Efficient Parallel Output-Sensitive Edit Distance

- ² Xiangyun Ding ⊠[©]
- ³ University of California, Riverside
- ₄ Xiaojun Dong ⊠©
- ⁵ University of California, Riverside
- ₀ Yan Gu ⊠©
- 7 University of California, Riverside
- 🔋 Youzhe Liu 🖂 🗈
- 9 University of California, Riverside

10 Yihan Sun 🖂 回

- ¹¹ University of California, Riverside
- 12 Abstract -

In this paper, we study efficient parallel edit distance algorithms, both in theory and in practice. 13 Given two strings A[1..n] and B[1..m], and a set of operations allowed to edit the strings, the 14 edit distance between A and B is the minimum number of operations required to transform A15 into B. In this paper, we use edit distance to refer to the Levenshtein distance, which allows 16 for unit-cost single-character edits (insertions, deletions, substitutions). Sequentially, a standard 17 Dynamic Programming (DP) algorithm solves edit distance with $\Theta(nm)$ cost. In many real-world 18 applications, the strings to be compared are similar to each other and have small edit distances. 19 To achieve highly practical implementations, we focus on output-sensitive parallel edit-distance 20 algorithms, i.e., to achieve asymptotically better cost bounds than the standard $\Theta(nm)$ algorithm 21 when the edit distance is small. We study four algorithms in the paper, including three algorithms 22 based on Breadth-First Search (BFS), and one algorithm based on Divide-and-Conquer (DaC). Our 23 BFS-based solution is based on the Landau-Vishkin algorithm. We implement three different data 24 structures for the longest common prefix (LCP) queries needed in the algorithm: the classic solution 25 using parallel suffix array, and two hash-based solutions proposed in this paper. Our DaC-based 26 solution is inspired by the output-insensitive solution proposed by Apostolico et al., and we propose 27 a non-trivial adaption to make it output-sensitive. All of the algorithms studied in this paper have 28 good theoretical guarantees, and they achieve different tradeoffs between work (total number of 29 operations), span (longest dependence chain in the computation), and space. 30 We test and compare our algorithms on both synthetic data and real-world data, including DNA 31

We test and compare our algorithms on both synthetic data and real-world data, including DNA sequences, Wikipedia texts, GitHub repositories, etc. Our BFS-based algorithms outperform the existing parallel edit-distance implementation in ParlayLib in all test cases. On cases with fewer than 10⁵ edits, our algorithm can process input sequences of size 10⁹ in about ten seconds, while ParlayLib can only process sequences of sizes up to 10⁶ in the same amount of time. By comparing our algorithms, we also provide a better understanding of the choice of algorithms for different input patterns. We believe that our paper is the first systematic study in the theory and practice of parallel edit distance.

³⁹ 2012 ACM Subject Classification Theory of computation \rightarrow Parallel algorithms

Keywords and phrases Edit Distance, Parallel Algorithms, String Algorithms, Dynamic Program ming, Pattern Matching

- 42 Digital Object Identifier 10.4230/LIPIcs.ESA.2023.46
- 43 Related Version Full Version: https://arxiv.org/abs/2306.17461 [16]
- 44 Supplementary Material Source Code: https://github.com/ucrparlay/Edit-Distance [17]
- 45 Funding This work is supported by NSF grants CCF-2103483, CCF-2238358, and IIS-2227669, and
- ⁴⁶ UCR Regents Faculty Fellowships.



licensed under Creative Commons License CC-BY 4.0 31st Annual European Symposium on Algorithms (ESA 2023). Editors: Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman; Article No. 46; pp. 46:1–46:18



Leibniz International Proceedings in Informatics

© Xiangyun Ding, Xiaojun Dong, Yan Gu, Youzhe Liu, Yihan Sun;

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

47 **1** Introduction

Given two strings (sequences) A[1..n] and B[1..m] over an alphabet Σ and a set of operations 48 allowed to edit the strings, the *edit distance* between A and B is the minimum number of 49 operations required to transform A into B. WLOG, we assume $m \leq n$. The most commonly 50 used metric is the *Levenshtein distance* which allows for unit-cost single-character edits 51 (insertions, deletions, substitutions). In this paper, we use *edit distance* to refer to the 52 Levenshtein distance. We use k to denote the edit distance for strings A and B throughout 53 this paper. Edit distance is usually used to measure the similarity of two strings (a smaller 54 distance means higher similarity). 55

Edit distance is a fundamental problem in computer science, and is introduced in most 56 algorithm textbooks (e.g., [14, 15, 23]). In practice, it is widely used in version-control 57 software [54], computational biology [12, 31, 39], natural language processing [10, 29], and 58 spell corrections [28]. It is also closely related to other important problems such as longest 59 common subsequence (LCS) [50], longest increasing subsequence (LIS) [34], approximate 60 string matching [56], and multi-sequence alignment [59]. The classic dynamic programming 61 (DP) solution can compute edit distance in O(nm) work (number of operations) between 62 two strings of sizes n and m. This complexity is impractical if the input strings are large. 63 One useful observation is that, in real-world applications, the strings to be compared are 64 usually *reasonably similar*, resulting in a relatively small edit distance. For example, in many 65 version-control softwares (e.g., Git), if the two committed versions are similar (within a 66 certain number of edits), the "delta" file is stored to track edits. Otherwise, if the difference 67 is large, the system directly stores the new version. Most of the DNA or genome sequence 68 alignment applications also only focus on when the number of edits is small [39]. We say 69 an edit distance algorithm is *output-sensitive* if the work is o(nm) when k = o(n). Many 70 more efficient and/or practical algorithms were proposed in this setting with cost bounds 71 parameterized by k [19, 20, 21, 22, 26, 35, 36, 37, 46, 47, 49]. 72

Considering the ever-growing data size and plateaued single-processor performance, it is 73 crucial to consider parallel solutions for edit distance. Although the problem is simple and 74 well-studied in the sequential setting, we observe a huge gap between theory and practice 75 in the parallel setting. The few implementations we know of [7, 55, 58] simply parallelize 76 the O(nm)-work sequential algorithm and require O(n) span (longest dependence chain), 77 which indicates low-parallelism and redundant work when $k \ll n$. Meanwhile, numerous 78 theoretical parallel algorithms exist [1, 3, 20, 37, 41, 48], but it remains unknown whether 79 these algorithms are practical (i.e., can be implemented with reasonable engineering effort), 80 and if so, whether they can yield high performance. The goal of this paper is to formally 81 study parallel solutions for edit distance. By carefully studying existing theoretical solutions, 82 we develop new output-sensitive parallel solutions with good theoretical guarantees 83 and high performance in practice. We also conduct in-depth experimental studies on 84 existing and our new algorithms. 85

The classic dynamic programming (DP) algorithm solves edit distance by using the states G[i, j] as the edit distance of transforming A[1..i] to B[1..j]. G[i, j] can be computed as:

$$G[i,j] = \begin{cases} G[i-1,j-1] & \text{if } A[i] = B[j] \text{ and } i > 0, j > 0 \\ 1 + \min(G[i-1,j], G[i-1,j-1], G[i,j-1]) & \text{otherwise} \end{cases}$$

$$G[i,j] = \max(i,j) & \text{if } i = 0 \text{ or } j = 0 \end{cases}$$

⁹¹ A simple parallelization of this computation is to compute all states with the same i + j

value in parallel, and process all i + j values in an incremental order [7, 55, 58]. However,

46:2 Efficient Parallel Output-Sensitive Edit Distance

Algorithm	Work	Span	\mathbf{Space}^*	Algorithm	Work	Span	\mathbf{Space}^*
BFS-SA	$O(n+k^2)$	$\tilde{O}(k)$	O(n)	$\mathbf{BFS} extsf{-}\mathbf{Hash}^*$	$O(n + k^2 \log n)$	$\tilde{O}(k)$	O(n)
DaC-SD	$O(nk\log k)$	$\tilde{O}(1)$	O(nk)	$\mathbf{BFS}\text{-}\mathbf{B}\text{-}\mathbf{Hash}^*$	$O(n+k^2b\log n)$	$\tilde{O}(kb)$	O(n/b+k)

Table 1 Algorithms in this paper. k is the edit distance. b is the block size. *: Monte Carlo algorithms due to the use of hashing. "Space*" means auxiliary space used in addition to the input. Here we assume constant alphabet size for BFS-SA.

⁹³ this approach has low parallelism as it requires n + m rounds to finish. Later work [1, 3, 41]⁹⁴ improved parallelism using a *divide-and-conquer* (*DaC*) approach and achieved $\tilde{O}(n^2)$ ⁹⁵ work and polylog(n) span. These algorithms use the monotonicity of the DP recurrence, and ⁹⁶ are complicated. There are two critical issues in the DaC approaches. First, to the best of ⁹⁷ our knowledge, there exist no implementations given the sophistication of these algorithms. ⁹⁸ Second, they are not output-sensitive ($\tilde{O}(nm)$ work), which is inefficient when $k \ll n$.

