>
- Le Gall, F., Seddighin, S.: Quantum meets fine-grained complexity: Sublinear time quantum algorithms for string problems. In: *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, pp. 97:1–97:23 (2022). <https://doi.org/10.4230/LIPIcs.ITCS.2022.97>
- Ambainis, A.: Quantum walk algorithm for element distinctness. *SIAM J. Comput.* **37**(1), 210–239 (2007). <https://doi.org/10.1137/S0097539705447311>
- Wang, Q., Ying, M.: Quantum algorithm for lexicographically minimal string rotation. *CoRR* (2020). [arXiv:2012.09376](https://arxiv.org/abs/2012.09376)
- Durr, C., Høyer, P.: A quantum algorithm for finding the minimum. Preprint (1996). [arXiv:quant-ph/9607014](https://arxiv.org/abs/quant-ph/9607014)
- Apostolico, A., Iliopoulos, C.S., Paige, R.: On $O(n \log n)$ cost parallel algorithm for the single function coarsest partition problem. In: *Parallel Algorithms and Architectures, International Workshop, 1987, Proceedings*, pp. 70–76 (1987). https://doi.org/10.1007/3-540-18099-0_30
- Iliopoulos, C.S., Smyth, W.F.: Optimal algorithms for computing the canonical form of a circular string. *Theor. Comput. Sci.* **92**(1), 87–105 (1992). [https://doi.org/10.1016/0304-3975\(92\)90137-5](https://doi.org/10.1016/0304-3975(92)90137-5)

18. Aaronson, S., Shi, Y.: Quantum lower bounds for the collision and the element distinctness problems. *J. ACM* **51**(4), 595–605 (2004). <https://doi.org/10.1145/1008731.1008735>
19. Kutin, S.: Quantum lower bound for the collision problem with small range. *Theory Comput.* **1**(1), 29–36 (2005). <https://doi.org/10.4086/toc.2005.v001a002>
20. Ambainis, A.: Polynomial degree and lower bounds in quantum complexity: collision and element distinctness with small range. *Theory Comput.* **1**(1), 37–46 (2005). <https://doi.org/10.4086/toc.2005.v001a003>
21. Bennett, C.H., Bernstein, E., Brassard, G., Vazirani, U.V.: Strengths and weaknesses of quantum computing. *SIAM J. Comput.* **26**(5), 1510–1523 (1997). <https://doi.org/10.1137/S0097539796300933>
22. Starikovskaya, T., Vildhøj, H.W.: Time-space trade-offs for the longest common substring problem. In: Proceedings of the 24th Annual Symposium on Combinatorial Pattern Matching (CPM 2013), pp. 223–234 (2013). https://doi.org/10.1007/978-3-642-38905-4_22
23. Charalampopoulos, P., Crochemore, M., Iliopoulos, C.S., Kociumaka, T., Pissis, S.P., Radoszewski, J., Rytter, W., Waleń, T.: Linear-time algorithm for long LCF with k mismatches. In: Proceedings of the 29th Annual Symposium on Combinatorial Pattern Matching (CPM 2018), pp. 23:1–23:16 (2018). <https://doi.org/10.4230/LIPIcs.CPM.2018.23>
24. Amir, A., Charalampopoulos, P., Pissis, S.P., Radoszewski, J.: Longest common substring made fully dynamic. In: Proceedings of the 27th Annual European Symposium on Algorithms (ESA 2019), pp. 6:1–6:17 (2019). <https://doi.org/10.4230/LIPIcs.ESA.2019.6>
25. Amir, A., Charalampopoulos, P., Pissis, S.P., Radoszewski, J.: Dynamic and internal longest common substring. *Algorithmica* **82**(12), 3707–3743 (2020). <https://doi.org/10.1007/s00453-020-00744-0>
26. Ben-Nun, S., Golan, S., Kociumaka, T., Kraus, M.: Time-space tradeoffs for finding a long common substring. In: Proceedings of the 31st Annual Symposium on Combinatorial Pattern Matching (CPM 2020), pp. 5:1–5:14 (2020). <https://doi.org/10.4230/LIPIcs.CPM.2020.5>
27. Charalampopoulos, P., Gawrychowski, P., Pokorski, K.: Dynamic longest common substring in poly-logarithmic time. In: Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP 2020), pp. 27:1–27:19 (2020). <https://doi.org/10.4230/LIPIcs.ICALP.2020.27>
28. Charalampopoulos, P., Kociumaka, T., Pissis, S.P., Radoszewski, J.: Faster algorithms for longest common substring. In: Proceedings of the 29th Annual European Symposium on Algorithms (ESA 2021), pp. 30:1–30:17 (2021). <https://doi.org/10.4230/LIPIcs.ESA.2021.30>
