

# Longest Gapped Repeats and Palindromes <sup>\*</sup>

Marius Dumitran<sup>1</sup>    Paweł Gawrychowski<sup>2</sup>    Florin Manea<sup>3</sup>

<sup>1</sup> Faculty of Mathematics and Computer Science, University of Bucharest

<sup>2</sup> Institute of Computer Science, University of Wrocław

<sup>3</sup> Department of Computer Science, Kiel University

received 24<sup>th</sup> Nov. 2015, revised 12<sup>th</sup> Oct. 2016, 25<sup>th</sup> Sep. 2017, accepted 3<sup>rd</sup> Oct. 2017.

---

A gapped repeat (respectively, palindrome) occurring in a word  $w$  is a factor  $uvu$  (respectively,  $u^Rvu$ ) of  $w$ . In such a repeat (palindrome)  $u$  is called the arm of the repeat (respectively, palindrome), while  $v$  is called the gap. We show how to compute efficiently, for every position  $i$  of the word  $w$ , the longest gapped repeat and palindrome occurring at that position, provided that the length of the gap is subject to various types of restrictions. That is, that for each position  $i$  we compute the longest prefix  $u$  of  $w[i..n]$  such that  $uv$  (respectively,  $u^Rv$ ) is a suffix of  $w[1..i-1]$  (defining thus a gapped repeat  $uvu$  – respectively, palindrome  $u^Rvu$ ), and the length of  $v$  is subject to the aforementioned restrictions.

**Keywords:** combinatorial pattern matching, gapped repeats, gapped palindromes

---

## 1 Introduction

Gapped repeats and palindromes have been investigated for a long time (see, e.g., Gusfield (1997); Brodal et al. (1999); Kolpakov and Kucherov (2000, 2009a); Kolpakov et al. (2014); Crochemore et al. (2010); Crochemore and Tischler (2011) and the references therein), with motivation coming especially from the analysis of DNA and RNA structures, where tandem repeats or hairpin structures play important roles in revealing structural and functional information of the analysed genetic sequence (see Gusfield (1997); Brodal et al. (1999); Kolpakov and Kucherov (2009a) and the references therein).

A gapped repeat (respectively, palindrome) occurring in a word  $w$  is a factor  $uvu$  (respectively,  $u^Rvu$ ) of  $w$ . The middle part  $v$  of such a structure is called gap, while the two factors  $u$  (respectively, the factors  $u^R$  and  $u$ ) are called left and right arms. Generally, the previous works were interested in finding all the gapped repeats and palindromes, under certain numerical restrictions on the length of the gap or on the relation between the length of the arm of the repeat or palindrome and the length of the gap.

A classical problem for palindromes asks to find the longest palindromic factor of a word (see, Manacher (1975)). This is our first inspiration point in proposing an alternative point of view in the study of gapped repeats and palindromes. As a second inspiration point for our work, we refer to the longest

---

<sup>\*</sup>This is an extension of the conference papers of Gawrychowski and Manea (2015), presented at the 20th International Symposium on Fundamentals of Computation Theory, FCT 2015, and of Dumitran and Manea (2015), presented at the 40th International Symposium on Mathematical Foundations of Computer Science, MFCS 2015.

previous factor table (LPF) associated to a word. This data structure was introduced and considered in the context of efficiently computing Lempel-Ziv-like factorisations of words (see Crochemore et al. (2013, 2010)). Such a table provides for each position  $i$  of the word the longest factor occurring both at position  $i$  and once again on a position  $j < i$ . Several variants of this table were also considered by Crochemore et al. (2010): the longest previous reverse factor (*LPrF*), where we look for the longest factor occurring at position  $i$  and whose mirror image occurs in the prefix of  $w$  of length  $i - 1$ , or the longest previous non-overlapping factor, where we look for the longest factor occurring both at position  $i$  and somewhere inside the prefix of length  $i - 1$  of  $w$ . Such tables may be seen as providing a comprehensive image of the long repeats and symmetries occurring in the analysed word.

According to the above, in our work we approach the construction of longest previous gapped repeat or palindrome tables: for each position  $i$  of the word we want to compute the longest factor occurring both at position  $i$  and once again on a position  $j < i$  (or, respectively, whose mirror image occurs in the prefix of length  $i - 1$  of  $w$ ) such that there is a gap (subject to various restrictions) between  $i$  and the previous occurrence of the respective factor (mirrored factor). Similar to the original setting, this should give us a good image of the long gapped repeats and symmetries of a word.

A simple way to restrict the gap is to lower bound it by a constant; i.e., we look for factors  $uvu$  (or  $u^Rvu$ ) with  $|v| > g$  for some  $g \geq 0$ . The techniques of Crochemore et al. (2010) can be easily adapted to compute for each position  $i$  the longest prefix  $u$  of  $w[i..n]$  such that there exists a suffix  $uv$  (respectively,  $u^Rv$ ) of  $w[1..i - 1]$ , forming thus a factor  $uvu$  (respectively,  $u^Rvu$ ) with  $|v| > g$ . Here we consider three other different types of restricted gaps.

We first consider the case when the length of the gap is between a lower bound  $g$  and an upper bound  $G$ , where  $g$  and  $G$  are given as input (so, may depend on the input word). This extends naturally the case of lower bounded gaps.

**Problem 1** Given  $w$  of length  $n$  and two integers  $g$  and  $G$ , such that  $0 \leq g < G \leq n$ , construct the arrays  $LPrF_{g,G}[\cdot]$  and  $LPF_{g,G}[\cdot]$  defined for  $1 \leq i \leq n$ :

- a.  $LPrF_{g,G}[i] = \max\{|u| \mid \text{there exists } v \text{ such that } u^Rv \text{ is a suffix of } w[1..i - 1] \text{ and } u \text{ is prefix of } w[i..n], \text{ with } g \leq |v| < G\}$ .
- b.  $LPF_{g,G}[i] = \max\{|u| \mid \text{there exists } v \text{ such that } uv \text{ is a suffix of } w[1..i - 1] \text{ and } u \text{ is prefix of } w[i..n], \text{ with } g \leq |v| < G\}$ .

We are able to solve Problem 1(a) in linear time  $\mathcal{O}(n)$ . Problem 1(b) is solved here in  $\mathcal{O}(n \log n)$  time. Intuitively, in the case of gapped palindromes, when trying to compute the longest prefix  $u$  of  $w[i..n]$  such that  $u^Rv$  is a suffix of  $w[1..i - 1]$  with  $g < |v| \leq G$ , we just have to compute the longest common prefix between  $w[i..n]$  and the words  $w[1..j]^R$  with  $g < i - j \leq G$ . The increased difficulty in solving the problem for repeats (reflected in the increased complexity of our algorithm) seems to come from the fact that when trying to compute the longest prefix  $u$  of  $w[i..n]$  such that  $uv$  is a suffix of  $w[1..i - 1]$  with  $g < |v| \leq G$ , it is hard to see where the  $uv$  factor may start, so we have to somehow try more variants for the length of  $u$ . Brodal et al. (1999) give an algorithm that finds all maximal repeats (i.e., repeats whose arms cannot be extended) with gap between a lower and an upper bound, running in  $\mathcal{O}(n \log n + z)$  time, where  $z$  is the number of such repeats. It is worth noting that there are words (e.g.,  $(a^2b)^{n/3}$ , from Brodal et al. (1999)) that may have  $\Theta(nG)$  maximal repeats  $uvu$  with  $|v| < G$ . Hence, for  $G > \log n$  and  $g = 0$ , for instance, our algorithm is faster than an approach that would first use the algorithms of Brodal et al. (1999) to get all maximal repeats, and then process them somehow to solve Problem 1(b).

The data structures we construct allow us to trivially find in linear time the longest gapped palindrome having the length of the gap between  $g$  and  $G$ , and in  $\mathcal{O}(n \log n)$  time the longest gapped repeat with the length of the gap between the bounds  $g$  and  $G$ .

In the second case, the gaps of the repeats and palindromes we investigate are only lower bounded; however, the bound on the gap allowed at each position is defined by a function depending on the position.

**Problem 2** Given  $w$  of length  $n$  and the values  $g(1), \dots, g(n)$  of  $g : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ , construct the arrays  $LPrF_g[\cdot]$  and  $LPF_g[\cdot]$  defined for  $1 \leq i \leq n$ :

- a.  $LPrF_g[i] = \max\{|u| \mid \text{there exists } v \text{ such that } u^R v \text{ is a suffix of } w[1..i-1] \text{ and } u \text{ is prefix of } w[i..n], \text{ with } g(i) \leq |v|\}$ .
- b.  $LPF_g[i] = \max\{|u| \mid \text{there exists } v \text{ such that } uv \text{ is a suffix of } w[1..i-1] \text{ and } u \text{ is prefix of } w[i..n], \text{ with } g(i) \leq |v|\}$ .

The setting of this problem can be seen as follows. An expert preprocesses the input word (in a way specific to the framework in which one needs this problem solved), and detects the length of the gap occurring at each position (so, computes  $g(i)$  for all  $i$ ). These values and the word are then given to us, to compute the arrays defined in our problems. We solve both problems in linear time. Consequently, we can find in linear time the longest gapped palindrome or repeat whose gap fulfils the length restriction defined by the position where this palindrome or repeat occurs (as above).

Finally, following Kolpakov and Kucherov (2009a); Kolpakov et al. (2014), we analyse gapped repeats  $uvu$  or palindromes  $u^Rvu$  where the length of the gap  $v$  is upper bounded by the length of the arm  $u$  multiplied by some factor. More precisely, Kolpakov et al. (2014) investigate  $\alpha$ -gapped repeats: words  $uvu$  with  $|uv| \leq \alpha|u|$ . Similarly, Kolpakov and Kucherov (2009a) analyse  $\alpha$ -gapped palindromes, i.e., words  $u^Rvu$  with  $|uv| \leq \alpha|u|$ . For  $\alpha = 2$ , these structures are called long armed repeats (or pairs) and palindromes, respectively; for  $\alpha = 1$ , they are squares and palindromes of even length, respectively. Intuitively, one is interested in repeats or palindromes whose arms are roughly close one to the other; therefore, the study of  $\alpha$ -gapped repeats and palindromes was rather focused on the cases with small  $\alpha$ . Here, we address the general case, of searching in a word  $w$   $\alpha$ -gapped repeats or palindromes for  $\alpha \leq |w|$ .

**Problem 3** Given  $w$  of length  $n$  and a number  $\alpha \leq n$ , construct the arrays  $LPal_\alpha[\cdot]$  and  $LRep_\alpha[\cdot]$ , defined for  $1 \leq i \leq n$ :

- a.  $LPal_\alpha[i] = \max\{|u| \mid \text{there exists } v \text{ such that } u^R v \text{ is a suffix of } w[1..i-1], u \text{ is a prefix of } w[i..n], \text{ and } |uv| \leq \alpha|u|\}$ .
- b.  $LRep_\alpha[i] = \max\{|u| \mid \text{there exists } v \text{ such that } uv \text{ is a suffix of } w[1..i-1], u \text{ is a prefix of } w[i..n], \text{ and } |uv| \leq \alpha|u|\}$ .

The problem of constructing the set  $S$  of all factors of a word of length  $n$  which are maximal  $\alpha$ -gapped repeats or palindromes (i.e., the arms cannot be extended simultaneously with one symbol to the right or to the left to get a longer similar structure) was thoroughly considered, and finally settled by Crochemore et al. (2016); Gawrychowski et al. (2016) (see also the references therein). In both these papers, it is shown that the number of  $\alpha$ -gapped repeats or palindromes a word of length  $n$  may contain is  $\Theta(\alpha n)$ . Using as starting point the algorithm presented in Gawrychowski and Manea (2015), that finds the longest  $\alpha$ -gapped repeat or palindrome (without constructing the set of all such structures), Gawrychowski et al.

(2016) give an algorithm finding all maximal  $\alpha$ -gapped repeats and  $\alpha$ -gapped palindromes in optimal time  $\Theta(\alpha n)$ . Here, we first present the algorithm of Gawrychowski and Manea (2015) for the identification of the longest  $\alpha$ -gapped repeat or palindrome contained in a word, and briefly explain how it was extended to output all maximal  $\alpha$ -gapped repeats and palindromes in a word. Then we use the algorithm of Gawrychowski et al. (2016) and a linear time algorithm finding the longest square/palindrome centred at each position of a word to solve Problem 3 in linear time.

Our algorithms are generally based on efficient data-structures. On one hand, we use efficient word-processing data structures like suffix arrays, longest common prefix structures, or dictionaries of basic factors. On the other hand, we heavily use specific data-structures for maintaining efficiently collections of disjoint sets, under union and find operations. Alongside these data-structures, we make use of a series of remarks of combinatorial nature, providing insight in the repetitive structure of the words.


## 2 Preliminaries

The computational model we use to design and analyze our algorithms is the standard unit-cost RAM with logarithmic word size, which is generally used in the analysis of algorithms. In this model, the memory word size is logarithmic in the size of the input.

Let  $V$  be a finite alphabet;  $V^*$  denotes the set of all finite words over  $V$ . In the upcoming algorithmic problems, we assume that the words we process are sequences of integers (i.e., over integer alphabets). In general, if the input word has length  $n$  then we assume its letters are in  $\{1, \dots, n\}$ , so each letter fits in a single memory-word. This is a common assumption in stringology (see, e.g., the discussion by Kärkkäinen et al. (2006)).