Alternatively, many existing solutions, both sequentially [19, 20, 21, 22, 26, 35, 36, 46, 47] 99 and in parallel [20, 37] use output-sensitive algorithms, and achieve O(nk) or $O(n+k^2)$ work 100 and $\hat{O}(k)$ span. These algorithms view DP table as a grid-like DAG, where each state (cell) 101 (x, y) has three incoming edges from (x-1, y), (x, y-1), and (x-1, y-1) (if they exist). The 102 edge weight is 0 from (x - 1, y - 1) to (x, y), when A[x] = B[y], and 1 otherwise. Then edit 103 distance is equivalent to the shortest path from (0,0) to (n,m). An example is given in Fig. 1. 104 Since the edge weights can only be 0 or 1, we can use **breadth-first search** (BFS) from the 105 cell (0,0) until (n,m) is reached. Ukkonen [56] further showed that using **longest common** 106 prefix (LCP) queries based on suffix trees or suffix arrays, the work can be improved to 107 $O(n + k^2)$. Landau and Vishkin [37] parallelized this algorithm (see Sec. 3). While the 108 sequential output-sensitive algorithms have been widely used in practice [21, 26, 36, 46, 47], 109 we are unaware of any existing implementations for the parallel version. 110

We systematically study parallel output-sensitive edit distance, using both the BFS-based 111 and the DaC-based approaches. Our first effort is to implement the BFS-based Landau-112 Vishkin algorithm with our carefully-engineered parallel suffix array (SA) implementation, 113 referred to as BFS-SA. Although suffix array is theoretically efficient with O(n) construction 114 work, the hidden constant is large. Thus, we use hashing-based solutions to replace SA for 115 LCP queries to improve the performance in practice. We first present a simple approach 116 BFS-HASH in Sec. 3.2 that stores a hash value for all prefixes of the input. This approach 117 has O(n) construction work, $O(\log n)$ per LCP query, and O(n) auxiliary space. While both 118 BFS-SA and BFS-HASH take O(n) extra space, such space overhead can be significant in 119 practice—for example, BFS-HASH requires n 64-bit hash values, which is $4 \times$ the input size 120 considering characters as inputs, and $32 \times$ with even smaller alphabet such as molecule bases 121 (alphabet as $\{A, C, G, T\}$). To address the space issue, we proposed BFS-B-HASH using 122 blocking. Our solution takes a user-defined parameter b as the block size, which trades off 123 between space usage and query time. BFS-B-HASH limits extra space in O(n/b) by using 124 $O(b \log n)$ LCP query time. Surprisingly, despite a larger LCP cost, our hash-based solutions 125 are consistently faster than BFS-SA in all real-world test cases, due to cheaper construction. 126 All of our BFS-based solutions are simple to program. 127

We also study the DaC-based approach and propose a parallel output-sensitive solution. We propose a non-trivial adaption for the AALM algorithm [1] to make it output-sensitive. Our algorithm is inspired by the BFS-based approaches, and improves the work from $\tilde{O}(nm)$ to $\tilde{O}(nk)$, with polylogarithmic span. The technical challenge is that the states in the computation are no longer a rectangle, but an irregular shape (see Fig. 1 and 3). We then present a highly non-trivial implementation of this algorithm. Among many key challenges, we highlight our solution to avoid dynamically allocating arrays in the recursive execution.

X. Ding, X. Dong, Y. Gu, Y. Liu, and Y. Sun

While memory allocation is mostly ignored theoretically, in practice it can easily be the performance bottleneck in the parallel setting. We refer to this implementation as DAC-SD, with details given in Sec. 4 and 5.2 and the full version of this paper [16].

The bounds of our algorithms (BFS-SA, BFS-HASH, BFS-B-HASH, and DAC-SD) are 138 presented in Tab. 1. We implemented them and show an experimental study in Sec. 6. We 139 tested both synthetic and real-world datasets, including DNA, English text from Wikipedia, 140 and code repositories from GitHub, with string lengths in 10^{5} - 10^{9} and varying edit distances, 141 many of them with real edits (e.g., edit history from Wikipedia and commit history on 142 GitHub). In most tests, our new BFS-B-HASH or BFS-HASH performs the best, and 143 their relative performance depends on the value of k and the input patterns. Our BFS-144 based algorithms are faster than the existing parallel output-insensitive implementation in 145 ParlayLib [7], even with a reasonably large $k \approx 10^5$. We believe that our paper is the first 146 systematic study in theory and practice of parallel edit distance, and we give the first publicly 147 available parallel edit distance implementation that can process *billion-scale strings* with 148 small edit distance and our code at [17]. Due to page limit, some details are provided in the 149 full version of this paper [16]. We summarize our contributions as follows: 150

151 1. Two new BFS-based edit distance solutions BFS-HASH and BFS-B-HASH using hash-

- ¹⁵² based LCP queries. Compared to the existing SA-based solution in Landau-Vishkin,
 ¹⁵³ our hash-based solutions are simpler and more practical. BFS-B-HASH also allows for
 ¹⁵⁴ tradeoffs between time and auxiliary space.
- ¹⁵⁵ **2.** A new DaC-based edit distance solution DAC-SD with $O(nk \log k)$ work and polylogar-¹⁵⁶ ithmic span.
- New implementations for four output-sensitive edit distance algorithms: BFS-SA, BFS HASH, BFS-B-HASH and DAC-SD. Our code is publicly available[17].
- ¹⁵⁹ 4. Experimental study of the existing and our new algorithms on different input patterns.

¹⁶⁰ **2** Preliminaries

We use O(f(n)) with high probability (whp) (in n) to mean O(cf(n)) with probability at least $1 - n^{-c}$ for $c \ge 1$. We use $\tilde{O}(f(n))$ to denote $O(f(n) \cdot \text{polylog}(n))$. For a string A, we use A[i] as the *i*-th character in A. We use string and sequence interchangeably. We use A[i..j] to denote the *i*-th to the *j*-th characters in A, and A[i..j) the *i*-th to the (j-1)-th characters in A. Throughout the paper, we use "auxiliary space" to mean space used in addition to the input.

¹⁶⁷ String Edit Distance. Given two strings A[1..n] and B[1..m], Levenshtein's Edit Dis-¹⁶⁸ tance [38] between A and B is the minimum number of operations needed to convert A to B ¹⁶⁹ by using insertions, deletions, and substitutions. We also call the operations *edits*. In this ¹⁷⁰ paper, we use *edit distance* to refer to Levenshtein's Edit Distance. The classic dynamic ¹⁷¹ programming (DP) algorithm for edit distance uses DP recurrence shown in Sec. 1 with ¹⁷² O(mn) work and space.

Hash Functions. For the simplicity of algorithm descriptions, we assume a perfect hash 173 function for string comparisons, i.e., a function $h: S \to [1, O(|S|)]$ such that $h(x) = h(y) \iff$ 174 x = y. For any alphabet Σ with size $|\alpha|$, we use a hash function $h(A[l..r]) = \sum_{i=l}^{r} A[i] \times p^{r-i}$ 175 for some prime numbers $p > |\alpha|$, which returns a unique hash value of the substring A[l...r]. 176 The hash values of two consecutive substrings S_1 and S_2 can be concatenated as $h([S_1, S_2]) =$ 177 $h(S_1) \cdot p^{|S_2|} + h(S_2)$, and the inverse can also be computed as $h(S_2) = h(|S_1, S_2|) - p^{|S_2|} \cdot h(S_1)$. 178 For simplicity, we denote concatenation and its inverse operation as \oplus and \ominus , respectively, as 179 $h([S_1, S_2]) = h(S_1) \oplus h(S_2)$ and $h(S_2) = h([S_1, S_2]) \oplus h(S_1)$. We assume perfect hashing for 180

46:4 Efficient Parallel Output-Sensitive Edit Distance

theoretical analysis. In practice, we use p as a large prime and modular arithmetic to keep the word-size hash values. In our experiment, we compare different approaches and validate that our implementations are correct in all test cases. However, collisions are possible for other datasets, since different strings may be mapped to the same hash value. If such cases arise, one can either use multiple hash functions for a better success rate in practice, or use the idea of Hirschberg's algorithm [27] to generate the edit sequence and run a correctness check (and restart with another hash function if failed).