29. Burkhardt, S., Kärkkäinen, J.: Fast lightweight suffix array construction and checking. In: Proceedings of the 14th Annual Symposium on Combinatorial Pattern Matching (CPM 2003), pp. 55–69 (2003). https://doi.org/10.1007/3-540-44888-8_5
30. Maekawa, M.: A \sqrt{N} algorithm for mutual exclusion in decentralized systems. *ACM Trans. Comput. Syst.* **3**(2), 145–159 (1985). <https://doi.org/10.1145/214438.214445>
31. Birenzweige, O., Golan, S., Porat, E.: Locally consistent parsing for text indexing in small space. In: Proceedings of the 31st ACM-SIAM Symposium on Discrete Algorithms (SODA 2020), pp. 607–626 (2020). <https://doi.org/10.1137/1.9781611975994.37>
32. Kempa, D., Kociumaka, T.: String synchronizing sets: sublinear-time BWT construction and optimal LCE data structure. In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019), pp. 756–767. ACM (2019). <https://doi.org/10.1145/3313276.3316368>
33. Magniez, F., Nayak, A., Roland, J., Santha, M.: Search via quantum walk. *SIAM J. Comput.* **40**(1), 142–164 (2011). <https://doi.org/10.1137/090745854>
34. Willard, D.E., Lueker, G.S.: Adding range restriction capability to dynamic data structures. *J. ACM* **32**(3), 597–617 (1985). <https://doi.org/10.1145/3828.3839>
35. Christian Worm Mortensen: Fully dynamic orthogonal range reporting on RAM. *SIAM J. Comput.* **35**(6), 1494–1525 (2006). <https://doi.org/10.1137/s0097539703436722>
36. Chan, T.M., Tsakalidis, K.: Dynamic orthogonal range searching on the RAM, revisited. In: Proceedings of the 33rd International Symposium on Computational Geometry (SoCG 2017), vol. 77, pp. 28:1–28:13 (2017). <https://doi.org/10.4230/LIPIcs.SocG.2017.28>
37. Masek, W.J., Paterson, M.: A faster algorithm computing string edit distances. *J. Comput. Syst. Sci.* **20**(1), 18–31 (1980). [https://doi.org/10.1016/0022-0000\(80\)90002-1](https://doi.org/10.1016/0022-0000(80)90002-1)
38. Backurs, A., Indyk, P.: Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). *SIAM J. Comput.* **47**(3), 1087–1097 (2018). <https://doi.org/10.1137/15M1053128>
39. Boroujeni, M., Ehsani, S., Ghodsi, M., HajiAghayi, M.T., Seddighin, S.: Approximating edit distance in truly subquadratic time: quantum and MapReduce. *J. ACM* **68**(3), 1–41 (2021). <https://doi.org/10.1145/3456807>

40. Chakraborty, D., Das, D., Goldenberg, E., Koucký, M., Saks, M.E.: Approximating edit distance within constant factor in truly sub-quadratic time. *J. ACM* **67**(6), 36:1–36:22 (2020). <https://doi.org/10.1145/3422823>
41. Naumovitz, T., Saks, M.E., Seshadhri, C.: Accurate and nearly optimal sublinear approximations to ulam distance. In: *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017)*, pp. 2012–2031 (2017). <https://doi.org/10.1137/1.9781611974782.131>
42. Montanaro, A.: Quantum pattern matching fast on average. *Algorithmica* **77**(1), 16–39 (2017). <https://doi.org/10.1007/s00453-015-0060-4>
43. Ambainis, A., Balodis, K., Iraids, J., Khadiev, K., Kļevickis, V., Prūsīs, K., Shen, Y., Smotrovs, J., Vihrovs, J.: Quantum lower and upper bounds for 2D-grid and Dyck language. In: *Proceedings of the 45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*, pp. 8:1–8:14 (2020). <https://doi.org/10.4230/LIPIcs.MFCS.2020.8>
44. Ambainis, A., Montanaro, A.: Quantum algorithms for search with wildcards and combinatorial group testing. *Quant. Inf. Comput.* **14**(5–6), 439–453 (2014). <https://doi.org/10.26421/QIC14.5-6-4>
45. Cleve, R., Iwama, K., Le Gall, F., Nishimura, H., Tani, S., Teruyama, J., Yamashita, S.: Reconstructing strings from substrings with quantum queries. In: *Proceedings of the 13th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2012)*, pp. 388–397 (2012). https://doi.org/10.1007/978-3-642-31155-0_34