The *length* of a word  $w \in V^*$  is denoted by  $|w|$ . The *empty word* is denoted by  $\lambda$ . A word  $u \in V^*$  is a *factor* of  $v \in V^*$  if  $v = xuy$ , for some  $x, y \in V^*$ ; we say that  $u$  is a *prefix* of  $v$ , if  $x = \lambda$ , and a *suffix* of  $v$ , if  $y = \lambda$ . We denote by  $w[i]$  the symbol occurring at position  $i$  in  $w$ , and by  $w[i..j]$  the factor of  $w$  starting at position  $i$  and ending at position  $j$ , consisting of the catenation of the symbols  $w[i], \dots, w[j]$ , where  $1 \leq i \leq j \leq n$ ; we define  $w[i..j] = \lambda$  if  $i > j$ . The powers of a word  $w$  are defined recursively by  $w^0 = \lambda$  and  $w^n = ww^{n-1}$  for  $n \geq 1$ . If  $w$  cannot be expressed as a nontrivial power (i.e.,  $w$  is not a repetition) of another word, then  $w$  is *primitive*. A *period* of a word  $w$  over  $V$  is a positive integer  $p$  such that  $w[i] = w[j]$  for all  $i$  and  $j$  with  $i \equiv j \pmod{p}$ . Let  $per(w)$  be the smallest period of  $w$ . A word  $w$  with  $per(w) \leq \frac{|w|}{2}$  is called *periodic*; a periodic  $w[i..j]$  (with  $p = per(w[i..j]) < \frac{j-i+1}{2}$ ) is a *run* if it cannot be extended to the left or right to get a word with the same period  $p$ , i.e.,  $i = 1$  or  $w[i-1] \neq w[i+p-1]$ , and,  $j = n$  or  $w[j+1] \neq w[j-p+1]$ . Kolpakov and Kucherov (1999) showed that the number of runs of a word is linear and their list (with a run  $w[i..j]$  represented as the triple  $(i, j, per(w[i..j]))$ ) can be computed in linear time. The exponent of a run  $w[i..j]$  occurring in  $w$  is defined as  $\frac{j-i+1}{per(w[i..j])}$ ; the sum of the exponents of all runs in a word of length  $n$  is  $\mathcal{O}(n)$  (see Kolpakov and Kucherov (1999)).

For a word  $u$ ,  $|u| = n$ , over  $V \subseteq \{1, \dots, n\}$  we build in  $\mathcal{O}(n)$  time the suffix array as well as data structures allowing us to retrieve in constant time the length of the longest common prefix of any two suffixes  $u[i..n]$  and  $u[j..n]$  of  $u$ , denoted  $LCP_u(i, j)$  (the subscript  $u$  is omitted when there is no danger of confusion). Such structures are called *LCP* data structures in the following. For details, see, e.g., Kärkkäinen et al. (2006); Gusfield (1997), and the references therein. In the solutions of the problems dealing with gapped palindromes inside a word  $w$  (Problems 1(a), 2(a), and 3(a)) we construct the suffix array and *LCP* data structures for the word  $u = w0w^R$ , where 0 is a new symbol lexicographically



**Fig. 1:** Occurrences of the basic factors  $y = w[i..i + 2^k - 1]$  in  $z = w[j..j + 8 \cdot 2^k - 1]$ . The overlapping occurrences are part of runs, and they can be returned as the pair formed of the first occurrence of  $y$  from each run and the period of  $y$ . The representation of the occurrences of  $y$  in  $z$  will return 4 elements: 3 runs and one separate occurrence.

smaller than all the symbols of  $V$ ; this takes  $\mathcal{O}(|w|)$  time. To check whether  $w[i..j]$  occurs at position  $\ell$  in  $w$  (respectively,  $w[i..j]^R$  occurs at position  $\ell$  in  $w$ ) we check whether  $\ell + (j - i + 1) \leq n$  and  $LCP_u(i, \ell) \geq j - i + 1$  (respectively,  $LCP_u(\ell, 2|w| - j + w) \geq j - i + 1$ ). To keep the notation free of subscripts, when we measure the longest common prefix of a word  $w[1..j]^R$  and word  $w[i..n]$  we write  $LCP(w[1..j]^R, w[1..i])$ , and this is in fact an  $LCP$ -query on  $u = w0w^R$ ; when we measure the longest common prefix of a word  $w[j..n]$  and word  $w[i..n]$  we write  $LCP(j, i)$ , and this is in fact an  $LCP$ -query on  $w$ .

The suffix array of  $w0w^R$  allows us to construct in linear time a list  $\mathcal{L}$  of the suffixes  $w[i..n]$  of  $w$  and of the mirror images  $w[1..i]^R$  of the prefixes of  $w$  (which correspond to the suffixes of length less than  $|w|$  of  $w0w^R$ ), ordered lexicographically. Generally, we denote by  $Rank[i]$  the position of  $w[i..n]$  in the ordered list  $\mathcal{L}$  of these factors, and by  $Rank_R[i]$  the position of  $w[1..i]^R$  in  $\mathcal{L}$ .

The dictionary of basic factors (introduced by Crochemore and Rytter (1991)) of a word  $w$  (DBF for short) is a data structure that labels the factors  $w[i..i + 2^k - 1]$  (called basic factors), for  $k \geq 0$  and  $1 \leq i \leq n - 2^k + 1$ , such that every two identical factors of  $w$  get the same label and we can retrieve the label of any basic factor in  $\mathcal{O}(1)$  time. The DBF of a word of length  $n$  is constructed in  $\mathcal{O}(n \log n)$  time.

Note that a basic factor  $w[i..i + 2^k - 1]$  occurs either at most twice in any factor  $w[j..j + 2^{k+1} - 1]$  or the occurrences of  $w[i..i + 2^k - 1]$  in  $w[j..j + 2^{k+1} - 1]$  form a run of period  $per(w[i..i + 2^k - 1])$  (so the corresponding positions where  $w[i..i + 2^k - 1]$  occurs in  $w[j..j + 2^{k+1} - 1]$  form an arithmetic progression of ratio  $per(w[i..i + 2^k - 1])$ , see Kociumaka et al. (2012)). Hence, the occurrences of  $w[i..i + 2^k - 1]$  in  $w[j..j + 2^{k+1} - 1]$  can be presented in a compact manner: either at most two positions, or the starting position of the progression and its ratio. For  $c \geq 2$ , the occurrences of the basic factor  $w[i..i + 2^k - 1]$  in  $w[j..j + c2^k - 1]$  can be also presented in a compact manner: the positions (at most  $c$ ) where  $w[i..i + 2^k - 1]$  occurs isolated (not inside a run) and/or at most  $c$  maximal runs that contain the overlapping occurrences of  $w[i..i + 2^k - 1]$ , each run having period  $per(w[i..i + 2^k - 1])$ .

**Remark 1** Using the DBF of a word  $w$  of length  $n$ , given a number  $\ell > 0$  we can produce in  $\mathcal{O}(n \log n)$  time a data structure answering the following type of queries in  $\mathcal{O}(1)$  time: “Given  $i$  and  $k$  return the compact representation of the occurrences of the basic factor  $w[i..i + 2^k - 1]$  inside the basic factor  $w[i - \ell - 2^{k+1}..i - \ell - 1]$ ”. Similarly, given  $\ell > 0$  and a constant  $c > 0$  (e.g.,  $c = 10$ ), we can produce in  $\mathcal{O}(n \log n)$  time a data structure answering the following type of queries in  $\mathcal{O}(1)$  time: “Given

$i$  and  $k$  return the compact representation of the occurrences of the basic factor  $w[i..i + 2^k - 1]$  in  $w[i - \ell - c2^k..i - \ell - 1]$ ”.

Indeed, once we construct the dictionary of basic factors of  $w$ , we reorganise it such that for each distinct basic factor we have an array with all the positions where it occurs, ordered increasingly. Now, we traverse each such array, keeping track of the current occurrence  $w[i..i + 2^k - 1]$  and a window containing its occurrences from the range between  $i - \ell - c2^k$  and  $i - \ell$  (and a compact representation of these occurrences); when we move in our traversal of this array to the next occurrence of the current basic factors, we also slide the window in the array, and, looking at the content of the previous window and keeping track of the occurrences that were taken out and those that were added to its content, we can easily obtain in constant time a representation of the occurrences of the considered basic factors inside the new window.

The previous remark is extended by the following lemma in a more general setting:  $c$  is no longer a constant, and in a query we look, this time, for the occurrences of a basic factor  $w[i..i + 2^k - 1]$  in factors  $z = w[j..j + c2^k - 1]$ , where there is no relation between  $i$  and  $j$ .

**Lemma 1** *Given a word  $w$  of length  $n$  and a number  $c \geq 2$ , we can preprocess  $w$  in time  $\mathcal{O}(n \log n)$  such that given any basic factor  $y = w[i..i + 2^k - 1]$  and any factor  $z = w[j..j + c2^k - 1]$ , with  $k \geq 0$ , we can compute in  $\mathcal{O}(\log \log n + c)$  time a compact representation of all the occurrences of  $y$  in  $z$ .*

**Proof:** We construct the dictionary of basic factors of a word of length  $n$  in  $\mathcal{O}(n \log n)$  time and reorganise it such that for each basic factor we have an array with all its occurrences, ordered by their starting positions. For each such array we construct data structures that allow predecessor/successor search in  $\mathcal{O}(\log \log n)$  time with linear time preprocessing w.r.t. its length (see, e.g., van Emde Boas (1975)). When we have to return the occurrences of  $y = w[i..i + 2^k - 1]$  in  $z = w[j..j + c2^k - 1]$ , we search in the basic factors-array corresponding to  $w[i..i + 2^k - 1]$  the successor of  $j$  and, respectively, the predecessor of  $j + c2^k - 1$  and then return a compact representation of the occurrences of  $w[i..i + 2^k - 1]$  between these two values. This representation can be obtained in  $\mathcal{O}(c)$  time. We just have to detect those occurrences that form a run; this can be done with a constant number of *LCP* queries. Indeed, for two consecutive occurrences, we compute the length of their overlap, which gives us a period of  $w[i..i + 2^k - 1]$ . Then we look in  $w$  to see how long can the run with this period be extended to the right, which gives us the number of occurrences of  $w[i..i + 2^k - 1]$  in that run. As their starting positions form an arithmetic progression, we can represent them compactly. So, we return the representation of the occurrences of  $w[i..i + 2^k - 1]$  from this run, and then move directly to the first occurrence of  $w[i..i + 2^k - 1]$  appearing after this run and still in the desired range; as there are at most  $\mathcal{O}(c)$  runs and separate occurrences of the given basic factor in the desired range, the conclusion follows.  $\square$

In the following we make use of a result by Gawrychowski (2011), where it is shown that one can also construct data structures that allow the fast identification of the suffixes of a word that start with a given basic factor.

**Lemma 2** (Gawrychowski (2011)) *A word  $w$  of length  $n$  can be processed in  $\mathcal{O}(n)$  time such that given any basic factor  $w[i..i + 2^k - 1]$ , with  $k \geq 0$ , we can retrieve in  $\mathcal{O}(1)$  time the range of the suffix array of  $w$  of suffixes starting with  $w[i..i + 2^k - 1]$ . Equivalently, we can find the node of the suffix tree of  $w$  closest to the root, such that the label of the path from the root to that node has the prefix  $w[i..i + 2^k - 1]$ .*

We can now show the following lemma.

**Lemma 3** *Given a word  $v$ ,  $|v| = \alpha \log n$ , we can process  $v$  in time  $\mathcal{O}(\alpha \log n)$  time such that given any basic factor  $y = v[j \cdot 2^k + 1..(j+1)2^k]$ , with  $j, k \geq 0$  and  $j2^k + 1 > (\alpha - 1) \log n$ , we can find in  $\mathcal{O}(\alpha)$  time  $\mathcal{O}(\alpha)$  bit-sets, each storing  $\mathcal{O}(\log n)$  bits, characterising all the occurrences of  $y$  in  $v$ .*

**Proof:** We first show how the proof works for  $\alpha = 1$ .

We first build the suffix tree for  $v$  in  $\mathcal{O}(\alpha \log n)$  time (see Farach (1997)). We further process this suffix tree such that we can find in constant time, for each factor  $v[j \cdot 2^k + 1..(j+1)2^k]$ , the node of the suffix tree which is closest to the root with the property that the label of path from the root to it starts with  $v[j \cdot 2^k + 1..(j+1)2^k]$ . According to Lemma 2, this can be done in linear time.

Now, we augment the suffix tree in such a manner that for each node we store an additional bit-set, indicating the positions of  $v$  where the word labelling the path from the root to the respective node occurs. Each of these bit-sets, of size  $\mathcal{O}(\log n)$ , can be stored in constant space; indeed, each  $\log n$  block of bits from the bit-set can be seen in fact as a number between 1 and  $n$  so we only need to store a constant number of numbers smaller than  $n$ ; in our model, each such number fits in a memory word. Computing the bit-sets can be done in a bottom up manner in linear time: for a node, we need to make the union of the bit-sets of its children, and this can be done by doing a bitwise **or** operation between all the bit-sets corresponding to the children. So, now, checking the bit-set associated to the lowest node of the suffix tree such that the label of the path from the root to that node starts with  $v[j \cdot 2^k + 1..(j+1)2^k]$  we can immediately output a representation of this factor's occurrences in  $v$ .

This concludes the proof for  $\alpha = 1$ .

For  $\alpha > 1$ , we just have to repeat the algorithm in the previous proof for the words  $v[i \log n + 1..(i+2) \log n]v[(\alpha - 1) \log n + 1..\alpha \log n]$ , for  $0 \leq i \leq \alpha - 2$ , which allows us to find all the occurrences of the basic factors of  $v[(\alpha - 1) \log n + 1..\alpha \log n]$  in  $v$ . The time is clearly  $\mathcal{O}(\alpha \log n)$ . □

Note that each of the bit-sets produced in the above lemma can be stored in a constant number of memory words in our model of computation. Essentially, this lemma states that we can obtain in  $\mathcal{O}(\alpha \log n)$  time a representation of size  $\mathcal{O}(\alpha)$  of all the occurrences of  $y$  in  $v$ .