Longest Common Prefix (LCP). For two sequences A[1..n] and B[1..m], the Longest Common Prefix (LCP) query at position x in A[1..n] and position y in B[1..m] is the longest substring starting from A[x] that match a prefix starting from B[y]. With clear context, we also use the term "LCP" to refer to the length of the LCP, i.e., LCP(A, B, x, y) is the length of the longest common prefix substring starting from A[x] and B[y] for A and B.

Computational Model. We use the *work-span model* in the classic multithreaded model 193 with **binary-forking** [2, 8, 9]. We assume a set of threads that share the memory. Each 194 thread acts like a sequential RAM plus a fork instruction that forks two child threads running 195 in parallel. When both child threads finish, the parent thread continues. A parallel-for is 196 simulated by fork for a logarithmic number of steps. A computation can be viewed as a 197 DAG (directed acyclic graph). The work W of a parallel algorithm is the total number 198 of operations, and the span (depth) S is the longest path in the DAG. The randomized 199 work-stealing scheduler can execute such a computation in W/P + O(S) time whp in W on 200 P processors [2, 9, 25]. 201

²⁰² **Suffix Array.** The suffix array (SA) [42] is a lexicographically sorted array of the suffixes of ²⁰³ a string, usually used together with the longest common prefix (LCP) array, which stores ²⁰⁴ the length of LCP between every adjacent pair of suffixes. The SA and LCP array can be ²⁰⁵ built in parallel in O(n) work and $O(\log^2 n)$ span whp [32, 53].

In edit distance, we need the LCP query between A[x..n] and B[y..m] for any x and y. This can be computed by building the SA and LCP arrays for a new string C[1..n + m]that concatenates A[1..n] and B[1..m]. The LCP between any pair of suffixes in C can be computed by a range minimum query (RMQ) on the LCP array, which can be built in O(n+m) work and $O(\log(n+m))$ span [8]. Combining all pieces gives the following theorem:

▶ Lemma 1. Given two strings A[1..n] and B[1..m], using a suffix array, the longest common prefix (LCP) between any two substrings A[x..n] and B[y..m] can be reported in O(1) work and span, with O(n + m) preprocessing work and $O(\log^2(n + m))$ span whp.

²¹⁴ **3** BFS-based Algorithms

3.1 Overview of Existing Sequential and Parallel BFS-based Algorithms

Many existing output-sensitive algorithms [19, 20, 21, 22, 26, 35, 36, 37, 46, 47] are based on breadth-first search (BFS). These algorithms view the DP matrix for edit distance as a DAG, as shown in Fig. 1. In this section, we use x and y to denote the row and column ids of the cells in the DP matrix, respectively. Each state (cell) (x, y) has three incoming edges from (x - 1, y), (x, y - 1), and (x - 1, y - 1) (if they exist). The edge weight is 0 from (x - 1, y - 1) to (x, y) when A[x] = B[y], and 1 otherwise. Then edit distance is equivalent to the shortest distance from (0, 0) to (n, m).

Since the edge weights are 0 or 1, we can use a special breadth-first search (BFS) to compute the shortest distance. In round t, we process states with edit distance t. The



algorithm terminates when we reach cell (n, m). First observed by Ukkonen [56], in the BFS-based approach, not all states need to be visited. For example, all states with |x-y| > kwill not be reached before we reach (n, m) with edit distance k, since they require more than k edits. Thus, this BFS will touch at most O(kn) cells, leading to O(kn) work.

Another key observation is that starting from any cell (x, y), if there are diagonal edges 229 with weight 0, we should always follow the edges until a unit-weight edge is encountered. 230 Namely, we should always find the longest common prefix (LCP) from A[x+1] and B[y+1], 231 and skip to the cell at (x + p, y + p) with no edit, where p is the LCP length. This idea is 232 used in Landau and Vishkin [37] on parallel approximate string matching, and we adapt this 233 idea to edit distance here. Using the modified parallel BFS algorithm by Landau-Vishkin [37] 234 (shown in Alg. 1), only $O(k^2)$ states need to be processed—on each diagonal and for each 235 edit distance t, only the last cell with t edits needs to be processed (see Fig. 1). Hence, 236 the BFS runs for k rounds on 2k + 1 diagonals, which gives the $O(k^2)$ bound above. In 237 the BFS algorithm, we can label each diagonal by the value of x - y. In round t, the BFS 238 visits a *frontier* of cells $f_t[\cdot]$, where $f_t[i]$ is the cell with edit distance t on diagonal i, for 239 $-t \leq i \leq t$. We present the algorithm in Alg. 1 and an illustration in Fig. 1. Note that in 240 the implementation, we only need to maintain two frontiers (the previous and the current 241 one), which requires O(k) space. We provide more details about this algorithm in the full 242 version [16]. If the LCP query is supported by suffix arrays, we can achieve $O(n+k^2)$ work 243 and $O(\log n + k \log k)$ span for the edit distance algorithm. 244

Algorithm Based on Suffix Array (BFS-SA). Using the SA algorithm in [32] and the LCP
algorithm in [53] for Landau-Vishkin gives the claimed bounds in Tab. 1. We present details
about our SA implementation in Sec. 5.1.

²⁴⁸ 3.2 Algorithm Based on String Hashing (BFS-Hash)

Although BFS-SA is theoretically efficient with O(n) preprocessing work to construct the SA, the hidden constant is large. For better performance, we consider string hashing as an alternative for SA. Similar attempts (e.g., locality-sensitive hashing) have also been used in approximate pattern matching problems [43, 44]. In our pursuit of exact output-sensitive

46:6 Efficient Parallel Output-Sensitive Edit Distance

edit distance computation, we draw inspiration from established string hashing algorithms, such as the Rabin-Karp algorithm (also known as rolling hashing) [33]. We will first present a simple hash-based solution BFS-HASH with O(n) preprocessing cost and O(n) auxiliary space. Then later in Sec. 3.3, we will present BFS-B-HASH, which saves auxiliary space by trading off more work in LCP queries.

As mentioned in Sec. 2, the hash function $h(\cdot)$ maps any substring A[l..r] to a unique hash value, which provides a fingerprint for this substring in the LCP query. The high-level idea is to binary search the query length, using the hash value as validation. We precompute the hash values for all prefixes, i.e., $T_A[x] = h(A[1..x])$ for the prefix substring A[1..x] (similar for B). They can be computed in parallel by using any scan (prefix-sum) operation [6] with O(n) work and $O(\log n)$ span. We can compute h(A[l..r]) by $T_A[r] \ominus T_A[l-1]$.

With the preprocessed hash values, we dual binary search the LCP of A[x..n] and B[y..m]. We compare the hash values starting from A[x] and B[y] with chunk sizes of 1, 2, 4, 8, ...,until we find value l, such that $A[x..x+2^l) = B[y..y+2^l)$, but $A[x..x+2^{l+1}) \neq B[y..y+2^{l+1})$. By doing this with $O(\log n)$ work, we know that the LCP of A[x..n] and B[y..m] must have a length in the range $[2^l, 2^{l+1})$. We then perform a regular binary search in this range, which costs another $O(\log n)$ work. This indicates $O(\log n)$ work in total per LCP query. Combining the preprocessing and query costs, we present the cost bounds of BFS-HASH:

▶ **Theorem 2.** BFS-HASH computes the edit distance between two sequences of length n and $m \le n$ in $O(n + k^2 \log n)$ work, $\tilde{O}(k)$ span, and O(n) auxiliary space, where k is the output size (fewest possible edits).

BFS-HASH is simple and easy to implement. Our experimental results indicate that its 274 simplicity also allows for a reasonably good performance in practice for most real-world 275 input instances. However, this algorithm uses n 64-bit integers as hash values, and such 276 space overhead may be a concern in practice. This is more pronounced when the input is 277 large and/or the alphabet is small (particularly when each input element can be represented 278 with smaller than byte size), as the auxiliary space can be much larger than the input 279 size. This concern also holds for BFS-SA as several O(n)-size arrays are needed during SA 280 construction. Note that for shared-memory parallel algorithms, space consumption is also a 281 key constraint—if an algorithm is slow, we can wait for longer; but if data (and auxiliary 282 data) do not fit into the memory, then this algorithm is not applicable to large input at 283 all. In this case, the problem size that is solvable by the algorithm is limited by the space 284 overhead, which makes the improvement from parallelism much narrower. Below we will 285 discuss how to make our edit distance algorithms more space efficient. 286

²⁸⁷ 3.3 Algorithm Based on Blocked-Hashing (BFS-B-Hash)

In this section, we introduce our BFS-B-HASH algorithm that provides a more space-efficient solution by trading off worst-case time (work and span). Interestingly, we observed that on many data sets, BFS-B-HASH can even outperform BFS-HASH and other opponents due to faster construction time, and we will analyze that in Sec. 6.