46. Szegedy, M.: Quantum speed-up of Markov chain based algorithms. In: *Proceedings of the 45th Symposium on Foundations of Computer Science (FOCS 2004)*, pp. 32–41 (2004). <https://doi.org/10.1109/FOCS.2004.53>
47. Magniez, F., Santha, M., Szegedy, M.: Quantum algorithms for the triangle problem. *SIAM J. Comput.* **37**(2), 413–424 (2007). <https://doi.org/10.1137/050643684>
48. Jeffery, S., Kothari, R., Magniez, F.: Nested quantum walks with quantum data structures. In: *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2013)*, pp. 1474–1485 (2013). <https://doi.org/10.1137/1.9781611973105.106>
49. Le Gall, F.: Improved quantum algorithm for triangle finding via combinatorial arguments. In: *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS 2014)*, pp. 216–225 (2014). <https://doi.org/10.1109/FOCS.2014.31>
50. Belovs, A., Childs, A.M., Jeffery, S., Kothari, R., Magniez, F.: Time-efficient quantum walks for 3-distinctness. In: *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP 2013), Part I*, pp. 105–122 (2013). https://doi.org/10.1007/978-3-642-39206-1_10
51. Bernstein, D.J., Jeffery, S., Lange, T., Meurer, A.: Quantum algorithms for the subset-sum problem. In: *Proceedings of the 5th International Workshop on Post-Quantum Cryptography (PQCrypto 2013)*, pp. 16–33 (2013). https://doi.org/10.1007/978-3-642-38616-9_2
52. Jeffery, S.: Frameworks for quantum algorithms. PhD thesis, University of Waterloo (2014). <http://hdl.handle.net/10012/8710>
53. Aaronson, S., Chia, N.-H., Lin, H.-H., Wang, C., Zhang, R.: On the quantum complexity of closest pair and related problems. In: *Proceedings of the 35th Computational Complexity Conference (CCC 2020)*, pp. 16:1–16:43 (2020). <https://doi.org/10.4230/LIPIcs.CCC.2020.16>
54. Buhrman, H., Patro, S., Speelman, F.: A framework of quantum strong exponential-time hypotheses. In: *Proceedings of the 38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021)*, pp. 19:1–19:19 (2021). <https://doi.org/10.4230/LIPIcs.STACS.2021.19>
55. Buhrman, H., Loff, B., Patro, S., Speelman, F.: Limits of quantum speed-ups for computational geometry and other problems: Fine-grained complexity via quantum walks. In: *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, pp. 31:1–31:12 (2022). <https://doi.org/10.4230/LIPIcs.ITCS.2022.31>
56. Ambainis, A., Larka, N.: Quantum algorithms for computational geometry problems. In: *Proceedings of the 15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, pp. 9:1–9:10 (2020). <https://doi.org/10.4230/LIPIcs.TQC.2020.9>
57. Gusfield, D.: *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology*. Cambridge University Press (1997). <https://doi.org/10.1017/CBO9780511574931>
58. Crochemore, M., Rytter, W.: *Jewels of Stringology*. World Scientific (2002). <https://doi.org/10.1142/4838>
59. Crochemore, M., Hancart, C., Lecroq, T.: *Algorithms on Strings*. Cambridge University Press (2007). <https://doi.org/10.1017/CBO9780511546853>

60. Kociumaka, T., Starikovskaya, Ta., Vildhøj, H.W.: Sublinear space algorithms for the longest common substring problem. In: Proceedings of the 22th Annual European Symposium on Algorithms (ESA 2014), pp. 605–617 (2014). https://doi.org/10.1007/978-3-662-44777-2_50
61. Abboud, A., Williams, R.R., Yu, H.: More applications of the polynomial method to algorithm design. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015), pp. 218–230 (2015). <https://doi.org/10.1137/1.9781611973730.17>
62. Flouri, T., Giaquinta, E., Kobert, K., Ukkonen, E.: Longest common substrings with k mismatches. Inf. Process. Lett. **115**(6–8), 643–647 (2015). <https://doi.org/10.1016/j.ipl.2015.03.006>
