

1 **ELASTIC-DEGENERATE STRING MATCHING**
2 **VIA FAST MATRIX MULTIPLICATION***

3 GIULIA BERNARDINI[†], PAWEŁ GAWRYCHOWSKI[‡], NADIA PISANTI[§], SOLON P.
4 PISSIS[¶], AND GIOVANNA ROSONE^{||}

5 **Abstract.** An elastic-degenerate (ED) string is a sequence of n sets of strings of total length
6 N , which was recently proposed to model a set of similar sequences. The ED string matching
7 (EDSM) problem is to find all occurrences of a pattern of length m in an ED text. The EDSM
8 problem has recently received some attention in the combinatorial pattern matching community,
9 and an $\mathcal{O}(nm^{1.5}\sqrt{\log m} + N)$ -time algorithm is known [Aoyama et al., CPM 2018]. The standard
10 assumption in the prior work on this question is that N is substantially larger than both n and
11 m , and thus we would like to have a linear dependency on the former. Under this assumption, the
12 natural open problem is whether we can decrease the 1.5 exponent in the time complexity, similarly
13 as in the related (but, to the best of our knowledge, not equivalent) *word break* problem [Backurs
14 and Indyk, FOCS 2016].

15 Our starting point is a conditional lower bound for the EDSM problem. We use the popular combi-
16 binatorial Boolean Matrix Multiplication (BMM) conjecture stating that there is no truly subcubic
17 *combinatorial* algorithm for BMM [Abboud and Williams, FOCS 2014]. By designing an appropriate
18 reduction we show that a combinatorial algorithm solving the EDSM problem in $\mathcal{O}(nm^{1.5-\epsilon} + N)$
19 time, for any $\epsilon > 0$, refutes this conjecture. Our reduction should be understood as an indication
20 that decreasing the exponent requires fast matrix multiplication.

21 String periodicity and fast Fourier transform are two standard tools in string algorithms. Our
22 main technical contribution is that we successfully combine these tools with fast matrix multipli-
23 cation to design a non-combinatorial $\tilde{\mathcal{O}}(nm^{\omega-1} + N)$ -time algorithm for EDSM, where ω denotes
24 the matrix multiplication exponent and the $\tilde{\mathcal{O}}(\cdot)$ notation suppresses polylog factors. To the best of
25 our knowledge, we are the first to combine these tools. In particular, using the fact that $\omega < 2.373$
26 [Alman and Williams, SODA 2021; Le Gall, ISSAC 2014; Williams, STOC 2012], we obtain an
27 $\mathcal{O}(nm^{1.373} + N)$ -time algorithm for EDSM. An important building block in our solution, that might
28 find applications in other problems, is a method of selecting a small set of length- ℓ substrings of the
29 pattern, called anchors, so that any occurrence of a string from an ED text set contains at least one
30 but not too many (on average) such anchors inside.

31 **Key words.** string algorithms, pattern matching, elastic-degenerate string, matrix multiplica-
32 tion, fast Fourier transform

33 **AMS subject classifications.** 68W01, 68W32, 68Q25, 68Q17

*Submitted to the editors DATE. A preliminary version appeared in ICALP 2019 [14].

Funding: GB is supported by the Netherlands Organisation for Scientific Research (NWO) under project OCENW.GROOT.2019.015 “Optimization for and with Machine Learning (OPTIMAL)”. PG is partially supported by the Bekker programme of the Polish National Agency for Academic Exchange (PPN/BEK/2020/1/00444). NP and GR are supported by the University of Pisa under the “PRA Progetti di Ricerca di Ateneo” (Institutional Research Grants) - Project no. PRA_20202021_26 “Metodi Informatici Integrati per la Biomedica” and partially supported by MIUR-SIR project CMACBioSeq “Combinatorial methods for analysis and compression of biological sequences” grant no. RBSI146R5L. NP, SPP and GR are partially supported by the ALPACA project that has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 956229. SPP is partially supported by the PANGAIA project that has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 872539.

[†]CWI, Amsterdam, The Netherlands (giulia.bernardini@cw.nl).

[‡]University of Wrocław, Wrocław, Poland (gawry@cs.uni.wroc.pl).

[§]University of Pisa, Pisa, Italy (nadia.pisanti@unipi.it).

[¶]CWI, Amsterdam, The Netherlands and Vrije Universiteit, Amsterdam, The Netherlands (solon.pissis@cw.nl).

^{||}University of Pisa, Pisa, Italy (giovanna.rosone@unipi.it).

34 **1. Introduction.** Boolean matrix multiplication (BMM) is one of the most fun-
 35 damental computational problems. Apart from its theoretical interest, it has a wide
 36 range of applications [34, 36, 44, 55, 64]. BMM is also the core combinatorial part of
 37 integer matrix multiplication. In both problems, we are given two $\mathcal{N} \times \mathcal{N}$ matrices
 38 and we are to compute \mathcal{N}^2 values. Integer matrix multiplication can be performed
 39 in *truly subcubic* time, i.e., in $\mathcal{O}(\mathcal{N}^{3-\epsilon})$ operations over the field, for some $\epsilon > 0$. The
 40 fastest known algorithms for this problem run in $\mathcal{O}(\mathcal{N}^{2.373})$ time [4, 51, 66]. These
 41 algorithms are known as algebraic: they rely on the ring structure of matrices over
 42 the field.

43 There also exists a different family of algorithms for the BMM problem known as
 44 combinatorial. Their focus is on unveiling the combinatorial structure in the Boolean
 45 matrices to reduce redundant computations. A series of results [9, 11, 20] culminating
 46 in an $\tilde{\mathcal{O}}(\mathcal{N}^3 / \log^4 \mathcal{N})$ -time algorithm [70, 71] (the $\tilde{\mathcal{O}}(\cdot)$ notation suppresses polylog
 47 factors) has led to the popular combinatorial BMM conjecture stating that there is no
 48 combinatorial algorithm for BMM working in time $\mathcal{O}(\mathcal{N}^{3-\epsilon})$, for any $\epsilon > 0$ [2]. There
 49 has been ample work on applying this conjecture to obtain BMM hardness results:
 50 see, e.g., [2, 22, 40, 49, 50, 52, 60].

51 String matching is another fundamental problem, asking to find all fragments of
 52 a string text of length n that match a string pattern of length m . This problem
 53 has several linear-time solutions [28]. In many real-world applications, it is often
 54 the case that letters at some positions are either unknown or uncertain. A way of
 55 representing these positions is with a subset of the alphabet Σ . Such a representation
 56 is called *degenerate string*. A special case of a degenerate string is when at such
 57 unknown or uncertain positions the only subset of the alphabet allowed is the whole
 58 alphabet. These special degenerate strings are more commonly known as strings
 59 with wildcards. The first efficient algorithm for a text and a pattern, where both
 60 may contain wildcards, was published by Fischer and Paterson in 1974 [35]. It has
 61 undergone several improvements since then [25, 26, 43, 46]. The first efficient algorithm
 62 for a standard text and a degenerate pattern, which may contain any non-empty
 63 subset of the alphabet, was published by Abrahamson in 1987 [3], followed by several
 64 practically efficient algorithms [41, 56, 69].

65 Degenerate letters are used in the IUPAC notation [45] to represent a position
 66 in a DNA sequence that can have multiple possible alternatives. These are used
 67 to encode the consensus of a population of sequences [5, 6, 37, 57, 63] in a multiple
 68 sequence alignment (MSA). In the presence of insertions or deletions in the MSA,
 69 we may need to consider alternative representations. Consider the following MSA of
 70 three closely-related sequences (on the left):

$$\begin{array}{l}
 \text{GCAACGGGTA--TT} \\
 \text{GCAACGGGTATATT} \\
 \text{GCACCTGG----TT}
 \end{array}
 \quad \tilde{T} = \{ \text{GCA} \} \cdot \left\{ \begin{array}{c} \text{A} \\ \text{C} \end{array} \right\} \cdot \{ \text{C} \} \cdot \left\{ \begin{array}{c} \text{G} \\ \text{T} \end{array} \right\} \cdot \{ \text{GG} \} \cdot \left\{ \begin{array}{c} \text{TA} \\ \text{TATA} \\ \varepsilon \end{array} \right\} \cdot \{ \text{TT} \}$$

72 These sequences can be compacted into a single sequence \tilde{T} of sets of strings (on
 73 the right) containing some deterministic and some non-deterministic segments. A
 74 non-deterministic segment is a finite set of deterministic strings and may contain the
 75 empty string ε corresponding to a deletion. The total number of segments is the
 76 *length* of \tilde{T} and the total number of letters is the *size* of \tilde{T} . We denote the length by
 77 $n = |\tilde{T}|$ and the size by $N = \|\tilde{T}\|$.

78 This representation has been defined in [42] by Iliopoulos et al. as an *elastic-*
 79 *degenerate* (ED) string. Being a sequence of subsets of Σ^* , it can be seen as a general-
 80 ization of a degenerate string. The natural problem that arises is finding all matches

81 of a deterministic pattern P in an ED text \tilde{T} . This is the *elastic-degenerate string*
 82 *matching* (EDSM) problem. Since its introduction in 2017 [42], it has attracted some
 83 attention in the combinatorial pattern matching community [58], and a series of re-
 84 sults have been published. The simple algorithm by Iliopoulos et al. [42] for EDSM
 85 was first improved by Grossi et al. in the same year, who showed that, for a pattern of
 86 length m , the EDSM problem can be solved *on-line* in $\mathcal{O}(nm^2 + N)$ time [39]; on-line
 87 means that it reads the text segment-by-segment and reports an occurrence as soon
 88 as this is detected. This result was improved by Aoyama et al. [8] who presented
 89 an $\mathcal{O}(nm^{1.5}\sqrt{\log m} + N)$ -time algorithm. An important feature of these bounds is
 90 their *linear dependency* on N . A different branch of on-line algorithms waiving the
 91 linear-dependency restriction exists [23, 24, 39, 59]. Moreover, the EDSM problem has
 92 been considered under Hamming and edit distance [16]. Recent results on founder
 93 block graphs [53] can also be casted on elastic-degenerate strings.

94 A question with a somewhat similar flavor is the *word break* problem. We are given
 95 a dictionary \mathcal{D} , $m = |\mathcal{D}|$, and a string S , $n = |S|$, and the question is whether we can
 96 split S into fragments that appear in \mathcal{D} (the same element of \mathcal{D} can be used multiple
 97 times). Backurs and Indyk [10] designed an $\tilde{\mathcal{O}}(nm^{1/2-1/18} + m)$ -time algorithm for
 98 this problem¹. Bringmann et al. [18] improved this to $\tilde{\mathcal{O}}(nm^{1/3} + m)$ and showed
 99 that this is optimal for combinatorial algorithms by a reduction from k -Clique. Their
 100 algorithm uses fast Fourier transform (FFT), and so it is not clear whether it should
 101 be considered combinatorial. While this problem seems similar to EDSM, there does
 102 not seem to be a direct reduction and so their lower bound does not immediately
 103 apply.

104 **Our Results.** It is known that BMM and triangle detection (TD) in graphs either
 105 both have truly subcubic combinatorial algorithms or none of them do [68]. Recall
 106 also that the currently fastest algorithm with linear dependency on N for the EDSM
 107 problem runs in $\mathcal{O}(nm^{1.5}\sqrt{\log m} + N)$ time [8]. In this paper we prove the following
 108 two theorems.

109 **THEOREM 1.1.** *If the EDSM problem can be solved in $\mathcal{O}(nm^{1.5-\epsilon} + N)$ time,*
 110 *for any $\epsilon > 0$, with a combinatorial algorithm, then there exists a truly subcubic*
 111 *combinatorial algorithm for TD.*

112 Arguably, the notion of combinatorial algorithms is not clearly defined, and The-
 113 orem 1.1 should be understood as an indication that in order to achieve a better
 114 complexity one should use fast matrix multiplication. Indeed, there are examples
 115 where a lower bound conditioned on BMM was helpful in constructing efficient algo-
 116 rithms using fast matrix multiplication [1, 17, 21, 30, 54, 67, 72]. We successfully design
 117 such a non-combinatorial algorithm by combining three ingredients: a string periodic-
 118 ity argument, FFT, and fast matrix multiplication. While periodicity is the usual tool
 119 in combinatorial pattern matching [29, 47, 48] and using FFT is also not unusual (for
 120 example, it often shows up in approximate string matching [3, 7, 25, 38]), to the best
 121 of our knowledge, we are the first to combine these with fast matrix multiplication.
 122 Specifically, we show the following result for the EDSM problem, where ω denotes the
 123 matrix multiplication exponent.

124 **THEOREM 1.2.** *The EDSM problem can be solved on-line in $\tilde{\mathcal{O}}(nm^{\omega-1} + N)$ time.*

125 In order to obtain a faster algorithm for the EDSM problem, we focus on the
 126 *active prefixes* (AP) problem that lies at the heart of all current solutions [8, 39]. In

¹ The $\tilde{\mathcal{O}}(\cdot)$ notation suppresses polylog factors.

127 the AP problem, we are given a string P of length m and a set of arbitrary prefixes
 128 $P[1..i]$ of P , called *active prefixes*, stored in a bit vector U so that $U[i] = 1$ if $P[1..i]$
 129 is active. We are further given a set \mathcal{S} of strings of total length N and we are asked to
 130 compute a bit vector V which stores the new set of active prefixes of P . A new active
 131 prefix of P is a concatenation of $P[1..i]$ (such that $U[i] = 1$) and some element of \mathcal{S} .

132 Using the algorithmic framework introduced in [39], EDSM is addressed by solving
 133 an instance of the AP problem per each segment i of the ED text corresponding to set
 134 \mathcal{S} of the AP problem. Hence, an $\mathcal{O}(f(m) + N_i)$ solution for the AP problem (with N_i
 135 being the size of a single segment of the ED text) implies an $\mathcal{O}(nf(m) + N)$ solution
 136 of EDSM, as $f(m)$ is repeated n times and $N = \sum_{i=1}^n N_i$. The algorithm of [8] solves
 137 the AP problem in $\mathcal{O}(m^{1.5}\sqrt{\log m} + N_i)$ time leading to $\mathcal{O}(nm^{1.5}\sqrt{\log m} + N)$ time
 138 for the EDSM problem. Our algorithm partitions the strings of each segment i of
 139 the ED text into three types according to a periodicity criterion, and then solves
 140 a restricted instance of the AP problem for each of the types. In particular, we
 141 solve the AP problem in $\tilde{\mathcal{O}}(m^{\omega-1} + N_i)$ time leading to $\tilde{\mathcal{O}}(nm^{\omega-1} + N)$ time for the
 142 EDSM problem. Given this connection between the two problems and, in particular,
 143 between their size parameter N , in the rest of the paper we will denote with N also
 144 the parameter N_i of the AP problem.

145 An important building block in our solution that might find applications in other
 146 problems is a method of selecting a small set of length- ℓ substrings of the pattern,
 147 called *anchors*, so that any relevant occurrence of a string from an ED text set contains
 148 at least one but not too many such anchors inside. This is obtained by rephrasing the
 149 question in a graph-theoretical language and then generalizing the well-known fact
 150 that an instance of the hitting set problem with m sets over $[n]$, each of size at least
 151 k , has a solution of size $\mathcal{O}(n/k \cdot \log m)$. While the idea of carefully selecting some
 152 substrings of the same length is not new (for example Kociumaka et al. [48] used it
 153 to design a data structure for pattern matching queries on a string), our setting is
 154 different and hence so is the method of selecting these substrings.

155 In addition to the conditional lower bound for the EDSM problem (Theorem 1.1),
 156 we also exhibit a reduction from BMM to AP that leads to the following conditional
 157 lower bound for AP.

158 **THEOREM 1.3.** *If the AP problem can be solved in $\mathcal{O}(m^{1.5-\epsilon} + N)$ time, for any*
 159 *$\epsilon > 0$, with a combinatorial algorithm, then there exists a truly subcubic combinatorial*
 160 *algorithm for the BMM problem.*

161 We remark that Theorem 1.3 is also implied by Theorem 1.1, as described at the
 162 end of Section 4, but we believe that a direct reduction from BMM to AP serves as a
 163 good starting point for the more complicated reduction from BMM to EDSM.

164 **Roadmap.** Section 2 provides the necessary definitions and notation as well as the
 165 algorithmic toolbox used throughout the paper. In Section 3 we prove our lower
 166 bound result for the AP problem (Theorem 1.3). The lower bound result for the
 167 EDSM problem is proved in Section 4 (Theorem 1.1). In Section 5 we present our
 168 algorithm for EDSM (Theorem 1.2); this is the most technically involved part of the
 169 paper.

170 **2. Preliminaries.** Let $T = T[1]T[2] \dots T[n]$ be a string of length $|T| = n$ over a
 171 finite ordered alphabet Σ of size $|\Sigma| = \sigma$. For two positions i and j on T , we denote by
 172 $T[i..j] = T[i] \dots T[j]$ the substring of T that starts at position i and ends at position
 173 j (it is of length 0 if $j < i$). By ε we denote the empty string of length 0. A prefix of T
 174 is a substring of the form $T[1..j]$, and a suffix of T is a substring of the form $T[i..n]$.

