

# Optimal pattern matching in LZW compressed strings <sup>\*</sup>

Paweł Gawrychowski <sup>†</sup>

## Abstract

We consider the following variant of the classical pattern matching problem: given an uncompressed pattern  $s[1..m]$  and a compressed representation of a string  $t[1..N]$ , does  $s$  occur in  $t$ ? When  $t$  is compressed using the LZW method, we are able to detect the occurrence in optimal linear time, thus answering a question of Amir, Benson, and Farach [2]. Previous results implied solutions with complexities  $\mathcal{O}(n \log m + m)$  [2],  $\mathcal{O}(n + m^{1+\epsilon})$  [14], or (randomized)  $\mathcal{O}(n \log \frac{N}{n} + m)$  [8], where  $n$  is the size of the compressed representation of  $t$ . Our algorithm is simple and fully deterministic.

**Key-words:** pattern matching, compression, Lempel-Ziv-Welch

## 1 Introduction

Even though the processing capabilities of a modern hardware are growing very rapidly, so is the amount of the data gathered and stored by both individual users and all types of organizations. While inventing efficient compressing methods to store this data has been a very active research area since decades, it seems that with more and more aspects of our lives being processed digitally, the necessity of developing very efficient compression methods is becoming even more important, where efficient should be understood as both fast and achieving a good compression ratio. But even the best compression scheme will not get us much gain if performing any operation on the stored information requires uncompressing it anyway. Thus we should aim to develop not only good compression methods but also efficient algorithms working on the compressed representation alone, preferably with the running times depending only (or mainly) on the size of this representation, not the original input.

The most ubiquitous problem concerning processing information is the *pattern matching* for strings, in which we are given a pattern  $p[1..m]$  and a text  $t[1..N]$ , and the goal is to find if there is an occurrence of  $p$  in  $t$ . There are many efficient solutions to this

problem, starting with the first linear time solution of [15], then optimized for the delay between reading two consecutive letter of the text in [10, 13], the amount of additional memory used [11], and the total number of comparisons [5], to name just a few improvements, with the two most well-known algorithms being Knuth-Morris-Pratt [13] and Boyer-Moore [4].

If the string is given in a compressed representation, we get the *compressed pattern matching problem*. Given a text  $t$  and a pattern  $p$  we need to check if there is an occurrence of  $p$  in  $t$  without decompressing  $t$ , i.e., the running time should depend on the size  $n$  of the compressed representation of  $T$  alone rather than on the original length  $N$ . According to Amir and Benson [1], a solution to such problem is *efficient* if its time complexity is  $\mathcal{O}(N)$ , *almost optimal* if its time complexity is  $\mathcal{O}(n \log m + m)$ , and *optimal* if the complexity is  $\mathcal{O}(n + m)$ .

Complexity of the compressed pattern matching problem clearly depends on the particular compression method chosen. If we use any nonadaptive compression, the problem becomes rather simple. On the other hand, adaptive compression methods based on the Lempel-Ziv method [20] seems to be quite challenging to deal with efficiently, mostly because of the fact that the same fragment of a text might be encoded differently depending on its exact location. There has been a substantial amount of work devoted to developing fast algorithms for pattern matching in both Lempel-Ziv and its simplified version Lempel-Ziv-Welch [19] compressed texts. In [2] the authors introduced two algorithms with time complexities  $\mathcal{O}(n + m^2)$  and  $\mathcal{O}(n \log m + m)$  for Lempel-Ziv-Welch compressed texts. The pattern preprocessing time has been improved in [14] to get  $\mathcal{O}(n + m^{1+\epsilon})$  time complexity. Then [8] developed algorithm working for any Lempel-Ziv compressed text in (randomized) time  $\mathcal{O}(n \log^2 \frac{N}{n} + m)$ , with a version simplified for Lempel-Ziv-Welch working in (also randomized) time  $\mathcal{O}(n \log \frac{N}{n} + m)$ . [16] developed a practical  $\mathcal{O}(n \frac{m}{w})$  algorithm exploiting bit-parallelism. There has been also a substantial amount of research devoted to a more general problem of fully compressed pattern matching, where both the text and the pattern are compressed. For the case of LZW compression, Gasieniec and Rytter [12] developed an  $\mathcal{O}((n + m) \log(n + m))$  algorithm,

<sup>\*</sup>Supported by MNiSW grant number N N206 492638, 2010–2012

<sup>†</sup>Institute of Computer Science, University of Wrocław, 50–383 Wrocław, Poland. [gawry@cs.uni.wroc.pl](mailto:gawry@cs.uni.wroc.pl)

where  $n$  and  $m$  are the compressed sizes of the text and the pattern, respectively. Here we are concerned only with the case of uncompressed patterns, though.

In this paper we show that the compressed pattern matching problem can be solved in optimal linear time for Lempel-Ziv-Welch compressed texts. Variants of this compression method are used, for example, in Unix `compress` utility, in GIF images, and (optionally) in TIFF and PDF files so this result is not only of purely theoretical interest. Our algorithm is simple and fully deterministic.

## 2 Preliminaries

Let  $\Sigma$  be a fixed finite alphabet (later we will also deal with the more general case of polynomial size integer alphabet). We consider strings over  $\Sigma$  given in a Lempel-Ziv-Welch compressed form, which is a simplified version of Lempel-Ziv compression defined as follows: given an uncompressed string  $t[1..N]$  we iteratively construct its representation by looking up and removing the longest prefix of the current suffix  $t[k..N]$  which already occurs in  $t[1..k-1]$ . The output is a sequence of triples  $(start_i, len_i, next_i)$  with  $t[start_i..start_i + len_i - 1]$  being this longest prefix and  $next_i$  being the proceeding character. In Lempel-Ziv-Welch compression the string is represented as a sequence of *codewords* where a codeword is either a single letter, or a previously occurring codeword concatenated with a single character. This additional character is not given explicitly: we define it as the first character of the next codeword, and initialize the set of codewords to contain all single characters in the very beginning. The resulting compression method enjoys a particularly simple encoding/decoding process, but unfortunately requires outputting at least  $\Omega(\sqrt{N})$  codewords. Still, its simplicity and good compression ratio achieved on real life instances make it an interesting model to work with. For the rest of the paper we will use LZW when referring to Lempel-Ziv-Welch compression.

We are interested in a variation of the classical pattern matching problem: given a pattern  $s[1..m]$  and a text  $t[1..N]$ , does  $s$  occur in  $t$ ? In our case  $t$  is given in a compressed form. We wish to achieve a running time depending on the size  $n$  of this compressed representation, not the length of  $t$  itself. Additionally, we would like to keep the memory complexity as low as  $\mathcal{O}(m)$ , preferably with the algorithm reading the compressed representation of  $t$  just once from left to right. If the pattern does occur in the text, we would like to get the position of its first occurrence, and if there are *occ* occurrences, we would like to output all of them while adding just  $\mathcal{O}(occ)$  to the overall time complexity.