**Remark 2** *By Lemma 2, given a word  $v$ ,  $|v| = \alpha \log n$ , and a basic factor  $y = v[j \cdot 2^k + 1, (j+1)2^k]$ , with  $j, k \geq 0$  and  $j2^k + 1 > (\alpha - 1) \log n$ , we can produce  $\mathcal{O}(\alpha)$  bit-sets, each containing exactly  $\mathcal{O}(\log n)$  bits, characterising all the occurrences of  $y$  in  $v$ . Let us also assume that we have access to all values  $\log x$  with  $x \leq n$  (which can be ensured by a  $\mathcal{O}(n)$  preprocessing). Now, using the bit-sets encoding the occurrences of  $y$  in  $v$  and given a factor  $z$  of  $v$ ,  $|z| = c|y|$  for some  $c \geq 1$ , we can obtain in  $\mathcal{O}(c)$  time the occurrences of  $y$  in  $z$ : the positions (at most  $c$ ) where  $y$  occurs outside a run and/or at most  $c$  runs containing the occurrences of  $y$ . Indeed, the main idea is to select by bitwise operations on the bit-sets encoding the factors of  $v$  that overlap  $z$  the positions where  $y$  occurs (so the positions with an 1). For each two consecutive such occurrences of  $y$  we detect whether they are part of a run in  $v$  (by LCP-queries on  $v$ ) and then skip over all the occurrences of  $y$  from that run (and the corresponding parts of the bit-sets) before looking again for the 1-bits in the bit-sets.*

Some of our solutions rely on an efficient solution for the *disjoint set union-find* problem. This problem asks to maintain a family consisting initially of  $d$  disjoint singleton sets from the universe  $U = [1, n + 1)$  (shorter for  $\{1, \dots, n\}$ ) so that given any element we can locate its current set and return the minimal (and/or the maximal) element of this set (operation called *find-query*) and we can merge two disjoint sets into one (operation called *union*). In our framework, we know from the beginning the pairs of elements

whose corresponding sets can be joined. Under this assumption, a data-structure fulfilling the above requirements can be constructed in  $\mathcal{O}(d)$  time such that performing a sequence of  $m$  find and union operations takes  $\mathcal{O}(m)$  time in the computation model we use (see Gabow and Tarjan (1983)). As a particular case, this data structure solves with  $\mathcal{O}(n)$  preprocessing time and  $\mathcal{O}(1)$  amortised time per operation the *interval union-find* problem, which asks to maintain a partition of the universe  $U = [1, n+1)$  into a number of  $d$  disjoint intervals, so that given an element of  $U$  we can locate its current interval, and we can merge two adjacent intervals of the partition.

**Remark 3** *As a first consequence of the algorithms of Gabow and Tarjan (1983), we recall a folklore result, stating that we can process in  $\mathcal{O}(n)$  time a weighted tree with  $\mathcal{O}(n)$  nodes and all weights in  $\mathcal{O}(n)$  so that we can compute off-line, also in linear time, the answer to  $\mathcal{O}(n)$  weighted level ancestor queries on the nodes of this tree (where such a query asked for the first node on a path from a given node to the root such that the sum of the weights of the edges between the given and the returned node is greater than a given weight).*

In the solutions of both tasks of Problem 1, we use the following variant of the interval union-find problem.

**Remark 4** *Let  $U = [1, n+1)$ . We are given integers  $d, \ell > 0$  and for each  $i \leq \ell$  a partition  $\mathcal{P}_i$  of  $U$  in  $d$  intervals, a sequence of  $d$  find-queries and a sequence of  $d$  union-operations to be performed alternatively on  $\mathcal{P}_i$  (that is, we perform one find query then the one union operation, and so on, in the order specified in the sequences of union and find-queries). We can maintain the structure and obtain the answers to all find-queries in  $\mathcal{O}(n + \ell d)$  total time.*

*To see why this holds, assume that we are given the sequence  $Q_i = \langle x_1^i, \dots, x_d^i \rangle$  of  $d$  find-queries and the sequence  $U_i = \langle y_1^i, \dots, y_d^i \rangle$  of  $d$  union-operations to be performed alternatively on this partition. As defined above, first we perform the first find query in which we search for the interval containing  $x_1$ , then the first union between the interval ending on  $y_1$  and the one starting on  $y_1$ , then the second find query on the updated partition, in which we search the interval containing  $x_2$ , and so on. In the following we show how to obtain the answers to all find-queries in  $\mathcal{O}(n + \ell d)$  time.*

*As a first step, we sort in time  $\mathcal{O}(n + \ell d)$  all the ends of the intervals in  $\mathcal{P}_i$  and the elements of  $Q_i$  for all  $i$  at once, using radix-sort. Then we separate them (now sorted) according to the partition they correspond to: so, we have for each  $i$  an ordered list of the ends of the intervals of  $\mathcal{P}_i$  and the elements of  $Q_i$ . Now we basically have all the data needed to be able to run the algorithms of Gabow and Tarjan (1983) in  $\mathcal{O}(d)$  time for each partition (e.g., the time used in our setting by the sub-routines defined and used by Gabow and Tarjan (1983) is exactly just the one the required to apply the results of that paper, for each of the partitions). In conclusion, the total time needed to process one partition and to perform all the find queries and union operations on it is  $\mathcal{O}(d)$ . This adds up to a total time of  $\mathcal{O}(n + \ell d)$ .*

We further give another lemma related to the union-find data structure.

**Lemma 4** *Let  $U = [1, n+1)$ . We are given  $k$  intervals  $I_1, \dots, I_k$  included in  $U$ . Also, for each  $j \leq k$  we are given a positive integer  $g_j \leq n$ , the weight of the interval  $I_j$ . We can compute in  $\mathcal{O}(n + k)$  time the values  $H[i] = \max\{g_j \mid i \in I_j\}$  (or, alternatively,  $h[i] = \min\{g_j \mid i \in I_j\}$ ) for all  $i \leq n$ .*

**Proof:** Let us assume that  $I_j = [a_j, b_j)$  where  $1 \leq a_j \leq n$  and  $1 \leq b_j \leq n+1$ , for all  $1 \leq j \leq k$ . We first show how the array  $H[\cdot]$  is computed.



We sort the intervals  $I_1, I_2, \dots, I_k$  with respect to their starting positions  $a_j$ , for  $1 \leq j \leq k$ . Then, we produce for each  $g$  from 1 to  $n$  the list of the intervals  $I_j$  that have  $g_j = g$  (again, sorted by their starting positions). Using radix-sort we can achieve this in time  $\mathcal{O}(n+k)$ . Further, we set up a disjoint set union-find data structure for the universe  $U = [1, n]$ . Initially, the sets in our structure are the singletons  $\{1\}, \dots, \{n\}$ ; the only unions that we can make while maintaining this structure are between the set containing  $i$  and the set containing  $i+1$ , for all  $1 \leq i \leq n-1$ . Therefore, we can think that our structure only contains intervals  $[a, b)$ ; initially, we have in our union-find data structures, for all  $i \leq n$ , the sets  $\{i\} = [i, i+1)$ . Now, we process the intervals that have weight  $g$ , for each  $g$  from  $n$  to 1 in decreasing order; that is, the intervals are considered in decreasing order of their weight. So, assume that we process the input intervals of weight  $g$ . Assume that at this moment,  $U$  is partitioned in some intervals (some of them just singletons); initially, as mentioned above, we only have singletons  $[i, i+1)$ . Let  $I_j = [a, b)$  be an input interval of weight  $g$ . Let now  $\ell = a$  (in the following,  $\ell$  will take different values between  $a$  and  $b$ ). We locate the interval  $[c_1, c_2)$  (of the union-find structure we maintain) where  $\ell$  is located; if this is a singleton and  $H[\ell]$  was not already set, then we set  $H[\ell] = g$ . Further, unless  $\ell$  is  $a$ , we merge the interval containing  $\ell$  to that containing  $\ell-1$  and set  $\ell = c_2$ ; if  $\ell = a$  we just set  $\ell = c_2$ . In both cases, we repeat the procedure while  $\ell < b$ . We process in this manner all the intervals  $I_j$  with  $g_j = g$ , in the order given by their starting positions, then continue and process the intervals with weight  $g-1$ , and so on.

The algorithm computes  $H[\cdot]$  correctly. Indeed, we set  $H[\ell] = g$  when we process an interval  $I_j$  of weight  $g$  that contains  $\ell$ , and no other interval of greater weight contained  $\ell$ . To see the complexity of the algorithm we need to count the number of union and find operations we make. First, we count the number of union operations. For this, it is enough to note that for each element  $\ell$  we might make at most 2 unions: one that unites the singleton interval  $[\ell, \ell+1)$  to the interval of  $\ell-1$ , which has the form  $[a, \ell)$  for some  $a$ , and another one that unites the interval of  $\ell+1$  to the one of  $\ell$ . So, this means that we make at most  $\mathcal{O}(n)$  union operations. For the find operations, we just have to note that when an interval  $I_j$  is processed the total number of finds is  $\mathcal{O}(|I_j \setminus U| + 2)$ , where  $U$  is the union of the intervals that were processed before  $I_j$ . This shows that the total number of find operations is  $\mathcal{O}(k + |\cup_{j=1,k} I_j|) = \mathcal{O}(n+k)$ .

By the results of Gabow and Tarjan (1983), as we know the structure of the unions we can make forms a tree, our algorithm runs in  $\mathcal{O}(n+k)$  time.

The computation of  $h[\cdot]$  is similar. The only difference is that we consider the intervals in increasing order of their weight.  $\square$

We conclude this section with a lemma that will be used in the solutions of Problem 3. It shows how to compute in linear time the length of the longest square centred at each position of a given word.

**Lemma 5** *Given a word  $w$  of length  $n$  we can compute in  $\mathcal{O}(n)$  time the values  $SC[i] = \max\{|u| \mid u \text{ is both a suffix of } w[1..i-1] \text{ and a prefix of } w[i..n]\}$ .*

**Proof:** Note that each square  $u^2$  occurring in a word  $w$  is part of a maximal run  $w[i'..j'] = p^\alpha p'$ , where  $p$  is primitive and  $p'$  is a prefix of  $p$ , and  $u = q^\ell$ , where  $q$  is a cyclic shift of  $p$  (i.e.,  $q$  is a factor of  $p^2$  of length  $|p|$ ) and  $\ell \leq \frac{\alpha}{2}$ .

So, if we consider a maximal run  $r = p^\alpha p'$  and some  $\ell \leq \frac{\alpha}{2}$ , we can easily detect the possible centre positions of the squares having the form  $(q^\ell)^2$  contained in this run, with  $q$  a cyclic shift of  $p$ . These positions occur consecutively in the word  $w$ : the first is the  $(|p|\ell+1)^{th}$  position of the run, and the last is the one where the suffix of length  $|p|\ell$  of the run starts. So they form an interval  $I_{r,\ell}$  and we associate to this interval the weight  $g_{r,\ell} = |p|\ell$  (i.e., the length of an arm of the square). In this way, we define  $\mathcal{O}(n)$

intervals (as their number is upper bounded by the sum of the exponents of the maximal runs of  $w$ ), all contained in  $[1, n+1)$ , and each interval having a weight between 1 and  $n$ . By Lemma 4, we can process these intervals so that we can determine for each  $i \in [1, n+1)$  the interval of maximum weight containing  $i$ , or, in other words, the maximum length  $SC[i]$  of a square centred on  $i$ . This procedure runs in  $\mathcal{O}(n)$  time.  $\square$

It is worth noting that using the same strategy as in the proof of Lemma 5, one can detect the minimum length of a square centred at each position of a word (also called the local period at that position) in linear time. Indeed, in this case we note that the square of minimum length centred at some position of the input word must be primitively rooted. Then, we repeat the same strategy as in the proof of Lemma 5, but only construct the intervals  $I_{r,\ell}$  for the case when  $\ell = 1$  (i.e., the intervals  $I_{r,1}$ ) for all runs  $r$ , with the corresponding weights. Then we use Lemma 4 for these intervals to determine the interval of minimum weight that contains each position of the word. This leads to an alternative solution to the problem of computing the local periods of a word, solved by Duval et al. (2003). Compared to the solution of Duval et al. (2003), ours uses a relatively involved data structures machinery (disjoint sets union-find structures, as well as an algorithm finding all runs in a word), but is much shorter and seems conceptually simpler as it does not require a very long and detailed combinatorial analysis of the properties of the input word.

The same strategy allows solving the problem of computing in linear time, for integer alphabets, the length of the shortest (or longest) square ending (or starting) at each position of a given word; this improves the results by Kosaraju (1994); Xu (2010), where such a result was only shown for constant size alphabets. Let us briefly discuss the case of finding the shortest square ending at each position of the word; such squares are, clearly, primitively rooted. So, just like in the case of searching squares centred at some position, we can determine for each run an interval of positions where a square having exactly the same period as the run may end. The period of each run becomes the weight of the interval associated to that run. Then we use again Lemma 4 for these intervals to determine the interval of minimum weight that contains each position of the word. Thus, we determine for each position the shortest square ending on it. In the case when we want the longest square ending at each position, just like in the case of centred squares, we define for each run several intervals. Fortunately, this does not increase asymptotically the complexity.

In the following we apply the machinery developed in this section to identify the longest gapped repeats and palindromes occurring in a word.

### 3 Lower and upper bounded gap


In this section we present the solutions of Problems 1(a) and 1(b). With these problems solved, we can immediately retrieve the longest gapped repeat and palindrome contained in the input word of these problems, such that the length gap is between a given lower bound and a given upper bound.