To achieve better space usage, we divide the strings into blocks of size b. As such, we only need to store the hash values for prefixes of the entire blocks $h(A[1..b]), h(A[1..2b]), \dots, h(A[1..[(n/b)] \cdot b])$. Our idea of blocking is inspired by many string algorithms (e.g., [4]). Using this approach, we only need auxiliary space to store O(n/b) hash values, and thus we can control the space usage using the parameter b. To compute these hash values, we will first compute the hash value for each block, and run a parallel scan (prefix sum on



 \oplus on the hash values for all the blocks. Similar to the above, we refer to these arrays as $T_A[i] = h(A[1..ib])$ (and $T_B[i]$ accordingly), and call them **prefix tables**.

We now discuss how to run LCP with only partial hash values available. The LCP300 function in Alg. 2 presents the process to find the LCP of A[x..n] and B[y..m] using the 301 prefix tables. We present an illustration in Fig. 2. We will use the same dual binary search 302 approach to find the LCP of two strings. Since we do not store the hash values for all 303 prefixes, we use a function $GetHash(A, T_A, x)$ to compute h(A[1..x]). We can locate the 304 closest precomputed hash value and use r as the previous block id before x. Then the hash 305 value up to block r is simply $\bar{h} = T_A[r]$. We then concatenate the rest characters to the hash 306 value (i.e., return $\overline{h} \oplus h(A[rb+1]) \oplus \cdots \oplus h(A[x]))$). In this way, we can compute the hash 307 value of any prefixes for both A and B, and plug this scheme into the dual binary search in 308 BFS-HASH. In each step of dual binary search, the concatenation of hash value can have at 309 most b steps, and thus leads to a factor of b overhead in query time than BFS-HASH. 310

▶ **Theorem 3.** BFS-B-HASH computes the edit distance between two sequences of length $n \text{ and } m \leq n \text{ in } O(n + k^2 \cdot b \log n) \text{ work and } \tilde{O}(kb) \text{ span, using } O(n/b+k) \text{ auxiliary space,}$ where k is the output size (fewest possible edits).

The term k in space usage is from the BFS (each frontier is at most size O(k)). $O(b \log n)$ is the work for each LCP query. Note that this is an upper bound—if the LCP length L is small, the cost can be significantly smaller (a tighter bound is $O(\min(L, b \log L)))$). Sec. 6 will show that for normal input strings where the LCP lengths are small in most queries,



Figure 3 The illustrations of the key concepts and notation in the AALM algorithm described in Sec. 4.

the performance of BFS-B-HASH is indeed the fastest, although for certain input instances when the worst case is reached, the performance is not as good.

³²⁰ **4** The Divide-and-Conquer Algorithms

Our parallel output-sensitive algorithm DAC-SD is inspired by the AALM algorithm [1], and also uses it as a subroutine. We first overview the AALM algorithm, and introduce our algorithm in details. We assume m = n is a power of 2 in this section for simple descriptions, but both our algorithm and AALM work for any n and m.

The AALM Algorithm. As described above, the edit distance problem can be considered as 325 a shortest distance (SD) problem from the top-left cell (0,0) to the bottom-right cell (n,n) in 326 the DP matrix G. Instead of directly computing the SD from (0,0) to (n,n), AALM computes 327 pairwise SD between any cell on the left/top boundaries and the bottom/right boundaries 328 (i.e., those on $L \cup U$ to $W \cup R$ in Fig. 3(a)). We relate all cells in $L \cup U$ as a sequence 329 $v = \{v_0, v_1, \dots, v_{2n}\}$ (resp., $W \cup R$ as $u = \{u_0, u_1 \cdots, u_{2n}\}$), as shown in Fig. 3. Therefore, 330 for the DP matrix G, the pairwise SD between v and u forms a $(2n+1) \times (2n+1)$ matrix. We 331 call it the **SD** matrix of G, and denote it as D_G . AALM uses a divide-and-conquer approach. 332 It first partitions G into four equal submatrices G_1 , G_2 , G_3 , and G_4 (See Fig. 3(b)), and 333 recursively computes the SD matrices for all G_i . We use D_i to denote the SD matrix for G_i . 334 In the "conquer" step, the AALM algorithm uses a *Combine* subroutine to combine two SD 335 matrices into one if they share a common boundary (our algorithm also uses this subroutine). 336 For example, consider combining G_1 and G_2 . We still use v_i and u_j to denote the cells on the 337 left/top and bottom/right boundaries of $\begin{pmatrix} G_1 \\ G_2 \end{pmatrix}$ (see Fig. 3(c)), and denote the cells on the 338 common boundary of G_1 and G_2 as $w_1, \dots, w_{n/2}$, ordered from left to right. For any pair v_i 339 and u_i , if they are in the same submatrix, we can directly get the SD from the corresponding 340 SD matrix. Otherwise, WLOG assume $v_i \in G_1$ and $u_j \in G_2$, then we compute the SD 341 between them by finding $\min_l D_1[i, l] + D_2[l, j]$, i.e., for all w_l on the common boundary, we 342 attempt to use the SD between v_i to w_l , and w_l to u_j , and find the minimum one. Similarly, 343 we can combine D_3 with D_4 , and $D_{1\cup 2}$ with $D_{3\cup 4}$, and eventually get D_G . We note that the 344 Combine algorithm, even theoretically, is highly involved. At a high level, it uses the Monge 345 property of the shortest distance (the monotonicity of the DP recurrence), and we refer 346 the readers to [1] for a detailed algorithm description and theoretical analysis. In Sec. 5.2, 347 we highlight a few challenges and our solutions for implementing this highly complicated 348 algorithm. Theoretically, combining two $n \times n$ SD matrices can be performed in $O(n^2)$ work 349 and $O(\log^2 n)$ span, which gives $O(n^2 \log n)$ work and $O(\log^3 n)$ span for AALM. 350

X. Ding, X. Dong, Y. Gu, Y. Liu, and Y. Sun

	and $D[1n]$.
1 Notes: We assume both A and B has size	Function $GETDISTANCE(i, j, n, t)$
$n = 2^c$ for simple description. Our	if $n/2 < t$ then
algorithm also works for strings with	\square Computed D by the AALM algorithm
different lengths with minor changes.	return D
2 Function DAC-SD	// Compute the SD matrices for
$3 \mid t \leftarrow 1$	G_1, G_2, G_3, G_4 (as shown in Fig. 3 (d)).
4 while true do	11 $D_1 \leftarrow \text{GETDISTANCE}(i, j, n/2, t)$
$5 \mid D \leftarrow Check(t)$	12 Compute D_2 and D_3 by the AALM al-
6 if $D[t][t] \leq t$ then break	gorithm
$ au \qquad t \leftarrow \min(2t, n)$	13 $D_4 \leftarrow \text{GetDistance}(i+n/2, j+n/2, n-$
s return $D[t][t]$	n/2,t)
// Find the SD in the DP matrix from $(0,0)$	// use the same Combine function as AALM
with width t	14 $D_{1\cup 2} \leftarrow Combine(D_1, D_2)$
a Function $Check(t)$	15 $D_{3\cup4} \leftarrow Combine(D_3, D_4)$
10 $D \leftarrow GETDISTANCE(0, 0, n, t)$	$16 D \leftarrow (D_{1\cup 2}, D_{3\cup 4})$
	$17 \mid \text{return } D$

Algorithm 3 Divide-and-Conquer edit distance algorithm on A[1..n] and B[1..n].

Our algorithm. The AALM algorithm has $\tilde{O}(n^2)$ work ($\tilde{O}(nm)$ if $n \neq m$) and polylogar-351 ithmic span, which is inefficient in the output-sensitive setting. As mentioned in Sec. 3.1, 352 only a narrow width-O(k) diagonal area in G is useful (Fig. 3(d)). We thus propose an 353 output-sensitive DAC-SD algorithm adapted from the AALM algorithm. We follow the 354 same steps in AALM, but restrict the paths to the diagonal area, although the exact size is 355 unknown ahead of time. We first present the algorithm to compute the shortest distance 356 on the diagonal region with width 2t + 1 as function Check(t) in Alg. 3, which restricts the 357 search in diagonals -t to t. First, we divide such a region into four sub-regions (see Fig. 3(d)). 358 Two of them $(G_1 \text{ and } G_4)$ are of the same shape, and the other two of them $(G_2 \text{ and } G_3)$ 359 are triangles. For G_2 and G_3 , we use the AALM algorithm to compute their SD matrices by 360 aligning them to squares. For G_1 and G_4 , we process them recursively, until the base case 361 where the edge length of the matrix is smaller than t and they degenerate to squares, in 362 which case we apply the AALM algorithm. Note that even though the width-(2t+1) diagonal 363 stripe is not a square (G_1 and G_4 are also of the same shape), the useful boundaries are still 364 the left/top and bottom/right boundaries $(L \cup U \text{ and } W \cup R \text{ in Fig. 3(d)})$. Therefore, we 365 can use the same *Combine* algorithm as in AALM to combine the SD matrices. For example, 366 in Fig. 3(d), when combining G_1 with G_2 , we obtain the pairwise distance between $L \cup U$ 367 and $R \cup R'$ using the common boundary W. We can similarly combine all G_1, G_2, G_3 , and 368 G_4 to get the SD matrix for G. 369