The very first thing we do is preprocessing  $s$ . Using

standard tools (suffix trees [7, 18] built for both the pattern and the reversed pattern, and lowest common ancestor queries [3]) we get the following:

**LEMMA 2.1.** *Pattern  $s$  can be preprocessed in linear time so that given any two fragments  $s[i..i+k]$  and  $s[j..j+k]$  we can find their longest common prefix (suffix) in constant time.*

**LEMMA 2.2.** *Pattern  $s$  can be preprocessed in linear time so that given any fragment  $s[i..j]$  we can find its longest prefix which is a suffix of the whole pattern in constant time, assuming we know the (explicit or implicit) vertex corresponding to  $s[i..j]$  in the suffix tree.*

*Proof.* We assume that the suffix tree is built for  $s$  concatenated with a special terminating character, say  $\$$ . Each leaf in the suffix tree corresponds to some suffix of  $s$ , and is connected to its parent with an edge labeled with a single letter. If we mark all those parents, finding the longest prefix which is a suffix of the whole  $s$  reduces to finding the lowest marked vertex on a given root-to-vertex path, which can be precomputed for all vertices in linear time.  $\square$

At a very high level our procedures resemble the old and well-known Knuth-Morris-Pratt pattern matching algorithm: we compute the longest prefix of  $s$  which is a suffix of the current prefix of the text. Recall that if we already have such longest prefix of a given  $t[1..i]$  and need to compute the longest match for  $t[1..i+1]$ , there are two possibilities. Either we extend the current longest prefix by  $t[i+1]$ , or we can safely replace it with the longest proper *border*, where border of a string is its fragment which is a suffix and a prefix at the same time (usually we identify such fragment with its length and say that  $border(t) = \{i_1, \dots, i_k\}$  is the set of all borders of  $t$ ). A closely related concept is the notion of *periods*:  $p$  is a period of  $w$  if  $w[i] = w[i+p]$  whenever both  $w[i]$  and  $w[i+p]$  are defined. It is easy to see that  $p$  is a period of  $w$  if and only if  $|w| - p$  is a border, and the following property holds:

**LEMMA 2.3. (PERIODICITY LEMMA)** *If both  $p$  and  $q$  are periods  $w$ , and  $p+q \leq |w| + \gcd(p, q)$ , then  $\gcd(p, q)$  is a period as well.*

This lemma helps us to avoid performing the computation for every possible prefix of the text. We will try to restrict attention to just  $\mathcal{O}(n)$  of them without losing any potential occurrence, and to this aim we will need the following property:




Figure 1: Detecting an occurrence in a concatenation of two snippets.

LEMMA 2.4. *If the longest border of  $t$  is of length  $b \geq \frac{|t|}{2}$  then all borders of length at least  $\frac{|t|}{2}$  create one arithmetic progression. More specifically,  $\text{border}(t) \cap \left\{ \frac{|t|}{2}, \dots, |t| \right\} = \left\{ |t| - \alpha p : 0 \leq \alpha \leq \frac{|t|}{2p} \right\}$ , where  $p = |t| - b$  is the period of  $t$ . We call this set the long borders of  $t$ .*

*Proof.* First note that  $p = |t| - b$  is the period of  $|t|$  so any  $\alpha p$  is also a period, as long as  $\alpha \leq \frac{|t|}{p}$ . Thus all elements of the arithmetic progression in the lemma are borders. We must show that there are no other borders of length at least  $\frac{|t|}{2}$ . Let  $b'$  be any such border. Observe that  $p' = |t| - b'$  is a period and  $p + p' \leq |t|$  so from the periodicity lemma  $\gcd(p, p')$  is a period as well. But  $p$  is the shortest period so  $p$  divides  $p'$  and  $b'$  belongs to the progression.  $\square$

Thanks to the above lemma instead of processing borders one-by-one we can split them into  $\log |t|$  groups with all borders in a single group creating one arithmetic progression, and hope to process them all at once. We will need the borders of both prefixes and suffixes of  $s$ .

LEMMA 2.5. *Pattern  $s$  can be preprocessed in linear time so that we can find the border of each its prefix (suffix) in constant time.*

*Proof.* Simply do the standard preprocessing from the Knuth-Morris-Pratt algorithm for both  $s$  and the reversed  $s$ . This preprocessing results in computing the border of each possible prefix.  $\square$

### 3 Pattern matching in a sequence of snippets

To detect an occurrence of the pattern in a LZW compressed string we will construct a collection of strings constructed by concatenating a number of substrings of the pattern one after another. In this section we show that after such reduction we can solve the problem in optimal linear time.

DEFINITION 3.1. *A snippet is any substring of the pattern  $s[i..j]$ . If  $i = 1$  we call it a prefix snippet, if  $j = m$  a suffix snippet. A sequence of snippets of size  $k$  is a concatenation of  $k$  substrings of the pattern  $s[i_1..j_1] \dots s[i_k..j_k]$ .*

We assume that for each snippet we store not only the indices  $i, j$  but also a link to the corresponding node (either explicit or implicit) in the suffix tree and the longest suffix which is a prefix of the whole pattern (in case of a prefix snippet this is of course the whole fragment). Given a sequence of such snippets we would like to simulate the Knuth-Morris-Pratt algorithm efficiently. To this aim we need the following two lemmas.

LEMMA 3.1. *Given any two snippets we can detect an occurrence of the pattern in their concatenation in constant time.*

*Proof.* Let the two snippets be  $s[i..j]$  and  $s[k..l]$ . If one of them is the whole pattern, we are done. Otherwise we are looking for  $x$  and  $y$  such that  $i \leq x \leq j$ ,  $k \leq y \leq l$  and  $s = s[x..j]s[k..y]$ . We can replace  $s[i..j]$  with its longest suffix which is a prefix of the whole  $s$ , and  $s[k..l]$  with its longest prefix which is a suffix of the whole  $s$ , without losing any possible occurrence. This takes just constant time thanks to Lemma 2.2. Then we are left with the following question: does  $s$  occur in  $s[1..i]s[j..m]$ ? Or, in other words, is there  $x \in \text{border}(s[1..i])$  and  $y \in \text{border}(s[j..m])$  such that  $x + y = m$ ? Note that either  $x \geq \frac{|s[1..i]|}{2}$  or  $y \geq \frac{|s[j..m]|}{2}$ , and without losing the generality assume the former. From Lemma 2.4 we know that all such possible values of  $x$  create one arithmetic progression. More specifically,  $x = i - \alpha p$ , where  $p \leq \frac{i}{2}$  is the period of  $s[1..i]$ . We need to check if there is an occurrence of  $s$  in  $s[1..i]s[j..m]$  starting after the  $\alpha p$ -th character, for some  $0 \leq \alpha \leq \frac{i}{p}$ . For any such possible interesting shift, there will be no mismatch in  $s[1..i]$ . There might be a mismatch in  $s[j..m]$ , though.