175 T^r denotes the reverse of T , that is, $T[n]T[n-1]\dots T[1]$. We say that a string X is
 176 a power of a string Y if there exists an integer $k > 1$, such that X is expressed as k
 177 consecutive concatenations of Y , denoted by $X = Y^k$. A period of a string X is any
 178 integer $p \in [1, |X|]$ such that $X[i] = X[i+p]$ for every $i = 1, 2, \dots, |X| - p$, and the
 179 *period*, denoted by $\text{per}(X)$, is the smallest such p . We call a string X *strongly periodic*
 180 if $\text{per}(X) \leq |X|/4$.

181 LEMMA 2.1 ([33]). *If p and q are both periods of the same string X , and addi-*
 182 *tionally $p + q \leq |X| + 1$, then $\text{gcd}(p, q)$ is also a period of X .*

183 A *trie* is a tree in which every edge is labeled with a single letter, and every two
 184 edges outgoing from the same node have different labels. The label of a node u in
 185 such a tree T , denoted by $\mathcal{L}(u)$, is defined as the concatenation of the labels of all
 186 the edges on the path from the root of T to u . By replacing each path p consisting
 187 of nodes with exactly one child by an edge labeled by the concatenation of the labels
 188 of the edges of p we obtain a *compact trie*. The nodes of the trie that are removed
 189 after this transformation are called *implicit*, while the remaining ones are referred to
 190 as *explicit*. The suffix tree of a string S is the compact trie representing all suffixes of
 191 S , $\$ \notin \Sigma$, where instead of explicitly storing the label $S[i..j]$ of an edge we represent
 192 it by the pair (i, j) .

193 A *heavy path decomposition* of a tree T is obtained by selecting, for every non-
 194 leaf node $u \in T$, its child v such that the subtree rooted at v is the largest. This
 195 decomposes the nodes of T into node-disjoint paths, with each such path p (called a
 196 heavy path) starting at some node, called the *head* of p , and ending at a leaf. An
 197 important property of such a decomposition is that the number of distinct heavy
 198 paths above any leaf (that is, intersecting the path from a leaf to the root) is only
 199 logarithmic in the size of T [62].

200 Let $\tilde{\Sigma}$ denote the set of all finite non-empty subsets of Σ^* . Previous works (cf. [8,
 201 15, 39, 42, 59]) define $\tilde{\Sigma}$ as the set of all finite non-empty subsets of Σ^* excluding $\{\varepsilon\}$
 202 but we waive here the latter restriction as it has no algorithmic implications. An
 203 *elastic-degenerate string* $\tilde{T} = \tilde{T}[1] \dots \tilde{T}[n]$, or ED string, over alphabet Σ , is a string
 204 over $\tilde{\Sigma}$, i.e., an ED string is an element of $\tilde{\Sigma}^*$, and hence each $\tilde{T}[i]$ is a set of strings.

205 Let \tilde{T} denote an ED string of length n , i.e. $|\tilde{T}| = n$. We assume that for any
 206 $1 \leq i \leq n$, the set $\tilde{T}[i] \in \tilde{\Sigma}$ is implemented as an array and can be accessed by an
 207 index, i.e., $\tilde{T}[i] = \{\tilde{T}[i][k] \mid k = 1, \dots, |\tilde{T}[i]|\}$. For any $\tilde{\sigma} \in \tilde{\Sigma}$, $|\tilde{\sigma}|$ denotes the total
 208 length of all strings in $\tilde{\sigma}$, and for any ED string \tilde{T} , $|\tilde{T}|$ denotes the total length of all
 209 strings in all $\tilde{T}[i]$ s. We will denote $N_i = \sum_{k=1}^{|\tilde{T}[i]|} |\tilde{T}[i][k]|$ the total length of all strings
 210 in $\tilde{T}[i]$ and $N = \sum_{i=1}^n |\tilde{T}[i]|$ the *size* of \tilde{T} . An ED string \tilde{T} can be thought of as a
 211 compact representation of the set of strings $\mathcal{A}(\tilde{T})$ which is the Cartesian product of
 212 all $\tilde{T}[i]$ s; that is, $\mathcal{A}(\tilde{T}) = \tilde{T}[1] \times \dots \times \tilde{T}[n]$ where $A \times B = \{xy \mid x \in A, y \in B\}$ for any
 213 sets of strings A and B .

214 For any ED string \tilde{X} and a pattern P , we say that P *matches* \tilde{X} if:

- 215 1. $|\tilde{X}| = 1$ and P is a substring of some string in $\tilde{X}[1]$, or,
- 216 2. $|\tilde{X}| > 1$ and $P = P_1 \dots P_{|\tilde{X}|}$, where P_1 is a suffix of some string in $\tilde{X}[1]$, $P_{|\tilde{X}|}$
 217 is a prefix of some string in $\tilde{X}[|\tilde{X}|]$, and $P_i \in \tilde{X}[i]$, for all $1 < i < |\tilde{X}|$.

218 We say that an occurrence of a string P ends at position j of an ED string \tilde{T} if
 219 there exists $i \leq j$ such that P matches $\tilde{T}[i] \dots \tilde{T}[j]$. We will refer to string P as the
 220 *pattern* and to ED string \tilde{T} as the *text*. We define the main problem considered in
 221 this paper.

ELASTIC-DEGENERATE STRING MATCHING (EDSM)

INPUT: A string P of length m and an ED string \tilde{T} of length n and size $N \geq m$.

OUTPUT: All positions in \tilde{T} where at least one occurrence of P ends.

EXAMPLE 1. $P = \text{GTAT}$ ends at positions 2, 6, and 7 of the following text \tilde{T} .

$$\tilde{T} = \{ \text{ATGTA} \} \cdot \left\{ \begin{array}{c} \text{A} \\ \text{T} \end{array} \right\} \cdot \{ \text{C} \} \cdot \left\{ \begin{array}{c} \text{G} \\ \text{T} \end{array} \right\} \cdot \{ \text{CG} \} \cdot \left\{ \begin{array}{c} \text{TA} \\ \text{TATA} \\ \varepsilon \end{array} \right\} \cdot \left\{ \begin{array}{c} \text{TATGC} \\ \text{TTTTA} \end{array} \right\}$$

Whenever $|\tilde{T}| = 1$, the problem reduces to Case 1 only (searching for P in all strings of $\tilde{T}[1]$), which can be done in $\mathcal{O}(N)$ time using any linear-time pattern-matching algorithm. In the general case of $|\tilde{T}| > 1$, at a high-level, previous on-line solutions to EDSM consist of the following steps: (i) For each $\tilde{T}[i]$, for each $S \in \tilde{T}[i]$ that is long enough, search for occurrences of the whole of P in S (this corresponds to Case 1 of the definition of a match of P given above). Then (Case 2 of the definition of a match of P , in which an occurrence of P spans over several sets of strings), (ii) find the prefixes of P that match any suffix of some $S \in \tilde{T}[i]$, (iii) try to extend at $\tilde{T}[i]$ every partial occurrence of P , which has started earlier in \tilde{T} , by solving an instance of AP, and (iv) if a full occurrence of P also ends at $\tilde{T}[i]$, then output position i ; otherwise store the prefixes of P extended at $\tilde{T}[i]$, which will be further extended at $\tilde{T}[i+1]$.

Aoyama et al. [8] obtained an on-line $\mathcal{O}(nm^{1.5}\sqrt{\log m} + N)$ -time algorithm by identifying Step (iii) as the bottleneck in this approach, observing that all other steps can be implemented in $\mathcal{O}(n + M)$ time, and designing an improved solution for Step (iii). We formally define the task that needs to be solved in Step (iii) as the ACTIVE PREFIXES problem:

ACTIVE PREFIXES (AP)

INPUT: A string P of length m , a bit vector U of size m , a set \mathcal{S} of strings of total length N .

OUTPUT: A bit vector V of size m with $V[j] = 1$ if and only if there exists $S \in \mathcal{S}$ and $i \in [1, m]$, $U[i] = 1$, such that $P[1..i] \cdot S = P[1..i+|S|]$ and $j = i+|S|$.

In particular, given an ED text $\tilde{T} = \tilde{T}[1] \dots \tilde{T}[n]$, one should consider an instance of the AP problem per each $\tilde{T}[i]$. Hence, an $\mathcal{O}(f(m) + N_i)$ solution for AP (N_i being the size of $\tilde{T}[i]$) implies an $\mathcal{O}(n \cdot f(m) + N)$ solution for EDSM, as $f(m)$ is repeated n times and $N = \sum_{i=1}^n N_i$. We provide an example of the AP problem.

EXAMPLE 2. Let $P = \text{ababbababab}$ of length $m = 11$, $U = 01000100000$, and $\mathcal{S} = \{\varepsilon, \text{ab}, \text{abb}, \text{ba}, \text{baba}\}$. We have that $V = 01011101010$.

For our lower bound results we rely on BMM and the following closely related problem.

BOOLEAN MATRIX MULTIPLICATION (BMM)

INPUT: Two $\mathcal{N} \times \mathcal{N}$ Boolean matrices A and B .

OUTPUT: $\mathcal{N} \times \mathcal{N}$ Boolean matrix C , where $C[i, j] = \bigvee_k (A[i, k] \wedge B[k, j])$.

TRIANGLE DETECTION (TD)

INPUT: Three $\mathcal{N} \times \mathcal{N}$ Boolean matrices A, B and C .

OUTPUT: Are there i, j, k such that $A[i, j] = B[j, k] = C[k, i] = 1$?

253 An algorithm is called *truly subcubic* if it runs in $\mathcal{O}(\mathcal{N}^{3-\epsilon})$ time, for some $\epsilon > 0$.
 254 TD and BMM either both have truly subcubic combinatorial algorithms, or none of
 255 them do [68].

256 **3. AP Conditional Lower Bound.** As a warm-up, in order to investigate the
 257 hardness of the EDSM problem, we first show that an $\mathcal{O}(m^{1.5-\epsilon} + N)$ -time solution
 258 to the active prefixes problem, that constitutes the core of the solutions proposed
 259 in [8, 39], would imply a truly subcubic combinatorial algorithm for Boolean matrix
 260 multiplication (BMM). We recall that in the AP problem, we are given a string P
 261 of length m and a set of prefixes $P[1..i]$ of P , called *active prefixes*, stored in a bit
 262 vector U ($U[i] = 1$ if and only if $P[1..i]$ is active). We are further given a set \mathcal{S} of
 263 strings of total length N and we are asked to compute a bit vector V storing the new
 264 set of active prefixes of P : a prefix of P that extends $P[1..i]$ (such that $U[i] = 1$)
 265 with some element of \mathcal{S} . Of course, we can solve BMM by working over integers and
 266 using one of the fast matrix multiplication algorithms; plugging in the best known
 267 bounds results in an $\mathcal{O}(\mathcal{N}^{2.373})$ -time algorithm [4]. However, such an algorithm is
 268 not *combinatorial*, i.e., it uses *algebraic* methods. In comparison, the best known
 269 combinatorial algorithm for BMM works in $\hat{\mathcal{O}}(\mathcal{N}^3/\log^4 \mathcal{N})$ time [71]. This leads to
 270 the following popular conjecture.

271 CONJECTURE 1 ([2]). *There is no combinatorial algorithm for the BMM problem*
 272 *working in time $\mathcal{O}(\mathcal{N}^{3-\epsilon})$, for any $\epsilon > 0$.*

273 Aoyama et al. [8] showed that the AP problem can be solved in $\mathcal{O}(m^{1.5}\sqrt{\log m} + N)$
 274 time for constant-sized alphabets. Together with some standard string-processing
 275 techniques applied similarly as in [39], this is then used to solve the EDSM problem
 276 by creating an instance of the AP problem for every set $\tilde{T}[i]$ of \tilde{T} , i.e., with $\mathcal{S} = \tilde{T}[i]$.

277 We argue that, unless Conjecture 1 is false, the AP problem cannot be solved in
 278 time $\mathcal{O}(m^{1.5-\epsilon} + N)$, for any $\epsilon > 0$, with a combinatorial algorithm (note that the
 279 algorithm of Aoyama et al. [8] uses FFT, and so it is not completely clear whether it
 280 should be considered to be combinatorial). We show this by a reduction from combi-
 281 natorial BMM. Assume that, for the AP problem, we seek combinatorial algorithms
 282 with the running time $\mathcal{O}(m^{1.5-\epsilon} + N)$, i.e., with linear dependency on the total length
 283 of the strings. We need to show that such an algorithm implies that the BMM prob-
 284 lem can be solved in $\mathcal{O}(\mathcal{N}^{3-\epsilon'})$ time, for some $\epsilon' > 0$, with a combinatorial algorithm,
 285 thus implying that Conjecture 1 is false.

286 THEOREM 1.3. *If the AP problem can be solved in $\mathcal{O}(m^{1.5-\epsilon} + N)$ time, for any*
 287 *$\epsilon > 0$, with a combinatorial algorithm, then there exists a truly subcubic combinatorial*
 288 *algorithm for the BMM problem.*

289 *Proof.* Recall that in the BMM problem the matrices are denoted by A and B .
 290 In order to compute $C = A \times B$, we need to find, for every $i, j = 1, \dots, \mathcal{N}$, an index k
 291 such that $A[i, k] = 1$ and $B[k, j] = 1$. To this purpose, we split matrix A into blocks
 292 of size $\mathcal{N} \cdot L$ and B into blocks of size $L \cdot L$. This corresponds to considering values of
 293 j and k in intervals of size L , and clearly there are \mathcal{N}/L such intervals. Matrix B is
 294 thus split into $(\mathcal{N}/L)^2$ blocks, giving rise to an equal number of instances of the AP
 295 problem, each one corresponding to an interval of j and an interval of k . We will now
 296 describe the instance corresponding to the (K, J) -th block, where $1 \leq K, J \leq \mathcal{N}/L$.

We build the string P of the AP problem, for any block, as a concatenation of
 \mathcal{N} gadgets corresponding to $i = 1, \dots, \mathcal{N}$, and we construct the bit vector $U^{(K, J)}$ of

the AP problem as a concatenation of \mathcal{N} bit vectors, one per gadget. Each gadget consists of the same string $\mathbf{a}^L \mathbf{b} \mathbf{a}^L$; we set to 1 the k' -th bit of the i -th gadget bit vector if $A[i, (K-1)L + k'] = 1$. The solution of the AP problem $V^{(K,J)}$ will allow us to recover the solution of BMM, as we will ensure that the bit corresponding to the j' -th \mathbf{a} in the second half of the gadget is set to 1 if and only if, for some $k' \in [L]$, $A[i, (K-1)L + k'] = 1$ and $B[(K-1)L + k', (J-1)L + j'] = 1$. In order to enforce this, we will include the following strings in set $\mathcal{S}^{(K,J)}$:

$$\mathbf{a}^{L-k'} \mathbf{b} \mathbf{a}^{j'}, \text{ for every } k', j' \in [L] \text{ such that } B[(K-1)L + k', (J-1)L + j'] = 1.$$

297 This guarantees that after solving the AP problem we have the required property,
 298 and thus, after solving all the instances, we have obtained matrix $C = A \times B$. Indeed,
 299 consider values j , i.e., the index that runs on the columns of C , in intervals of size L .
 300 By construction and by definition of BMM, the i -th line of the J -th column interval
 301 of C is obtained by taking the disjunction of the second half of the i -th interval of
 302 each (K, J) -th bit vector for every $K = 1, 2, \dots, \mathcal{N}/L$.