**Theorem 1** *Problem 1(a) can be solved in linear time.*

**Proof:** Let  $\delta = G - g$ ; for simplicity, assume that  $n$  is divisible by  $\delta$ . Further, for  $1 \leq i \leq n$ , let  $u$  be the longest factor which is a prefix of  $w[i..n]$  such that  $u^R$  is a suffix of some  $w[1..k]$  with  $g < i - k \leq G$ ; then,  $B[i]$  is the rightmost position  $j$  such that  $g < i - j \leq G$  and  $u^R$  is a suffix of  $w[1..j]$ . Knowing  $B[i]$  means knowing  $LPrF_{g,G}[i]$ : we just have to return  $LPrF(w[i..n], w[1..j]^R)$ .

We split the set  $\{1, \dots, n\}$  into  $\frac{n}{\delta}$  ranges of consecutive numbers:  $I_0 = \{1, 2, \dots, \delta\}$ ,  $I_1 = \{\delta + 1, \delta + 2, \dots, 2\delta\}$ , and so on. Note that for all  $i$  in some range  $I_k = \{k\delta + 1, k\delta + 2, \dots, (k+1)\delta\}$  from those

defined above, there are at most three consecutive ranges where some  $j$  such that  $g < i - j \leq G$  may be found: the range containing  $k\delta - G + 1$ , the range containing  $(k + 1)\delta - g - 1$ , and the one in between these two (i.e., the range containing  $k\delta + 1 - g$ ). Moreover, for some fixed  $i$ , we know that we have to look for  $B[i]$  in the interval  $\{i - G, \dots, i - g - 1\}$ , of length  $\delta$ ; when we search  $B[i + 1]$  we look at the interval  $\{i + 1 - G, \dots, i - g\}$ . So, basically, when trying to find  $B[i]$  for all  $i \in I_k$ , we move a window of length  $\delta$  over the three ranges named above, and try to find for every content of the window (so for every  $i$ ) its one element that fits the description of  $B[i]$ . The difference between the content of the window in two consecutive steps is not major: we just removed an element and introduced a new one. Also, note that at each moment the window intersects exactly two of the aforementioned three ranges. We try to use these remarks, and maintain the contents of the window such that the update can be done efficiently, and the values of  $B[i]$  (and, implicitly,  $LPrF_{g,G}[i]$ ) can be, for each  $i \in I$ , retrieved very fast. Intuitively, grouping the  $i$ 's on ranges of consecutive numbers allows us to find the possible places of the corresponding  $B[i]$ 's for all  $i \in I_k$  in  $\mathcal{O}(\delta)$  time.



**Fig. 2:** Proof of Theorem 1: Construction of the ranges  $J_1, J_2, J_3$  for some range  $I_k$ , and of the sliding window  $J'_i \cup J''_i$  for some  $i \in I_k$

We now go into more details. As we described in the preliminaries, let  $\mathcal{L}$  be the lexicographically ordered list of the suffixes  $w[i..n]$  of  $w$  and of the mirror images  $w[1..i]^R$  of the prefixes of  $w$  (which correspond to the suffixes of  $w^R$ ). For the list  $\mathcal{L}$  we compute the arrays  $Rank[\cdot]$  and  $Rank_R[\cdot]$ . We use the suffix array for  $w0w^R$  to produce for each of the ranges  $I_k$  computed above the set of suffixes of  $w$  that start in the respective range (sorted lexicographically, in the order they appear in the suffix array) and the set of prefixes of  $w$  ending in  $I$ , ordered lexicographically with respect to their mirror image.

We consider now one of the ranges of indexes  $I_k = \{k\delta + 1, k\delta + 2, \dots, (k + 1)\delta\}$  from the above, and show how we can compute the values  $B[i]$ , for all  $i \in I_k$ . For some  $i \in I_k$  we look for the maximum  $\ell$  such that there exists a position  $j$  with  $w[j - \ell + 1..j]^R = w[i..i + \ell - 1]$ , and  $i - G \leq j \leq i - g$ . As already explained, for the  $i$ 's of  $I_k$  there are three consecutive ranges where the  $j$ 's corresponding to the  $i$ 's of  $I_k$  may be found; we denote them  $J_1, J_2, J_3$ .

Now, for an  $i \in I_k$  we have that  $B[i]$  (for which  $g < i - B[i] \leq G$ ) belongs to  $J'_i \cup J''_i$ , where  $J'_i$  is an interval starting with  $i - G$  and extending until the end of the range that contains  $i - G$  (which is one of the ranges  $J_1, J_2, J_3$ ) and  $J''_i$  is an interval ending with  $i - g - 1$ , which starts at the beginning of the range that contains  $i - g - 1$  (which is the range that comes next after the one containing  $i - G$ ). Referencing back to the intuitive explanation we gave at the beginning of this proof,  $J'_i \cup J''_i$  is the window we use to locate the value of  $B[i]$  for an  $i \in I_k$ .

To compute  $B[i]$  for some  $i$ , take  $f_i \in J'_i$  such that  $LCP(w[1..f_i]^R, w[i..n]) \geq LCP(w[1..j]^R, w[i..n])$  for all  $j' \in J'_i$ . Similarly, take  $s_i \in J''_i$  such that for all  $j' \in J''_i$  we have  $LCP(w[1..s_i]^R, w[i..n]) \geq$

$LCP(w[1..j]^R, w[i..n])$ . Once  $s_i$  and  $f_i$  computed, we set  $B[i] = s_i$  if  $LCP(w[1..s_i]^R, w[i..n]) \geq LCP(w[1..f_i]^R, w[i..n])$ ; we set  $B[i] = f_i$ , otherwise. So, in order to compute, for some  $i$ , the value  $B[i]$ , that determines  $LPrF_{g,G}[i]$ , we first compute  $f_i$  and  $s_i$ .

We compute for all the indices  $i \in I_k$ , considered in increasing order, the values  $f_i$ . We consider for each  $i \in I_k$  the interval  $J'_i$  and note that  $J'_i \setminus J'_{i+1} = \{i - G\}$ , and, if  $J'_i$  is not a singleton (i.e.,  $J'_i \neq \{i - G\}$ ) then  $J'_{i+1} \subset J'_i$ . If  $J'_i$  is a singleton, then  $J'_{i+1}$  is, in fact, one of the precomputed range  $I_p$ , namely the one which starts on position  $i + 1 - G$  (so,  $p = \frac{i-G}{\delta}$ ).

These relations suggest the following approach. We start with  $i = k\delta + 1$  and consider the set of words  $w[1..j]^R$  with  $j \in J'_i$ ; this set can be easily obtained in  $\mathcal{O}(\delta)$  time by finding first the range  $J_1$  in which  $i - G$  is contained (which takes  $\mathcal{O}(1)$  time, as  $J_1 = I_p$  for  $p = \lfloor \frac{i-G}{\delta} \rfloor$ ), and then selecting from  $J_1$  of the set of prefixes  $w[1..d]$  of  $w$ , ending in  $J_1$  with  $d \geq i - G$  (ordered lexicographically with respect to their mirror image). The ranks corresponding to these prefixes in the ordered list  $\mathcal{L}$  (i.e., the set of numbers  $Rank_R[d]$ ) define a partition of the universe  $U = [0, 2n + 1)$  in at most  $\delta + 2$  disjoint intervals. So, we can maintain an interval union-find data structures like in Remark 4, where the ranks are seen as limits of the intervals in this structure. We assume that the intervals in our data structure are of the form  $[a, b)$ , with  $a$  and  $b$  equal to some  $Rank_R[d_a]$  and  $Rank_R[d_b]$ , respectively. The first interval in the structure is of the form  $[0, a)$ , while the last is of the form  $[b, 2n + 1)$ . We now find the interval to which  $Rank[i]$  belongs; say that this is  $[a, b)$ . This means that the words  $w[1..d_a]^R$  and  $w[1..d_b]^R$  are the two words of  $\{w[1..d]^R \mid d \in J'_i\}$  which are closest to  $w[i..n]$  lexicographically ( $w[1..d_a]^R$  is lexicographically smaller,  $w[1..d_b]^R$  is greater). Clearly,  $f_i = d_a$  if  $LCP(w[1..d_a]^R, w[i..n]) \geq LCP(w[1..d_b]^R, w[i..n])$  and  $f_i = d_b$ , otherwise (in case of a tie, we take  $f_i$  to be the greater of  $d_a$  and  $d_b$ ). So, to compute  $f_i$  we query once the union-find data structure to find  $a$  and  $b$ , and the corresponding  $d_a$  and  $d_b$ , and then run two more  $LCP$  queries.

When moving on to compute  $f_{i+1}$ , we just have to update our structure and then run the same procedure. Now,  $i - G$  is no longer a valid candidate for  $f_{i+1}$ , and it is removed from  $J'_i$ . So we just delete it from the interval union-find data structure, and merge the interval ending right before  $Rank_R[i - G]$  and the one starting with  $Rank_R[i - G]$ . This means one union operation in our interval union-find structure. Then we proceed to compute  $f_{i+1}$  as in the case of  $f_i$ .

The process continues until  $J'_i$  is a singleton, so  $f_i$  equals its single element.

Now,  $i - G$  is the last element of one of the ranges  $J_1, J_2$ , or  $J_3$ ; assume this range is  $I_p$ . So far, we performed alternatively at most  $\delta$  find queries and  $\delta$  union operations on the union-find structure. Now, instead of updating this structure, we consider a new interval partition of  $[0, 2n + 1)$  induced by the ranks of the  $\delta$  prefixes ending in  $I_{p+1}$ . When computing the values  $f_i$  for  $i \in I_k$  we need to consider a new partition of  $U$  at most once: at the border between  $J_1$  and  $J_2$ .

It is not hard to see from the comments made in the above presentation that our algorithm computes  $f_i \in I_k$  correctly. In summary, in order to compute  $f_i$ , we considered the elements  $i \in I_k$  from left to right, keeping track of the left part  $J'_i$  of the window  $J'_i \cup J''_i$  while it moved from left to right through the ranges  $J_1, J_2$  and  $J_3$ . Now, to compute the values  $s_i$  for  $i \in I_k$ , we proceed in a symmetric manner: we consider the values  $i$  in decreasing order, moving the window from right to left, and keep track of its right part  $J''_i$ .

As already explained, by knowing the values  $f_i$  and  $s_i$  for all  $i \in I_k$  and for all  $k$ , we immediately get  $B[i]$  (and, consequently  $LPrF_{g,G}[i]$ ) for all  $i$ .

We now evaluate the running time of our approach. We can compute, in  $\mathcal{O}(n)$  time, from the very beginning of our algorithm the partitions of  $[1, 2n - 1)$  we need to process (basically, for each  $I_k$  we

find  $J_1$ ,  $J_2$  and  $J_3$  in constant time, and we get the three initial partitions we have to process in  $\mathcal{O}(\delta)$  time), and we also know the union-operations and find-queries that we will need to perform for each such partition (as we know the order in which the prefixes are taken out of the window, so the order in which the intervals are merged). In total we have  $\mathcal{O}(n/\delta)$  partitions, each having initially  $\delta + 2$  intervals, and on each we perform  $\delta$  find-queries and  $\delta$ -union operations. So, by Remark 4, we can preprocess this data (once, at the beginning of the algorithm) in  $\mathcal{O}(n)$  time, to be sure that the time needed to obtain the correct answers to all the find queries is  $\mathcal{O}(n)$ . So, the total time needed to compute the values  $f_i$  for all  $i \in I_k$  and for all  $k$  is  $\mathcal{O}(n)$ . Similarly, the total time needed to compute the values  $s_i$  for all  $i$  is  $\mathcal{O}(n)$ . Then, for each  $i$  we get  $B[i]$  and  $LPrF_{g,G}[i]$  in  $\mathcal{O}(1)$  time.

Therefore, Problem 1(a) can be solved in linear time.  $\square$

Before giving the formal solution for Problem 1(b), we give a short intuitive description of how this algorithm is different from the previous one. In this case, when trying to construct the repeat  $uvu$  with the longest arm occurring at a certain position, we need to somehow restrict the range where its left arm may occur. To this end, we restrict the length of the arm, and, consequently, search for  $u$  with  $2^k \leq |u| \leq 2^{k+1}$ , for each  $k \leq \log n$ . Such a factor always start with  $w[i + 1..i + 2^k]$  and may only occur in the factor  $w[i - G - 2^{k+1}..i - g]$ . If  $2^k \geq G - g$ , using the dictionary of basic factors data structures (from Crochemore and Rytter (1991)) and the results by Kociumaka et al. (2012), we get a constant size representation of the occurrences of  $w[i + 1..i + 2^k]$  in that range (which are a constant number of times longer than the searched factor), and then we detect which one of these occurrences produces the repeat  $uvu$  with the longest arm. If  $2^k < G - g$ , we use roughly the same strategy to find the repeat  $uvu$  with the longest arm and  $u$  starting in a range of radius  $2^{k+1}$  centred around  $i - G$  or in a range of length  $2^{k+1}$  ending on  $i - g$  (again, these ranges are just a constant number of times longer than the searched factor). To detect a repeat starting between  $i - G + 2^{k+1}$  and  $i - g - 2^{k+1}$  we use the strategy from the solution of Problem 1(a); in that case, we just have to return the longest common prefix of  $w[i..n]$  and the words  $w[j..n]$  with  $i - G + 2^{k+1} \leq j \leq i - g - 2^{k+1}$ . Overall, this approach can be implemented to work in  $\mathcal{O}(n \log n)$  time.