Let  $k \geq i$  be the longest prefix of  $s$  for which  $p$  is a period (such  $k$  can be calculated efficiently by looking up the longest common prefix of  $s[p+1..m]$  and the whole  $s$ ). We shift  $s[1..k]$  by  $\left\lfloor \frac{\min(i, i-j+1)}{p} \right\rfloor p$  characters (the minimum is chosen so that extending  $s[1..k]$  to the whole  $s$  does not result in sticking out of the right end of  $s[j..m]$ ) and compute its leftmost mismatch with  $s[j..m]$ , see Figure 1. Note that such shift is the largest possible among all shifts of the form

---

**Algorithm 1** NAIVE-PATTERN-MATCHING( $s_1, s_2, \dots, s_n$ )

---

```
1:  $q \leftarrow 0$ 
2:  $\ell \leftarrow$  longest prefix of  $s$  ending  $s_1$  ▷ Lemma 2.2
3:  $k \leftarrow 2$ 
4: while  $k \leq n$  and  $\ell + \sum_{i=k}^n |s_i| \geq m$  do
5:   check for an occurrence of  $s$  in  $s[1.. \ell]s_k$  ▷ Lemma 3.1
6:   if  $s[1.. \ell]s_k$  is a prefix of  $s$  then
7:      $\ell \leftarrow \ell + |s_k|$ 
8:      $k \leftarrow k + 1$ 
9:   continue
10:  end if
11:   $p \leftarrow$  period( $s[1.. \ell]$ ) ▷ Lemma 2.5
12:   $b \leftarrow$  longest long border of  $s[1.. \ell]$  such that  $s[1.. b]s_k$  is a prefix of  $s$  ▷ Lemma 3.2
13:  if  $b$  is undefined then
14:     $\ell \leftarrow$  longest prefix of  $s$  ending  $s[\lceil \frac{\ell}{2} \rceil .. \ell]$  ▷ Lemma 3.3
15:    continue
16:  end if
17:   $\ell \leftarrow b + |s_k|$ 
18:   $k \leftarrow k + 1$ 
19: end while
```

---

$\alpha p$  which start in  $s[1..i]$  and end in  $s[j..m]$ . Position of the first mismatch, or its nonexistence, allows us to eliminate all but one interesting shift. More precisely, we have two cases to consider.

1. There is no mismatch. If  $k = m$  we are done, otherwise  $s[k+1] \neq s[k+1-p]$ , meaning that choosing any smaller interesting shift results in a mismatch.
2. There is a mismatch. Let the conflicting characters be  $a$  and  $b$  and call the position at which  $a$  occurs the *obstacle*. Observe that we must choose a shift  $\alpha p$  so that  $s[1..k]$  shifted by  $\alpha p$  is completely on the left of the obstacle. On the other hand, if  $s[1..k]$  shifted by  $(\alpha+1)p$  is completely on the left as well, shifting  $s[1..k]$  by  $\alpha p$  results in a mismatch because  $s[k+1] \neq s[k+1-p]$  and  $s[k+1-p]$  matches with the corresponding character in  $s[j..m]$ . Thus we may restrict our attention to the largest shift for which  $s[1..k]$  is on the left of the obstacle.

Having identified the only interesting shift, we verify if there is a match using one longest common prefix query on  $s$ . Overall, the whole procedure takes constant time.  $\square$

**LEMMA 3.2.** *Given a prefix snippet  $s[1..i]$  and a snippet  $s[j..k]$  we can find the longest long border  $b$  of  $s[1..i]$  such that  $s[1..b]s[j..k]$  is a prefix of the whole  $s$  in constant time.*

*Proof.* As in the previous proof, we begin with computing  $k \geq i$ , the longest prefix of  $s$  for which  $p$  is a period, where  $p$  is the period of  $s[1..i]$ . We need to choose the smallest  $\alpha$  so that shifting  $s$  by  $\alpha p$  characters results in no mismatches with  $s[1..i]s[j..k]$ , and the shifted  $s$  does not end before the right end of  $s[1..i]s[j..k]$ . We begin with shifting  $s[1..k]$  by  $\left\lceil \frac{i+|s[j..k]|-m}{p} \right\rceil p$  characters, which is the smallest possible shift, and compute its leftmost mismatch with  $s[1..i]s[j..k]$ , which must be on the right of  $s[1..i]$ . We have two cases to consider.

1. There is no mismatch. We check if this shift corresponds to a valid long border in constant time using at most one longest common prefix query. If so, this is the longest long border. Otherwise  $k < m$  and  $s[k+1] \neq s[k+1-p]$ , and choosing any larger shift results in a mismatch, so there is no such long border.
2. There is a mismatch. Again, call the position in  $s[1..i]s[j..k]$  at which the mismatch occurs the *obstacle*. Even if we choose a larger shift,  $s[1..k]$  will expect the same character at the obstacle, so we will get a mismatch as well. Thus there is no such long border.  $\square$

Now we are ready to present the first algorithm NAIVE-PATTERN-MATCHING. It simulates the Knuth-Morris-Pratt algorithm, adding just two optimizations: we extend the current match by whole snippets, not




Figure 2: Shifting  $s$  to get an occurrence in  $t$ .

single letters, and process all long borders of the current match at once. This is not enough to achieve a linear complexity yet: it might happen that we need to spend  $\Omega(\log m)$  time at each snippet. Nevertheless, such method will serve as a basis for developing the optimal solution.

While we assume that for all input snippets the information about corresponding vertex in the suffix tree and longest suffix being a suffix of  $s$  is already known, during the execution we might create some new snippets. Fortunately, they are either prefix snippets, or snippets of the form  $s[\lceil \frac{\ell}{2} \rceil .. \ell]$  which will be called the *half snippets*. Before we analyze the running time of NAIVE-PATTERN-MATCHING we need a small technical lemma concerning those fresh snippets.