We have a total of $(\mathcal{N}/L)^2$ instances. In each of them, the total length of all strings is $\mathcal{O}(L^3)$, and the length of the input string P is $(2L+1)\mathcal{N} = \mathcal{O}(L \cdot \mathcal{N})$. Using our assumed algorithm for each instance, we obtain the following total time:

$$\mathcal{O}((\mathcal{N}/L)^2 \cdot (L^3 + (\mathcal{N} \cdot L)^{1.5-\epsilon})) = \mathcal{O}(\mathcal{N}^2 \cdot L + \mathcal{N}^{3.5-\epsilon}/L^{0.5+\epsilon}).$$

303 If we set $L = \mathcal{N}^{(1.5-\epsilon)/(1.5+\epsilon)}$, then the total time becomes:

$$\begin{aligned} & \mathcal{O}(\mathcal{N}^{2+(1.5-\epsilon)/(1.5+\epsilon)} + \mathcal{N}^{3.5-\epsilon-(0.5+\epsilon)(1.5-\epsilon)/(1.5+\epsilon)}) \\ &= \mathcal{O}(\mathcal{N}^{2+(1.5-\epsilon)/(1.5+\epsilon)} + \mathcal{N}^{2+(1.5-\epsilon)-(1.5-\epsilon)(0.5+\epsilon)/(1.5+\epsilon)}) \\ &= \mathcal{O}(\mathcal{N}^{2+(1.5-\epsilon)/(1.5+\epsilon)} + \mathcal{N}^{2+(1.5-\epsilon)(1.5+\epsilon-0.5-\epsilon)/(1.5+\epsilon)}) \\ &= \mathcal{O}(\mathcal{N}^{2+(1.5-\epsilon)/(1.5+\epsilon)}). \end{aligned}$$

308 Hence we obtain a combinatorial BMM algorithm with complexity $\mathcal{O}(\mathcal{N}^{3-\epsilon'})$, where
 309 $\epsilon' = 1 - (1.5 - \epsilon)/(1.5 + \epsilon) > 0$. \square

310 **EXAMPLE 3.** Consider the following instance of the BMM problem with $\mathcal{N} = 6$
 311 and $L = 3$.

$$\begin{array}{c} 312 \\ 313 \\ 314 \\ 315 \\ 316 \end{array} \begin{array}{ccc} & A & B & C \\ \left[\begin{array}{ccc|ccc} 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ \hline 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{array} \right] & \times & \left[\begin{array}{ccc|ccc} 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ \hline 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 \end{array} \right] & = & \left[\begin{array}{ccc|ccc} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{array} \right] \end{array}$$

317 From matrices A and B , we now show how the resulting matrix C can be found
 318 by building and solving 4 instances of the AP problem constructed as follows. The
 319 pattern is

$$320 \quad P = \mathbf{aaabaaa} \cdot \mathbf{aaabaaa} \cdot \mathbf{aaabaaa} \cdot \mathbf{aaabaaa} \cdot \mathbf{aaabaaa} \cdot \mathbf{aaabaaa}$$

321 where the six gadgets are separated by a '.' to be highlighted. For the AP instances,
 322 the vectors $U^{(K,J)}$ shown below are the input bit vectors, and the sets $\mathcal{S}^{(K,J)}$ are the
 323 input set of strings. For each instance, the bit vector $V^{(K,J)}$ shown below is the output
 324 of the AP problem.

i	1	2	3	4	5	6
$U^{(1,1)}$: [0100000	1010000	0000000	1000000	0000000	0100000]
$S^{(1,1)}$: { aba ,baaa}					
$V^{(1,1)}$: [0000 100	0000001	0000000	0000000	0000000	0000100]
$U^{(1,2)}$: [0100000	1010000	0000000	1000000	0000000	0100000]
$S^{(1,2)}$: {aabaaa, baa }					
$V^{(1,2)}$: [0000000	0000011	0000000	0000001	0000000	0000000]
$U^{(2,1)}$: [0100000	0000000	0010000	0100000	1000000	0000000]
$S^{(2,1)}$: {aabaa, ba }					
$V^{(2,1)}$: [0000 000	0000000	0000100	0000000	0000010	0000000]
$U^{(2,2)}$: [0100000	0000000	0010000	0100000	1000000	0000000]
$S^{(2,2)}$: { aba ,baa}					
$V^{(2,2)}$: [0000100	0000000	0000010	0000100	0000000	0000000]

325

326

327

328

As an example on how to obtain matrix C , consider the bold part of C above (i.e., the first line of block (1,1) of C). This is obtained by taking the disjunction of the bold parts of $V^{(1,1)}$ and $V^{(2,1)}$.

329

330

331

332

333

334

335

336

337

338

339

340

341

342

343

344

345

346

347

4. EDSM Conditional Lower Bound. Since the lower bound for the AP problem does not imply *per se* a lower bound for the whole EDSM problem, in this section we show a conditional lower bound for the EDSM problem. Specifically, we perform a reduction from Triangle Detection to show that, if the EDSM problem could be solved in $\mathcal{O}(nm^{1.5-\epsilon} + N)$ time, this would imply the existence of a truly subcubic algorithm for TD. We show that TD can be reduced to the decision version of the EDSM problem: the goal is to detect whether there exists at least one occurrence of P in \tilde{T} . To this aim, given three matrices A, B, C , we first decompose matrix B into blocks of size $\mathcal{N}/s \times \mathcal{N}/s$, where s is a parameter to be determined later; the pattern P is obtained by concatenating a number (namely $z = \mathcal{N}s^2$) of constituent parts P_i of length $\mathcal{O}(\mathcal{N}/s)$, each one built with five letters from disjoint subalphabets. The ED text \tilde{T} is composed of three parts: the central part consists of three degenerate segments, the first one encoding the 1s of matrix A , the second one those of matrix B and the third one those of matrix C . These segments are built in such a way that the concatenation of strings of subsequent segments is of the same form as the pattern's building blocks. This central part is then padded to the left and to the right with sets containing appropriately chosen concatenations of substrings P_i of P , so that an occurrence of the pattern in the text implies that one of its building blocks matches the central part of the text, thus corresponding to a triangle. Formally:

348

349

THEOREM 1.1. *If the EDSM problem can be solved in $\mathcal{O}(nm^{1.5-\epsilon} + N)$ time, for any $\epsilon > 0$, with a combinatorial algorithm, then there exists a truly subcubic*

350 *combinatorial algorithm for TD.*

351 *Proof.* Consider an instance of TD, where we are given three $\mathcal{N} \times \mathcal{N}$ Boolean
 352 matrices A, B, C , and the question is to check if there exist i, j, k such that $A[i, j] =$
 353 $B[j, k] = C[k, i] = 1$. Let s be a parameter, to be determined later, that corresponds
 354 to decomposing B into blocks of size $(\mathcal{N}/s) \times (\mathcal{N}/s)$. We reduce to an instance of
 355 EDSM over an alphabet Σ of size $\mathcal{O}(\mathcal{N})$. Let us remark that, since we search for exact
 356 occurrences of the pattern, it would also be possible to assume that the instance of
 357 EDSM we reduce to is over a constant-sized (binary) alphabet. We could in fact
 358 replace each letter of the $\mathcal{O}(\mathcal{N})$ -sized alphabet with its binary encoding, increasing
 359 the length of the involved strings by only a logarithmic factor.

Pattern P . We construct P by concatenating, in some fixed order, the following strings:

$$P(i, x, y) = v(i)xa^{\mathcal{N}/s}x\$ya^{\mathcal{N}/s}yv(i)$$

360 for every $i = 1, 2, \dots, \mathcal{N}$ and $x, y = 1, 2, \dots, s$, where $a \in \Sigma_1, \$ \in \Sigma_2, x \in \Sigma_3, y \in \Sigma_4,$
 361 $v(i) \in \Sigma_5$, and $\Sigma_1, \Sigma_2, \dots, \Sigma_5$ are disjoint subsets of Σ .

362 **ED text \tilde{T} .** The text \tilde{T} consists of three parts. Its middle part encodes all the entries
 363 equal to 1 in matrices A, B and C , and consists of three string sets $\mathcal{X} = \mathcal{X}_1 \cdot \mathcal{X}_2 \cdot \mathcal{X}_3,$
 364 where:

- 365 1. \mathcal{X}_1 contains all strings of the form $v(i)xa^j$, for some $i \in [\mathcal{N}], x \in [s]$ and
 366 $j \in [\mathcal{N}/s]$ such that $A[i, (x-1) \cdot (\mathcal{N}/s) + j] = 1$;
- 367 2. \mathcal{X}_2 contains all strings of the form $a^{\mathcal{N}/s-j}x\$ya^{\mathcal{N}/s-k}$, for some $x, y \in [s]$ and
 368 $j, k \in [\mathcal{N}/s]$ such that $B[(x-1) \cdot (\mathcal{N}/s) + j, (y-1) \cdot (\mathcal{N}/s) + k] = 1$, i.e., if
 369 the corresponding entry of B is 1;
- 370 3. \mathcal{X}_3 contains all strings of the form $a^k yv(i)$, for some $i \in [\mathcal{N}], y \in [s]$ and
 371 $k \in [\mathcal{N}/s]$ such that $C[(y-1) \cdot (\mathcal{N}/s) + k, i] = 1$.

372 It is easy to see that $|P(i, x, y)| = \mathcal{O}(\mathcal{N}/s)$. This implies the following:

- 373 1. The length of the pattern is $m = \mathcal{O}(\mathcal{N} \cdot s^2 \cdot \mathcal{N}/s) = \mathcal{O}(\mathcal{N}^2 \cdot s)$;
- 374 2. The total length of \mathcal{X} is $|\mathcal{X}| = \mathcal{O}(\mathcal{N} \cdot s \cdot \mathcal{N}/s \cdot \mathcal{N}/s + s^2 \cdot (\mathcal{N}/s)^2 \cdot \mathcal{N}/s + \mathcal{N} \cdot$
 375 $s \cdot \mathcal{N}/s \cdot \mathcal{N}/s) = \mathcal{O}(\mathcal{N}^3/s)$.

376 By the above construction, we obtain the following fact.

377 **FACT 1.** $P(i, x, y)$ matches \mathcal{X} if and only if, for some $j, k = 1, 2, \dots, \mathcal{N}/s$, we
 378 have $A[i, (x-1) \cdot (\mathcal{N}/s) + j] = 1, B[(x-1) \cdot (\mathcal{N}/s) + j, (y-1) \cdot (\mathcal{N}/s) + k] = 1$ and
 379 $C[(y-1) \cdot (\mathcal{N}/s) + k, i] = 1$.

380 Solving the TD problem thus reduces to taking the disjunction of all such con-
 381 ditions. Let us write down all strings $P(i, x, y)$ in some arbitrary but fixed order to
 382 obtain $P = P_1 P_2 \dots P_z$ with $z = \mathcal{N}s^2$ being a power of 2, where every $P_t = P(i, x, y),$
 383 for some i, x, y . We aim to construct a small number of sets of strings that, when
 384 considered as an ED text, match any prefix $P_1 P_2 \dots P_t$ of the pattern, $1 \leq t \leq z-1$;
 385 a similar construction can be carried on to obtain sets of strings that match any suffix
 386 $P_k \dots P_{z-1} P_z, 2 \leq k \leq z$. These sets will then be added to the left and to the right
 387 of \mathcal{X} , respectively, to obtain the ED text \tilde{T} .

388 **ED Prefix.** We construct $\log z$ sets of strings as follows. The first one contains
 389 the empty string ε and $P_1 P_2 \dots P_{z/2}$. The second one contains $\varepsilon, P_1 P_2 \dots P_{z/4}$
 390 and $P_{z/2+1} \dots P_{z/2+z/4}$. The third one contains $\varepsilon, P_1 P_2 \dots P_{z/8}, P_{z/4+1} \dots P_{z/4+z/8},$
 391 $P_{z/2+1} \dots P_{z/2+z/8}$ and $P_{z/2+z/4+1} \dots P_{z/2+z/4+z/8}$.

392 Formally, for every $i = 1, 2, \dots, \log z$, the i -th of such sets is:

$$393 \tilde{T}_i^p = \varepsilon \cup \{P_{j_{\frac{z}{2^{i-1}}+1}} \dots P_{j_{\frac{z}{2^{i-1}}+2^i}} \mid j = 0, 1, \dots, 2^{i-1} - 1\}.$$

394 **ED Suffix.** We similarly construct $\log z$ sets to be appended to \mathcal{X} :

$$395 \quad \tilde{T}_i^s = \varepsilon \cup \{P_{z-j\frac{z}{2^{i-1}}-\frac{z}{2^i}+1} \dots P_{z-j\frac{z}{2^{i-1}}} \mid j = 0, 1, \dots, 2^{i-1} - 1\}.$$

396 The total length of all the ED prefix and ED suffix strings is $\mathcal{O}(\log z \cdot \mathcal{N}^2 \cdot s) =$
 397 $\mathcal{O}(\mathcal{N}^2 \cdot s \cdot \log \mathcal{N})$. The whole ED text \tilde{T} is thus: $\tilde{T} = \tilde{T}_1^p \dots \tilde{T}_{\log z}^p \cdot \mathcal{X} \cdot \tilde{T}_{\log z}^s \dots \tilde{T}_1^s$.
 398 We next show how a solution of such instance of EDSM corresponds to the solution
 399 of TD.

400 **LEMMA 4.1.** *The pattern P occurs in the ED text \tilde{T} if and only if there exist i, j, k*
 401 *such that $A[i, j] = B[j, k] = C[k, i] = 1$.*

402 *Proof.* By Fact 1, if such i, j, k exist then P_t matches \mathcal{X} , for some $t \in \{1, \dots, z\}$.
 403 Then, by construction of the sets \tilde{T}_i^p and \tilde{T}_i^s , the prefix $P_1 \dots P_{t-1}$ matches the ED
 404 prefix (this can be proved by induction), and similarly the suffix $P_{t+1} \dots P_z$ matches
 405 the ED suffix, so the whole P matches \tilde{T} , and so P occurs therein. In the other
 406 direction, assume that there is an occurrence of the pattern P in \tilde{T} . Because the
 407 letter $\$$ appears only in the center of every P_i and in the strings from \mathcal{X}_2 , and it can
 408 be verified that in any string from $\tilde{T}_1^p \dots \tilde{T}_{\log z}^p$ or $\tilde{T}_{\log z}^s \dots \tilde{T}_1^s$ there are fewer than
 409 z such letters, it must be the case that for some P_t its $\$$ is aligned with a $\$$ from some
 410 $X_2 \in \mathcal{X}_2$. But then by the subalphabets being disjoint we must have $X_1 X_2 X_3 = P_t$
 411 for some $X_1 \in \mathcal{X}_1, X_2 \in \mathcal{X}_2, X_3 \in \mathcal{X}_3$, and by Fact 1 there exists a triangle. \square

412 Note that for the EDSM problem we have $m = \mathcal{N}^2 \cdot s$, $n = 1 + 2\log z$ and $N =$
 413 $\|\mathcal{X}\| + \mathcal{O}(\mathcal{N}^2 \cdot s \cdot \log \mathcal{N})$. Thus if we had a solution running in $\mathcal{O}(\log z \cdot m^{1.5-\epsilon} + \|\mathcal{X}\| +$
 414 $\mathcal{N}^2 \cdot s \cdot \log \mathcal{N}) = \mathcal{O}(\log \mathcal{N} \cdot (\mathcal{N}^2 \cdot s)^{1.5-\epsilon} + \mathcal{N}^3/s)$ time, for some $\epsilon > 0$, by choosing
 415 a sufficiently small $\alpha > 0$ and setting $s = \mathcal{N}^\alpha$ we would obtain, for some $\delta > 0$, an
 416 $\mathcal{O}(\mathcal{N}^{3-\delta})$ -time algorithm for TD. This ends the proof of Theorem 1.1. \square

417 In order to show that AP cannot be solved in time $\mathcal{O}(m^{1.5-\epsilon} + N)$ with a combi-
 418 natorial algorithm unless there is a truly subcubic combinatorial algorithm for BMM
 419 (Theorem 1.3), in Section 3, we have exhibited a fully detailed reduction from BMM
 420 to the AP problem. However, now that we have proved a lower bound for EDSM,
 421 we remark that Theorem 1.1 also implies Theorem 1.3. Indeed, assuming that the
 422 AP problem can be solved in $\mathcal{O}(m^{1.5-\epsilon} + N)$ time, then by calling the AP problem n
 423 times (as described in Section 2 under the definition of the EDSM problem), we could
 424 solve the EDSM problem in $\mathcal{O}(nm^{1.5-\epsilon} + N)$ time. At that point, we could apply
 425 Theorem 1.1 and obtain a truly subcubic combinatorial algorithm for BMM.

426 **5. An $\tilde{\mathcal{O}}(nm^{\omega-1} + N)$ -time Algorithm for EDSM.** Our goal is to design a
 427 non-combinatorial $\tilde{\mathcal{O}}(nm^{\omega-1} + N)$ -time algorithm for EDSM, which in turn can be
 428 achieved with a non-combinatorial $\tilde{\mathcal{O}}(m^{\omega-1} + N)$ -time algorithm for the AP problem,
 429 that is the bottleneck of EDSM (cf. [39]).

430 We reduce AP to a logarithmic number of restricted instances of the same prob-
 431 lem, based on the length of the strings in \mathcal{S} . We start by giving a lemma that we will
 432 use to process naively the strings of length up to a constant c , to be determined later,
 433 in $\mathcal{O}(m \log m + N)$ time.