**Theorem 2** *Problem 1(b) can be solved in  $\mathcal{O}(n \log n)$  time.*


**Proof:** For  $1 \leq i \leq n$ , let  $B[i]$  denote the value  $j$  such that  $w[j..j + LPrF_{g,G}[i] - 1] = w[i..i + LPrF_{g,G}[i] - 1]$  and  $g < i - (j + LPrF_{g,G}[i] - 1) \leq G$ . In other words, to define  $B[i]$ , let  $u$  be the longest factor which is both a prefix of  $w[i..n]$  and a suffix of some  $w[1..k]$  with  $g < i - k \leq G$ ; then,  $B[i]$  is the rightmost position  $j$  such that  $g < i - (j + |u| - 1) \leq G$  and  $u$  is a suffix of  $w[1..j + |u| - 1]$ . Clearly, knowing  $B[i]$  means knowing  $LPrF_{g,G}[i]$ .

Intuitively, computing  $B[i]$  is harder in this case than it was in the case of the solution of Problem 1(a). Now we have no real information where  $B[i]$  might be, we just know the range of  $w$  where the longest factor that occurs both at  $B[i]$  and at  $i$  ends. So, to determine  $B[i]$  we try different variants for the length of this factor, and see which one leads to the right answer.

Let  $\delta = G - g$  and  $k_0 = \lceil \log \delta \rceil$ ; assume w.l.o.g. that  $n$  is divisible by  $\delta$ .

As already noted, the solution used in the case of gapped palindromes in the proof of Theorem 1 does not work anymore in this case: we do not know for a given  $i$  a range in which  $B[i]$  is found. So, we try to restrict the places where  $B[i]$  can occur. We split our discussion in two cases.

In the first case, we try to find, for each  $i$  and each  $k \geq k_0$ , the factor  $uvu$  with the longest  $u$ , such that the second  $u$  occurs at position  $i$  in  $w$ ,  $2^k \leq |u| < 2^{k+1}$  and  $g < |v| \leq G$ . Clearly,  $u$  should start with the



**Fig. 3:** Proof of Theorem 2: Occurrences of  $w[i..i + 2^k - 1]$  inside  $w[i - G - 2^{k+1}..i - g - 1]$ : one separate occurrence, and a run containing three occurrences. The separate occurrence can be prolonged to produce a gapped palindrome  $uvu$ , which is not, however, long enough as it does not reach the range between  $i - G$  and  $i - g$ . The rightmost occurrence in the run produces the gapped repeat  $w[i..i + 2^k - 1]v'w[i..i + 2^k - 1]$ , which fulfills the conditions imposed on the gap.

basic factor  $w[i..i + 2^k - 1]$ , and the left arm  $u$  in the factor  $uvu$  we look for should have its prefix of length  $2^k$  (so, a factor equal to the basic factor  $w[i..i + 2^k - 1]$ ) contained in the factor  $w[i - G - 2^{k+1}..i - g - 1]$ , whose length is

$$G - g + 2^{k+1} \leq 2^{k_0+1} + 2^{k+1} - 1 \leq 2^{k+2} - 1 \leq 4 \cdot 2^k.$$

So, by Remark 1, using the dictionary of basic factors we can retrieve a compact representation of the occurrences of  $w[i..i + 2^k - 1]$  in  $w[i - G - 2^{k+1}..i - g - 1]$ : these consist in a constant number of isolated occurrences and a constant number of runs. For each of the isolated occurrences  $w[j'..j' + 2^k - 1]$  we compute  $LCP(j', i)$  and this gives us a possible candidate for  $u$ ; we measure the gap between the two occurrences of  $u$  (the one at  $j'$  and the one at  $i$ ) and if it is between smaller than  $G$ , we store this  $u$  as a possible solution to our problem (we might have to cut a suffix of  $u$  so that the gap is also longer than  $g$ ). Further, each of the runs has the form  $p^\alpha p'$ , where  $|p|$  is the period of  $w[i..i + 2^k - 1]$  and  $p'$  is a prefix of  $p$  (which varies from run to run); such a run cannot be extended to the right: either the period breaks, or it would go out of the factor  $w[i - G - 2^{k+1}..i - g - 1]$ . Using a  $LCP$  query we find the longest factor  $p^\beta p''$  with period  $|p|$ , occurring at position  $i$ . For a run  $p^\alpha p'$ , the longest candidate for the first  $u$  of the factor  $uvu$  we look for starts with  $p^\gamma p'''$  where  $\gamma = \min\{\alpha, \beta\}$  and  $p'''$  is the shortest of  $p'$  and  $p''$ ; if  $p' = p''$ , then this factor is extended with the longest factor that occurs both after the current run and after  $p^\beta p''$  and does not overlaps the minimal gap (i.e., end at least  $g$  symbols before  $i$ ). This gives us a candidate for the factor  $u$  we look for (provided that the gap between the two candidates for  $u$  we found is not too large). In the end, we just take the longest of the candidates we identified (in case of ties, we take the one that starts on the rightmost position), and this gives the factor  $uvu$  with the longest  $u$ , such that the second  $u$  occurs at position  $i$  in  $w$ ,  $2^k \leq |u| < 2^{k+1}$  and  $g < |v| \leq G$ .

Iterating this process for all  $k$  and  $i$  as above, we obtain for each  $i$  the factor  $uvu$  with the longest  $u$ , such that the second  $u$  occurs at position  $i$  in  $w$ ,  $2^{k_0} \leq |u|$  and  $g < |v| \leq G$ . By Remark 1, the time spent in this computation for some  $i$  and  $k$  is constant, so the overall time used in the above computation is  $\mathcal{O}(n \log n)$ . Clearly, if for some  $i$  we found such a factor  $u$ , then this gives us both  $B[i]$  and  $LPF_{g,G}[i]$ ; if not, we continue as follows.

In the second case, we try to identify for each  $i$  the factor  $uvu$  with the longest  $u$ , such that the second  $u$  factor occurs at position  $i$  in  $w$ ,  $u < 2^{k_0}$  and  $g < |v| \leq G$ . Again, we consider each  $k \leq k_0 - 1$  separately and split the discussion in three cases.

Firstly, for each  $i$ , we find the factor  $uvu$  with the longest  $u$ , such that the second  $u$  occurs at position  $i$  in  $w$ ,  $2^k \leq |u| < 2^{k+1}$ ,  $g < |v| \leq G$ , and the first  $u$  has its prefix of length  $2^k$  in the factor  $w[i - G -$

$2^{k+1}..i - G + 2^k]$ , whose length is  $2^k + 2^{k+1} + 1 \leq 4 \cdot 2^k$ . This can be done similarly to the above (report all the occurrences of  $w[i..i + 2^k - 1]$  in that range, and try to extend them to get  $u$ ); for all  $i$  and all  $k$  takes  $\mathcal{O}(n \log \delta)$  time.

Secondly, for each  $i$ , we find the factor  $uvu$  with the longest  $u$ , such that the second  $u$  occurs at position  $i$  in  $w$ ,  $2^k \leq u < 2^{k+1}$ ,  $g < |v| \leq G$ , and the first  $u$  has its prefix of length  $2^k$  in the factor  $w[i - g - 2^{k+1}..i - g - 1]$ , whose length is  $2^{k+1}$ . Again, this can be done like above, and for all  $i$  and all  $k$  takes  $\mathcal{O}(n \log \delta)$  time.

The third and more complicated subcase is when the first  $u$  starts in the factor  $w[i - G..i - g - 2^{k+1} - 1]$ , of length  $\delta_k = \delta - 2^{k+1}$ . Let  $g_k = g + 2^{k+1} + 1$ ; obviously,  $\delta_k = G - g_k$ . Note that, in this case, every factor of length at most  $2^{k+1}$  starting in  $w[i - G..i - g - 2^{k+1} - 1]$  ends before  $i - g$ , so it is a valid candidate for the  $u$  we look for. In this case we can follow the algorithm from the proof of Theorem 1. We split the set  $\{1, \dots, n\}$  into ranges of consecutive numbers:  $I_0 = \{1, 2, \dots, \delta_k\}$ ,  $I_1 = \{\delta_k + 1, \delta_k + 2, \dots, 2\delta_k\}$ , and so on. For some  $I_\ell = \{\ell\delta_k + 1, \dots, (\ell + 1)\delta_k\}$ , considering all  $i \in I_\ell$  there are three consecutive ranges from the ones defined above, where the first  $u$  of the factor  $uvu$  we look for may occur. The first (leftmost, with respect to its starting position) such range is the one containing  $\ell\delta_k - G$ , and let us denote it  $J_1$ ; the last one (rightmost) is the one that contains  $(\ell + 1)\delta_k - g_k - 1$ , and we denote it by  $J_3$ . Clearly, between the ranges containing  $\ell\delta_k - G$  and  $(\ell + 1)\delta_k - g_k - 1$ , respectively, there is exactly one complete range, call it  $J_2$ .

Moreover, for a precise  $i \in I_\ell$  the possible starting positions of the left arm  $u$  of the repeat  $uvu$  of the type we are searching for (i.e., with the gap between  $g$  and  $G$ ), with the second  $u$  starting on  $i$ , form a contiguous range  $J'_i \cup J''_i$  where  $J'_i$  is an interval starting with  $i - G$  and extending until the end of the range  $J_p$  that contains  $i - G$ , and  $J''_i$  is an interval ending with  $i - g_k - 1$ , contained in  $J_{p+1}$  (when  $p < 3$ ). Like before,  $J'_i \cup J''_i$  can be seen as the content of a window that slides through  $J_1, J_2, J_3$  while searching for  $B[i]$ . We denote by  $f_i$  the position of  $J'_i$  such that  $LCP(f_i, i)$  is maximum among all such positions; we denote by  $s_i$  the position of  $J''_i$  such that  $LCP(s_i, i)$  is maximum among all such positions. Then we just have to check where, at  $f_i$  or  $s_i$ , occurs a longer factor that also occurs at  $i$ . We just explain how to compute  $f_i$ .

We start with  $i = \ell\delta_k + 1$  and consider the set of words  $w[j..n]$  with  $j \in J'_i$ ; this set can be easily obtained in  $\mathcal{O}(\delta)$  time by finding first the range  $J_1$  in which  $i - G$  is contained and then selecting from  $J_1$  of the set of words  $w[d..n]$  of  $w$  starting in  $J_1$  with  $d \geq i - G$  (ordered lexicographically). The ranks corresponding to these suffixes in the suffix array of  $w$  (i.e., the set of numbers  $Rank[d]$ ) define a partition of the universe  $U = [0, n + 1)$  in at most  $\delta_k + 2$  disjoint intervals. So, we can maintain an interval union-find data structures like in Remark 4, where the ranks are seen as limits of the intervals in this structure. We assume that the intervals in our data structure are of the form  $[a, b)$ , with  $a$  and  $b$  integers that are equal to some  $Rank[d_a]$  and  $Rank[d_b]$ . The first interval in the structure is of the form  $[0, a)$ , while the last is of the form  $[b, n + 1)$ . Recall that we want to compute  $f_i$ . To this end, we just have to find the interval to which  $Rank[i]$  belongs; say that this is  $[a, b)$ . This means that the words  $w[d_a..n]$  and  $w[d_b..n]$  are the two words of the set  $\{w[d..n] \mid d \in J'_i\}$  which are closest to  $w[i..n]$  lexicographically ( $w[d_a..n]$  is lexicographically smaller, while  $w[d_b..n]$  is lexicographically greater). Clearly,  $f_i = d_a$  if  $LCP(d_a, i) \geq LCP(d_b, i)$  and  $f_i = d_b$ , otherwise (in case of a tie, we take  $d_a$  if  $d_b < d_a$ , and  $d_b$  otherwise). So, to compute  $f_i$  we have to query once the union-find data structure to find  $a$  and  $b$ , and the corresponding  $d_a$  and  $d_b$ , and then compute the answer to two  $LCP$  queries.

When moving on to compute  $f_{i+1}$ , we just have to update our structure. In this case,  $i - G$  is no longer a valid candidate for  $f_{i+1}$ , as it should be removed from  $J'_i$ . So we just delete it from the interval

union-find data structure, and merge the interval ending right before  $\text{Rank}[i - G]$  and the one starting with  $\text{Rank}[i - G]$ . This means one union operation in our interval union-find structure. Then we proceed to compute  $f_{i+1}$  just in the same manner as in the case of  $f_i$ .

The process continues in this way until we try to compute  $f_i$  for  $J'_i$  being a singleton. This means that  $i - G$  is the last element of one of the ranges  $J_1, J_2$ , or  $J_3$ ; let us assume that this range is  $I_p$  from the ranges defined above. Clearly, this time  $f_i$  is the single element of  $J'_i$ . Until now, we performed alternatively at most  $\delta_k$  find operations and  $\delta_k$  union operations on the union-find data structure. Further, instead of updating the union-find data structure, we consider a new interval partition of  $[0, n+1)$  induced by the ranks of the  $\delta_k$  suffixes starting in  $I_{p+1}$ . Note that when computing the values  $f_i$  for  $i \in I_\ell$  we need to consider a new partition of  $U$  once: at the border between  $J_1$  and  $J_2$ .

Now, by the same arguments as in the proof for gapped palindromes, the process takes  $\mathcal{O}(n)$  in total for all intervals  $I_\ell$  (so, when we iterate  $\ell$ ), for each  $k$ . In this way we find the factor  $uvu$  with the longest  $u$ , such that the second  $u$  occurs at position  $i$  in  $w$ ,  $2^k \leq u < 2^{k+1}$ ,  $g < |v| \leq G$ , and the first  $u$  starts in  $w[i - G..i - g - 2^{k+1} - 1]$ . Iterating for all  $k$ , we complete the computation of  $B[i]$  and  $LPF_{g,G}[i]$  for all  $i$ ; the needed time is  $\mathcal{O}(n \log \delta)$  in this case.