**LEMMA 3.3.** *Information for all prefix and half snippets can be computed in linear time.*

*Proof.* In case of prefix snippets the computation is trivial. The case of all half snippets is not completely trivial, though. To compute the suffixes we loop through the possible values of  $\ell$ . Assuming we have the longest suffix of  $s[\lceil \frac{\ell}{2} \rceil .. \ell]$  which is a prefix  $s[1..i]$ , we try to compute such longest suffix of  $s[\lceil \frac{\ell+1}{2} \rceil .. \ell+1]$ . For that we consider the borders of  $s[1..i]$  looking for the one which can be extended with  $s[\ell+1]$ . Then we check if the extended border is longer than  $s[\lceil \frac{\ell+1}{2} \rceil .. \ell+1]$  and if so, take its border. The total complexity of this procedure is linear because at each step we increase the length of the current prefix by at most 1. Now consider locating the vertices in the suffix tree. Assume that we know the corresponding vertex for a given value of  $\ell$  and need to update the information after increasing  $\ell$  by one. Extending the current word by one letter requires traversing at most one edge in the suffix tree, then we might need to remove the first letter. If the current vertex is explicit, we can use its suffix link. Otherwise we use the suffix link of its deepest explicit ancestor, and then traverse edges down from the vertex found. This traversing might require more than constant time, but can be amortized by noting that the number of explicit ancestors of the current vertex cannot exceed  $n$ , and decreases by at most one at for each  $\ell$ . To finish the

proof note that looking up the edge can be performed very quickly even when the alphabet is not constant: there are at most two half snippets of a given length and so we can afford to simply iterate through all outgoing edges, then the total complexity of all lookups will be just linear in  $|s|$ .  $\square$

**THEOREM 3.1.** NAIVE-PATTERN-MATCHING works in time  $\mathcal{O}(n \log m)$ .

*Proof.* Observe that each iteration of the **while** loop results in either increasing  $k$  by 1 or decreasing  $\ell \leq m$  at least twice. Thus the total number of iterations is  $\mathcal{O}(n \log m)$ . In each iteration we spend just a constant time due to the above lemmas. Note that we require that for any snippet we know not only its start and end in  $s$  but also the corresponding vertex in the suffix tree (this will be very important for the improved method). While we assume that we get such information for all the input snippets, we might create some new snippets during the execution of the algorithm. Fortunately, the only possible non-input snippets we create are prefix and half snippets, and so we can use Lemma 3.3 to preprocess them.  $\square$

To accelerate the above basic method we use the concept of *levers*:

**DEFINITION 3.2.** We call  $s_i$  a *lever* of a sequence of snippets  $s_1 s_2 \dots s_k$  if  $\sum_{j=1}^{i-1} |s_j| \leq \frac{|s_i|}{2}$ .

Given such a lever we can eliminate many potential occurrences at once. This is formalized in Lemma 3.4 and Lemma 3.5.

**LEMMA 3.4.** Given a sequence of snippets  $s_1 s_2 \dots s_i$  in which  $s_i$  is a lever, we can detect an occurrence of  $s$  in  $s_1 s_2 \dots s_i$  in time  $\mathcal{O}(i)$ .

*Proof.* Let  $L = |s_i|$  and  $S = \sum_{j=1}^{i-1} |s_j|$ . Because  $S \leq \frac{L}{2}$  and  $L \leq m$ ,  $S$  is at most  $\frac{m}{2}$ , meaning that any occurrence of  $s$  must end in  $s_i$ . Note that we can safely replace  $s_i$  with the longest suffix  $s[k..m]$  which is its prefix, by Lemma 2.2 this requires just constant time. First we check if  $s$  is a suffix of the whole sequence




Figure 3: All occurrences of  $s_i$  in  $s[1..S+L]$ .

in time  $\mathcal{O}(i)$ . Now, any possible occurrence of  $s$  cannot end earlier than after the  $m-S \geq \frac{L}{2}$ -th character and so in fact we are looking for a long border  $b$  of  $s[k..m]$  such that  $s[1..m-b] = s_1s_2\dots s_{i-1}$ . From Lemma 2.4 we know that any such  $b$  must be of the form  $|s[k..m]|-\alpha p$ , where  $p \leq \frac{|s[k..m]|}{2}$  is the period of  $s[k..m]$  calculated in constant time by Lemma 2.5. It turns out that we can restrict the set of possible values of  $\alpha$  to just one, which can be then checked naively in time  $\mathcal{O}(i)$ .

First we compute how far the period of  $s[k..m]$  continues in  $s[1..k-1]$  and  $s_1s_2\dots s_{i-1}$ . In order to perform both those computations efficiently we only have to develop a constant time procedure which, given two snippets  $s[i..j]$  and  $s[k..\ell]$ , finds the longest suffix of the former which is also a suffix of some power of the latter. Such procedure works in two steps:

1. find the longest common suffix of  $s[i..j]$  and  $s[k..\ell]$ , return if it is shorter than  $|s[k..\ell]|$ ,
2. find and return the longest common suffix of  $s[i..j]$  and  $s[i..j - |s[k..\ell]|]$ .

By repeating this procedure at most  $i$  times we can assume that we know how far the period continues in both strings. If  $p$  is a period of the whole  $s$ , check if  $s$  is a suffix of the whole  $s_1s_2\dots s_i$  and note that if it is not, shifting  $s$  to the left by any multiple of  $p$  cannot result in a match. Otherwise we have found the rightmost position  $x$  such that  $s[x] \neq s[x+p]$ . If  $p$  is a period of the whole  $t = s_1s_2\dots s_i$ ,  $s$  does not occur there. Otherwise we have found the rightmost position  $y$  such that  $t[y] \neq t[y+p]$ . Now we claim that in order to get a match, we must shift  $s$  so that its  $x$ -th character is aligned with the  $y$ -th character of  $t$ , see Figure 2 (recall that we shift  $s$  by a multiple of  $p$ ). Indeed, choosing any smaller shift creates a mismatch concerning  $s[x]$  and choosing any bigger shift creates a mismatch concerning  $t[y]$ . This gives us a simple arithmetic condition on the only possible value of  $\alpha$ : if

$p$  does not divide  $|t| - |s| + x - y$  there is no occurrence, otherwise we check  $\alpha = \frac{|t| - |s| + x - y}{p}$ .  $\square$

LEMMA 3.5. *Given a sequence of snippets  $s_1s_2\dots s_i$  in which  $s_i$  is a lever, we can compute the longest prefix of  $s$  which is a suffix of  $s_1s_2\dots s_i$  in time  $\mathcal{O}(i)$ .*

*Proof.* As in the previous proof, let  $L = |s_i|$  and  $S = \sum_{j=1}^{i-1} |s_j|$ . We calculate the longest prefix of  $s$  which is a suffix of  $s_i$  in constant time by Lemma 2.2. Then we should check if there is any longer prefix. If so, it corresponds to an occurrence of  $s_i$  in  $s[1..S+L]$ . First we locate the node corresponding to  $s_i$  in the suffix tree built for  $s$ . The tree can be preprocessed (in linear time) so that having this node we can compute the first and second occurrence of  $s_i$  in  $s[1..S+L]$  (more precisely, we compute the first and second occurrences in the whole  $s$ , and check if they are inside  $s[1..S+L]$ ). If there is none, we terminate. If there is just one, we check naively in time  $\mathcal{O}(i)$  the corresponding prefix. If there are two, the situation is more complicated, as in fact there can be many more of them, and we cannot afford to iterate through all possibilities.