434 **LEMMA 5.1.** *For any integer t , all strings in \mathcal{S} of length at most t can be processed*
 435 *in $\mathcal{O}(m \log m + mt + N)$ time.*

436 *Proof.* We first construct the suffix tree ST of P in $\mathcal{O}(m \log m)$ time [65]. Let us
 437 remark that we are spending $\mathcal{O}(m \log m)$ time and not just $\mathcal{O}(m)$ so as to avoid any
 438 assumptions on the size of the alphabet. For every explicit node $u \in ST$, we construct
 439 a perfect hash function mapping the first letter on every edge outgoing from u to the

440 corresponding edge. This takes $\mathcal{O}(m \log m)$ total time [61] and allows us to navigate
 441 in ST in constant time per letter. For every $S \in \mathcal{S}$, find and mark its corresponding
 442 (implicit or explicit) node of ST . This takes $\mathcal{O}(N)$ time overall. For every possible
 443 length $t' \leq t$, scan P with a window of length t' while maintaining its corresponding
 444 (implicit or implicit) node of ST . To move the window to the right, we first follow
 445 the suffix link of the current node (if the node is implicit, we follow the suffix link
 446 of its nearest explicit ancestor, and then descend to find the node corresponding to
 447 the truncated window), and then follow the appropriate edge. This takes $\mathcal{O}(mt)$ total
 448 time by standard amortization based on counting the number of explicit ancestors of
 449 the current node. If the current window $P[i..(i+t'-1)]$ corresponds to a marked
 450 node of ST and additionally $U[i-1] = 1$, we set $V[i+t'-1] = 1$. \square

451 We build the restricted instances of the AP problem by considering only strings in
 452 $\mathcal{S}_k \subseteq \mathcal{S}$ of length in $[(19/18)^k, (19/18)^{k+1})$ for each integer k ranging from $\left\lceil \frac{\log c}{\log(19/18)} \right\rceil$
 453 to $\left\lfloor \frac{\log m}{\log(19/18)} \right\rfloor$. These sets form a partition of the set of all strings in \mathcal{S} of lengths up
 454 to m ; longer strings are not needed when solving the AP problem.

455 For each integer k from $\left\lceil \frac{\log c}{\log(19/18)} \right\rceil$ to $\left\lfloor \frac{\log m}{\log(19/18)} \right\rfloor$, let ℓ be an integer such that
 456 the length of every string in \mathcal{S}_k belongs to $[9/8 \cdot \ell, 5/4 \cdot \ell)$. Note that such an integer
 457 always exists for an appropriate choice of the integer constant c . In fact, it must hold
 458 that

$$459 \quad \frac{9}{8} \cdot \ell \leq \left(\frac{19}{18}\right)^k < \left(\frac{19}{18}\right)^{k+1} \leq \frac{5}{4} \cdot \ell \iff \frac{4}{5} \cdot \left(\frac{19}{18}\right)^{k+1} \leq \ell \leq \frac{8}{9} \cdot \left(\frac{19}{18}\right)^k.$$

460 To ensure that there exists an *integer* ℓ satisfying such conditions, we require that

$$461 \quad \frac{4}{5} \cdot \left(\frac{19}{18}\right)^{k+1} + 1 \leq \frac{8}{9} \cdot \left(\frac{19}{18}\right)^k \iff \frac{45}{2} \leq \left(\frac{19}{18}\right)^k.$$

462 The last equation holds for $k \geq 58$, implying that we will process naïvely the strings
 463 of length up to $c = 23$, and each \mathcal{S}_k , for k ranging from 58 to $\left\lfloor \frac{\log m}{\log(19/18)} \right\rfloor$, will be
 464 processed separately as described in the next paragraph.

465 *Remark 5.2.* The length of every string in \mathcal{S} belonging to $[9/8 \cdot \ell, 5/4 \cdot \ell)$ implies
 466 that every string in \mathcal{S} contains at most $\ell/4$ length- ℓ substrings (and at least $1 + \ell/8$
 467 of them).

468 Denoting by N_k the total size of strings in \mathcal{S}_k , we have that, if we solve every
 469 such instance of AP in $\mathcal{O}(N_k + f(m))$ time, then we can solve the original instance of
 470 AP in $\mathcal{O}(N + f(m) \log m)$ time by taking the disjunction of the results. Switching to
 471 $\tilde{\mathcal{O}}$ notation that disregards polylog factors, it thus suffices to solve each such instance
 472 of the AP problem in $\tilde{\mathcal{O}}(N + m^{\omega-1})$ time.

473 We further partition the strings in \mathcal{S}_k into three types, compute the corresponding
 474 bit vector V for each type separately and, finally, take the disjunction of the resulting
 475 bit vectors V to obtain the answer for each restricted instance.

476 **Partitioning \mathcal{S}_k .** Keeping in mind that from now on (until Section 5.4) we address
 477 the AP problem assuming that \mathcal{S} only contains strings of length in $[9/8 \cdot \ell, 5/4 \cdot \ell)$,
 478 and thus is in fact \mathcal{S}_k , to lighten the notation we now switch back to denote it simply
 479 with \mathcal{S} , and similarly write N to denote the total length of all strings given as the
 480 input to the AP problem. The three types of strings are as follows:

481 **Type 1:** Strings $S \in \mathcal{S}$ such that every length- ℓ substring of S is not strongly peri-
 482 odic.

483 **Type 2:** Strings $S \in \mathcal{S}$ containing at least one length- ℓ substring that is not strongly
 484 periodic and at least one length- ℓ substring that is strongly periodic.

485 **Type 3:** Strings $S \in \mathcal{S}$ such that every length- ℓ substring of S is strongly periodic
 486 (in Lemma 5.3 we show that in this case $\text{per}(S) \leq \ell/4$).

487 These three types are evidently a partition of \mathcal{S} . We start with showing that, in
 488 fact, strings of type 3 are exactly strings with period at most $\ell/4$. It is straightfor-
 489 ward to verify that strings with period at most $\ell/4$ are such that all their length- ℓ
 490 substrings have period at most $\ell/4$; the following lemma addresses the other (less
 491 obvious) direction.

492 **LEMMA 5.3.** *Let S be a string. If $\text{per}(S[j..j + \ell - 1]) \leq \ell/4$ for every j then*
 493 *$\text{per}(S) \leq \ell/4$.*

494 *Proof.* We first show that, for any string W and letters a, b , if $\text{per}(aW) \leq |aW|/4$
 495 and $\text{per}(Wb) \leq |Wb|/4$ then $\text{per}(aW) = \text{per}(Wb)$. This follows from Lemma 2.1: since
 496 $\text{per}(aW)$ and $\text{per}(Wb)$ are both periods of W and $(1 + |W|)/4 \leq |W|/2$, then we have
 497 that $d = \text{gcd}(\text{per}(aW), \text{per}(Wb))$ is a period of W . Assuming by contradiction that
 498 $\text{per}(aW) \neq \text{per}(Wb)$, then it must be that either $d < \text{per}(aW)$ or $d < \text{per}(Wb)$; by
 499 symmetry it is enough to consider the former possibility, and we claim that then d is a
 500 period of aW . Indeed, $a = W[\text{per}(aW)]$ (observe that, since $\text{per}(aW) \leq (1 + |W|)/4 \leq$
 501 $|W|/2$, in particular $\text{per}(aW) < |W|$) and $W[i] = W[i + d]$ for any $i = 1, 2, \dots, |W| - d$,
 502 so by $\text{per}(aW)$ being a multiple of d , we obtain that $a = W[\text{per}(aW)] = W[d]$, which is
 503 a contradiction because, by definition of $\text{per}(aW)$, we have that $d < \text{per}(aW)$ cannot
 504 be a period of aW .

505 If $\text{per}(S[j..j + \ell - 1]) \leq \ell/4$ for every j then by the above reasoning the periods of
 506 all substrings $S[j..j + \ell - 1]$ are all equal to the same $p \leq \ell/4$. But then $S[i] = S[i + p]$
 507 for every i , so $\text{per}(S) \leq \ell/4$. \square

508 Before proceeding with the algorithm, we show that, for each string $S \in \mathcal{S}$, we
 509 can determine its type in $\mathcal{O}(|S|)$ time.

510 **LEMMA 5.4.** *Given a string S we can determine its type in $\mathcal{O}(|S|)$ time.*

511 *Proof.* It is well-known that $\text{per}(T)$ can be computed in $\mathcal{O}(|T|)$ time for any string
 512 T (cf. [28]). We partition S into blocks $T_\alpha = S[\alpha \lfloor \ell/2 \rfloor .. (\alpha + 1) \lfloor \ell/2 \rfloor - 1]$ of size $\lfloor \ell/2 \rfloor$,
 513 and compute $\text{per}(T_\alpha)$ for every α in $\mathcal{O}(|S|)$ total time. Observe that every substring
 514 $S[i..i + \ell - 1]$ contains at least one whole block inside.

515 If $\text{per}(T_\alpha) > \ell/4$ then the period of any substring $S[i..i + \ell - 1]$ that contains T_α
 516 is also larger than $\ell/4$. Consequently, if $\text{per}(T_\alpha) > \ell/4$ for every α , then we declare S
 517 to be of type 1.

518 Consider any α such that $p = \text{per}(T_\alpha) \leq \ell/4$. If the period p' of a substring
 519 $S' = S[i..i + \ell - 1]$ that contains T_α is at most $\ell/4$, then in fact it must be equal to
 520 p , because $p' \geq p$ and so, by Lemma 2.1 applied on T_α , p' must be a multiple of p
 521 and, by repeatedly applying $S'[j] = S'[j + p']$ and $T_\alpha[j] = T_\alpha[j + p]$ and using the fact
 522 that T_α occurs inside S' , we conclude that in fact $S'[j] = S'[j + p]$ for any j , and thus
 523 $p' = p$. This allows us to check whether there exists a substring $S' = S[i..i + \ell - 1]$
 524 that contains T_α such that $\text{per}(S') \leq \ell/4$ by computing, in $\mathcal{O}(\ell)$ time, how far the
 525 period p extends to the left and to the right of T_α in $T_{\alpha-1}T_\alpha T_{\alpha+1}$ (if either $T_{\alpha-1}$ or
 526 $T_{\alpha+1}$ do not exist, then we do not extend the period in the corresponding direction).
 527 There exists such a substring S' if and only if the length of the extended substring
 528 with period p is at least ℓ . Therefore, for every α we can check in $\mathcal{O}(\ell)$ time if there

529 exists a length- ℓ substring S' containing T_α with $\text{per}(S') \leq \ell/4$. By repeating this
 530 procedure for every α , we can distinguish between S of type 2 and S of type 3 in
 531 $\mathcal{O}(|S|)$ total time. \square

532 Since we have shown how to efficiently partition the strings of S into the three
 533 types, in what follows we present our solution of the AP problem for each type of
 534 strings separately.

535 **5.1. Type 1 Strings.** In this section we show how to solve a restricted instance
 536 of the AP problem where every string $S \in \mathcal{S}$ is of type 1, that is, each of its length- ℓ
 537 substrings is not strongly periodic, and furthermore $|S| \in [9/8 \cdot \ell, 5/4 \cdot \ell)$ for some
 538 $\ell \leq m$. Observe that all (and hence at most $\ell/4$ by Remark 5.2) length- ℓ substrings of
 539 any $S \in \mathcal{S}$ must be distinct, as otherwise we would be able to find two occurrences of
 540 a length- ℓ substring at distance at most $\ell/4$ in S , making the period of the substring
 541 at most $\ell/4$ and contradicting the assumption that S is of type 1.

542 We start with constructing the suffix tree ST of P (our pattern in the EDSM
 543 problem) and storing, for every node, the first letters on its outgoing edges in a static
 544 dictionary with constant access time. Then, for every $S \in \mathcal{S}$, we check in $\mathcal{O}(|S|)$ time
 545 using ST if it occurs in P and, if not, we disregard it from further consideration.
 546 Therefore, from now on we assume that all strings S , and thus all their length- ℓ
 547 substrings, occur in P . We will select a set of length- ℓ substrings of P , called the
 548 *anchors*, each represented by one of its occurrences in P , such that:

- 549 1. The total number of occurrences of all anchors in P is $\mathcal{O}(m/\ell \cdot \log m)$.
- 550 2. For every $S \in \mathcal{S}$, at least one of its length- ℓ substrings is an anchor.
- 551 3. The total number of occurrences of all anchors in strings $S \in \mathcal{S}$ is $\mathcal{O}(|\mathcal{S}| \cdot$
 552 $\log m)$.

553 We formalize this using the following auxiliary problem, which is a strengthening of
 554 a well-known *Hitting Set* problem, which given a collection of m sets over $[n]$, each of
 555 size at least k , asks to choose a subset of $[n]$ of size $\mathcal{O}(n/k \cdot \log m)$ that nontrivially
 556 intersects every set.

NODE SELECTION (NS)

557 **INPUT:** A bipartite graph $G = (U, V, E)$ with $\deg(u) \in (d, 2d]$ for every $u \in U$
 and weight $w(v)$ for every $v \in V$, where $W = \sum_{v \in V} w(v)$.

OUTPUT: A set $V' \subseteq V$ of total weight $\mathcal{O}(W/d \cdot \log |U|)$ such that $N[u] \cap V' \neq \emptyset$
 for every node $u \in U$, and $\sum_{u \in U} |N[u] \cap V'| = \mathcal{O}(|U| \log |U|)$.

558 We reduce the problem of finding anchors to an instance of the NS problem, by
 559 building a bipartite graph G in which the nodes in U correspond to strings $S \in \mathcal{S}$,
 560 the nodes in V correspond to distinct length- ℓ substrings of P , and there is an edge
 561 (u, v) if the length- ℓ string corresponding to v occurs in the string S corresponding
 562 to u . Using suffix links, we can find the node of the suffix tree corresponding to
 563 every length- ℓ substring of S in $\mathcal{O}(|S|)$ total time, so the whole construction takes
 564 $\mathcal{O}(m \log m + \sum_{S \in \mathcal{S}} |S|) = \mathcal{O}(m \log m + N)$ time. The size of G is $\mathcal{O}(m + N)$, and the
 565 degree of every node in U belongs to $(\ell/8, \ell/4]$. We set the weight of a node $v \in V$ to
 566 be its number of occurrences in P , and solve the obtained instance of the NS problem
 567 to obtain the set of anchors. We remark that, because each string $S \in \mathcal{S}$ can be
 568 assumed to be a substring of P and we do not need to keep duplicate strings in \mathcal{S} ,
 569 we have $\log |U| = \Theta(\log m)$ and the three required properties indeed hold assuming
 570 that we have found a solution. However, it is not immediately clear that an instance
 571 of the NS problem always has a solution. We show that indeed it does, and that it

572 can be found in linear time.

573 LEMMA 5.5. *A solution to an instance of the NS problem always exists and can*
 574 *be found in linear time in the size of G .*

575 *Proof.* We first show a solution that uses the probabilistic method and leads us
 576 to an efficient Las Vegas algorithm; we will then derandomize the solution using the
 577 method of conditional expectations.

578 We independently choose each node of V with probability p to obtain the set V'
 579 of selected nodes. The expected total weight of V' is $\sum_{v \in V} p \cdot w(v) = p \cdot W$, so by
 580 Markov's inequality it exceeds $4p \cdot W$ with probability at most $1/4$. For every node
 581 $u \in U$, the probability that $N[u]$ does not intersect V' is at most $(1 - p)^d \leq e^{-pd}$.
 582 Finally, $\mathbb{E}[\sum_{u \in U} |N[u] \cap V'|] \leq |U| \cdot 2pd$, so by Markov's inequality $\sum_{u \in U} |N[u] \cap V'|$
 583 exceeds $|U| \cdot 8pd$ with probability at most $1/4$. We set $p = \ln(4|U|)/d$ (observe that
 584 if $p > 1$ then we can select all nodes in V). By union bound, the probability that V'
 585 is not a valid solution is at most $3/4$, so indeed a valid solution exists. Furthermore,
 586 this reasoning gives us an efficient Las Vegas algorithm that chooses V' randomly
 587 as described above and then verifies if it constitutes a valid solution. Each iteration
 588 takes linear time in the size of G , and the expected number of required iterations is
 589 constant.