63. Thankachan, S.V., Apostolico, A., Aluru, S.: A provably efficient algorithm for the k -mismatch average common substring problem. J. Comput. Biol. **23**(6), 472–482 (2016). <https://doi.org/10.1089/cmb.2015.0235>
64. Starikovskaya, T.: Longest common substring with approximately k mismatches. In: Proceedings of the 27th Annual Symposium on Combinatorial Pattern Matching (CPM 2016), pp. 21:1–21:11 (2016). <https://doi.org/10.4230/LIPIcs.CPM.2016.21>
65. Kociumaka, T., Radoszewski, J., Starikovskaya, T.: Longest common substring with approximately k mismatches. Algorithmica **81**(6), 2633–2652 (2019). <https://doi.org/10.1007/s00453-019-00548-x>
66. Gourdel, G., Kociumaka, T., Radoszewski, J., Starikovskaya, T.: Approximating longest common substring with k mismatches: Theory and practice. In: Proceedings of the 31st Annual Symposium on Combinatorial Pattern Matching (CPM 2020), pp. 16:1–16:15 (2020). <https://doi.org/10.4230/LIPIcs.CPM.2020.16>
67. Apostolico, A., Crochemore, M.: Optimal canonization of all substrings of a string. Inf. Comput. **95**(1), 76–95 (1991). [https://doi.org/10.1016/0890-5401\(91\)90016-U](https://doi.org/10.1016/0890-5401(91)90016-U)
68. Babenko, M.A., Kolesnichenko, I.I., Starikovskaya, T.: On minimal and maximal suffixes of a substring. In: Proceedings of the 24th Annual Symposium on Combinatorial Pattern Matching (CPM 2013), pp. 28–37, Springer (2013). https://doi.org/10.1007/978-3-642-38905-4_5
69. Babenko, M.A., Gawrychowski, P., Kociumaka, T., Kolesnichenko, I.I., Starikovskaya, T.: Computing minimal and maximal suffixes of a substring. Theor. Comput. Sci. **638**, 112–121 (2016). <https://doi.org/10.1016/j.tcs.2015.08.023>
70. Kociumaka, T.: Minimal suffix and rotation of a substring in optimal time. In: Proceedings of the 27th Annual Symposium on Combinatorial Pattern Matching (CPM 2016), pp. 28:1–28:12 (2016). <https://doi.org/10.4230/LIPIcs.CPM.2016.28>
71. Urabe, Y., Nakashima, Y., Inenaga, S., Bannai, H., Takeda, M.: Longest Lyndon substring after edit. In: Proceedings of the 29th Annual Symposium on Combinatorial Pattern Matching, (CPM 2018), pp. 19:1–19:10 (2018). <https://doi.org/10.4230/LIPIcs.CPM.2018.19>
72. Crochemore, M.: An optimal algorithm for computing the repetitions in a word. Inf. Process. Lett. **12**(5), 244–250 (1981). [https://doi.org/10.1016/0020-0190\(81\)90024-7](https://doi.org/10.1016/0020-0190(81)90024-7)
73. Main, M.G., Lorentz, R.J.: An $O(n \log n)$ algorithm for finding all repetitions in a string. J. Algorithms **5**(3), 422–432 (1984). [https://doi.org/10.1016/0196-6774\(84\)90021-X](https://doi.org/10.1016/0196-6774(84)90021-X)
74. Amir, A., Boneh, I., Charalampopoulos, P., Kondratovsky, E.: Repetition detection in a dynamic string. In: Proceedings of the 27th Annual European Symposium on Algorithms (ESA 2019), pp. 5:1–5:18 (2019). <https://doi.org/10.4230/LIPIcs.ESA.2019.5>
75. Bille, P., Gawrychowski, P.G., Inge, L., Landau, G.M., Weimann, O.: Longest common extensions in trees. In: Proceedings of the 26th Annual Symposium on Combinatorial Pattern Matching (CPM 2015), pp. 52–64 (2015). https://doi.org/10.1007/978-3-319-19929-0_5
76. Gawrychowski, P., Kociumaka, T., Rytter, W., Waleń, T.: Faster longest common extension queries in strings over general alphabets. In: Proceedings of the 27th Annual Symposium on Combinatorial Pattern Matching (CPM 2016), pp. 5:1–5:13 (2016). <https://doi.org/10.4230/LIPIcs.CPM.2016.5>
77. Alzamel, M., Crochemore, M., Iliopoulos, C.S., Kociumaka, T., Radoszewski, J., Rytter, W., Straszyski, J., Waleń, T., Zuba, W.: Quasi-linear-time algorithm for longest common circular factor. In: Proceedings of the 30th Annual Symposium on Combinatorial Pattern Matching (CPM 2019), pp. 25:1–25:14 (2019). <https://doi.org/10.4230/LIPIcs.CPM.2019.25>