Because  $s_i$  is a lever,  $L \geq 2|S|$  and  $L \geq \frac{2}{3}(S+L)$ . Thus the overlap of any two occurrences of  $s_i$  in  $s[1..S+L]$  is of length at least  $\frac{L}{2}$  and so any non-leftmost occurrence corresponds to a long border of  $s_i$ . Let  $p$  be the period of  $s_i$  (note that we do know  $p$  yet), the first occurrence starts at the  $k$ -th character of  $s$ , and the last occurrence starts at the  $k+\alpha p$ -th character of  $s$ . Then it is clear that there are occurrences starting at the  $k+\beta p$ -th characters, for any  $0 \leq \beta \leq \alpha$ , see Figure 3. In particular, the second occurrence starts at the  $k+p$ -th character, so knowing the positions of the first and second occurrence allows us to compute  $p$ . Having the value of  $p$  and  $k$ , we find the longest common prefix of  $s[k..S+L]$  and  $s[k+p..S+L]$ , which gives us the value of  $\alpha$  (more precisely, if the length is  $l$ ,  $\alpha = \left\lfloor \frac{l+p-L}{p} \right\rfloor$ ). Now we can use the same method as in the previous proof. First compute how far the period of  $s_i$  continues

---

**Algorithm 2** LEVERED-PATTERN-MATCHING( $s_1, s_2, \dots, s_n$ )

---

```
1:  $q \leftarrow 0$ 
2:  $\ell \leftarrow$  longest prefix of  $s$  ending  $s_1$  ▷ Lemma 2.2
3:  $k \leftarrow 2$ 
4: while  $k \leq n$  and  $\ell + \sum_{i=k}^n |s_i| \geq m$  do
5:   choose  $t \geq k$  minimizing  $|s_k| + |s_{k+1}| + \dots + |s_{t-1}| - \frac{|s_t|}{2}$ 
6:   if  $\ell + |s_k| + |s_{k+1}| + \dots + |s_{t-1}| \leq \frac{|s_t|}{2}$  then
7:     check for occurrence of  $s$  in  $s[1.. \ell]s_k s_{k+1} \dots s_t$  ▷ Lemma 3.4
8:      $\ell \leftarrow$  longest prefix of  $s$  ending  $s[1.. \ell]s_k s_{k+1} \dots s_t$  ▷ Lemma 3.5
9:      $k \leftarrow t + 1$ 
10:  else
11:    execute lines 5–18 of NAIVE-PATTERN-MATCHING
12:  end if
13: end while
```

---

in both  $s[1..k]$  and  $s_1 s_2 \dots s_{i-1}$ . Then consider the following three cases.

1. If  $p$  is a period of  $s_1 s_2 \dots s_i$ , for a nonempty suffix which is a prefix of  $s$  to exist  $p$  must be a period of  $s[1..k+p]$  as well. Then the longest such suffix corresponds to the biggest  $\beta \leq \alpha$  for which  $S \geq k - 1 + \beta p$ .
2. If  $p$  is not a period of the whole  $s_1 s_2 \dots s_i$  but is a period of  $s[1..k+p]$ , any prefix of  $s$  which is a suffix of  $s_1 s_2 \dots s_i$  must be completely contained in the periodic suffix of  $s_1 s_2 \dots s_i$ . We can select the longest such suffix in constant time.
3. If  $p$  is neither a period of the whole  $s[1..k+p]$  nor  $s_1 s_2 \dots s_i$ , we get a simple arithmetic condition on the only possible value of  $\alpha$ . Then we verify this value in time  $\mathcal{O}(i)$ .

□

We are ready to present the improved (and final) algorithm. The idea is that whenever the sequence contains a lever, we can use Lemma 3.4 and Lemma 3.5 to quickly process a bunch of snippets. Otherwise we stick to the same method as in NAIVE-PATTERN-MATCHING. We call the resulting method LEVERED-PATTERN-MATCHING.

While the idea is rather simple, it turns out that the notion of levers captures all inputs on which NAIVE-PATTERN-MATCHING runs in superlinear time. To formalize this claim we will need a few technical lemmas, but before that we show how to execute line 5 efficiently.

**LEMMA 3.6.** *Line 5 of LEVERED-PATTERN-MATCHING can be executed in amortized constant time, and  $\mathcal{O}(m)$  additional memory.*

*Proof.* If we are allowed to use as much as  $\Theta(n)$  additional memory, we can preprocess all minima going right to left: the best possible choice of  $t$  for a given value of  $k$  is either the same as for  $k+1$ , or it is equal  $k$ . If we would like to optimize the amount of additional memory used, and avoid the necessity of reading all input snippets when there is a match somewhere near the very beginning, we can use a slightly more complicated method. First note that we are interested only in  $t \leq k + m$ . For each current value of  $k$  we keep an increasing list of *candidates*  $k \leq t_1 < t_2 < \dots < t_c \leq k + m$ . Let  $f(i) = |s_1| + |s_2| + \dots + |s_{i-1}| - \frac{|s_i|}{2}$ , then  $t_1$  is the position with the minimum value of  $f$  in the interval  $[k, k + m]$ ,  $t_2$  is the position with the minimum value of  $f$  in the interval  $[t_1 + 1, k + m]$ , and so on. Before increasing  $k$  by one we need to remove  $t_1$  from the list, if  $t_1 = k$ , and consider a new candidate  $k + 1 + m$ : remove all  $t_i$  with  $f(t_i) > f(k + 1 + m)$ , and add  $k + 1 + m$  to the list. The amortized cost of this update is clearly constant, and the maximum number of stored candidates  $\mathcal{O}(m)$ . To find  $t$ , simply take the first candidate  $t_1$ . □

With each sequence of snippets  $s_1 s_2 \dots s_k$  we associate its potential  $\Phi(|s_1|, |s_2|, \dots, |s_k|)$ , which roughly corresponds to the amount of work we still need to perform if we use the concept of levers. Before we use it to bound the running time of LEVERED-PATTERN-MATCHING, we need a few observations.