590 To derandomize the above procedure we apply the method of conditional expect-
 591 ations. Let X_1, X_2, \dots be the binary random variables corresponding to the nodes
 592 of V . Recall that in the above proof we set $X_i = 1$ with probability p . Now we
 593 will choose the values of X_1, X_2, \dots one-by-one. Define a function $f(X_1, X_2, \dots)$ that
 594 bounds the probability that X_1, X_2, \dots corresponds to a valid solution as follows:

$$595 \quad f(X_1, X_2, \dots) = \frac{\sum_v X_v \cdot w(v)}{4W/d \cdot \ln(4|U|)} + \sum_{u \in U} \prod_{v \in N[u]} (1 - X_v) + \frac{\sum_{u \in U} \sum_{v \in N[u]} X_v}{8|U| \ln(4|U|)}.$$

596 As explained above, we have $\mathbb{E}[f(X_1, X_2, \dots)] = 3/4$. Assume that we have already
 597 fixed the values $X_1 = x_1, \dots, X_i = x_i$. Then there must be a choice of $X_{i+1} = x_{i+1}$
 598 that does not increase the expected value of $f(X_1, X_2, \dots)$ conditioned on the already
 599 chosen values. We want to compare the following two quantities:

$$600 \quad \mathbb{E}[f(X_1, X_2, \dots) \mid X_1 = x_1, \dots, X_i = x_i, X_{i+1} = 0]$$

$$601 \quad \mathbb{E}[f(X_1, X_2, \dots) \mid X_1 = x_1, \dots, X_i = x_i, X_{i+1} = 1]$$

603 and choose x_{i+1} corresponding to the smaller one. Canceling out the shared terms,
 604 we need to compare the expected values of:

$$605 \quad 0 \quad + \quad \sum_{u \in N[i+1]} \prod_{v \in N[u]} (1 - X_v) \quad + \quad 0 \quad \text{and}$$

$$606 \quad \frac{w(i+1)}{4W/d \cdot \ln(4|U|)} \quad + \quad 0 \quad + \quad \frac{\deg(i+1)}{8|U| \ln(4|U|)}.$$

608 The second quantity can be computed in constant time. We claim that (ignoring the
 609 issue of numerical precision) the first quantity can be computed in time $\mathcal{O}(\deg(i+1))$
 610 after a linear-time preprocessing as follows. In the preprocessing we compute and
 611 store $E[i] = (1 - p)^i$, for every $i = 0, 1, \dots, |V|$ in $\mathcal{O}(|V|)$ total time. Then, during
 612 the computation we maintain, for every $u \in U$, the number $c[u]$ of $v \in N[u]$ for which
 613 we still need to choose the value X_v , and a single bit $b[u]$ denoting whether for some

614 $v \in N[u] \cap \{1, \dots, i\}$ we already have $x_v = 1$. This information can be updated in
 615 $\mathcal{O}(\deg(i+1))$ time after selecting x_{i+1} . Now to compute the first quantity, we iterate
 616 over $u \in N[i+1]$ and, if $b[u] = 0$ then we add $E[c[u]]$ to the result. Finally, we
 617 claim that it is enough to implement all calculations with precision of $\Theta(\log |V|)$ bits.
 618 This is because such precision allows us to calculate both quantities with relative
 619 accuracy $1/(8|V|)$, so the expected value of $f(X_1, X_2, \dots)$ might increase by a factor
 620 of $(1 + 1/(4|V|))$ in every step, which is at most $(1 + 1/(4|V|))^{|V|} \leq e^{1/4}$ overall. This
 621 still guarantees that the final value is at most $3/4 \cdot e^{1/4} < 1$, so we obtain a valid
 622 solution. \square

623 In the rest of this section we explain how to compute the bit vector V from the bit
 624 vector U , and thus solve the AP problem, after having obtained a set \mathcal{A} of anchors.
 625 For any $S \in \mathcal{S}$, since S contains an occurrence of at least one anchor $H \in \mathcal{A}$, say
 626 $S[j..(j+|H|-1)] = H$, so any occurrence of S in P can be generated by choosing
 627 some occurrence of H in P , say $P[i..(i+|H|-1)] = H$, and then checking that
 628 $S[1..(j-1)] = P[(i-j+1)..(i-1)]$ and $S[(j+|H|)..|S|] = P[(i+|H|)..(i+|S|-j)]$.
 629 In other words, $S[1..(j-1)]$ should be a suffix of $P[1..(i-1)]$ and $S[(j+|H|)..|S|]$
 630 should be a prefix of $P[(i+|H|)..|P|]$. In such case, we say that the occurrence of S in
 631 P is generated by H . By the properties of \mathcal{A} , any occurrence of $S \in \mathcal{S}$ is generated by
 632 $occ_S \geq 1$ occurrences of anchors, where $\sum_{S \in \mathcal{S}} occ_S = \mathcal{O}(|\mathcal{S}| \log m)$. For every $H \in \mathcal{A}$
 633 we create a separate data structure $D(H)$ responsible for setting $V[i+|S|-1] = 1$,
 634 when $U[i-1] = 1$ and $P[i..(i+|S|-1)] = S$ is generated by H . We now first describe
 635 what information is used to initialize each $D(H)$, and how this is later processed to
 636 update V .

637 **Initialization.** $D(H)$ consists of two compact tries $T(H)$ and $T^r(H)$. For every
 638 occurrence of H in P , denoted by $P[i..(i+|H|-1)] = H$, $T(H)$ should contain a leaf
 639 corresponding to $P[(i+|H|)..|P|]\$$ and $T^r(H)$ should contain a leaf corresponding
 640 to $(P[1..(i-1)])^r\$$, both decorated with position i . Additionally, $D(H)$ stores a list
 641 $L(H)$ of pairs of nodes (u, v) , where $u \in T^r(H)$ and $v \in T(H)$ (both nodes might be
 642 implicit or explicit). Each such pair corresponds to an occurrence of H in a string
 643 $S \in \mathcal{S}$, $S[j..(j+|H|-1)] = H$, where u is the node of $T^r(H)$ corresponding to
 644 $(S[1..(j-1)])^r\$$ and v is the node of $T(H)$ corresponding to $S[(j+|H|)..|S|]\$$.
 645 We claim that $D(H)$, for all H , can be constructed in $\mathcal{O}(m \log m + N)$ total time.

646 We first construct the suffix tree ST of $P\$$ and the suffix tree ST^r of $P^r\$$ (again in
 647 $\mathcal{O}(m \log m)$ time not to make assumptions on the alphabet). We augment both trees
 648 with data for answering both *weighted ancestor* (WA) and *lowest common ancestor*
 649 (LCA) queries, that are defined as follows. For a rooted tree T on n nodes with an
 650 integer weight $\mathcal{D}(v)$ assigned to every node u , such that the weight of the root is
 651 zero and $\mathcal{D}(u) < \mathcal{D}(v)$ if u is the parent of v , we say that a node v is a weighted
 652 ancestor of a node u at depth ℓ , denoted by $WA_T(u, \ell)$, if v is the highest ancestor
 653 of u with weight at least ℓ . Such queries can be answered in $\mathcal{O}(\log n)$ time after an
 654 $\mathcal{O}(n)$ preprocessing [32]. For a rooted tree T , $LCA_T(u, v)$ is the lowest node that is an
 655 ancestor of both u and v . Such queries can be answered in $\mathcal{O}(1)$ time after an $\mathcal{O}(n)$
 656 preprocessing [12]. Recall that every anchor H is represented by one of its occurrences
 657 in P . Using WA queries, we can access in $\mathcal{O}(\log m)$ time the nodes corresponding to H
 658 and H^r , respectively, and extract a lexicographically sorted list of suffixes following an
 659 occurrence of H in $P\$$ and a lexicographically sorted list of reversed prefixes preceding
 660 an occurrence of H in $P^r\$$ in time proportional to the number of such occurrences.
 661 Then, by iterating over the lexicographically sorted list of suffixes and using LCA
 662 queries on ST we can build $T(H)$ in time proportional to the length of the list, and




FIG. 1. An occurrence of S starting at position i in P is generated by H : (u, v) corresponds to $S[j..(j+|H|-1)] = H$ and i appears in the subtree of $T^r(H)$ rooted at u , as well as in the subtree of $T(H)$ rooted at v .

663 similarly we can build $T^r(H)$. To construct $L(H)$ we start by computing, for every
 664 $S \in \mathcal{S}$ and $j = 1, \dots, |S|$, the node of ST^r corresponding to $(S[1..j])^r$ and the node
 665 of ST corresponding to $S[(j+1)..|S|]$ (the nodes might possibly be implicit). This
 666 takes only $\mathcal{O}(|S|)$ time, by using suffix links. We also find, for every length- ℓ substring
 667 $S[j..(j+\ell-1)]$ of S , an anchor $H \in \mathcal{A}$ such that $S[j..(j+\ell-1)] = H$, if any exists.
 668 This can be done by finding the nodes (implicit or explicit) of ST that correspond to
 669 the anchors, and then scanning over all length- ℓ substrings while maintaining the node
 670 of ST corresponding to the current substring using suffix links in $\mathcal{O}(|S|)$ total time.
 671 After having determined that $S[j..(j+\ell-1)] = H$ we retrieve the previously found
 672 nodes u of ST^r and v of ST corresponding to $(S[1..(j-1)])^r$ and $S[(j+\ell)..|S|]$,
 673 respectively. Then we look up the node $u' \in T^r(H)$ corresponding to u and the node
 674 $v' \in T(H)$ corresponding to v , and if they both exist we add (u, v) to $L(H)$. This
 675 lookup can be implemented in $\mathcal{O}(\log m)$ time by binary searching over the leaves of
 676 the compact tries. By construction, we have the following property, also illustrated
 677 in Figure 1.

678 **FACT 2.** A string $S \in \mathcal{S}$ starts at position $i-j+1$ in P if and only if, for some
 679 anchor $H \in \mathcal{A}$, $L(H)$ contains a pair (u, v) corresponding to $S[j..(j+|H|-1)] = H$,
 680 such that the subtree of $T^r(H)$ rooted at u and that of $T(H)$ rooted at v contain a leaf
 681 decorated with i .

682 Note that the overall size of all lists $L(H)$, when summed up over all $H \in \mathcal{A}$, is
 683 $\sum_{S \in \mathcal{S}} |L(H)| = \mathcal{O}(|\mathcal{S}| \log m)$, and since each S is of length at least ℓ this is $\mathcal{O}(N/\ell \cdot \log m)$.

684 **Processing.** The goal of processing $D(H)$ is to efficiently process all occurrences
 685 generated by H . As a preliminary step, we decompose $T^r(H)$ and $T(H)$ into heavy
 686 paths. Then, for every pair of leaves $u \in T^r(H)$ and $v \in T(H)$ decorated by the same
 687 i , we consider all heavy paths above u and v . Let $p = u_1 - u_2 - \dots$ be a heavy path
 688 above u in $T^r(H)$ and $q = v_1 - v_2 - \dots$ be a heavy path above v in $T(H)$, where
 689 u_1 is the head of p and v_1 is the head of q , respectively. Further, choose the largest
 690 x such that u is in the subtree rooted at u_x , and the largest y such that v is in the
 691 subtree rooted at v_y (this is well-defined by the choice of p and q , as u is in the subtree
 692 rooted at u_1 and v is in the subtree rooted at v_1). We add $(i, |\mathcal{L}(u_x)|, |\mathcal{L}(v_y)|)$ to an
 693 auxiliary list associated with the pair of heavy paths (p, q) , where $\mathcal{L}(u)$ denotes the
 694 concatenation of the edge labels on the path from the root to node u . In the rest
 695 of the processing we work with each such list separately. Notice that the overall size
 696 of all auxiliary lists, when summed up over all $H \in \mathcal{A}$, is $\mathcal{O}(m/\ell \cdot \log^3 m)$, because
 697 there are at most $\log^2 m$ pairs of heavy paths above u and v decorated by the same i ,




FIG. 2. An occurrence of S starting at position i in P corresponds to a triple $(i, \mathcal{L}(u_x), \mathcal{L}(v_y))$ on some auxiliary list.

698 and the total number of leaves in all trees $T^r(H)$ and $T(H)$ is bounded by the total
 699 number of occurrences of all anchors in P , which is $\mathcal{O}(m/\ell \cdot \log m)$. By Fact 2, there
 700 is an occurrence of a string S generated by H and starting at position $i - j + 1$ in P if
 701 and only if $L(H)$ contains a pair (u, v) corresponding to $S[j..(j + |H| - 1)] = H$ such
 702 that, denoting by p the heavy path containing u in $T^r(H)$ and by q the heavy path
 703 containing v in $T(H)$, the auxiliary list associated with (p, q) contains a triple (i, x, y)
 704 such that $x \geq |\mathcal{L}(u)|$ and $y \geq |\mathcal{L}(v)|$. This is illustrated in Figure 2. Henceforth,
 705 we focus on the problem of processing a single auxiliary list associated with (p, q) ,
 706 together with a list of pairs (u, v) , such that u belongs to p and v belongs to q .

707 Processing an auxiliary list can be interpreted geometrically as follows: for every
 708 (i, x, y) we create a red point (x, y) , and for every (u, v) we create a blue point
 709 $(|\mathcal{L}(u)|, |\mathcal{L}(v)|)$. Then, each occurrence of $S \in \mathcal{S}$ generated by H corresponds to a
 710 pair of points (p_1, p_2) such that p_1 is red, p_2 is blue, and p_1 dominates p_2 . We further
 711 reduce this to a collection of simpler instances in which all red points already dominate
 712 all blue points. This can be done with a divide-and-conquer procedure which
 713 is essentially equivalent to constructing a 2D range tree [13]: we first apply a divide-
 714 and-conquer that splits the current set of points along the median x coordinate, and
 715 inside every each obtained subproblem consisting of the left and the right part we apply
 716 another divide-and-conquer that splits the current set of points along the median y
 717 coordinate. The total number of points in all obtained instances increases by a factor
 718 of $\mathcal{O}(\log^2 m)$, making the total number of red points in all instances $\mathcal{O}(m/\ell \cdot \log^5 m)$,
 719 while the total number of blue points is $\mathcal{O}(N/\ell \cdot \log^3 m)$. There is an occurrence of
 720 a string $S \in \mathcal{S}$ generated by H and starting at position $i - j + 1$ in P if and only if
 721 some simpler instance contains a red point created for some (i, x, y) and a blue point
 722 created for some (u, v) corresponding to $S[j..(j + |H| - 1)] = H$. In the following we
 723 focus on processing a single simpler instance.

724 To process a simpler instance we need to check if $U[i - j] = 1$, for a red point
 725 created for some (i, x, y) and a blue point created for some (u, v) corresponding to
 726 $S[j..(j + |H| - 1)] = H$, and if so set $V[i - j + |S|] = 1$. This has a natural
 727 interpretation as an instance of BMM: we create a $\lceil 5/4 \cdot \ell \rceil \times \lceil 5/4 \cdot \ell \rceil$ matrix M such
 728 that $M[\lceil 5/4 \cdot \ell \rceil - j, \lceil 5/4 \cdot \ell \rceil + 1 - j] = 1$ if and only if there is a blue point created for some
 729 (u, v) corresponding to $S[j..(j + |H| - 1)] = H$; then for every red point created for
 730 some (i, x, y) we construct a bit vector $U_i = U[(i - \lceil 5/4 \cdot \ell \rceil) .. (i - 1)]$ (if $i < \lceil 5/4 \cdot \ell \rceil$,
 731 we pad U_i with 0s to make its length always equal to $\lceil 5/4 \cdot \ell \rceil$); calculate $V_i = M \times U_i$;
 732 and finally set $V[i + j] = 1$ whenever $V_i[j] = 1$ (and $i + j \leq m$).

733 LEMMA 5.6. $V_i[k] = 1$ if and only if there is a blue point created for some (u, v)
 734 corresponding to $S[j..(j + |H| - 1)] = H$ such that $U[i - j] = 1$ and $k = |S| - j$.

735 *Proof.* By definition of $V_i = M \times U_i$, we have that $V_i[k] = 1$ if and only if
 736 $M[k, t] = 1$ for some t such that $U_i[t] = 1$. By definition of U_i , we have that $U_i[t] = 1$
 737 if and only if $U[i - \lceil 5/4 \cdot \ell \rceil + t - 1] = 1$, and hence the previous condition can be
 738 rewritten as $M[k, t] = 1$ and $U[i - \lceil 5/4 \cdot \ell \rceil + t - 1] = 1$, or equivalently, by substituting
 739 $j = \lceil 5/4 \cdot \ell \rceil + 1 - t$, $M[k, \lceil 5/4 \cdot \ell \rceil + 1 - j] = 1$ and $U[i - j] = 1$. By definition of M ,
 740 we have that $M[k, \lceil 5/4 \cdot \ell \rceil + 1 - j] = 1$ if and only if there is a blue point created for
 741 some (u, v) corresponding to $S[j \dots (j + |H| - 1)] = H$ with $k = |S| - j$, which proves
 742 the lemma. \square

743 The total length of all vectors U_i and V_i is $\mathcal{O}(m \log^5 m)$, so we can afford to
 744 extract the appropriate fragment of U and then update the corresponding fragment
 745 of V . The bottleneck is computing the matrix-vector product $V_i = M \times U_i$. Since the
 746 total number of 1s in all matrices M is bounded by the total number of blue points,
 747 a naïve method would take $\mathcal{O}(N/\ell \cdot \log^3 m)$ time; we overcome this by processing
 748 together all multiplications concerning the same matrix M , thus amortizing the costs.
 749 Let $U_{i_1}, U_{i_2}, \dots, U_{i_s}$ be all bit vectors that need to be multiplied with M , and let z
 750 a parameter to be determined later. We distinguish between two cases: (i) if $s < z$,
 751 then we compute the products naïvely by iterating over all 1s in M , and the total
 752 computation time, when summed up over all such matrices M , is $\mathcal{O}(N/\ell \cdot \log^3 m \cdot z)$;
 753 (ii) if $s \geq z$, then we partition the bit vectors into $\lceil s/z \rceil \leq s/z + 1$ groups of z
 754 (padding the last group with bit vectors containing all 0s) and, for every group, we
 755 create a single matrix whose columns contain all the bit vectors belonging to the
 756 group. Thus, we reduce the problem of computing all matrix-vector products $M \times U_i$
 757 to that of computing $\mathcal{O}(s/z)$ matrix-matrix products of the form $M \times M'$, where M'
 758 is an $\lceil 5/4 \cdot \ell \rceil \times z$ matrix. Even if M' is not necessarily a square matrix, we can still
 759 apply the fast matrix multiplication algorithm to compute $M \times M'$ using the standard
 760 trick of decomposing the matrices into square blocks.