78. Kempa, D., Kociumaka, T.: Breaking the $O(n)$ -barrier in the construction of compressed suffix arrays. CoRR, (2021). To appear in SODA 2023. [arXiv:2106.12725](https://arxiv.org/abs/2106.12725)
79. Kociumaka, T., Radoszewski, J., Rytter, W., Waleń, T.: Internal pattern matching queries in a text and applications. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015), pp. 532–551 (2015). <https://doi.org/10.1137/1.9781611973730.36>

80. Kociumaka, T.: Efficient data structures for internal queries in texts. PhD thesis, University of Warsaw (2018). <https://depotuw.ceon.pl/handle/item/3614>
81. Jin, C., Nogler, J.: Quantum speed-ups for string synchronizing sets, longest common substring, and k-mismatch matching. CoRR (2022). To appear in SODA 2023. [arXiv:2211.15945](https://arxiv.org/abs/2211.15945)
82. Childs, A.M., Kothari, R., Kovacs-Deak, M., Sundaram, A., Wang, D.: Quantum divide and conquer. CoRR 2022. [arXiv:2210.06419](https://arxiv.org/abs/2210.06419)
83. Kent, C., Lewenstein, M., Sheinwald, D.: On demand string sorting over unbounded alphabets. Theor. Comput. Sci. **426**, 66–74 (2012). <https://doi.org/10.1016/j.tcs.2011.12.001>
84. Fine, N.J., Wilf, H.S.: Uniqueness theorems for periodic functions. Proc. Am. Math. Soc. **16**(1), 109–114 (1965). <https://doi.org/10.2307/2034009>
85. Plandowski, W., Rytter, W.: Application of Lempel-Ziv encodings to the solution of words equations. In: Proceedings of the 25th International Colloquium on Automata, Languages and Programming (ICALP 1998), pp. 731–742 (1998). <https://doi.org/10.1007/BFb0055097>
86. Ambainis, A.: Quantum query algorithms and lower bounds. In: Classical and New Paradigms of Computation and their Complexity Hierarchies, pp. 15–32. Springer (2004). https://doi.org/10.1007/978-1-4020-2776-5_2
87. Buhrman, H., de Wolf, R.: Complexity measures and decision tree complexity: a survey. Theor. Comput. Sci. **288**(1), 21–43 (2002). [https://doi.org/10.1016/S0304-3975\(01\)00144-X](https://doi.org/10.1016/S0304-3975(01)00144-X)
88. Barenco, A., Bennett, C.H., Cleve, R., DiVincenzo, D.P., Margolus, N., Shor, P., Sleator, T., Smolin, J.A., Weinfurter, H.: Elementary gates for quantum computation. Phys. Rev. A **52**, 3457–3467 (1995). <https://doi.org/10.1103/PhysRevA.52.3457>
89. Brassard, G., Høyer, P., Mosca, M., Tapp, A.: Quantum amplitude amplification and estimation. Preprint (2000). [arXiv:quant-ph/0005055](https://arxiv.org/abs/quant-ph/0005055)
90. Høyer, P., Mosca, M., de Wolf, R.: Quantum search on bounded-error inputs. In: Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP 2003), pp. 291–299 (2003). https://doi.org/10.1007/3-540-45061-0_25
91. de Wolf, R.: Quantum computing: Lecture notes. CoRR (2019). [arXiv:1907.09415v2](https://arxiv.org/abs/1907.09415v2)
92. Blelloch, G.E., Golovin, D., Vassilevska, V.: Uniquely represented data structures for computational geometry. In: Proceedings of the 11th Scandinavian Workshop on Algorithm Theory (SWAT 2008), pp. 17–28 (2008). https://doi.org/10.1007/978-3-540-69903-3_4
93. Pugh, W.: Skip lists: a probabilistic alternative to balanced trees. Commun. ACM **33**(6), 668–676 (1990). <https://doi.org/10.1145/78973.78977>
94. Pugh, W.: A skip list cookbook. Technical Report CS-TR-2286.1, University of Maryland at College Park, USA (1990). <http://hdl.handle.net/1903/544>
95. Indyk, P.: A small approximately min-wise independent family of hash functions. J. Algorithms **38**(1), 84–90 (2001). <https://doi.org/10.1006/jagm.2000.1131>
96. Chen, K.T., Fox, R.H., Lyndon, R.C.: Free differential calculus, IV: the quotient groups of the lower central series. Ann. Math. (1958). <https://doi.org/10.2307/1970044>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.