Considering the three cases above leads to finding for each  $i$  the factor  $uvu$  with the longest  $u$ , such that the second  $u$  factor occurs at position  $i$  in  $w$ ,  $u < 2^{k_0}$  and  $g < |v| \leq G$ . The complexity of this analysis is  $\mathcal{O}(n \log \delta)$ . After concluding this, we get the value  $B[i]$  for each  $i$ . Therefore, the entire process of computing  $B[i]$  and  $LPF_{g,G}[i]$  for all  $i$  takes  $\mathcal{O}(n \log n)$  time.  $\square$

We conclude this section with the following consequence of Theorems 1 and 2.

**Theorem 3** *Given  $w$  of length  $n$  and two integers  $g$  and  $G$ , such that  $0 \leq g < G \leq n$ , we can find in linear time the gapped palindrome  $u^Rvu$  occurring in that word with the longest arm  $u$  and  $g \leq |v| < G$ . Given  $w$  of length  $n$  and two integers  $g$  and  $G$ , such that  $0 \leq g < G \leq n$ , we can find in  $\mathcal{O}(n \log n)$  time the gapped repeat  $uvu$  occurring in that word with the longest arm  $u$  and  $g \leq |v| < G$ .*

**Proof:** It is immediate that, after solving Problem 1 for the word  $w$ , we just have to check which is the longest gapped palindrome (respectively, repeat) stored in the array  $LPrF_{g,G}[\cdot]$  (respectively,  $LPF_{g,G}[\cdot]$ ).  $\square$

## 4 Lower bounded gap

To solve Problem 2(a) we need to find, for some position  $i$  of  $w$ , the factor  $w[1..j]^R$  with  $j \leq i - g(i)$  that occurs closest to  $w[1..i]$  in the lexicographically ordered list  $\mathcal{L}$  of all the suffixes  $w[k..n]$  of  $w$  and of the mirror images  $w[1..k]^R$  of its prefixes. In the following we show how to do this for all  $i$  in  $\mathcal{O}(n)$  time, by reducing it to answering  $n$  find queries in an extended interval union-find data structure.

**Theorem 4** *Problem 2(a) can be solved in linear time.*

**Proof:** In a preprocessing step of our algorithm, we produce the suffix array of  $w0w^R$  and the lexicographically ordered list  $\mathcal{L}$  of the suffixes of  $w[i..n]$  of  $w$  and of the mirror images  $w[1..i]^R$  of the prefixes of  $w$  (which correspond to the suffixes of  $w^R$ ). For the list  $\mathcal{L}$  we compute the arrays  $\text{Rank}[\cdot]$  and  $\text{Rank}_R[\cdot]$ .

We first want to find, for each  $i$ , the prefix  $w[1..j]$ , such that  $i - j > g(i)$ ,  $w[1..j]^R$  occurs before  $w[i..n]$  in  $\mathcal{L}$  (i.e.,  $\text{Rank}_R[j] < \text{Rank}[i]$ ), and the length of the common prefix of  $w[1..j]^R$  and  $w[i..n]$  is greater or equal to the length of the common prefix of  $w[1..j']^R$  and  $w[i..n]$  for  $j'$  such that  $\text{Rank}_R[j'] < \text{Rank}[i]$ ;



for the prefix  $w[1..j]$  as above, the length of the common prefix of  $w[1..j]^R$  and  $w[i..n]$  is denoted by  $LPrF_g^<[i]$ , while  $B_{<}[i]$  denotes  $j$ . If  $w[i..n]$  has no common prefix with any factor  $w[1..j]^R$  with  $i - j > g(i)$  and  $Rank_R[j] < Rank[i]$ , then  $LPrF_g^<[i] = 0$  and  $B_{<}[i]$  is not defined; as a convention, we set  $B_{<}[i]$  to  $-1$ .

Afterwards, we compute for each  $i$  the prefix  $w[1..j]$ , such that  $i - j > g(i)$ ,  $w[1..j]^R$  occurs after  $w[i..n]$  in  $\mathcal{L}$  (i.e.,  $Rank_R[j] > Rank[i]$ ), and the length of the common prefix of  $w[1..j]^R$  and  $w[i..n]$  is greater or equal to the length of the common prefix of  $w[1..j']^R$  and  $w[i..n]$  for  $j'$  such that  $Rank_R[j'] > Rank[i]$ ; for the prefix  $w[1..j]$  as above, the length of the common prefix of  $w[1..j]^R$  and  $w[i..n]$  is denoted by  $LPrF_g^>[i]$ , while  $B_{>}[i]$  denotes  $j$ . Clearly,  $LPrF_g[i] = \max\{LPrF_g^<[i], LPrF_g^>[i]\}$ . If  $w[i..n]$  has no common prefix with any factor  $w[1..j]^R$  with  $i - j > g(i)$  and  $Rank_R[j] > Rank[i]$ , then  $LPrF_g^>[i] = 0$  and  $B_{>}[i]$  is set to  $-1$ .

For simplicity, we just present an algorithm computing  $LPrF_g^<[\cdot]$  and  $B_{<}[\cdot]$ . The computation of  $LPrF_g^>[\cdot]$  is performed in a similar way.

The main idea behind the computation of  $LPrF_g^<[i]$ , for some  $1 \leq i \leq n$ , is that if  $w[1..j_1]$  and  $w[1..j_2]$  are such that  $Rank_R[j_2] < Rank_R[j_1] < Rank[i]$  and  $j_1 < j_2 < i$  then definitely  $B_{<}[i] \neq j_2$ . Indeed,  $LCP(w[1..j_2]^R, w[i..n]) < LCP(w[1..j_1]^R, w[i..n])$ , and, moreover,  $i - j_2 < i - j_1$ . So, if  $i - j_2 > g(i)$  then also  $i - j_1 > g(i)$ , and it follows that  $LPrF_g^<[i] \geq LCP(w[1..j_1]^R, w[i..n]) > LCP(w[1..j_2]^R, w[i..n])$ , so  $B_{<}[i]$  cannot be  $j_2$ . This suggests that we could try to construct for each  $i$  an ordered list  $\mathcal{A}_i$  of all the integers  $j \leq n$  such that  $Rank_R[j] < i$  and moreover, if  $j_1$  and  $j_2$  are in  $\mathcal{A}_i$  and  $j_1 < j_2$  then also  $Rank_R[j_1] < Rank_R[j_2]$ .

Now we describe how to implement this. Let us now consider  $i_1$  and  $i_2$  which occur on consecutive positions of the suffix array of  $w$ , such that  $Rank[i_1] < Rank[i_2]$ . The list  $\mathcal{A}_{i_2}$  can be obtained from  $\mathcal{A}_{i_1}$  as follows. We consider one by one, in the order they appear in  $Rank_R[\cdot]$ , the integers  $j$  such that  $Rank[i_1] < Rank_R[j] < Rank[i_2]$ , and for each of them update a temporary list  $\mathcal{A}$ , which initially is equal to  $\mathcal{A}_{i_1}$ . When a certain  $j$  is considered, we delete from the right end of the list  $\mathcal{A}$  (where  $\mathcal{A}$  is ordered increasingly from left to right) all the values  $j' > j$ ; then we insert  $j$  in  $\mathcal{A}$ . When there are no more indices  $j$  that we need to consider, we set  $\mathcal{A}_{i_2}$  to be equal to  $\mathcal{A}$ . It is clear that the list  $\mathcal{A}_{i_2}$  is computed correctly.

Now, for each  $i$  we need to compute the greatest  $j \in \mathcal{A}_i$  such that  $j < i - g(i)$ . As  $\mathcal{A}_i$  is ordered increasingly, we could obtain  $j$  by performing a predecessor search on  $\mathcal{A}_i$  (that is, binary searching the greatest  $j$  of the list, which is smaller than  $i - g(i)$ ), immediately after we computed it, and save the answer in  $B_{<}[i]$ . However, this would be inefficient. Before proceeding, we note that if we compute the lists  $\mathcal{A}_i$  for the integers  $i$  in the order they appear in the suffix array of  $w$ , then it is clear that the time needed to compute all these lists is linear. Indeed, each  $j \leq n$  is introduced exactly once in the temporary list, and then deleted exactly once from it. Doing the above mentioned binary searches would add up to a total of  $\mathcal{O}(n \log n)$ . We can do better than that.

Now we have reduced the original problem to a data-structures problem. We have to maintain an increasingly ordered (from left to right) list  $\mathcal{A}$  of numbers (at most  $n$ , in the range  $\{1, \dots, 2n\}$ , each two different), subject to the following update and query operations. This list can be updated by the following procedure: we are given a number  $j$ , we delete from the right end of  $\mathcal{A}$  all the numbers greater than  $j$ , then we append  $j$  to  $\mathcal{A}$ . By this update, the list remains increasingly ordered. The following queries can be asked: for a given  $\ell$ , which is the rightmost number of  $\mathcal{A}$ , smaller than  $\ell$ ? We want to maintain this list while  $n$  updates are executed, and  $n$  queries are asked at different moments of time. Ideally, the total

time we can spend in processing the list during all the updates should be  $\mathcal{O}(n)$ , and, after all the updates are processed, we should be able to provide in  $\mathcal{O}(n)$  time the correct answer for all the queries (i.e., if a certain query was asked after  $k$  update operations were performed on the list, we should return the answer to the query with respect to the state of the list after those  $k$  update operations were completed).

Next we describe our solution to this problem.

We use a dynamic tree data-structure to maintain the different stages of  $\mathcal{A}$ . Initially, the tree contains only one path: the root 0 and the leaf  $2n + 1$ . When an update is processed, in which a number  $j$  is added to the list, we go up the rightmost path of the tree (from leaf to root) until we find a node with a value smaller than  $j$ . Then  $j$  becomes the rightmost child of that node (i.e.,  $j$  is a leaf). Basically, the rightmost path of a tree after  $k$  updates contains the elements of  $\mathcal{A}$  after those  $k$  updates, preceded by 0. When a query is asked we associate that query with the leaf corresponding to the rightmost leaf of the tree at that moment. In this way, we will be able to identify, after all updates were processed, the contents of the list at the moments of time the queries were, respectively, asked: we just have to traverse the path from the node of the tree associated to that query (this node was a leaf when the query is asked, but after all the updates were processed might have become an internal node) to the root.

The tree can be clearly constructed in linear time: each node is inserted once on the rightmost tree, and it disappears from this rightmost tree (and will not be reinserted there) when a smaller value is inserted in the tree.

In this new setting, the queries can be interpreted as weighted level ancestor queries on the nodes of the constructed tree (where the weight of an edge is the difference between the two nodes bounding it). Considering that the size of the tree is  $\mathcal{O}(n)$ , all weights are also  $\mathcal{O}(n)$ , there are  $\mathcal{O}(n)$  queries, and these queries are to be answered off-line, it follows (see Remark 3) that we can return the answers to all these queries in  $\mathcal{O}(n)$  time.

This completes the linear solution to our problem.  $\square$

To solve Problem 2(b) we use the following lemma.

**Lemma 6** *Given a word  $w$ , let  $L[i] = \min\{j \mid j < i, LCP(j, i) \geq LCP(k, i) \text{ for all } k < i\}$ . The array  $L[\cdot]$  can be computed in linear time.*

**Proof:** We first produce the suffix array of  $w$ . We denote by  $Rank[i]$  the position of  $w[i..n]$  in the suffix array of  $w$ .

We first want to find, for each  $i$ , the suffix  $w[j..n]$ , such that  $j < i$  is minimum with  $Rank[j] < Rank[i]$  and  $LCP(j, i) \geq LCP(j', i)$  for all  $j' < i$  such that  $Rank[j'] < Rank[i]$ . For the suffix  $w[j..n]$  as above, let  $B_{<}[i] = j$ . If no such  $j$  exists, we set  $B_{<}[i] = -1$ .

Similarly, we compute, for each  $i$ , the suffix  $w[j..n]$ , such that  $j < i$  is minimum with  $Rank[j] > Rank[i]$  and  $LCP(j, i) \geq LCP(j', i)$  for all  $j' < i$  such that  $Rank[j'] > Rank[i]$ . For the suffix  $w[j..n]$  as above, let  $B_{>}[i] = j$ . Again, if no such  $j$  exists, we set  $B_{>}[i] = -1$ .

For simplicity, we just present an algorithm computing  $B_{<}[\cdot]$ . The computation of  $B_{>}[\cdot]$  is performed in a similar way. Then,  $L[i] = B_{>}[i]$  if  $LCP(B_{>}[i], i) > LCP(B_{<}[i], i)$  or  $LCP(B_{>}[i], i) = LCP(B_{<}[i], i)$  and  $B_{>}[i] > B_{<}[i]$ ; similarly,  $L[i] = B_{<}[i]$  if  $LCP(B_{>}[i], i) < LCP(B_{<}[i], i)$  or  $LCP(B_{>}[i], i) = LCP(B_{<}[i], i)$  and  $B_{>}[i] < B_{<}[i]$ .

The main idea behind the computation of  $B_{<}[i]$ , for some  $1 \leq i \leq n$ , is that if  $w[j_1..n]$  and  $w[j_2..n]$  are such that  $Rank_R[j_2] < Rank_R[j_1] < Rank[i]$  and  $j_1 < j_2 < i$  then definitely  $B_{<}[i] \neq j_2$ . Indeed,  $LCP(j_2, i) \leq LCP(j_1, i)$ , and, moreover,  $j_2 > j_1$ , so  $j_1$  is a better candidate than  $j_2$  for  $B_{<}[i]$ .