761 **LEMMA 5.7.** *If two $N \times N$ matrices can be multiplied in $\mathcal{O}(N^\omega)$ time, then, for*
 762 *any $N \geq N'$, an $N \times N$ and an $N \times N'$ matrix can be multiplied in $\mathcal{O}((N/N')^2 N'^\omega)$*
 763 *time.*

764 *Proof.* We partition both matrices into blocks of size $N' \times N'$. There are $(N/N')^2$
 765 such blocks in the first matrix and N/N' in the second matrix. Then, to compute
 766 the product we multiply each block from the first matrix by the appropriate block in
 767 the second matrix in $\mathcal{O}(N'^\omega)$ time, resulting in the claimed complexity. \square

768 By applying Lemma 5.7, we can compute $M \times M'$ in $\mathcal{O}(\ell^2 z^{\omega-2})$ time (as long
 769 as we later verify that $5/4 \cdot \ell \geq z$), so all products $M \times U_i$ can be computed in
 770 $\mathcal{O}(\ell^2 z^{\omega-2} \cdot (s/z + 1))$ time. Note that this case can occur only $\mathcal{O}(m/(\ell \cdot z) \cdot \log^5 m)$ times,
 771 because all values of s sum up to $\mathcal{O}(m/\ell \cdot \log^5 m)$. This makes the total computation
 772 time, when summed up over all such matrices M , $\mathcal{O}(\ell^2 z^{\omega-2} \cdot m/(\ell \cdot z) \cdot \log^5 m) =$
 773 $\mathcal{O}(\ell z^{\omega-3} \cdot m \log^5 m)$. We can now prove our final result for strings of type 1.

774 **THEOREM 5.8.** *An instance of the AP problem where all strings are of type 1 can*
 775 *be solved in $\tilde{\mathcal{O}}(m^{\omega-1} + N)$ time.*

776 *Proof.* The total time complexity is first $\mathcal{O}(m + N)$ to construct the graph G ,
 777 then $\mathcal{O}(m \log m + N)$ to solve its corresponding instances of the NODESELECTION
 778 problem and obtain the set of anchors H . The time to initialize all structures $D(H)$
 779 is $\mathcal{O}(m \log m + N)$. For every $D(H)$, we obtain in $\mathcal{O}(m/\ell \cdot \log^5 m + N/\ell \cdot \log^3 m)$ time a
 780 number of simpler instances, and then construct the corresponding Boolean matrices
 781 M and bit vectors U_i in additional $\mathcal{O}(m \log^5 m)$ time. Note that some M might be
 782 sparse, so we need to represent them as a list of 1s. Then, summing up over all matrices

783 M and both cases, we spend $\mathcal{O}(N/\ell \cdot \log^3 m \cdot z + \ell z^{\omega-3} \cdot m \log^5 m)$ time. We would like
 784 to assume that $\ell \geq \log^3 m$ so that we can set $z = \ell/\log^3 m$. This is indeed possible,
 785 because for any t we can switch to a more naïve approach to process all strings of length
 786 at most t in $\mathcal{O}(m \log m + mt + N)$ time as described in Lemma 5.1. After applying
 787 it with $t = \log^3 m$ in $\mathcal{O}(m \log^3 m + N)$ time, we can set $z = \ell/\log^3 m$ (so that indeed
 788 $5/4 \cdot \ell \geq z$ as required in case $s \geq z$) and the overall time complexity for all matrices
 789 M and both cases becomes $\mathcal{O}(N + \ell^{\omega-2} \cdot m \log^{5+3(3-\omega)} m)$. Taking the initialization
 790 into account we obtain $\mathcal{O}(m \log^5 m + \ell^{\omega-2} \cdot m \log^{5+3(3-\omega)} m + N) = \tilde{\mathcal{O}}(m^{\omega-1} + N)$
 791 total time. \square

792 **5.2. Type 2 Strings.** In this section we show how to solve a restricted instance
 793 of the AP problem where every string $S \in \mathcal{S}$ is of type 2, that is, S contains a length- ℓ
 794 substring that is not strongly periodic as well as a length- ℓ substring that is strongly
 795 periodic, and furthermore $|S| \in [9/8 \cdot \ell, 5/4 \cdot \ell)$ for some $\ell \leq m$.

796 Similarly as in Section 5.1, we select a set of anchors. In this case, instead of the
 797 NODESELECTION problem we need to exploit periodicity. We call a string T ℓ -periodic
 798 if $|T| \geq \ell$ and $\text{per}(T) \leq \ell/4$. We consider all maximal ℓ -periodic substrings of S , that
 799 is, ℓ -periodic substrings $S[i..j]$ such that either $i = 1$ or $\text{per}(S[(i-1)..j]) > \ell/4$,
 800 and $j = |S|$ or $\text{per}(S[i..(j+1)]) > \ell/4$. We know that S contains at least one such
 801 substring (because there exists a length- ℓ substring that is strongly periodic), and
 802 that the whole S is not such a substring (because otherwise S would be of type 3).
 803 Further, two maximal ℓ -periodic substrings cannot overlap too much, as formalized
 804 in the following lemma.

805 **LEMMA 5.9.** *Any two distinct maximal ℓ -periodic substrings of the same string S*
 806 *overlap by less than $\ell/2$ letters.*

807 *Proof.* Assume (by contradiction) the opposite; then we have two distinct ℓ -
 808 periodic substrings $S[i..j]$ and $S[i'..j']$ such that $i < i' \leq j < j'$ and $j - i' + 1 \geq \ell/2$.
 809 Then, both $p = \text{per}(S[i..j])$ and $p' = \text{per}(S[i'..j'])$ are periods of $S[i'..j]$, and hence
 810 by Lemma 2.1 we have that $\text{gcd}(p, p')$ is a period of $S[i'..j]$. If $p \neq p'$ then, because
 811 $S[i'..j]$ contains an occurrence of both $S[i..(i+p-1)]$ and $S[i'..(i'+p'-1)]$, we
 812 obtain that one of these two substrings is a power of a shorter string, thus contradict-
 813 ing the definition of p or p' . So $p = p'$, but then $p \leq \ell/4$ is actually a period of the
 814 whole $S[i..j']$, meaning that $S[i..j]$ and $S[i'..j']$ are not maximal, a contradiction. \square

815 By Lemma 5.9, every $S \in \mathcal{S}$ contains exactly one maximal ℓ -periodic substring,
 816 and by the same argument P contains $\mathcal{O}(m/\ell)$ such substrings. The set of anchors
 817 will be generated by considering the unique maximal ℓ -periodic substring of every
 818 $S \in \mathcal{S}$, so we first need to show how to efficiently generate such substrings.

819 **LEMMA 5.10.** *Given a string S of length at most $5/4 \cdot \ell$, we can generate its*
 820 *(unique) maximal ℓ -periodic substring in $\mathcal{O}(|S|)$ time.*

821 *Proof.* We start with observing that any length- ℓ substring of S must contain
 822 $S[(\lfloor \ell/2 \rfloor + 1).. \ell]$ inside. Consequently, we can proceed similarly as in the proof of
 823 Lemma 5.4. We compute $p = \text{per}(S[(\lfloor \ell/2 \rfloor + 1).. \ell])$ in $\mathcal{O}(|S|)$ time. If $p > \ell/4$ then
 824 S does not contain any ℓ -periodic substrings. Otherwise, we compute in $\mathcal{O}(|S|)$ time
 825 how far the period p extends to the left and to the right; that is, we compute the
 826 smallest $i \leq \lfloor \ell/2 \rfloor + 1$ such that $S[k] = S[k+p]$ for every $k = i, i+1, \dots, \lfloor \ell/2 \rfloor$
 827 and the largest $j \geq \ell$ such that $S[k] = S[k-p]$ for every $k = \ell+1, \ell+2, \dots, j$. If
 828 $j - i + 1 \geq \ell$ then $S[i..j]$ is a maximal ℓ -periodic substring of S , and, as shown earlier
 829 by Lemma 5.9, S cannot contain any other maximal ℓ -periodic substrings. We return

830 $S[i..j]$ as the (unique) maximal ℓ -periodic substring of S . \square

831 For every $S \in \mathcal{S}$, we apply Lemma 5.10 on S to find its (unique) maximal ℓ -
 832 periodic substring $S[i..j]$ in $\mathcal{O}(|S|)$ time. If $i > 1$ then we designate $S[(i-1)..(i-1+\ell)]$
 833 as an anchor, and similarly if $j < |S|$ we designate $S[(j+1-\ell)..(j+1)]$ as an
 834 anchor. Observe that because S is of type 2 (and not of type 3) either $i > 1$ or $j < |S|$,
 835 so for every $S \in \mathcal{S}$ we designate at least one of its length- $(\ell+1)$ substrings as an anchor.
 836 As in Section 5.1, we represent each anchor by one of its occurrences in P , and so
 837 need to find its corresponding node in the suffix tree of P (if any). This can be done
 838 in $\mathcal{O}(|S|)$ time, so $\mathcal{O}(N)$ overall. During this process we might designate the same
 839 string as an anchor multiple times, but we can easily remove the possible duplicates
 840 to obtain the set \mathcal{A} of anchors in the end. Then, we generate the occurrences of
 841 all anchors in P by accessing their corresponding nodes in the suffix tree of P and
 842 iterating over all leaves in their subtrees. We claim that the total number of all these
 843 occurrences is only $\mathcal{O}(m/\ell)$. This follows from the following characterization.

844 LEMMA 5.11. *If $P[x..(x+\ell)]$ is an occurrence of an anchor then either $P[(x+1)..y]$ is a maximal ℓ -periodic substring of P , for some $y \geq x+\ell$, or $P[x'..(x+\ell-1)]$ is a maximal ℓ -periodic substring of P , for some $x' \leq x$.*

847 *Proof.* By symmetry, it is enough to consider an anchor H created because of a
 848 maximal ℓ -periodic substring $S[i..j]$ such that $i > 1$, when we add $S[(i-1)..(i-1+\ell)]$
 849 to \mathcal{A} . Thus, $\text{per}(H[2..|H|]) \leq \ell/4$ and if $P[x..(x+\ell)] = H$ then $\text{per}(P[(x+1)..(x+\ell)]) \leq \ell/4$, making $P[(x+1)..(x+\ell)]$ a substring of some maximal ℓ -periodic substring of $P[(x'+1)..y]$, where $x' \leq x$ and $y \geq x+\ell$. If $x' < x$ then $\text{per}(H) \leq \ell/4$. But then $H = S[(i-1)..(i-1+\ell)]$ can be extended to some maximal ℓ -periodic substring $S[i'..j']$ such that $i' \leq i-1$ and $j' \geq i-1+\ell$. The overlap between $S[i..j]$ and $S[i'..j']$ is at least ℓ , so by Lemma 5.9 $i = i'$ and $j = j'$, which is a contradiction. Consequently, $x' = x$ and we obtain the lemma. \square

856 By Lemma 5.11, the number of occurrences of all anchors in P is at most two
 857 per each maximal ℓ -periodic substrings, so $\mathcal{O}(m/\ell)$ in total. We thus obtain a set of
 858 length- $(\ell+1)$ anchors with the following properties:

- 859 1. The total number of occurrences of all anchors in P is $\mathcal{O}(m/\ell)$.
- 860 2. For every $S \in \mathcal{S}$, at least one of its length- $(\ell+1)$ substrings is an anchor.
- 861 3. For every $S \in \mathcal{S}$, at most two of its length- $(\ell+1)$ substrings are anchors.

862 These properties are even stronger than what we had used in Section 5.1 (except that
 863 now we are working with length- $(\ell+1)$ substrings, which is irrelevant), we can now
 864 prove our final result also for strings of type 2.

865 THEOREM 5.12. *An instance of the AP problem where all strings are of type 2 can be solved in $\tilde{\mathcal{O}}(m^{\omega-1} + N)$ time.*

867 **5.3. Type 3 Strings.** In this section we show how to solve a restricted instance
 868 of the AP problem where every string $S \in \mathcal{S}$ is of type 3, and furthermore $|S| \in [9/8 \cdot \ell, 5/4 \cdot \ell)$ for some $\ell \leq m$. Recall that strings $S \in \mathcal{S}$ are such that every length- ℓ substring of S is strongly periodic and, by Lemma 5.3, in this case, $\text{per}(S) \leq \ell/4$. An occurrence of such S in P must be contained in a maximal ℓ -periodic substring of P . Recall that a string T is called ℓ -periodic if $|T| \geq \ell$ and $\text{per}(T) \leq \ell/4$. For an ℓ -periodic string T , let its *root*, denoted by $\text{root}(T)$, be the lexicographically smallest cyclic shift of $T[1.. \text{per}(T)]$. Because $\text{per}(T) \leq \ell/4$ and $|T| \geq \ell$ by definition, there are at least four repetitions of the period in T , so we can write $T = R[i..|R|]R^\alpha R[1..j]$, where $R = \text{root}(T)$, for some $i, j \in [1, |R|]$ and $\alpha \geq 2$. It is well known that $\text{root}(T)$

877 can be computed in $\mathcal{O}(|T|)$ time [31].

878 **EXAMPLE 4.** Let $T = \mathbf{babababab}$ and $\ell = 8$. We have $|T| = 9 \geq \ell = 8$ and
 879 $\text{per}(T) = 2 \leq \ell/4 = 2$, so T is ℓ -periodic. We have $\text{root}(T) = R = \mathbf{ab}$, and T can be
 880 written as $T = \mathbf{b} \cdot (\mathbf{ab})^3 \cdot \mathbf{ab}$, for $i = 2$ and $j = 2$.

881 We will now make a partition of type 3 strings based on their roots. We start
 882 with extracting all maximal ℓ -periodic substrings of P by proceeding similarly as in
 883 the proof of Lemma 5.10, and then compute the root of every such substring in $\mathcal{O}(m)$
 884 total time. In more detail, we partition P into blocks of length $\ell/2$, and compute the
 885 period of each such block. Any maximal ℓ -periodic substring of P needs to contain
 886 at least one such block inside. Therefore, for each block with period at most $\ell/2$ we
 887 can compute how far its period extends to the left and to the right, and output the
 888 corresponding substring if it is long enough. The only difficulty is that we should not
 889 extend the period beyond the preceding block. Two maximal ℓ -periodic substrings
 890 cannot overlap by more than $\ell/2$ letters, hence their total length is $\mathcal{O}(m)$ and we can
 891 compute the root of each such substring in $\mathcal{O}(m)$ total time. We also extract the root
 892 of every $S \in \mathcal{S}$ in $\mathcal{O}(N)$ total time. We then partition maximal ℓ -periodic substrings
 893 of P and strings $S \in \mathcal{S}$ into groups that have the same root. In the remaining part
 894 we describe how to process each such group corresponding to root R in which all
 895 maximal ℓ -periodic substrings of P have total length m' , and the strings $S \in \mathcal{S}$ have
 896 total length N' .

897 Recall that the bit vector U stores the active prefixes input to the AP problem,
 898 and the bit vector V encodes the new active prefixes we aim to compute. For every
 899 maximal ℓ -periodic substring of P with root R we extract the corresponding fragment
 900 of the bit vector U and need to update the corresponding fragment of the bit vector
 901 V . To make the description less cluttered, we assume that each such substring of P
 902 is a power of R , that is, R^α for some $\alpha \geq 4$. This can be assumed without loss of
 903 generality as it can be ensured by appropriately padding the extracted fragment of
 904 U and then truncating the results, while increasing the total length of all considered
 905 substrings of P by at most half of their length. In the description below, for simplicity
 906 of presentation, U and V denote these padded fragments of the original U and V .
 907 When computing V from U we use two different methods for processing the elements
 908 $S = R[i..|R|]R^\beta R[1..j]$ of \mathcal{S} depending on their length: either $\beta \geq t$ (large β) or
 909 $\beta < t$ (small β), for some parameter t to be chosen later. In both cases, we rely on
 910 the observation that $S = R[i..|R|]R^\beta R[1..j]$ occurs R^α at positions $i + \gamma \cdot |R|$, for
 911 $\gamma = 0, \dots, \alpha - \beta - 2$. This follows from R being the root and $\beta \geq 1$.

912 **Large β .** We proceed in phases corresponding to $\beta = t, \dots, \alpha$. In each single phase,
 913 we consider all strings $S \in \mathcal{S}$ with $S = R[i..|R|]R^\beta R[1..j]$, for some i and j . Let $C(\beta)$
 914 be the set of the corresponding pairs (i, j) , and observe that $\sum_\beta |C(\beta)| \cdot |R^\beta| \leq N'$.
 915 This is because the length of R^β is not greater than that of $S = R[i..|R|]R^\beta R[1..j]$,
 916 there are $|C(\beta)|$ distinct strings of the latter form in \mathcal{S} , and the total length of all $S \in \mathcal{S}$
 917 is N' . The total number of occurrences of a string $S = R[i..|R|]R^\beta R[1..j]$ in R^α is
 918 bounded by $\mathcal{O}(\alpha)$, and all such occurrences can be generated in time proportional to
 919 their number. Thus, for every $(i, j) \in C(\beta)$, we can generate all occurrences of the
 920 corresponding string and appropriately update V in $\mathcal{O}(\alpha \cdot |C(\beta)|)$ total time.