**DEFINITION 3.3.** *Let  $x_1, x_2, \dots, x_k \leq m$  be a sequence of natural numbers. Consider a segment of length  $\sum_{i=1}^k x_i$  split into  $k$  blocks of length  $x_i$ , for  $i = 1, 2, \dots, k$ . First mark its suffix of length  $m$ . Then, for each  $i$ , mark a fragment of length  $\frac{x_i}{2}$  ending just before the part corresponding to  $x_i$ . Let  $y_i$  be the length of the marked suffix in the block corresponding to  $x_i$ , and define the potential  $\Phi(x_1, x_2, \dots, x_k)$  as  $\sum_{i=1}^k 2 + \log \frac{x_i}{y_i}$ .*




Figure 4: Inserting  $x_i$  between  $x_{i-1}$  and  $x_{i+1}$ .

Note than all  $y_i$  are strictly positive so the potential is well-defined for any sequence of snippets.

LEMMA 3.7.  $\Phi(x_1, x_2, \dots, x_k) \leq 7k$ .

*Proof.* We apply induction on  $k$ . If  $k = 1$  the whole segment is marked so the potential is 2 and the claim trivially holds. Let  $k > 1$ , choose  $x_i = \min\{x_1, x_2, \dots, x_k\}$  and assume that the claim holds for  $\Phi(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$ . We compute the maximum possible increase in the potential after inserting a block of length  $x_i$  between  $x_{i-1}$  and  $x_{i+1}$ , see Figure 4. There are two reasons the potential might increase.

1. We create a new block of  $x_i$  cells. Note that either  $i = k$  and they are all marked, or  $i < k$  and  $x_{i+1} \geq x_i$  so at least  $\frac{x_i}{2}$  of them are marked. In either case, the resulting increase in the potential is at most  $2 + \log \frac{x_i}{y_i} \leq 2 + \log 2 = 3$ .
2. We move all blocks corresponding to  $x_{i+1}, \dots, x_k$  further to the right. It might result in unmarking some cells in the blocks corresponding to  $x_1, x_2, \dots, x_{i-1}$ . Because we shift all those blocks by  $x_i$  to the right, the unmarked blocks are contained in a segment of such length. Because  $x_i$  is smaller than any  $x_j$  with  $j < i$ , any segment of length  $x_i$  intersects at most two blocks on the left of  $x_i$ . Consider the situation in one such block: the old potential was  $2 + \log \frac{x_j}{y_j}$ , the new potential is  $2 + \log \frac{x_j}{y'_j}$  with  $y'_j \geq \max(y_j - x_i, \frac{x_{j+1}}{2}) \geq \max(y_j - x_i, \frac{x_i}{2})$ . The increase is at most:

$$\begin{aligned} \log \frac{x_j}{y'_j} - \log \frac{x_j}{y_j} &= \log \frac{y_j}{y'_j} \leq \log \frac{y_j}{\max(y_j - x_i, \frac{x_i}{2})} \\ &= \begin{cases} \log \frac{y_j}{y_j - x_i} = \log \frac{y_j}{y_j - \frac{2}{3}y_j} = \log 3 & \text{if } x_i \leq \frac{2}{3}y_j \\ \log \frac{2y_j}{x_i} = \log \frac{2\frac{3}{2}x_i}{x_i} = \log 3 & \text{if } x_i > \frac{2}{3}y_j \end{cases} \end{aligned}$$

because there might be two such blocks, the maximum increase is twice as much,  $2 \log 3 \leq 4$ .

By summing the above cases  $\Phi(x_1, x_2, \dots, x_k) \leq \Phi(x_1, \dots, x_{i-1}, x_i, \dots, x_k) + 7 \leq 7k$ .  $\square$

LEMMA 3.8.  $\Phi(x_1, x_2, \dots, x_k) \geq 1 + \Phi(x_1 + x_2, \dots, x_k)$ .

*Proof.* Let  $y \geq y_2$  be the number of marked cells in the block corresponding to  $x_1 + x_2$  in  $\Phi(x_1 + x_2, x_3, \dots, x_k)$ , see Figure 5. The only change in the potential concerns the first  $x_1 + x_2$  cells:

$$\begin{aligned} \Delta &= \Phi(x_1, x_2, x_3, \dots, x_k) - \Phi(x_1 + x_2, x_3, \dots, x_k) \\ &= 2 + \log \frac{x_1}{y_1} + \log \frac{x_2}{y_2} - \log \frac{x_1 + x_2}{y} \end{aligned}$$

We have a few cases to consider.

1.  $x_1 \leq \frac{x_2}{2}$ . Then  $y_1 = x_1$  and the change in the potential is:

$$\begin{aligned} \Delta &\geq 2 + \log \frac{x_2}{y_2} - \log \frac{x_1 + x_2}{y_2} \\ &= 1 + \log \frac{2x_2}{x_1 + x_2} \geq 1 + \log \frac{4}{3} \geq 1 \end{aligned}$$

2.  $x_1 \geq \frac{x_2}{2}$  and  $y_1 = \frac{x_2}{2}$ . The change in the potential is:

$$\begin{aligned} \Delta &\geq 2 + \log \frac{x_1}{y_1} + \log \frac{x_2}{y_2} - \log \frac{x_1 + x_2}{y_2} \\ &\geq 3 + \log \frac{x_1}{x_2} + \log \frac{x_2}{y_2} - \log \frac{x_1 + x_2}{y_2} \\ &= 3 + \log \frac{x_1}{y_2} - \log \frac{x_1 + x_2}{y_2} = 3 + \log \frac{x_1}{x_1 + x_2} \\ &\geq 3 + \log \frac{1}{3} = 1 + \log \frac{4}{3} \geq 1 \end{aligned}$$

3.  $y_1 > \frac{x_2}{2}$ . Then all marked cells in the block corresponding to  $x_1$  are marked either because of some long  $x_i$  with  $i > 2$  or because they are among the  $m$  rightmost cells, so  $y_2 = x_2$  and merging two first blocks does not change the number of marked cell there,  $y = y_1 + y_2$ . We can bound  $\Delta$  as follows:

$$\begin{aligned} \Delta &= 2 + \log \frac{x_1}{y_1} - \log \frac{x_1 + x_2}{y_1 + y_2} \\ &= 2 + \log \frac{x_1}{y_1} - \log \frac{x_1 + x_2}{y_1 + x_2} \geq 2 \end{aligned}$$

where the last inequality follows from the fact that if  $p \geq q$  then  $\frac{p}{p} \geq \frac{p+r}{q+r}$ .  $\square$






Figure 5: Merging blocks corresponding to  $x_1$  and  $x_2$ .

LEMMA 3.9.  $\Phi(x_1, x_2, \dots, x_k) \geq \Phi(x'_1, x_2, \dots, x_k)$  if  $x_1 \geq x'_1$ .