This suggests that we should construct for each  $i$  an ordered list  $\mathcal{A}_i$  of all the integers  $j \leq n$  such that  $\text{Rank}_R[j] < i$  and moreover, if  $j_1$  and  $j_2$  are in  $\mathcal{A}_i$  and  $j_1 < j_2$  then also  $\text{Rank}_R[j_1] < \text{Rank}_R[j_2]$ . If  $j'$  is on top of  $\mathcal{A}_i$ , then  $B[i]$  is just the minimum  $j$  of  $\mathcal{A}_i$  such that  $LCP(j, i) = LCP(j', i)$ .

Let us now consider  $i_1$  and  $i_2$  which occur on consecutive positions of the suffix array of  $w$ , such that  $\text{Rank}[i_1] < \text{Rank}[i_2]$ . Assume that we already computed  $B[i_1]$ . The list  $\mathcal{A}_{i_2}$  can be obtained from  $\mathcal{A}_{i_1}$  by deleting from the right end of the list  $\mathcal{A}_{i_1}$  (where  $\mathcal{A}_{i_1}$  is ordered increasingly from left to right) all the values  $j > i_1$ . Also, to compute  $B[i_2]$  more efficiently, we also delete all the values of  $\mathcal{A}_{i_1}$  that are greater than  $B[i_1]$ . Indeed, for some  $k > B[i_1]$  we have that  $LCP(B[i_1], i_1) = LCP(B[i_1], i_2)$  and  $LCP(k, i_1) = LCP(k, i_2)$ , so  $B[i_1]$  is a better candidate for  $B[i_2]$  than  $k$ . Afterwards, we insert  $i_1$  in the updated list  $\mathcal{A}_{i_1}$ , to obtain a list  $\mathcal{A}_{i_2}$  from which  $B[i_2]$  can be computed. Clearly, this entire process takes linear time, and leads to a correct computation of  $B_{<}[\cdot]$ .

Afterwards, we compute  $B_{>}[\cdot]$  similarly, and get  $L[\cdot]$  in linear time.  $\square$

Another lemma shows how the computation of the array  $LPF_g[\cdot]$  can be connected to that of the array  $L[\cdot]$ . For an easier presentation, let  $B[i]$  denote the leftmost starting position of the longest factor  $x_i$  that occurs both at position  $i$  and at a position  $j$  such that  $j + |x_i| \leq i - g(i)$ ; if there is no such factor  $x_i$ , then  $B[i] = -1$ . In other words, the length of the factor  $x_i$  occurring at position  $B[i]$  gives us  $LPF_g[i]$ . In fact,  $LPF_g[i] = \min\{LCP(B[i], i), i - g(i) - B[i]\}$ .

Now, let  $L^1[i] = L[i]$  and  $L^k[i] = L[L^{k-1}[i]]$ , for  $k \geq 2$ ; also, we define

$$L^+[i] = \{L[i], L[L[i]], L[L[L[i]]], \dots\}.$$


The following lemma shows the important fact that  $B[i]$  can be obtained just by looking at the values of  $L^+[i]$ . More precisely,  $B[i]$  equals  $L^k[i]$ , where  $k$  is obtained by looking at the values  $L^j[i] \leq i - g(i)$  and taking the one such that the factor starting on it and ending on  $i - g(i) - 1$  has a maximal common prefix with  $w[i..n]$ . Afterwards, Theorem 5 shows that this check can be done in linear time for all  $i$ , thus solving optimally Problem 2.

**Lemma 7** *For a word  $w$  of length  $n$  and all  $1 \leq i \leq n$  such that  $B[i] \neq -1$ , we have that  $B[i] \in L^+[i]$ .*

**Proof:** Let us assume, for the sake of a contradiction, that  $B[i] = j \notin L^+[i]$ . This means that  $j < i - g(i)$  and  $\min\{LCP(j, i), i - g(i) - j\} \geq \min\{LCP(j', i), i - g(i) - j'\}$  for all  $j' < i - g(i)$ . We further consider two simple cases.

In the first case, there is  $j' < j$  such that  $j' \in L^+[i]$ ; take  $j'$  to be the greatest number less than  $j$  that belongs to  $L^+[i]$ . Then, it is clear that the longest factor that occurs both at  $j'$  and at  $i$  ends before  $i - g(i) - 1$ . Otherwise, we would have  $\min\{LCP(j, i), i - g(i) - j\} \leq i - g(i) - j < \min\{LCP(j', i), i - g(i) - j'\} = i - g(i) - j'$ , so we would have  $B[i] = j'$ , a contradiction. So, it follows immediately, that  $LCP(j, i) > LCP(j', i)$ . Now, if  $j' = L^k[i]$  (that is  $L[L[\dots L[i]\dots]]$ , where  $L$  is applied  $k$  times on  $i$ ), then there exists  $k' < k$  such that  $j'' = L^{k'}[i] > j$  and  $LCP(j'', i) > LCP(j', i)$  and  $LCP(j', i) = LCP(L[j''], i)$ . Clearly, this means that  $j$  occurs in the suffix array of  $w$  closer to the suffix  $w[i..n]$  than the suffix  $w[L[j'']..n]$ , but farther than the suffix  $w[j''..n]$ . So,  $LCP(j, j'') > LCP(L[j''], j)$ , which is a contradiction to the definition of  $L[j'']$ . In conclusion, we cannot have that there is  $j' < j$  such that  $j' \in L^+[i]$ .

Now, assume that there is no  $j' < j$  such that  $j' \in L^+[i]$ . As  $j \notin L^+[i]$ , this leads immediately to a contradiction. Since  $B[i] \neq -1$ , we have that the suffix starting at position  $j$  has at least the prefix of length one common with all the factors starting at positions from  $L^+[i]$ . Therefore, at least one element from  $L^+[i]$  should be less or equal to  $j$ . This concludes our proof.  $\square$



**Fig. 4:** Proof of Theorem 5: We measure the length of the common prefix between each factor starting on  $L^k[i]$  and ending on  $i - g(i)$  and the suffix  $w[i..n]$ . The starting position of the factor that produces the longest such common prefix is  $B[i]$ . Also, this longest common prefix defines the arm  $u$  of the gapped repeat  $uvu$ .

**Theorem 5** *Problem 2(b) can be solved in linear time.*

**Proof:** The main idea in this proof is that, to compute  $LPF_g[i]$ , it is enough to check the elements  $j \in L^+[i]$  with  $j \leq i - g(i)$ , and choose from them the one for which  $\min\{LCP(j, i), i - g(i) - j\}$  is maximum; this maximum will be the value we look for. In the following, we show how we can obtain these values efficiently.

First, if  $j = L[i]$  for some  $i$ , we define the value  $end[j] = k \leq LCP(L[j], i)$ , where  $|w[j..k]| \leq \min\{LCP(L[j], i), k - L[j]\}$ . Basically, for each position  $k' \leq k$ , the longest factor  $x$  starting at  $L[j]$  and ending on  $k'$  which also occurs at position  $i$  is longer than any factor starting on position  $j$  and ending on  $k'$  which also occurs at position  $i$ . Now we note that if  $j \in L^+[i]$  is the greatest element of this set such that  $end[j] \leq i - g(i) - 1$ , then  $LPF_g[i] = \min\{LCP(j, i), i - g(i) - j\}$ . Clearly,  $end[j]$  can be computed in constant time for each  $j$ .

To be able to retrieve efficiently for some  $i$  the greatest element of this set such that  $end[j] \leq i - g(i) - 1$  we proceed as follows.

First we define a disjoint-set union-find data structure on the universe  $U = [1, n + 1]$ , where the unions can be only performed between the set containing  $i$  and that containing  $L[i]$ , for all  $i$ . Initially, each number between 1 and  $n$  is a singleton set in this structure. Moreover, our structure fulfils the conditions that the efficient union-find data structure of Gabow and Tarjan (1983) should fulfil: the unions we make form a tree.

Further, we sort in linear time the numbers  $i - g(i)$ , for all  $i$ ; we also sort in linear time the numbers  $end[k]$  for all  $k \leq n$ . We now traverse the numbers from  $n$  to 1, in decreasing order. When we reach position  $j$  we check whether  $j$  equals  $end[k]$  for some  $k$ ; if yes, we unite the set containing  $k$  with the set containing  $end[k]$  for all  $k$  such that  $end[k] = j$ . Then, if  $j = i - g(i)$  for some  $i$ , we just have to return the minimum of the set containing  $i$ ; this value gives exactly the greatest element  $j \in L^+[i]$  such that  $end[j] \leq i - g(i) - 1$ . So, as described above, we can obtain from it the value of  $LPF_g[i]$ . The computation of this array follows from the previous remarks.

To evaluate the complexity of our approach, note that we do  $\mathcal{O}(n)$  union operations and  $\mathcal{O}(n)$  find queries on the union-find data structure. By the results of Gabow and Tarjan (1983), the time needed to construct the union-find data structure and perform these operations on it is also  $\mathcal{O}(n)$ . From every find query we get in constant time the value of a element  $LPF_g[i]$ . So the solution of Problem 2 is linear.  $\square$

Just like in the previous section, the following consequence of the previous two theorems is immediate.

**Theorem 6** *Given  $w$  of length  $n$  and the values  $g(1), \dots, g(n)$  of  $g : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ , we can find in linear time the gapped palindrome (or repeat)  $u^Rvu$  (respectively,  $uvu$ ) occurring in  $w$ , with the longest arm  $u$ , such that if its right arm starts on position  $i$  then  $|v| \geq g(i)$ .*

## 5 $\alpha$ -gapped Repeats and Palindromes

Recall that an  $\alpha$ -gapped palindrome (respectively, repeat)  $w[i..j]vw[i'..j']$  is called maximal if the arms cannot be extended to the right or to the left: neither  $w[i..j+1]v'w[i'-1..j']$  nor  $w[i-1..j]vw[i'..j'+1]$  (respectively, neither  $w[i..j+1]v'w[i'..j'+1]$  nor  $w[i-1..j]v''w[i'-1..j']$ ) are  $\alpha$ -gapped palindromes (respectively, repeats). Gawrychowski and Manea (2015) defined algorithms that find the longest  $\alpha$ -gapped palindromes and repeats in  $\mathcal{O}(\alpha n)$  time; these algorithms do not compute the set of all  $\alpha$ -gapped palindromes or repeats, but just the ones with the longest arms. We present them below.


We first consider the case of  $\alpha$ -gapped repeats.

**Theorem 7** *Given a word  $w$  of length  $n$  and an integer  $\alpha \leq n$ , the longest  $\alpha$ -gapped repeat  $uvu$  contained in  $w$  can be found in  $\mathcal{O}(\alpha n)$  time.*

**Proof:** Informally, our approach works as follows (see also Figure 5). For each  $k$ , we try to find the longest  $\alpha$ -gapped repeat  $u_1vu_2 = uvu$ , with  $u_1 = u_2 = u$ , and  $2^{k+1} \log n \leq |u| \leq 2^{k+2} \log n$ . In each such repeat, the right arm  $u_2$  must contain a factor (called  $k$ -block)  $z$ , of length  $2^k \log n$ , starting on a position of the form  $j2^k \log n + 1$ . So, we try each such factor  $z$ , fixing in this way a range of the input word where  $u_2$  could appear. Now,  $u_1$  must also contain a copy of  $z$ . However, it is not mandatory that this copy of  $z$  occurs nicely aligned to its original occurrence; that is, the copy of  $z$  does not necessarily occur on a position of the form  $i \log n + 1$ . But, it is not hard to see that  $z$  has a factor  $y$  of length  $2^{k-1} \log n$ , starting in its first  $\log n$  positions and whose corresponding occurrence in  $u_1$  starts on a position of the form  $i \log n + 1$ . Further, we can use the fact that  $u_1vu_2$  is  $\alpha$ -gapped and apply Lemma 1 to a suitable encoding of the input word to locate in constant time for each  $y$  starting in the first  $\log n$  positions of  $z$  all possible occurrences of  $y$  on a position of the form  $i \log n + 1$ , occurring not more than  $(8\alpha + 2)|y|$  positions to the left of  $z$ . Intuitively, each occurrence of  $y$  found in this way fixes a range where  $u_1$  might occur in  $w$ , such that  $u_1vu_2$  is  $\alpha$ -gapped. So, around each such occurrence of  $y$  (supposedly, in the range corresponding to  $u_1$ ) and around the  $y$  from the original occurrence of  $z$  we try to effectively construct the arms  $u_1$  and  $u_2$ , respectively, and see if we really get an  $\alpha$ -gapped repeat. In the end, we just return the longest repeat we obtained, going through all the possible choices for  $z$  and the corresponding  $y$ 's. We describe in the following an  $\mathcal{O}(\alpha n)$  time implementation of this approach.

The first step of the algorithm is to construct a word  $w'$ , of length  $\frac{n}{\log n}$ , whose symbols, called *blocks*, encode  $\log n$  consecutive symbols of  $w$  grouped together. Basically, now we have two versions of the word  $w$ : the original one, and the one where it is split in blocks. Clearly, the blocks can be encoded into numbers between 1 and  $n$  in linear time. Indeed, we produce the suffix array of  $w$ , and group together the suffixes that share a common prefix of length  $\log n$ . Then, we label (injectively) the groups obtained in this way with numbers between 1 and  $n$ . Finally, a block is encoded as the label of the suffix that starts with that block. Consequently, we can construct in  $\mathcal{O}(n)$  time the suffix arrays and *LCP*-data structures for both  $w$  and  $w'$ . We can also build in  $\mathcal{O}(n)$  time the data structures of Lemma 1 for the word  $w'$ .