However, the output value k is unknown before we run the algorithm. To overcome this 370 issue, we use a strategy based on prefix doubling to "binary search" the value of k without 371 asymptotically increasing the work of the algorithm. We start with t = 1, and run the 372 Check(t) in Alg. 3 (i.e., restricting the search in a width-(2t+1) diagonal). Assume that 373 the *Check* function returns σ edits. If $\sigma < t$, we know that σ is the SD from (0,0) to (n,n), 374 since allowing the path to go out of the diagonal area will result in an answer greater than t. 375 Otherwise, we know $\sigma > t$, and σ is not necessarily the shortest distance from the (0,0) to 376 (n, n), since not restricting the path in the t-diagonal area may allow for a shorter path. If 377 so, we double t and retry. Although we need $O(\log k)$ searches before finding the final answer 378 k, we will show that the total search cost is asymptotically bounded by the last search. In 379 the last search, we have t < 2k. 380

We first analyze the cost for Check(t). It contains two recursive calls, two calls to AALM, and three calls to the *Combine* function. Therefore, the work for Check(t) is W(n) =

46:10 Efficient Parallel Output-Sensitive Edit Distance

³⁸³ $2W(n/2) + O(t^2 \log t)$, with base cases $W(t) = t^2 \log t$, which solves to $W(n) = O(nt \log t)$. ³⁸⁴ For span, note that there are $\log(n/t)$ levels of recursion before reaching the base cases. In ³⁸⁵ each level, the *Combine* function combines $t \times t$ SD matrices with $O(\log^2 t)$ span. In the leaf ³⁸⁶ level, the base case uses AALM with $O(\log^3 t)$ span. Therefore, the total span of a *Check* is:

 $O(\log n/t \cdot \log^2 t + (\log n/t + \log^3 t)) = O(\log^2 t \cdot (\log n/t + \log t)) = O(\log^2 t \log n)$ (1)

We will apply $Check(\cdot)$ for $O(\log k)$ times, with $t = 1, 2, 4, \ldots$ up to at most 2k. Therefore, the total work is dominated by the last Check, which is $O(nk \log k)$. The span is $O(\log n \log^3 k)$.

Theorem 4. The DAC-SD algorithm computes the edit distance between two sequences of length n and $m \le n$ in $O(nk \log k)$ work and $O(\log n \log^3 k)$ span, where k is the output size (fewest possible edits).

³⁹³ Compared to the BFS-based algorithms with $\tilde{O}(k)$ span, our DAC-SD is also output-sensitive ³⁹⁴ and achieves polylogarithmic span. However, the work is $\tilde{O}(kn)$ instead of $\tilde{O}(n + k^2)$, which ³⁹⁵ will lead to more running time in practice for a moderate size of k.

³⁹⁶ **5** Implementation Details

We provide all implementations for the four algorithms as well as testing benchmarks at [17]. In this section, we highlight some interesting and challenging parts of our implementations.

³⁹⁹ 5.1 Implementation Details of BFS-based Algorithms

For the suffix array construction in BFS-SA, we implemented a parallel version of the 400 DC3 algorithm [32]. We also compared our implementation with the SA implementation 401 in ParlayLib [7], which is a highly optimized version of the prefix doubling algorithm with 402 $O(n \log n)$ work and $O(\log^2 n)$ span. On average, our implementation is about $2 \times$ faster 403 than that in ParlayLib when applied to edit distance. We present some results for their 404 comparisons in the full version [16]. For LCP array construction and preprocessing RMQ 405 queries, we use the implementation in ParlayLib [7], which requires $O(n \log n)$ work and 406 $O(\log^2 n)$ span. With them, the query has O(1) cost. 407

In our experiments on both synthetic and real-world data, we observed that the LCP 408 length is either very large when we find two long matched chunks, or in most of the cases, 409 very short when they are not corresponding to each other. This is easy to understand—for 410 genomes, text or code with certain edit history, it is unlikely that two random starting 411 positions share a large common prefix. Based on this, we add a simple optimization for all 412 LCP implementations such that we first compare the leading eight characters, and only when 413 they all match, we use the regular LCP query. This simple optimization greatly improved 414 the performance of BFS-SA, and also slightly improved the hash-based solutions. 415

416 5.2 Implementation Details of the DaC-SD Algorithm

Although our DAC-SD algorithm given in Alg. 3 is not complicated, we note that imple-417 menting it is highly non-trivial in two aspects. First, in Sec. 4, we assume both strings A418 and B have the same length n, which is a power of two. However, handling two strings 419 with different lengths makes the matrix partition more complicated in practice. Another 420 key challenge is that the combining step in the AALM algorithm is recursive and needs to 421 allocate memory with varying sizes in the recursive execution. While memory allocation is 422 mostly ignored theoretically, frequent allocation in practice can easily be the performance 423 bottleneck in the parallel setting. We discuss our engineering efforts as follows. 424



Data	Alias	A	B	$m{k}$	$ \Sigma $
Wilripodio	Wiki v1	$0.56 \mathrm{M}$	$0.56 \mathrm{M}$	439	256
	Wiki v2	$0.56 \mathrm{M}$	$0.56 \mathrm{M}$	5578	256
pages [45]	Wiki v3	$0.56 \mathrm{M}$	$0.55 \mathrm{M}$	15026	256
Linux hornol	Linux v1	$6.47 \mathrm{M}$	$6.47 \mathrm{M}$	236	256
Linux kerner	Linux v2	$6.47 \mathrm{M}$	$6.47 \mathrm{M}$	1447	256
code [40]	Linux v3	$6.47 \mathrm{M}$	$6.46 \mathrm{M}$	9559	256
DNA	DNA 1	42.3M	42.3M	928	4
	DNA 2	42.3M	42.3M	9162	4
sequences [5]	DNA 3	42.3M	42.3M	91419	4

Figure 4 The illustrations of our output-sensitive DaC-SD algorithm. (a) Two parameters t_1 and t_2 are needed to denote the lengths of the diagonal area on each side. (b) The case that $t_2 = 0$ and G_3 degenerates.

	Table 2 Real-world datasets in our experiments,
in	cluding input sizes $ A $ and $ B $, number of edits
k,	and alphabet sizes $ \Sigma $.

Irregularity. The general case, when n and m are not powers of two and not the same, is 425 more complicated than the case in Alg. 3. In this case, all four subproblems G_1, G_2, G_3 , 426 and G_4 will have different sizes. While theoretically, we can always round up, for better 427 performance in practice, we need to introduce additional parameters to restrict the search 428 within the belt region as shown in Fig. 4. Therefore, we use two parameters t_1 and t_2 , to 429 denote the lengths of the diagonal area on each side. We show an illustration in Fig. 4(a) 430 along with how to compute the subproblem sizes. In extreme cases, t_1 or t_2 can degenerate 431 to 0, which results in three subproblems (Fig. 4(b)). In such cases, we will first merge G_2 432 and G_4 , then merge G_1 and $G_{2\cup 4}$. 433

The Combining Step. As mentioned in Sec. 4, achieving an efficient combining step is 434 highly non-trivial. The straightforward solution to combine two matrices is to use the 435 Floyd-Warshall algorithm [18], but it incurs $O(n^3)$ work and will be a bottleneck. The 436 AALM algorithm improves this step to $O(n^2)$ by taking advantage of the Monge property of 437 the two matrices. For page limit, we introduce the details of the combining algorithm in the 438 full version [16]. However, the original AALM algorithm is based on divide-and-conquer and 439 requires memory allocation for every recursive function call. This is impractical as frequent 440 parallel memory allocation is extremely inefficient. To overcome this challenge, we redesign 441 the recursive solution to an iterative solution, such that we can preallocate the memory 442 space before the combining step. No dynamic memory allocation is involved during the 443 computation. We provide the details of this approach in the full version [16]. 444

445 **6** Experiments

Setup. We implemented all algorithms in C++ using ParlayLib [7] for fork-join parallelism
and some parallel primitives (e.g., reduce). Our tests use a 96-core (192 hyperthreads)
machine with four Intel Xeon Gold 6252 CPUs, and 1.5 TB of main memory. We utilize
numactl -i all in tests with more than one thread to spread the memory pages across
CPUs in a round-robin fashion. We run each test three times and report the median.