*Proof.* Decreasing  $x_1$  cannot change any  $y_i$  with  $i > 1$ . Let  $y'_1$  be the number of marked cells in the block corresponding to  $x'_1$ . Then  $y_1 = y'_1$ , and from the definition of the potential we get the claim, or  $y'_1 = x'_1$ , and  $\log \frac{x'_1}{y'_1} = 0 \leq \log \frac{x_1}{y_1}$ .  $\square$

THEOREM 3.2. LEVERED-PATTERN-MATCHING can be implemented to work in time  $\mathcal{O}(n)$  and use  $\mathcal{O}(m)$  additional memory.

*Proof.* Consider any execution of the algorithm. In the very beginning we allocate  $\Phi(|s_1|, |s_2|, \dots, |s_n|)$  credits which we then use to pay for all executions of lines 5–11 of LEVERED-PATTERN-MATCHING. We keep the following invariant during the whole execution: we have  $\Phi(\ell, |s_k|, |s_{k+1}|, \dots, |s_n|)$  credits available. Consider a single pass through the body of the **while** loop. From Lemma 3.6 the amortized cost of executing 5 is constant. Consider the remaining lines.

1. There is a lever  $s_t$  in  $s[1 \dots \ell]_{s_k s_{k+1} \dots s_t}$ . We need as much as  $\mathcal{O}(t - k + 1)$  time to process it, but then by  $t - k + 1$  applications of Lemma 3.8 the required potential is at most:

$$\begin{aligned} \Phi(\ell + |s_k| + |s_{k+1}| + \dots + |s_t|, |s_{t+1}|, \dots, |s_n|) \\ \leq \Phi(\ell, |s_k|, |s_{k+1}|, \dots, |s_n|) - (t - k + 1) \end{aligned}$$

which leaves us with  $t - k$  free credits which we can use to pay for the processing time.

2. There is no lever. Then we execute the corresponding lines from NAIVE-PATTERN-MATCHING which takes just constant time and results in either increasing  $k$  by 1 and  $\ell$  by at most  $|s_k|$  or decreasing  $\ell$  at least twice. In the first case the required potential is by Lemma 3.8 and Lemma 3.9 at most:

$$\Phi(\ell + |s_k|, |s_{k+1}|, \dots, |s_n|) \leq \Phi(\ell, |s_k|, \dots, |s_n|) - 1$$

The second case is slightly more involved: if decreasing  $\ell$  creates a lever, we pay for the pass in

the next round. Otherwise we replace  $\log \frac{x}{y}$  with  $\log \frac{x}{2y}$  in the potential so the required amount of credits left is just:

$$\Phi\left(\frac{\ell}{2}, |s_k|, \dots, |s_n|\right) \leq \Phi(\ell, |s_k|, \dots, |s_n|) - 1$$

which leaves us with one spare credit.

The total amount of allocated credits is by Lemma 3.7 just  $\mathcal{O}(n)$  and so is the time complexity.  $\square$

#### 4 LZW compression

Recall that a LZW compressed string is a sequence of codewords  $t_1 t_2 \dots t_n$ , each codeword being either a single letter, or a previously occurring codeword concatenated with a single letter. First we check if the pattern  $s$  occurs inside one of the strings represented by the codewords. Then for each codeword we need to know if the string it represents occurs in the pattern  $s$ , and if it does, we need to find the corresponding vertex in the suffix tree. We also need to know its longest prefix which is a suffix of  $s$ , and the longest suffix which is a prefix of  $s$  (actually, knowing the prefix is necessary only when the string does not occur inside  $s$ ). All those informations can be found efficiently (in constant time for each codeword) assuming that the alphabet is of constant size.

LEMMA 4.1. *If the alphabet is of constant size, we can perform the preprocessing for all codewords in total linear time.*

*Proof.* The whole set of codewords should be viewed as a trie. For each vertex of this trie we are required to compute:

1. the longest prefix of the corresponding word which is a suffix of  $s$ ,
2. if the corresponding word occurs in  $s$ , locate its (implicit or explicit) vertex in the suffix tree,
3. the longest suffix of the corresponding word which is a prefix of  $s$ .

We build an automaton  $\mathcal{A}$  recognizing all prefixes of  $s$ , i.e., its states set is  $\{0, 1, \dots, |s|\}$  and after reading a word  $w$  we are in a state corresponding to the longest prefix of  $s$  ending  $w$ . Such automaton can be easily constructed in linear time if the alphabet is of constant size. Using this automaton we can compute the longest suffix of each codeword in constant time per codeword.

Then we build the suffix tree for  $s$  in linear time [18]. For each codeword we will find the corresponding vertex in the suffix tree, if any. Note that such information is actually enough to compute the longest prefix which is a suffix of  $s$  for each codeword: first apply Lemma 2.2 to all codewords with corresponding vertex in the suffix tree. Then for all other codewords, simply take the answer computed for its lowest ancestor with a known result. To locate all corresponding vertices in the suffix tree we traverse the trie in a top-bottom fashion. To find the vertex for  $v$ , take the vertex found for its parent and traverse at most one edge. This gives a total linear time.  $\square$

After applying the above lemma to the set of all codewords, for each maximal sequence of codewords representing substrings of  $s$  ( $s[i_1 \dots j_1]s[i_2 \dots j_2] \dots s[i_k \dots j_k]$ ) we take the longest suffix of the proceeding codeword which is a prefix of  $s$  ( $s[1 \dots i]$ ) and the longest prefix of the succeeding codeword which is a suffix of  $s$  ( $s[j \dots m]$ ), and run the algorithm from the previous section on the sequence of snippets  $s[1 \dots i]s[i_1 \dots j_1] \dots s[i_k \dots j_k]s[j \dots m]$ . Because each such run takes time  $\mathcal{O}(k + 2)$ , and all those values of  $k$  sum up to at most  $n$ , the total complexity including the pre-processing of the pattern will be  $\mathcal{O}(n + m)$ . Note that in fact we do not have to process all codewords in the very beginning: with a little care it is possible to process the input in an online fashion, which results in an algorithm stopping as soon as the first occurrence is found (or, more precisely, reading at most  $\mathcal{O}(m)$  succeeding codewords).

**THEOREM 4.1.** *Pattern matching for LZW compressed strings over a constant size alphabet can be solved in optimal linear time.*

## 5 Integer alphabets

The assumption of the constant size alphabet is not necessary to achieve the claimed linear running time. If the alphabet is of non-constant size but consists of integers which can be sorted in linear time, we can create all necessary snippets sequences in linear time as well.

**LEMMA 5.1.** *If the alphabet consists of integers which can be sorted in linear time, we can perform the pre-processing for all codewords in total linear time, assuming the RAM model of computation.*

*Proof.* We begin with constructing the suffix tree in linear time using the assumption of integer alphabet and applying the result of [7]. Then we are left with answering the following three questions for each different codeword:

1. find its longest prefix which is a suffix of  $s$ ,
2. check if it occurs in  $s$ , and if it does, find the corresponding vertex,
3. find its longest suffix which is a prefix of  $s$ .