921 **Small β .** We start by giving a technical lemma on the complexity of multiplying two
 922 $r \times r$ matrices whose cells are polynomials of degree up to d .

923 **LEMMA 5.13.** *If two $r \times r$ matrices over \mathbb{Z} can be multiplied in $\mathcal{O}(r^\omega)$ time, then*
 924 *two $r \times r$ matrices over $\mathbb{Z}[x]$ with degrees up to d can be multiplied in $\mathcal{O}(r^\omega d + r^2 d \log d)$*

925 *time.*

926 *Proof.* Let A and B be two $r \times r$ matrices over $\mathbb{Z}[x]$ with degrees up to d . We
 927 reduce the product $A \times B = C$ to $(2d + 1)$ products of $r \times r$ matrices over \mathbb{Z} as
 928 follows. We evaluate the polynomials of each matrix in the complex $(2d + 1)$ -th roots
 929 of unity: let A_i and B_i be the matrices obtained by evaluating the polynomials of
 930 A and B in the i -th such root, respectively. We then perform the $2d + 1$ products
 931 $A_1 \times B_1, \dots, A_{2d+1} \times B_{2d+1}$ to obtain matrices C_1, \dots, C_{2d+1} : the $2d + 1$ values
 932 $C_1[i, j], \dots, C_{2d+1}[i, j]$ are finally interpolated to obtain the coefficient representation
 933 of $C[i, j]$, for each $i, j = 1, \dots, r$, in $\mathcal{O}(d \log d)$ time for each polynomial [27]. Since
 934 we perform $2d + 1$ products of matrices in $\mathbb{Z}^{r \times r}$, and we evaluate and interpolate
 935 r^2 polynomials of degree up to $2d + 1$, the overall time complexity is $2d\mathcal{O}(r^\omega) +$
 936 $r^2\mathcal{O}(d \log d) = \mathcal{O}(r^\omega d + r^2 d \log d)$. \square

937 Unlike in the large β case, we process $\beta = 2, \dots, t - 1$ simultaneously as follows
 938 when $t \geq 3$.

939 We construct three-dimensional Boolean tables: M with dimensions $|R| \times |R| \times t$
 940 and M' with dimensions $\lceil \alpha/t \rceil \times |R| \times t$. We set $M[i, j, \beta + 1] = 1$ if and only if $(i, j) \in$
 941 $C(\beta)$. M can be constructed in time proportional to its size by first precomputing
 942 a lexicographically sorted list of triples (β, i, j) corresponding to $S \in \mathcal{S}$ such that
 943 $S = R[i \dots |R|]R^\beta R[1 \dots j]$. The lists corresponding to different roots are constructed
 944 in $\mathcal{O}(N')$ total time, and we sort them together with radix sort to avoid paying $\mathcal{O}(m)$
 945 per each root. Then, we construct M by considering the prefix of the list consisting of
 946 all triples with sufficiently small first coordinates. Next, we set $M'[k, i, \gamma + 1] = 1$ if and
 947 only if $U[((k - 1)t + \gamma)|R| + i - 1] = 1$. Finally, we interpret M' and M as matrices over
 948 $\mathbb{Z}[x]$ with degrees up to $t - 1$, and compute their product $M'' = M' \times M$. That is, we
 949 think that $M'[k, i] = \sum_{\gamma=0}^{t-1} M'[k, i, \gamma + 1]x^\gamma$ and $M[i, j] = \sum_{\beta=0}^{t-1} M[i, j, \beta + 1]x^\beta$, and
 950 compute $M''[k, j] = \sum_{i=1}^{|R|} M'[k, i] \cdot M[i, j]$ for every $k = 1, \dots, \lceil \alpha/t \rceil$ and $j = 1, \dots, |R|$
 951 (this will be eventually implemented with Lemma 5.13). Note that each $M''[k, j]$ is
 952 a polynomial with degree up to $2(t - 1)$. We claim that this allows us to recover the
 953 updates to V by setting $V[((k - 1)t + q + 1)|R| + j] = 1$ whenever x^q appears with non-
 954 zero coefficient in the polynomial at $M''[k, j]$, for all $k = 1, \dots, \lceil \alpha/t \rceil$, $j = 1, \dots, |R|$
 955 and $q = 0, \dots, 2(t - 1)$. Equivalently, we set $V[((k - 1)t + \gamma + \beta + 1)|R| + j] = 1$ whenever
 956 $M'[k, i, \gamma + 1] = 1$ and $M[i, j, \beta + 1] = 1$, for all $k = 1, \dots, \lceil \alpha/t \rceil$, $i, j = 1, \dots, |R|$ and
 957 $\gamma, \beta = 0, \dots, t - 1$. This can be rewritten as setting $V[((k - 1)t + \gamma + \beta + 1)|R| + j] = 1$
 958 whenever $U[((k - 1)t + \gamma)|R| + i - 1] = 1$ and there exists $S \in \mathcal{S}$ such that $S =$
 959 $R[i \dots |R|]R^\beta R[1 \dots j]$, for all $k = 1, \dots, \lceil \alpha/t \rceil$, $j = 1, \dots, |R|$ and $\gamma, \beta = 0, \dots, t - 1$,
 960 which is indeed correct as any $x \in \{0, \dots, \alpha - 1\}$ can be written as $x = (k - 1)t + \gamma$
 961 for $k \in \{1, \dots, \lceil \alpha/t \rceil\}$ and $\gamma \in \{0, \dots, t - 1\}$.

962 We are now in a position to prove the following result for type 3 strings.

963 **THEOREM 5.14.** *An instance of the AP problem where all strings are of type 3*
 964 *can be solved in $\tilde{\mathcal{O}}(m^{\omega-1} + N)$ time.*

965 *Proof.* Recall that we consider strings S of type 3 with root R and substrings
 966 of P with root R together. We first analyze the time to process a single group
 967 containing a number of substrings of P of total length m' and a number of strings
 968 $S \in \mathcal{S}$ of total length N' . Let us denote by R^{α_h} the h -th considered substring of
 969 P and by t_h the value of t used to distinguish between small and large value of β
 970 when processing this substring. We partition all substrings into $\log m$ levels, with the
 971 k -th level G_k containing h such that $\alpha_h \in [2^k, 2^{k+1})$. We define $\bar{\alpha}_k = \sum_{h \in G_k} \alpha_k$ and

972 choose $t_h = \min(2^{k+1}, \lceil \bar{\alpha}_k/|R| \cdot \log m \rceil)$ for every $h \in G_k$.

973 For each level k , $h \in G_k$ and $\beta = t_h, \dots, \alpha_h$, we use the first method and spend
 974 $\mathcal{O}(\alpha_h \cdot |C(\beta)|)$ time, where $C(\beta)$ is the set of (i, j) for this specific β . This needs to
 975 be done only when $t_h \leq \alpha_h$, that is, $t_h = \lceil \bar{\alpha}_k/|R| \cdot \log m \rceil$. The overall time used for
 976 all applications of the first method is thus at most:

$$\begin{aligned}
 977 \quad \sum_k \sum_{h \in G_k} \mathcal{O}(\alpha_h \cdot \sum_{\beta \geq t_h} |C(\beta)|) &= \mathcal{O}(\sum_k \sum_{h \in G_k} \alpha_h / |R|^{t_h} \sum_{\beta \geq t_h} |C(\beta)| \cdot |R|^{t_h}) \\
 978 &= \mathcal{O}(\sum_k \sum_{h \in G_k} a_h / (|R| \cdot t_h) \sum_{\beta \geq t_h} |C(\beta)| \cdot |R|^\beta) \\
 979 &= \mathcal{O}(\sum_k \bar{a}_k / (|R| \cdot \bar{\alpha}_k / |R| \cdot \log m) \cdot N') = \mathcal{O}(N'),
 \end{aligned}$$

980 using the fact that $\sum_\beta |C(\beta)| \cdot |R|^\beta \leq N'$ and there are $\log m$ values of k .

981 For each level k and $h \in G_k$, we process together all $\beta = 2, \dots, t_h - 1$ using the
 982 second method. This requires multiplying two matrices of polynomials of degree up to
 983 $t_h - 1$. We observe that the second matrix is in fact the same for all $h \in G_k$, and so we
 984 denote the first matrix by M'_h , the second by simply M , and think that the degree of
 985 each polynomial in M'_h and M is strictly upper bounded by $d_k = \min(2^{k+1}, \lceil \bar{\alpha}_k/|R| \cdot$
 986 $\log m \rceil)$. M'_h is of size $\lceil \alpha_h/d_k \rceil \times |R|$ while M is of size $|R| \times |R|$. Instead of computing
 987 each product $M'_h \times M$ separately, we vertically concatenate all matrices M'_h to obtain
 988 a single matrix M' . The number of rows in M' is $r = \sum_{h \in G_k} \lceil \alpha_h/d_k \rceil$. Next, we
 989 compute $M' \times M$ with $\lceil r/|R| \rceil$ invocations of Lemma 5.13. We separately analyse the
 990 overall time complexity for $d_k = 2^{k+1}$ and $d_k = \lceil \bar{\alpha}_k/|R| \cdot \log m \rceil$.
 991 $d_k = \lceil \bar{\alpha}_k/|R| \cdot \log m \rceil$: Using $\alpha_h \geq 2^k \geq d_k/2$ we bound r as follows:

$$\begin{aligned}
 992 \quad r &= \sum_{h \in G_k} \lceil \alpha_h/d_k \rceil \leq \sum_{h \in G_k} (\alpha_h + d_k)/d_k \leq \sum_{h \in G_k} (\alpha_h + 2\alpha_h)/d_k \\
 993 &\leq 3 \sum_{h \in G_k} \alpha_h / (\bar{\alpha}_k / |R| \cdot \log m) = 3|R|/\log m \leq |R|,
 \end{aligned}$$

994 for sufficiently large m . Thus, one invocation suffices and takes time

$$995 \quad \mathcal{O}(|R|^\omega d_k + |R|^2 d_k \log d_k) = \mathcal{O}(|R|^{\omega-1} \bar{\alpha}_k \log^2 m)$$

996 using $d_k \geq 3$ and $d_k \leq 2m$.

997 $d_k = 2^{k+1}$: Because $\alpha_h \in [2^k, 2^{k+1})$ for each $h \in G_k$, we have $r = |G_k| \leq \bar{\alpha}_k/2^k$.

998 The number of invocations is thus at most $\lceil \bar{\alpha}_k/(2^k \cdot |R|) \rceil \leq \bar{\alpha}_k/(2^k \cdot |R|) + 1$.

999 The total time used by all these invocations is

$$\begin{aligned}
 1000 &(\bar{\alpha}_k/(2^k \cdot |R|) + 1) \mathcal{O}(|R|^\omega 2^{k+1} + |R|^2 2^{k+1} (k+1)) \\
 1001 &= \mathcal{O}(|R|^{\omega-1} \bar{\alpha}_k \log m + |R|^\omega 2^{k+1} \log m)
 \end{aligned}$$

1002 using $2^{k+1} \leq 2m$. Next, because $2^{k+1} \leq \lceil \bar{\alpha}_k/|R| \cdot \log m \rceil$ and $2^{k+1} \geq 2$ we
 1003 have $2^{k+1} \leq 2\bar{\alpha}_k/|R| \cdot \log m$, so the total time can be further bounded by

$$\begin{aligned}
 1004 &\mathcal{O}(|R|^{\omega-1} \bar{\alpha}_k \log m + |R|^\omega 2^{k+1} \log m) \\
 1005 &= \mathcal{O}(|R|^{\omega-1} \bar{\alpha}_k \log m + |R|^\omega (\bar{\alpha}_k/|R| \cdot \log m) \log m) \\
 1006 &= \mathcal{O}(|R|^{\omega-1} \bar{\alpha}_k \log^2 m).
 \end{aligned}$$

1007 Hence in both cases the time used by all multiplications is $\mathcal{O}(|R|^{\omega-1} \bar{\alpha}_k \log^2 m)$.
 1008 Using $\sum_k \bar{\alpha}_k = m'/|R|$ and $|R| \leq m'$, when summed over all $\log m$ levels k this
 1009 is in fact $\mathcal{O}((m')^{\omega-1} \log^2 m)$. We remark that the matrix M can be built in time
 1010 proportional to its size assuming $\mathcal{O}(N')$ preprocessing, while the matrix M' can be
 1011 built in time proportional to its size by just scanning over the corresponding fragment
 1012 of U .

1013 Finally, summing possibly many groups corresponding to different roots R , be-
 1014 cause all values of N' sum up to N and all values of m' sum up to $\mathcal{O}(m)$, by convexity
 1015 of $x^{\omega-1}$ we obtain that the overall time complexity including the preprocessing is
 1016 $\tilde{\mathcal{O}}(m^{\omega-1} + N)$. \square

1017 **5.4. Wrapping Up.** In Sections 5.1, 5.2 and 5.3 we design three $\tilde{\mathcal{O}}(m^{\omega-1} + N)$ -
 1018 time algorithms for an instance of the AP problem where all strings are of type 1,
 1019 2 and 3, respectively. Summing up over all values of k and all the types, we thus
 1020 obtain Theorem 1.2. In every case, the complexity is actually $\tilde{\mathcal{O}}(nm^{\omega-1}) + \mathcal{O}(N)$, so
 1021 using the fact that $\omega < 2.373$ [51, 66] we can hide the polylog factors and obtain the
 1022 following corollary.

1023 **COROLLARY 5.15.** *The EDSM problem can be solved on-line in $\mathcal{O}(nm^{1.373} + N)$*
 1024 *time.*

1025 **6. Final Remarks.** Our contribution in this paper is twofold. First, we de-
 1026 signed an appropriate reduction showing that a combinatorial algorithm solving the
 1027 EDSM problem in $\mathcal{O}(nm^{1.5-\epsilon} + N)$ time, for any $\epsilon > 0$, refutes the well-known BMM
 1028 conjecture. Second, we designed a non-combinatorial $\tilde{\mathcal{O}}(nm^{\omega-1} + N)$ -time algorithm
 1029 to attack the same problem. By using the fact that $\omega < 2.373$, our algorithm runs in
 1030 $\mathcal{O}(nm^{1.373} + N)$ time thus circumventing the combinatorial conditional lower bound
 1031 for the EDSM problem. Let us point out that if $\omega = 2$ then our algorithm for the
 1032 AP problem is time-optimal up to polylog factors, as any algorithm needs to read
 1033 the input. As for the EDSM problem, such an argument only shows a lower bound
 1034 of $\Omega(N)$. However, at the same time we can show that there is no $\mathcal{O}((nm)^{1-\epsilon})$ -time
 1035 algorithm, assuming the Strong Exponential Time Hypothesis (SETH) [19], by the
 1036 following argument. By prepending and appending a unique letter to both the ED
 1037 text and the pattern, we can reduce checking membership for a regular expression
 1038 of type $\cdot|\cdot$, as defined by Backurs and Indyk [10]. Combining this with their reduc-
 1039 tion from SETH, we immediately obtain the claimed conditional lower bound for the
 1040 EDSM problem.

1041 We finally remark that, if we use the simple cubic-time matrix multiplication
 1042 algorithm in our solution then the total time complexity becomes $\tilde{\mathcal{O}}(nm^{\omega-1} + N) =$
 1043 $\tilde{\mathcal{O}}(nm^2 + N)$. At the same time, the solution by Aoyama et al. [8], which also does
 1044 not use fast matrix multiplication, runs in time $\mathcal{O}(nm^{1.5} + N)$. It is thus plausible
 1045 that one could obtain an $\tilde{\mathcal{O}}(nm^{\omega/2} + N)$ -time algorithm for the EDSM problem. We
 1046 leave this question open for future work.

1047

REFERENCES

- 1048 [1] A. ABBOUD, A. BACKURS, AND V. WILLIAMS, *If the current clique algorithms are optimal, so is*
 1049 *Valiant's parser*, in 56th IEEE Symposium on Foundations Of Computer Science (FOCS),
 1050 2015, pp. 98–117.
 1051 [2] A. ABBOUD AND V. WILLIAMS, *Popular conjectures imply strong lower bounds for dynamic*
 1052 *problems*, in 55th IEEE Symposium on Foundations Of Computer Science (FOCS), 2014,
 1053 pp. 434–443.
 1054 [3] K. ABRAHAMSON, *Generalized string matching*, SIAM J. Comput., 16 (1987), pp. 1039–1051.