⁴⁵¹ **Tested Algorithms and Datasets.** We tested five algorithms in total: four output-sensitive ⁴⁵² algorithms in this paper (BFS-SA, BFS-HASH, BFS-B-HASH, DAC-SD), and a baseline ⁴⁵³ algorithm from ParlayLib [7], which is a parallel output-insensitive implementation with ⁴⁵⁴ O(nm) work. The ParlayLib implementation is intended to showcase the simplicity of ⁴⁵⁵ parallel algorithms, and as a result, it may not be well-ptimized. We are unaware of other



Synthetic Datasets:

Figure 5 Running time (in seconds) of synthetic and real-world datasets for all algorithms. Lower is better. We put an " \times " if the algorithm does not finish within 1000 seconds. For BFS-based algorithms, we separate the time into building time (constructing the data structure for LCP queries) and query time (running BFS). All bars out of the range of the y-axis are annotated with numbers. The number is the total running time for DAC-SD and ParlayLib, and is in the format of a + b for BFS-SA, where a is the building time and b is the query time. Full results are presented in the full version [16].

parallel implementations that provide output-sensitive cost bounds. We use b = 32 for our 456 BFS-B-HASH. As we will show later, the running time is generally stable with $4 \le b \le 64$. 457 We tested the algorithms on both synthetic and real-world datasets. For synthetic datasets, 458 we generate random strings with different string lengths $n = 10^i$ for $6 \le i \le 9$ and k (number 459 of edits) varying from 1 to 10^5 , and set the size of the alphabet as 256. We create strings A 460 and B by generating n random characters, and applying k edits. The k edits are uniformly 461 random for insertion, deletion and substitution. For $k \ll n$, we have $m \approx n$. All values of 462 k shown in the figures and tables are approximate values. Our real-world datasets include 463 Wikipedia [45], Linux kernel [40], and DNA sequences [51]. We compare the edit distance 464 between history pages on Wikipedia and history commits of a Linux kernel file on GitHub. 465 We also compare DNA sequences by adding valid modifications to them to simulate DNA 466 damage or genome editing techniques, as is used in many existing papers [11, 13, 30, 57]. We 467 present the statistics of the real-world datasets in Tab. 2. 468

⁴⁶⁹ **Overall Performance on Synthetic Data.** We present our results on synthetic data in ⁴⁷⁰ the upper part of Fig. 5. We also present the complete results in the full version [16]. For ⁴⁷¹ BFS-based algorithms, we also separate the time for **building** the data structures for LCP ⁴⁷² queries, and the **query** time (the BFS process). ParlayLib cannot process instances with ⁴⁷³ $n > 10^6$ due to its O(nm) work bound.



Figure 6 Performance of BFS-based algorithms vs. average LCP length. Some building times are invisible because they are too small.



Figure 7 Time-space trade-off in BFS-B-Hash. The space shown is the memory required for the prefix tables. The dotted line is the input size. Note that by setting b = 1, the algorithm is equivalent to BFS-HASH.

We first compare our solutions with ParlayLib [7]. Since ParlayLib is not output-sensitive, its running time remains the same regardless of the value of k. Among the tests that ParlayLib can process $(n = 10^6)$, our output-sensitive algorithms are much faster than ParlayLib, especially when k is small (up to $10^5 \times$). For $n = 10^6$, all our BFS-based algorithms are at least $1.7 \times$ faster than ParlayLib even when $k \approx n/10$.

We then compare our DaC- and BFS-based solutions. DAC-SD has the benefit of 479 polylogarithmic span, compared to O(k) span for the BFS-based algorithm. Although this 480 seems to suggest that DAC-SD should have better performance when k is large, the result 481 shows the opposite. The reason is that DAC-SD has $\tilde{O}(nk)$ work, compared to $\tilde{O}(n+k^2)$ 482 cost of the BFS-based algorithms. When k becomes larger, the overhead in work is also 483 more significant. On the other hand, when k is small, the O(nk) work becomes linear, which 484 hides the inefficiency in work. Therefore, the gap between DAC-SD and other algorithms is 485 smaller when k is small, but DAC-SD is still slower than BFS-based algorithms in all test 486 cases, especially when k is large. This experiment reaffirms the *importance of work efficiency* 487 on practical performance for parallel algorithms. 488

Finally, we compare all our BFS-based solutions. Our hash-based solutions have significant 489 advantages over the other implementations when k is small, since the pre-processing time 490 for hash-based solutions is much shorter. When k is large, pre-processing time becomes 491 negligible, and BFS-HASH seems to be the ideal choice since its query is also efficient. In 492 particular, for $n \approx m \approx 10^9$, hash-based algorithms use about 1 second for pre-processing 493 while BFS-SA uses about 100 seconds. Although BFS-SA also has O(n) construction time, 494 the constant is much larger and its memory access pattern is much worse than the two 495 hash-based solutions. We note that in some cases, the query time of BFS-SA can still be 496 faster than BFS-HASH and BFS-B-HASH, especially when k is large, which is consistent 497 with the theory (O(1) vs. $O(\log n)$ or $O(b \log n)$ per LCP query). 498

In theory, BFS-B-HASH reduces space usage in BFS-HASH by increasing the query time. 499 Interestingly, when k is small, BFS-B-HASH can also be faster than BFS-HASH by up to 500 $2.5\times$. This is because BFS-B-HASH incurs fewer writes (and thus smaller memory footprints) 501 in preprocessing that leads to faster building time. When k is small, the running time is 502 mostly dominated by the building time, and thus BFS-B-HASH can perform better. When k503 is relatively large and k^2 is comparable to n, BFS-HASH becomes faster than BFS-B-HASH 504 due to better LCP efficiency. In fact, when k is large, the running time is mainly dominated 505 by the query (BFS), and all three algorithms behave similarly. It is worth noting that in 506 these experiments with $|\Sigma| = 256$ and random edits, in most of the cases, the queried LCP is 507

46:14 Efficient Parallel Output-Sensitive Edit Distance

⁵⁰⁸ small. Therefore, the $O(\log n)$ or $O(b \log n)$ query time for BFS-HASH and BFS-B-HASH ⁵⁰⁹ are not tight, and they have much better memory access patterns than BFS-SA in LCP ⁵¹⁰ queries. As a result, they can have matching or even better performance than BFS-SA. ⁵¹¹ Later we will show that under certain input distributions where the average LCP length is ⁵¹² large, BFS-SA can have some advantage over both BFS-HASH and BFS-B-HASH.

Real-World Datasets. We now analyze how our algorithms perform on real-world string and edit patterns. The results are shown in the lower part of Fig. 5. The results are mostly consistent with our synthetic datasets, where BFS-B-HASH is more advantageous when k is small, and BFS-HASH performs the best when k is large. When k is large, BFS-SA can also have comparable performance to the hash-based solutions.

LCP Length vs. Performance. It seems that for both synthetic and real-world data shown 518 above, our hash-based solutions are always better than BFS-SA. It is worth asking, whether 519 BFS-SA can give the best performance in certain cases, given that it has the best theoretical 520 bounds (see Tab. 1). By investigating the bounds carefully, BFS-SA has better LCP query 521 cost as O(1), while the costs for BFS-HASH and BFS-B-HASH are $O(\log L)$ and $O(b \log L)$, 522 respectively, where L is the LCP length. This indicates that BFS-SA should be advantageous 523 when k and L are both large. To verify this, we artificially created input instances with 524 medium to large values of k and controlled average LCP query lengths, and showed the 525 results in Fig. 6 on two specific settings. 526

The experimental result is consistent with the theoretical analysis. The running time 527 for BFS-HASH increases slowly with L, while the performance of BFS-B-HASH grows 528 much faster, since it is affected by a factor of O(b) more than BFS-HASH. The query time 529 for BFS-SA almost stays the same, but also increases slightly with increasing L. This is 530 because in general, with increasing L, the running time for all three algorithms may increase 531 slightly due to worse cache locality in BFS due to more long matches. In Figure 6(a), the 532 building time for both BFS-HASH and BFS-B-HASH are negligible, while BFS-SA still 533 incurs significant building time. Even in this case, with an LCP length of 300, the query 534 time of the hash-based solutions still becomes larger than the *total* running time of BFS-SA. 535 In Figure 6(b) with a larger k, the building time for all three algorithms is negligible. In 536 this case, BFS-SA always has comparable performance with BFS-HASH, and may perform 537 better when L > 20. However, such extreme cases (both k and L are large) should be very 538 rare in real-world datasets - when k is large enough so that the query time is large enough 539 to hide SA's building time, L is more likely to be small, which in turn is beneficial for the 540 query bounds in hash-based solutions. Indeed such cases did not appear in our 33 tests on 541 both synthetic and real data. 542

Parallelism. We test the self-relative speedup of all algorithms. We present speedup 543 numbers on two representative tests with different values of n and k in Tab. 3. For BFS-based 544 algorithms, we separate the speedup for building and query. All our algorithms are highly 545 parallelized. Even though BFS-SA and DAC-SD have a longer running time, they still 546 have a 48–68× speedup, indicating good scalability. Our BFS-HASH algorithm has about 547 $40-50\times$ speedup in building, and BFS-B-HASH has a lower but decent speedup of about 548 $20-40\times$. When k is small, the frontier sizes (and the total work) of BFS are small, and the 549 running time is also negligible. In this case, we cannot observe meaningful speedup. For 550 larger $k = 10^5$, three BFS-based algorithms achieve $27-48 \times$ speedup both in query and entire 551 edit distance algorithm. 552

⁵⁵³ **Space Usage.** We study the time-space tradeoff of our BFS-B-HASH with different block ⁵⁵⁴ sizes b. We present the *auxiliary space* used by the prefix table in BFS-B-HASH along with ⁵⁵⁵ running time in Fig. 7 using one test case with $n = 10^8$ and $k = 10^5$ in our synthetic dataset.