We answer those questions for all codewords at once, which requires reading the whole input even if there is an occurrence in the very beginning. Questions of each type are processed separately.

1. We use Lemma 2.2 and the information found for the second type questions.
2. We build a trie  $T$  containing all the codewords. Then answering the questions reduces to computing the intersection of this trie and the suffix tree  $S$  (which is a compressed representation of a trie containing all subwords of  $s$ ). This can be done in time  $\mathcal{O}(|S| + |T|)$  assuming the edges outgoing from each vertex (both in  $S$  and  $T$ ) are sorted according to their labels. We process the codewords in order of their lengths. Assuming that we know the corresponding vertices for all codewords of length  $L$ , we consider all codewords  $wx$  of length  $L+1$ , and group them according to the vertex corresponding to  $w$  (if there is none,  $wx$  clearly does not occur in  $s$ ). In each group the codewords are sorted using  $x$  as the key, which can be assured by sorting all codewords in the very beginning. Then we find the corresponding vertices for all codewords in a single group at once, using a left-to-right scan.
3. Finding such suffix for a single codeword can be performed by running the Knuth-Morris-Pratt algorithm. While the amortized complexity of processing a single letter is constant, we have a lot of different letters that can extend a given codeword, and the amortization argument does not give a linear bound in such case.

Consider the automaton recognizing all prefixes of  $s$  from Lemma 4.1, i.e., with the states set  $\{0, 1, \dots, m\}$  and the transitions  $\delta(i, a) = \max\{j : s[1 \dots j] = s[i - j + 1 \dots i]\}$ . If the alphabet is of unbounded size the simple linear time construction is no longer possible. Fortunately, the number of nonzero transitions outgoing from a single vertex is at most  $2 \log n$ , which follows from a well known fact [6, Lemma 2.32] that  $\pi'^{(2)}(k) < \frac{k}{2}$ , where

$\pi'$  is the so-called strong failure function, defined as follows:  $\pi'(k)$  is the longest border of  $s[1..k]$  such that  $s[k+1] \neq s[\pi'(k)+1]$ . In fact much more is known: due to a result of [17], the total number of nontrivial transitions is linear in  $m$ , regardless of the size of the alphabet. Using the atomic heaps of Fredman and Willard [9], we can maintain a number of collections of polylogarithmic size, with amortized constant time update and worst-case constant time look-up, which is enough to construct the automaton in linear time. Then we are able to compute the transition function in constant time per character, which allows us to process all codewords in linear time. Introducing the atomic heaps results in a rather complicated method, and in this particular case it is possible to significantly simplify this structure because the universe is of just polynomial size. It does not result in changing the claimed time bounds, though.

□

This gives us the final result. Note that in case of integer alphabets, it is not clear how to avoid reading the whole input even in the case when there is an occurrence somewhere in the very beginning.

**THEOREM 5.1.** *Pattern matching for LZW compressed strings over a polynomial size alphabet can be solved in optimal linear time assuming the word RAM model.*

## References

- [1] A. Amir and G. Benson. Efficient two-dimensional compressed matching. In *Data Compression Conference*, pages 279–288, 1992.
- [2] A. Amir, G. Benson, and M. Farach. Let sleeping files lie: pattern matching in z-compressed files. In *SODA '94: Proceedings of the fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 705–714, Philadelphia, PA, USA, 1994. Society for Industrial and Applied Mathematics.
- [3] M. A. Bender and M. Farach-Colton. The lca problem revisited. In *LATIN '00: Proceedings of the 4th Latin American Symposium on Theoretical Informatics*, pages 88–94, London, UK, 2000. Springer-Verlag.
- [4] R. S. Boyer and J. S. Moore. A fast string searching algorithm. *Commun. ACM*, 20(10):762–772, 1977.
- [5] D. Breslauer. Saving comparisons in the Crochemore-Perrin string matching algorithm. In *In Proc. of 1st European Symp. on Algorithms*, pages 61–72, 1995.
- [6] M. Crochemore, C. Hancart, and T. Lecroq. *Algorithms on Strings*. Cambridge University Press, 2007.
- [7] M. Farach. Optimal suffix tree construction with large alphabets. In *FOCS '97: Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, page 137, Washington, DC, USA, 1997. IEEE Computer Society.
- [8] M. Farach and M. Thorup. String matching in Lempel-Ziv compressed strings. In *STOC '95: Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, pages 703–712, New York, NY, USA, 1995. ACM.
- [9] M. L. Fredman and D. E. Willard. Trans-dichotomous algorithms for minimum spanning trees and shortest paths. *J. Comput. Syst. Sci.*, 48(3):533–551, 1994.
- [10] Z. Galil. String matching in real time. *J. ACM*, 28(1):134–149, 1981.
- [11] Z. Galil and J. Seiferas. Time-space-optimal string matching (preliminary report). In *STOC '81: Proceedings of the thirteenth annual ACM symposium on Theory of computing*, pages 106–113, New York, NY, USA, 1981. ACM.
- [12] L. Gasieniec and W. Rytter. Almost optimal fully LZW-compressed pattern matching. In *DCC '99: Proceedings of the Conference on Data Compression*, page 316, Washington, DC, USA, 1999. IEEE Computer Society.
- [13] D. E. Knuth, J. H. Morris, Jr., and V. R. Pratt. Fast pattern matching in strings. *SIAM J. Comput.*, 6(2):323–350, 1977.
- [14] S. R. Kosaraju. Pattern matching in compressed texts. In *Proceedings of the 15th Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 349–362, London, UK, 1995. Springer-Verlag.
- [15] J. H. Morris, Jr. and V. R. Pratt. A linear pattern-matching algorithm. Technical Report 40, University of California, Berkeley, 1970.
- [16] G. Navarro and M. Raffinot. A general practical approach to pattern matching over Ziv-Lempel compressed text. In *CPM '99: Proceedings of the 10th Annual Symposium on Combinatorial Pattern Matching*, pages 14–36, London, UK, 1999. Springer-Verlag.
- [17] I. Simon. String matching algorithms and automata. In *Proceedings of the Colloquium in Honor of Arto Salomaa on Results and Trends in Theoretical Computer Science*, pages 386–395, London, UK, 1994. Springer-Verlag.
- [18] E. Ukkonen. On-line construction of suffix trees. *Algorithmica*, 14(3):249–260, 1995.
- [19] T. A. Welch. A technique for high-performance data compression. *Computer*, 17(6):8–19, 1984.
- [20] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, 23:337–343, 1977.