- 1055 [4] J. ALMAN AND V. V. WILLIAMS, *A refined laser method and faster matrix multiplication*, in
1056 Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), 2021,
1057 pp. 522–539.
- 1058 [5] M. ALZAMEL, L. AYAD, G. BERNARDINI, R. GROSSI, C. ILIOPOULOS, N. PISANTI, S. PISSIS,
1059 AND G. ROSONE, *Degenerate string comparison and applications*, in 18th Workshop on
1060 Algorithms in Bioinformatics (WABI), vol. 113 of LIPIcs, 2018, pp. 21:1–21:14.
- 1061 [6] M. ALZAMEL, L. A. K. AYAD, G. BERNARDINI, R. GROSSI, C. S. ILIOPOULOS, N. PISANTI, S. P.
1062 PISSIS, AND G. ROSONE, *Comparing degenerate strings*, *Fundam. Informaticae*, 175 (2020),
1063 pp. 41–58.
- 1064 [7] A. AMIR, M. LEWENSTEIN, AND E. PORAT, *Faster algorithms for string matching with k mis-*
1065 *matches*, *J. Algorithms*, 50 (2004), pp. 257–275.
- 1066 [8] K. AOYAMA, Y. NAKASHIMA, T. I. S. INENAGA, H. BANNAL, AND M. TAKEDA, *Faster online*
1067 *elastic degenerate string matching*, in 29th Symposium on Combinatorial Pattern Matching
1068 (CPM), vol. 105 of LIPIcs, 2018, pp. 9:1–9:10.
- 1069 [9] V. ARLAZAROV, E. DINIC, M. KRONROD, AND I. FARADŽEV, *On economical construction of the*
1070 *transitive closure of a directed graph*, *Soviet Mathematics Doklady*, 11 (1970), pp. 1209–
1071 1210.
- 1072 [10] A. BACKURS AND P. INDYK, *Which regular expression patterns are hard to match?*, in 57th
1073 IEEE Symposium on Foundations Of Computer Science (FOCS), 2016, pp. 457–466.
- 1074 [11] N. BANSAL AND R. WILLIAMS, *Regularity lemmas and combinatorial algorithms*, in 50th IEEE
1075 Symposium on Foundations Of Computer Science (FOCS), 2009, pp. 745–754.
- 1076 [12] M. BENDER AND M. FARACH-COLTON, *The LCA problem revisited*, in 4th Latin American
1077 symposium on Theoretical Informatics (LATIN), vol. 1776 of Springer LNCS, 2000, pp. 88–
1078 94.
- 1079 [13] J. BENTLEY, *Multidimensional binary search trees used for associative searching*, *Commun.*
1080 *ACM*, 18 (1975), pp. 509–517.
- 1081 [14] G. BERNARDINI, P. GAWRYCHOWSKI, N. PISANTI, S. P. PISSIS, AND G. ROSONE, *Even Faster*
1082 *Elastic-Degenerate String Matching via Fast Matrix Multiplication*, in 46th International
1083 Colloquium on Automata, Languages, and Programming (ICALP), vol. 132 of LIPIcs,
1084 2019, pp. 21:1–21:15.
- 1085 [15] G. BERNARDINI, N. PISANTI, S. PISSIS, AND G. ROSONE, *Pattern matching on elastic-degenerate*
1086 *text with errors*, in 24th International Symposium on String Processing and Information
1087 Retrieval (SPIRE), 2017, pp. 74–90.
- 1088 [16] G. BERNARDINI, N. PISANTI, S. P. PISSIS, AND G. ROSONE, *Approximate pattern matching on*
1089 *elastic-degenerate text*, *Theor. Comput. Sci.*, 812 (2020), pp. 109–122.
- 1090 [17] K. BRINGMANN, F. GRANDONI, B. SAHA, AND V. WILLIAMS, *Truly sub-cubic algorithms for*
1091 *language edit distance and RNA-folding via fast bounded-difference min-plus product*, in
1092 56th IEEE Symposium on Foundations Of Computer Science (FOCS), 2016, pp. 375–384.
- 1093 [18] K. BRINGMANN, A. GRØNLUND, AND K. LARSEN, *A dichotomy for regular expression mem-*
1094 *bership testing*, in 58th IEEE Symposium on Foundations Of Computer Science (FOCS),
1095 2017, pp. 307–318.
- 1096 [19] C. CALABRO, R. IMPAGLIAZZO, AND R. PATURI, *The complexity of satisfiability of small depth*
1097 *circuits*, in Parameterized and Exact Computation, 4th International Workshop (IWPEC),
1098 vol. 5917 of Lecture Notes in Computer Science, 2009, pp. 75–85.
- 1099 [20] T. CHAN, *Speeding up the four Russians algorithm by about one more logarithmic factor*, in
1100 26th ACM-SIAM Symposium On Discrete Algorithms (SODA), 2015, pp. 212–217.
- 1101 [21] Y.-J. CHANG, *Hardness of RNA folding problem with four symbols*, in 27th Symposium on
1102 Combinatorial Pattern Matching (CPM), vol. 54 of LIPIcs, 2016, pp. 13:1–13:12.
- 1103 [22] K. CHATTERJEE, B. CHOUDHARY, AND A. PAVLOGIANNIS, *Optimal Dyck reachability for data-*
1104 *dependence and alias analysis*, in 45th ACM SIGPLAN Symposium on Principles of Pro-
1105 gramming Languages (POPL), 2018, pp. 30:1–30:30.
- 1106 [23] A. CISLAK AND S. GRABOWSKI, *Sopang 2: online searching over a pan-genome without false*
1107 *positives*, *CoRR*, abs/2004.03033 (2020).
- 1108 [24] A. CISLAK, S. GRABOWSKI, AND J. HOLUB, *SOPanG: online text searching over a pan-genome*,
1109 *Bioinformatics*, 34 (2018), pp. 4290–4292.
- 1110 [25] P. CLIFFORD AND R. CLIFFORD, *Simple deterministic wildcard matching*, *Inf. Process. Lett.*,
1111 101 (2007), pp. 53–54.
- 1112 [26] R. COLE AND R. HARIHARAN, *Verifying candidate matches in sparse and wildcard matching*,
1113 in 34th ACM Symposium on Theory Of Computing (STOC), 2002, pp. 592–601.
- 1114 [27] J. COOLEY AND J. TUKEY, *An algorithm for the machine calculation of complex Fourier series*,
1115 *Mathematics of Computation*, 19 (1965), pp. 297–301.
- 1116 [28] M. CROCHEMORE, C. HANCART, AND T. LECROQ, *Algorithms on strings*, Cambridge University

- 1117 Press, 2007.
- 1118 [29] M. CROCHEMORE AND D. PERRIN, *Two-way string matching*, J. ACM, 38 (1991), pp. 651–675.
- 1119 [30] A. CZUMAJ AND A. LINGAS, *Finding a heaviest vertex-weighted triangle is not harder than*
- 1120 *matrix multiplication*, SIAM J. Comput., 39 (2009), pp. 431–444.
- 1121 [31] J. DUVAL, *Factorizing words over an ordered alphabet*, J. Algorithms, 4 (1983), pp. 363–381.
- 1122 [32] M. FARACH-COLTON AND S. MUTHUKRISHNAN, *Perfect hashing for strings: formalization*
- 1123 *and algorithms*, in 7th Annual Symposium on Combinatorial Pattern Matching (CPM),
- 1124 vol. 1075 of Springer LNCS, 1996, pp. 130–140.
- 1125 [33] N. FINE AND H. WILF, *Uniqueness theorems for periodic functions*, Proceedings of the American
- 1126 Mathematical Society, 16 (1965), pp. 109–114.
- 1127 [34] M. FISCHER AND A. MEYER, *Boolean matrix multiplication and transitive closure*, in 12th IEEE
- 1128 Symposium on Switching and Automata Theory (SWAT/FOCS), 1971, pp. 129–131.
- 1129 [35] M. FISCHER AND M. PATERSON, *String matching and other products*, SIAM-AMS Proceedings,
- 1130 7 (1974), pp. 113–125. Complexity of Computation.
- 1131 [36] M. FURMAN, *Application of a method of fast multiplication of matrices in the problem of finding*
- 1132 *the transitive closure of a graph*, Soviet Mathematics Doklady, 11 (1970), p. 1252.
- 1133 [37] P. GAWRYCHOWSKI, S. GHAZAWI, AND G. M. LANDAU, *On indeterminate strings matching*,
- 1134 in 31st Symposium on Combinatorial Pattern Matching (CPM), vol. 161 of LIPIcs, 2020,
- 1135 pp. 23:1–23:12.
- 1136 [38] P. GAWRYCHOWSKI AND P. UZNANSKI, *Towards unified approximate pattern matching for Ham-*
- 1137 *ming and L_1 distance*, in 45th International Colloquium on Automata, Languages and
- 1138 Programming (ICALP), vol. 107 of LIPIcs, 2018, pp. 62:1–62:13.
- 1139 [39] R. GROSSI, C. ILIOPOULOS, C. LIU, N. PISANTI, S. PISSIS, A. RETHA, G. ROSONE, F. VAYANI,
- 1140 AND L. VERSARI, *On-line pattern matching on similar texts*, in 28th Symposium on Com-
- 1141 binatorial Pattern Matching (CPM), vol. 78 of LIPIcs, 2017, pp. 9:1–9:14.
- 1142 [40] M. HENZINGER, S. KRINNINGER, D. NANONGKAI, AND T. SARANURAK, *Unifying and strengthen-*
- 1143 *ing hardness for dynamic problems via the online matrix-vector multiplication conjecture*,
- 1144 in 47th ACM Symposium on Theory Of Computing (STOC), 2015, pp. 21–30.
- 1145 [41] J. HOLUB, W. SMYTH, AND S. WANG, *Fast pattern-matching on indeterminate strings*, J.
- 1146 Discrete Algorithms, 6 (2008), pp. 37–50.
- 1147 [42] C. ILIOPOULOS, R. KUNDU, AND S. PISSIS, *Efficient pattern matching in elastic-degenerate texts*,
- 1148 in 11th International Conference on Language and Automata Theory and Applications
- 1149 (LATA), vol. 10168 of Springer LNCS, 2017, pp. 131–142.
- 1150 [43] P. INDYK, *Faster algorithms for string matching problems: Matching the convolution bound*,
- 1151 in 39th Symposium on Foundations Of Computer Science (FOCS), 1998, pp. 166–173.
- 1152 [44] A. ITAI AND M. RODEH, *Finding a minimum circuit in a graph*, in 9th ACM Symposium on
- 1153 Theory Of Computing (STOC), 1977, pp. 1–10.
- 1154 [45] IUPAC-IUB COMMISSION ON BIOCHEMICAL NOMENCLATURE, *Abbreviations and symbols for*
- 1155 *nucleic acids, polynucleotides, and their constituents*, Biochemistry, 9 (1970), pp. 4022–
- 1156 4027.
- 1157 [46] A. KALAI, *Efficient pattern-matching with don't cares*, in 13th ACM-SIAM Symposium on
- 1158 Discrete Algorithms (SODA), 2002, pp. 655–656.
- 1159 [47] D. KNUTH, J. M. JR., AND V. PRATT, *Fast pattern matching in strings*, SIAM J. Comput., 6
- 1160 (1977), pp. 323–350.
- 1161 [48] T. KOCIUMAKA, J. RADOSZEWSKI, W. RYTTER, AND T. WALE, *Internal pattern matching*
- 1162 *queries in a text and applications*, in 26th ACM-SIAM Symposium on Discrete Algorithms
- 1163 (SODA), 2015, pp. 532–551.
- 1164 [49] T. KOPELOWITZ AND R. KRAUTHGAMER, *Color-distance oracles and snippets*, in 27th Sym-
- 1165 posium on Combinatorial Pattern Matching (CPM), vol. 54 of LIPIcs, 2016, pp. 24:1–24:10.
- 1166 [50] K. LARSEN, I. MUNRO, J. NIELSEN, AND S. THANKACHAN, *On hardness of several string indexing*
- 1167 *problems*, Theor. Comput. Sci., 582 (2015), pp. 74–82.
- 1168 [51] F. LE GALL, *Powers of tensors and fast matrix multiplication*, in 39th International Symposium
- 1169 on Symbolic and Algebraic Computation (ISSAC), 2014, pp. 296–303.
- 1170 [52] L. LEE, *Fast context-free grammar parsing requires fast boolean matrix multiplication*, J. ACM,
- 1171 49 (2002), pp. 1–15.
- 1172 [53] V. MÄKINEN, B. CAZAUX, M. EQUI, T. NORRI, AND A. I. TOMESCU, *Linear time construction of*
- 1173 *indecomposable founder block graphs*, in 20th Workshop on Algorithms in Bioinformatics (WABI),
- 1174 vol. 172 of LIPIcs, 2020, pp. 7:1–7:18.
- 1175 [54] J. MATOUŠEK, *Computing dominances in E^n* , Inf. Process. Lett., 38 (1991), pp. 277–278.
- 1176 [55] I. MUNRO, *Efficient determination of the transitive closure of a directed graph*, Inf. Process.
- 1177 Lett., 1 (1971), pp. 56–58.
- 1178 [56] G. NAVARRO, *Nr-grep: a fast and flexible pattern-matching tool*, Softw., Pract. Exper., 31

- 1179 (2001), pp. 1265–1312.
- 1180 [57] P. PETERLONGO, N. PISANTI, F. BOYER, A. P. DO LAGO, AND M. SAGOT, *Lossless filter for*
 1181 *multiple repetitions with hamming distance*, J. Discrete Algorithms, 6 (2008), pp. 497–509.
- 1182 [58] N. PISANTI, *On-line pattern matching on D-texts (invited talk)*, in Proceedings of the 32nd
 1183 Annual Symposium on Combinatorial Pattern Matching (CPM), vol. 191 of LIPIcs, 2021,
 1184 pp. 3:1–3:2.
- 1185 [59] S. PISSIS AND A. RETHA, *Dictionary matching in elastic-degenerate texts with applications in*
 1186 *searching VCF files on-line*, in 17th International Symposium on Experimental Algorithms
 1187 (SEA), vol. 103 of LIPIcs, 2018, pp. 16:1–16:14.
- 1188 [60] L. RODITTY AND U. ZWICK, *On dynamic shortest paths problems*, in 12th European Symposium
 1189 on Algorithms (ESA), vol. 3221 of Springer LNCS, 2004, pp. 580–591.
- 1190 [61] M. RUŽIĆ, *Constructing efficient dictionaries in close to sorting time*, in 35th International
 1191 Colloquium on Automata, Languages and Programming (ICALP), vol. 5125 of Springer
 1192 LNCS, 2008, pp. 84–95.
- 1193 [62] D. SLEATOR AND R. TARJAN, *A data structure for dynamic trees*, J. Comput. Syst. Sci., 26
 1194 (1983), pp. 362–391.
- 1195 [63] THE COMPUTATIONAL PAN-GENOMICS CONSORTIUM, *Computational pan-genomics: status,*
 1196 *promises and challenges*, Briefings in Bioinformatics, 19 (2018), pp. 118–135.
- 1197 [64] L. VALIANT, *General context-free recognition in less than cubic time*, J. Comput. Syst. Sci., 10
 1198 (1975), pp. 308–315.
- 1199 [65] P. WEINER, *Linear pattern matching algorithms*, in 14th IEEE Annual Symposium on Switch-
 1200 ing and Automata Theory (SWAT/FOCS), 1973, pp. 1–11.
- 1201 [66] V. WILLIAMS, *Multiplying matrices faster than Coppersmith-Winograd*, in 44th ACM Sympo-
 1202 sium on Theory Of Computing Conference (STOC), 2012, pp. 887–898.
- 1203 [67] V. WILLIAMS AND R. WILLIAMS, *Finding a maximum weight triangle in $n^{3-\delta}$ time, with ap-*
 1204 *plications*, in 38th ACM Symposium on Theory Of Computing Conference (STOC), 2006,
 1205 pp. 225–231.
- 1206 [68] V. WILLIAMS AND R. WILLIAMS, *Subcubic equivalences between path, matrix and triangle*
 1207 *problems*, in 51st IEEE Symposium on Foundations Of Computer Science (FOCS), 2010,
 1208 pp. 645–654.
- 1209 [69] S. WU AND U. MANBER, *Agrep – a fast approximate pattern-matching tool*, in USENIX Tech-
 1210 nical Conference, 1992, pp. 153–162.
- 1211 [70] H. YU, *An improved combinatorial algorithm for boolean matrix multiplication*, in 42nd Inter-
 1212 national Colloquium on Automata, Languages, and Programming (ICALP), vol. 9134 of
 1213 Springer LNCS, 2015, pp. 1094–1105.
- 1214 [71] H. YU, *An improved combinatorial algorithm for boolean matrix multiplication*, Inf. Comput.,
 1215 261 (2018), pp. 240–247.
- 1216 [72] U. ZWICK, *All pairs shortest paths using bridging sets and rectangular matrix multiplication*,
 1217 J. ACM, 49 (2002), pp. 289–317.