	${m k}$	BFS-B-Hash		BFS-Hash			BFS-SA			DaC-SD	
\boldsymbol{n}		Build	\mathbf{Query}	Total	Build	\mathbf{Query}	Total	Build	\mathbf{Query}	Total	Total
$\frac{10^8}{10^9}$	$10 \\ 10^5$	$\begin{vmatrix} 20.4 \\ 24.2 \end{vmatrix}$	- 36.4	$19.9 \\ 36.3$	$ \begin{array}{c} 46.6 \\ 42.7 \end{array} $	- 46.8	$\begin{array}{c} 46.5\\ 46.6\end{array}$	$49.6 \\ 51.2$	- 27.1	$49.4 \\ 48.3$	68.2 t.o.

Table 3 Self-relative speedup of each implementation in each step. "Build" = constructing the data structure for LCP queries. "Query" = the BFS process. "t.o." = timeout. We omit query speedup when k = 10 because there is little parallelism to be explored for BFS with small k, and the BFS time is also small and hardly affects the overall speedup. 192 hyperthreads are used for parallel executions.

The dotted line shows the input size. Note that when b = 1, it is exactly BFS-HASH. Since 556 the inputs are 8-bit characters and the hash values are 64-bit integers, BFS-HASH incurs $8\times$ 557 space overhead than the input size. Using blocking, we can avoid such overhead and keep 558 the auxiliary space even lower than the input. The auxiliary space decreases linearly with 559 the block size b. Interestingly, although blocking itself incurs time overhead, the impact in 560 time is small: the time grows by $1.19 \times$ from b = 1 to 2, and grows by $1.08 \times$ from b = 2 to 561 64. This is mostly due to two reasons: 1) as mentioned, with 8-bit character input type and 562 random edits, the average LCP length is likely short and within the first block, and therefore 563 the query costs in both approaches are close to O(L) for LCP length L, and 2) the extra 564 factor of b in queries (Line 17) is mostly cache hits (consecutive locations in an array). This 565 illustrates the benefit of using blocking in such datasets, since blocking saves much space 566 while only increasing the time by a small fraction. 567

7 Conclusion and Discussions

We proposed output-sensitive parallel algorithms for the edit-distance problem, as well as 569 careful engineering of them. We revisited the BFS-based Landau-Vishkin algorithm. In 570 addition to using SA as is used in Landau-Vishkin (our BFS-SA implementation), we 571 also designed two hash-based data structures to replace the SA for more practical LCP 572 queries (BFS-HASH and BFS-B-HASH). We also presented the first output-sensitive parallel 573 algorithm based on divide-and-conquer with O(nk) work and polylogarithmic span. We have 574 also shown the best of our engineering effort on this algorithm, although its performance 575 seems less competitive than other candidates due to work inefficiency. 576

We implemented all these algorithms and tested them on synthetic and real-world 577 datasets. In summary, our BFS-based solutions show the best overall performance on 578 datasets with real-world edits or random edits, due to faster preprocessing time and better 579 I/O-friendliness. BFS-HASH performs the best in time when k is large. BFS-B-HASH 580 has better performance when k is small. The blocking scheme also greatly improves space 581 efficiency without introducing much overhead in time. In very extreme cases where both 582 k and the LCP lengths are large, BFS-SA can have some advantages over the hash-based 583 solutions, while BFS-B-HASH can be much slower than BFS-HASH. However, such input 584 patterns seem rare in the real world. 585

All our BFS-based solutions perform better than the output-insensitive solution in ParlayLib, and the DaC-based solution with $\tilde{O}(nk)$ work and polylogarithmic span, even for large $k > \sqrt{n}$. The results also imply the importance of work efficiency in parallel algorithm designs, consistent with the common belief in the literature [52, 24]. Because the number of cores in modern multi-core machines is small (usually hundreds to thousands) compared to the problem size, an algorithm is less practical if it blows up the work significantly, as parallelism cannot compensate for the performance loss due to larger work.

593		References
594	1	Alberto Apostolico, Mikhail J Atallah, Lawrence L Larmore, and Scott McFaddin, Efficient
595		parallel algorithms for string editing and related problems. SIAM J. on Computing, 19(5):968–
596		988. 1990.
597	2	Nimar S Arora, Robert D Blumofe, and C Greg Plaxton. Thread scheduling for multipro-
598		grammed multiprocessors. Theory of Computing Systems (TOCS), 34(2):115–144, 2001.
599	3	K Nandan Babu and Sanjeev Saxena. Parallel algorithms for the longest common subsequence
600		problem. In IEEE International Conference on High Performance Computing (HiPC), pages
601		120–125. IEEE, 1997.
602	4	Michael A. Bender and Martin Farach-Colton. The lca problem revisited. In Latin American
603		Symposium on Theoretical Informatics (LATIN), pages 88–94. Springer, 2000.
604	5	Dennis A Benson, Mark Cavanaugh, Karen Clark, Ilene Karsch-Mizrachi, David J Lipman,
605		James Ostell, and Eric W Sayers. Genbank. Nucleic acids research, 41(D1):D36–D42, 2012.
606	6	Guy E. Blelloch. Scans as primitive parallel operations. IEEE Trans. on Comput., 38(11),
607		1989.
608	7	Guy E. Blelloch, Daniel Anderson, and Laxman Dhulipala. Parlaylib — a toolkit for parallel
609		algorithms on shared-memory multicore machines. In ACM Symposium on Parallelism in
610		Algorithms and Architectures (SPAA), pages 507–509, 2020.
611	8	Guy E. Blelloch, Jeremy T. Fineman, Yan Gu, and Yihan Sun. Optimal parallel algorithms in
612		the binary-forking model. In ACM Symposium on Parallelism in Algorithms and Architectures
613		(SPAA), pages 89–102, 2020.
614	9	Robert D. Blumofe and Charles E. Leiserson. Space-efficient scheduling of multithreaded
615		computations. SIAM J. on Computing, 27(1), 1998.
616	10	Nicholas Boucher, Ilia Shumailov, Ross Anderson, and Nicolas Papernot. Bad characters:
617		Imperceptible nlp attacks. In IEEE Symposium on Security and Privacy (SP), pages 1987–2004.
618		IEEE, 2022.
619	11	Dana Carroll. Focus: genome editing: genome editing: past, present, and future. The Yale
620		journal of biology and medicine, 90(4):653, 2017.
621	12	Jung Hee Cheon, Miran Kim, and Kristin Lauter. Homomorphic computation of edit distance.
622		In International Conference on Financial Cryptography and Data Security, pages 194–212.
623		Springer, 2015.
624	13	Kendell Clement, Holly Rees, Matthew C Canver, Jason M Gehrke, Rick Farouni, Jonathan Y
625		Hsu, Mitchel A Cole, David R Liu, J Keith Joung, Daniel E Bauer, et al. Crispresso2 provides
626		accurate and rapid genome editing sequence analysis. Nature biotechnology, 37(3):224–226,
627	14	
628	14	I nomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. Introduction
629	15	Sonior Descripto, Christian H. Banadimitrion, and Umash Virkuman Vazimani. Algorithms
630	15	MaCrow Hill Higher Education New York 2008
631	16	Viangyun Ding Vianjun Dong Van Cu Viban Sun and Vougho Liu Efficient parallel
632	10	output-sensitive edit distance arXiv preprint 9306 17/61 2023
634	17	Viangvun Ding Vianiun Dong Van Cu. Vihan Sun, and Youzhe Liu. Parallel implementations
635	11	for output-sensitive edit distance https://github.com/ucrparlay/Edit-Distance 2023
626	18	Robert W Floyd Algorithm 97: shortest path Commun ACM 5(6):345 1962
627	10	Zvi Galil and Raffaele Giancarlo. Improved string matching with k mismatches ACM SIGACT
638	10	News, 17(4):52–54, 1986.
639	20	Zvi Galil and Raffaele Giancarlo. Parallel string matching with k mismatches. Theoretical
640	20	Computer Science (TCS), 51(3):341–348, 1987.
641	21	Zvi Galil and Raffaele Giancarlo. Data structures and algorithms for approximate string
642	-	matching. Journal of Complexity, 4(1):33–72, 1988.
643	22	Zvi Galil and Kunsoo Park. An improved algorithm for approximate string matching. SIAM
644		Journal on Computing, 19(6):989–999, 1990.

X. Ding, X. Dong, Y. Gu, Y. Liu, and Y. Sun

- ⁶⁴⁵ 23 Michael T Goodrich and Roberto Tamassia. Algorithm design and applications. Wiley Hoboken,
 ⁶⁴⁶ 2015.
- Yan Gu, Ziyang Men, Zheqi Shen, Yihan Sun, and Zijin Wan. Parallel longest increasing
 subsequence and van emde boas trees. In ACM Symposium on Parallelism in Algorithms and
 Architectures (SPAA), 2023.
- Yan Gu, Zachary Napier, and Yihan Sun. Analysis of work-stealing and parallel cache
 complexity. In SIAM Symposium on Algorithmic Principles of Computer Systems (APOCS),
 pages 46–60. SIAM, 2022.
- ⁶⁵³ 26 Dov Harel and Robert Endre Tarjan. Fast algorithms for finding nearest common ancestors.
 ⁶⁵⁴ siam Journal on Computing, 13(2):338–355, 1984.
- Daniel S. Hirschberg. A linear space algorithm for computing maximal common subsequences.
 Commun. ACM, 18(6):341–343, 1975.
- ⁶⁵⁷ 28 Daniel Hládek, Ján Staš, and Matúš Pleva. Survey of automatic spelling correction. *Electronics*,
 ⁶⁵⁸ 9(10):1670, 2020.
- Md Mosabbir Hossain, Md Farhan Labib, Ahmed Sady Rifat, Amit Kumar Das, and Monira
 Mukta. Auto-correction of english to bengali transliteration system using levenshtein distance.
 In International Conference on Smart Computing & Communications (ICSCC), pages 1–5.
 IEEE, 2019.
- Yoon-Seong Jeon, Kihyun Lee, Sang-Cheol Park, Bong-Soo Kim, Yong-Joon Cho, Sung-Min
 Ha, and Jongsik Chun. Ezeditor: a versatile sequence alignment editor for both rrna-and
 protein-coding genes. International journal of systematic and evolutionary microbiology,
 64(Pt_2):689-691, 2014.
- Tao Jiang, Guohui Lin, Bin Ma, and Kaizhong Zhang. A general edit distance between RNA
 structures. Journal of Computational Biology, 9(2):371–388, 2002.
- Juha Kärkkäinen and Peter Sanders. Simple linear work suffix array construction. In Intl.
 Colloq. on Automata, Languages and Programming (ICALP), pages 943–955. Springer, 2003.
- ⁶⁷¹ 33 Richard M Karp and Michael O Rabin. Efficient randomized pattern-matching algorithms.
 ⁶⁷² IBM journal of research and development, 31(2):249–260, 1987.
- ⁶⁷³ 34 Peter Krusche and Alexander Tiskin. New algorithms for efficient parallel string comparison.
 ⁶⁷⁴ In ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), pages 209–216,
 ⁶⁷⁵ 2010.
- Gad M Landau and Uzi Vishkin. Efficient string matching with k mismatches. Theoretical
 Computer Science (TCS), 43:239–249, 1986.
- Gad M Landau and Uzi Vishkin. Fast string matching with k differences. J. Computer and
 System Sciences, 37(1):63–78, 1988.
- Gad M Landau and Uzi Vishkin. Fast parallel and serial approximate string matching. J.
 Algorithms, 10(2):157–169, 1989.
- ⁶⁸² 38 Vladimir I Levenshtein et al. Binary codes capable of correcting deletions, insertions, and
 ⁶⁸³ reversals. In *Soviet physics doklady*, volume 10, pages 707–710. Soviet Union, 1966.
- Heng Li and Nils Homer. A survey of sequence alignment algorithms for next-generation
 sequencing. *Briefings in bioinformatics*, 11(5):473–483, 2010.
- 40 Linux Kernel File dcn_1_0_sh_mask.h Commit History on GitHub. https:
 //github.com/torvalds/linux/blob/master/drivers/gpu/drm/amd/include/asic_reg/
 dcn/dcn_1_0_sh_mask.h.
- Mi Lu and Hua Lin. Parallel algorithms for the longest common subsequence problem. *IEEE Transactions on Parallel and Distributed Systems*, 5(8):835–848, 1994.
- 42 Udi Manber and Gene Myers. Suffix arrays: a new method for on-line string searches. SIAM
 592 J. on Computing, 22(5):935-948, 1993.
- Guillaume Marçais, Dan DeBlasio, Prashant Pandey, and Carl Kingsford. Locality-sensitive
 hashing for the edit distance. *Bioinformatics*, 35(14):i127–i135, 2019.

46:18 Efficient Parallel Output-Sensitive Edit Distance

- Samuel McCauley. Approximate Similarity Search Under Edit Distance Using Locality Sensitive Hashing. In 24th International Conference on Database Theory (ICDT 2021), volume
 186, pages 21:1–21:22. Schloss Dagstuhl Leibniz-Zentrum für Informatik, 2021.
- 45 Municipal history of Quebec on Wikipedia. https://en.wikipedia.org/wiki/Municipal_
 history_of_Quebec.
- 46 Eugene W Myers. An o (nd) difference algorithm and its variations. Algorithmica, 1(1-4):251–
 266, 1986.
- 47 Eugene Wimberly Myers. Incremental alignment algorithms and their applications. University
 of Arizona, Department of Computer Science, 1986.
- 48 Jean-Frédéric Myoupo and David Seme. Time-efficient parallel algorithms for the longest
 common subsequence and related problems. J. Parallel Distrib. Comput., 57(2):212–223, 1999.
- Gonzalo Navarro. A guided tour to approximate string matching. ACM computing surveys (CSUR), 33(1):31–88, 2001.
- ⁷⁰⁸ 50 Mike Paterson and Vlado Dančík. Longest common subsequences. In International Symposium
 ⁷⁰⁹ on Mathematical Foundations of Computer Science, pages 127–142. Springer, 1994.
- Jane Peterson, Susan Garges, Maria Giovanni, Pamela McInnes, Lu Wang, Jeffery A Schloss,
 Vivien Bonazzi, Jean E McEwen, Kris A Wetterstrand, Carolyn Deal, et al. The nih human
 microbiome project. *Genome research*, 19(12):2317–2323, 2009.
- ⁷¹³ 52 Zheqi Shen, Zijin Wan, Yan Gu, and Yihan Sun. Many sequential iterative algorithms can be
 ⁷¹⁴ parallel and (nearly) work-efficient. In ACM Symposium on Parallelism in Algorithms and
 ⁷¹⁵ Architectures (SPAA), 2022.
- Julian Shun. Fast parallel computation of longest common prefixes. In International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), pages 387–398.
 IEEE, 2014.
- 719 54 Diomidis Spinellis. Git. *IEEE software*, 29(3):100–101, 2012.
- Vianney Kengne Tchendji, Armel Nkonjoh Ngomade, Jerry Lacmou Zeutouo, and Jean Frédéric
 Myoupo. Efficient cgm-based parallel algorithms for the longest common subsequence problem
 with multiple substring-exclusion constraints. *Parallel Computing*, 91:102598, 2020.
- 56 Esko Ukkonen. Algorithms for approximate string matching. Information and Control,
 64(1-3):100-118, 1985.
- ⁷²⁵ 57 Liang-Jiao Xue and Chung-Jui Tsai. Ageseq: analysis of genome editing by sequencing.
 ⁷²⁶ Molecular plant, 8(9):1428–1430, 2015.
- Jiaoyun Yang, Yun Xu, and Yi Shang. An efficient parallel algorithm for longest common subsequence problem on GPUs. In *World Congress on Engineering*, volume 1, pages 499–504, 2010.
- F30 59 Hongyu Zhang. Alignment of blast high-scoring segment pairs based on the longest increasing
 r31 subsequence algorithm. *Bioinformatics*, 19(11):1391–1396, 2003.