Now, we try to find the longest  $\alpha$ -gapped repeat  $u_1vu_2 = uvu$  of  $w$ , with  $u_1 = u_2 = u$ , and  $2^{k+1} \log n \leq |u| \leq 2^{k+2} \log n$ , for each  $k \geq 1$  if  $\alpha > \log n$  or  $k \geq \log \log n$ , otherwise. Let us consider



**Fig. 5:** Proof of Theorem 7: Segment of  $w$ , split into blocks of length  $\log n$ . In this segment,  $z$  is a  $k$ -block of length  $2^k \log n$ . For each factor  $y$ , of length  $2^{k-1} \log n$ , occurring in the first  $\log n$  symbols of  $z$  (not necessarily a sequence of blocks), we find the occurrences of  $y$  that correspond to sequences of  $2^{k-1}$  blocks, and start at most  $(8\alpha + 2)|y| = (4\alpha + 1) \cdot 2^k \log n$  symbols (or, alternatively,  $(4\alpha + 1) \cdot 2^k$  blocks) to the left of the considered  $z$ . These  $y$  factors may appear as runs or as separate occurrences. Some of them can be extended to form an  $\alpha$ -gapped repeat  $u_1vu_2 = wvu$  such that the respective occurrence of  $y$  has the same offset in  $u_1$  as the initial  $y$  in  $u_2$ .

now one such  $k$ . We split again the word  $w$ , this time in factors of length  $2^k \log n$ , called  $k$ -blocks. For simplicity, assume that each split is exact.

Clearly, if an  $\alpha$ -gapped repeat  $u_1vu_2$  like above exists, then  $u_2$  contains at least one of the  $k$ -blocks. Consider such a  $k$ -block  $z$  and assume it is the leftmost  $k$ -block of  $u_2$ . On the other hand,  $u_1$  contains at least  $2^{k+1} - 1$  consecutive blocks from  $w'$ , so there should be a factor  $y$  of  $w$  corresponding to  $2^{k-1}$  of these  $(2^{k+1} - 1)$  blocks which is also a factor of  $z$ , and starts on one of the first  $\log n$  positions of  $z$ . Now, for each  $k$ -block  $z$  and each  $y$ , with  $|y| = 2^{k-1} \log n$  and starting in its prefix of length  $\log n$ , we check whether there are occurrences of  $y$  in  $w$  ending before  $z$  that correspond to exactly  $2^{k-1}$  consecutive blocks of  $w'$  (one of them should be the occurrence of  $y$  in  $u_1$ ); note that the occurrence of  $y$  in  $z$  may not necessarily correspond to a group of  $2^{k-1}$  consecutive blocks, but the one from  $u_1$  should. As  $u_1vu_2$  is  $\alpha$ -gapped and  $|u_1| \leq 2^{k+2} \log n$ , then the occurrence of  $y$  from  $u_1$  starts at most  $(4\alpha + 1)2^k \log n$  symbols before  $z$  (as  $|u_1v| \leq \alpha|u_2| \leq \alpha 2^{k+2} \log n$ , and  $z$  occurs with an offset of at most  $2^k \log n$  symbols in  $u_2$ ). So, the block-encoding of the occurrence of the factor  $y$  from the left arm  $u_1$  should occur in a factor of  $(4\alpha + 1)2^k$  blocks of  $w'$ , to the left of the blocks corresponding to  $z$ .

For the current  $z$  and an  $y$  as above, we check whether there exists a factor  $y'$  of  $w'$  whose blocks correspond to  $y$ , by binary searching the suffix array of  $w'$  (using *LCP*-queries on  $w$  to compare the factors of  $\log n$  symbols of  $y$  and the blocks of  $w'$ , at each step of the search). If not, we try another possible  $y$ . If yes, using Lemma 1 for  $w'$ , we retrieve (in  $\mathcal{O}(\log \log |w'| + \alpha)$  time) a representation of the occurrences of  $y'$  in the range of  $(4\alpha + 1)2^k$  blocks of  $w'$  occurring before the blocks of  $z$ ; this range corresponds to a range of length  $(4\alpha + 1)2^k \log n$  of  $w$ .

If  $y'$  is aperiodic then there are only  $\mathcal{O}(\alpha)$  such occurrences. Each factor of  $w$  corresponding to one of these occurrences might be the occurrence of  $y$  from  $u_1$ , so we try to extend both the factor corresponding to the respective occurrence of  $y'$  from  $w'$  and the factor  $y$  from  $z$  in a similar way to the left and right to see whether we obtain the longest  $\alpha$ -gapped repeat. If  $y'$  is periodic (so,  $y$  is periodic as well), we know that the representation of its occurrences consists of  $\mathcal{O}(\alpha)$  separate occurrences and  $\mathcal{O}(\alpha)$  runs in which  $y'$  occurs (see Preliminaries). The separate occurrences are treated as above. Each run  $r'$  of  $w'$  where  $y'$  occurs is treated differently, depending on whether its corresponding run  $r$  from  $w$  (made of the blocks corresponding to  $r'$ ) supposedly starts inside  $u_1$ , ends inside  $u_1$ , or both starts and ends inside  $u_1$ . We can check each of these three cases separately, each time trying to establish a correspondence between  $r$  and

the run containing the occurrence of  $y$  from  $z$ , which should also start, end, or both start or end inside  $u_2$ , respectively. Then we define  $u_1$  and  $u_2$  as the longest equal factors containing these matching runs on matching positions. Hence, for each separate occurrence of  $y'$  or run of such occurrences, we may find an  $\alpha$ -gapped repeat in  $w$ ; we just store the longest. This whole process takes  $\mathcal{O}(\alpha)$  time.

If  $\alpha > \log n$ , we run this algorithm for all  $k \geq 1$  and find the longest  $\alpha$ -gapped repeat  $uvu$ , with  $4 \log n \leq |u|$ , in  $\mathcal{O}(\alpha n)$  time.

If  $\alpha \leq \log n$ , we run this algorithm for all  $k \geq \log \log n$  and find the longest  $\alpha$ -gapped repeat  $uvu$ , with  $2^{\log \log n + 1} \log n \leq |u|$ , in  $\mathcal{O}(\alpha n)$  time. If our algorithm did not find such a repeat, we should look for  $\alpha$ -gapped repeats with shorter arm. Now,  $|u|$  is upper bounded by  $2^{\log \log n + 1} \log n = 2(\log n)^2$ , so  $|uvu| \leq \ell_0$ , for  $\ell_0 = \alpha \cdot 2(\log n)^2 + 2(\log n)^2 = (2\alpha + 2)(\log n)^2$ . Such an  $\alpha$ -gapped repeat  $uvu$  is, thus, contained in (at least) one factor of length  $2\ell_0$  of  $w$ , starting on a position of the form  $1 + m\ell_0$  for  $m \geq 0$ . So, we take the factors  $w[1 + m\ell_0..(m+2)\ell_0]$  of  $w$ , for  $m \geq 0$ , and apply for each such factor, separately, the same strategy as above. As an important detail, before running the algorithm presented above, we first encode the symbols of  $w[1 + m\ell_0..(m+2)\ell_0]$ , which were numbers between 1 and  $n$ , to numbers between 1 and  $2\ell_0$ ; again, this is done by looking at the suffix array of  $w$ , and it allows us to apply recursively the algorithm described before. The total time needed to do that is  $\mathcal{O}\left(\alpha \ell_0 \frac{n}{\ell_0}\right) = \mathcal{O}(\alpha n)$ . Hence, we found the longest  $\alpha$ -gapped repeats  $uvu$ , with  $2^{\log \log(2\ell_0) + 1} \log(2\ell_0) \leq |u|$ . If our search was still fruitless, we need to search  $\alpha$ -gapped repeats with  $|u| \leq 2^{\log \log(2\ell_0) + 1} \log(2\ell_0) \leq 16 \log n$  (a rough estimation, based on the fact that  $\alpha \leq \log n$ ).

So, in both cases,  $\alpha > \log n$  or  $\alpha \leq \log n$ , it is enough to find the longest  $\alpha$ -gapped repeats with  $|u| \leq 16 \log n$ . The right arm  $u_2$  of such a repeat is contained in a factor  $w[m \log n + 1..(m+17) \log n]$  of  $w$ , while  $u_1$  surely occurs in a factor  $x = w[m \log n - 16\alpha \log n + 1..(m+17) \log n]$  (or, if  $m \log n - 16\alpha \log n + 1 \leq 0$ , then in a factor  $x = w[1..(m+17) \log n]$ ); in total, there are  $\mathcal{O}(n/\log n)$  such  $x$  factors. As before, we can process them in linear time in order to re-encode each of them as a word over the alphabet consisting of numbers which are in  $\mathcal{O}(\alpha n)$ . In each of these factors, we look for  $\alpha$ -gapped repeats  $u_1 v u_2 = uvu$  with  $2^{k+1} \leq |u| \leq 2^{k+2}$ , where  $0 \leq k \leq \log \log n + 2$  (the case  $|u| < 2$  is trivial), and  $u_2$  occurs in the suffix of length  $17 \log n$  of this factor. Moreover,  $u_2$  contains a factor  $y$  of the form  $x[j2^k + 1..(j+1)2^k]$ . Using Lemma 3 and Remark 2, for each such possible  $y$  occurring in the suffix of length  $17 \log n$  of  $x$ , we assume it is the one contained in  $u_2$  and we produce in  $\mathcal{O}(\alpha)$  time a representation of the  $\mathcal{O}(\alpha)$  occurrences of  $y$  in the factor of length  $(4\alpha + 1)|y|$  preceding  $y$ . One of these should be the occurrence of  $y$  from  $u_1$ . Similarly to the previous cases, we check in  $\mathcal{O}(\alpha)$  time which is the longest  $\alpha$ -gapped repeat obtained by pairing one of these occurrences to  $y$ , and extending them similarly to the left and right. The time needed for this is  $\mathcal{O}(\alpha \log n)$  per each of the  $\mathcal{O}(\frac{n}{\log n})$  factors  $x$  defined above. This adds up to an overall complexity of  $\mathcal{O}(\alpha n)$ , again.

This was the last case we needed to consider. In conclusion, we can find the longest  $\alpha$ -gapped repeat  $uvu$  in  $\mathcal{O}(\alpha n)$  time.  $\square$

Further we discuss the case of  $\alpha$ -gapped palindromes.

**Theorem 8** *Given a word  $w$  of length  $n$ , the longest  $\alpha$ -gapped palindrome  $u^R v u$  contained in  $w$  can be determined in  $\mathcal{O}(\alpha n)$  time.*

**Proof:** Our approach is similar to the case of repeats, presented in the previous theorem. For each  $k$ , we try to find the longest  $\alpha$ -gapped palindrome  $u^R v u$ , with  $2^{k+1} \log n \leq |u| \leq 2^{k+2} \log n$ . In each such

$\alpha$ -gapped palindrome, the right arm  $u$  must contain a factor  $z$ , of length  $2^k \log n$ , starting on a position of the form  $j2^k \log n + 1$ . So, we try each such factor  $z$ , fixing in this way a range of the input word where  $u$  could appear. Now,  $u^R$  must contain a factor  $z^R$ ; however, it is not mandatory that this factor occurs at a position of the form  $i \log n + 1$ . But, just like before, it is not hard to see that  $z$  has a factor  $y$ , of length  $2^{k-1} \log n$ , that starts in its first  $\log n$  positions and whose corresponding occurrence  $y^R$  from  $u^R$  should start on a position of the form  $i \log n + 1$ . Further, we can use the fact that  $u^R v u$  is  $\alpha$ -gapped and apply Lemma 1 to an encoding of the input word to locate in constant time for each  $y$  starting in the first  $\log n$  positions of  $z$  all possible occurrences of  $y^R$  on a position of the form  $i \log n + 1$ , occurring not more than  $(8\alpha + 2)|y|$  positions to the left of  $z$ . Intuitively, each occurrence of  $y$  found in this way fixes a range where  $u^R$  might occur in  $w$ , such that  $u^R v u$  is  $\alpha$ -gapped. So, around each such occurrence of  $y^R$  (supposedly, in the range corresponding to  $u^R$ ) and around the  $y$  from  $z$  we try to effectively construct the arms  $u^R$  and  $u$ , respectively, and see if we get the longest  $\alpha$ -gapped palindrome. This approach can be implemented in  $\mathcal{O}(\alpha n)$  time, just like in the case of  $\alpha$ -gapped repeats.

The first step of the algorithm is to construct a word  $w'$ , of length  $\frac{n}{\log n}$ , whose symbols, called *blocks*, encode  $\log n$  consecutive symbols of  $w$  grouped together. Now we have two versions of the word  $w$ : the original one, and the one made of blocks. As before, the blocks can be encoded as numbers between 1 and  $n$  in linear time. We construct in  $\mathcal{O}(n)$  time the suffix arrays and *LCP*-data structures for both  $w$  and  $w'$ , and we build in  $\mathcal{O}(n)$  time the data structures of Lemma 1 for the word  $w'$ .

Considering the original word  $w$ , we find the  $\alpha$ -gapped palindrome  $u^R v u$  with  $2^{k+1} \log n \leq |u| \leq 2^{k+2} \log n$ , for each  $k \geq \log \log n$ . To this end, we split again the word  $w$ , this time in factors of length  $2^k \log n$ , called *k-blocks*. For simplicity, assume that each split we do is exact; to achieve this, we may have to pad the word with a new symbol in a suitable manner.


If an  $\alpha$ -gapped palindrome  $u^R v u$  of the kind we search for exists, then  $u$  contains at least one of the *k-blocks*. Consider such a *k-block*  $z$  and assume it is the leftmost *k-block* of  $u$ . On the other hand,  $u^R$  contains at least  $2^{k+1} - 1$  consecutive blocks from  $w'$ , so there should be a factor  $y$  of  $w$  with  $y^R$  corresponding to  $2^{k-1}$  of these  $(2^{k+1} - 1)$  blocks such that  $y$  is a factor of  $z$  that starts on one of its first  $\log n$  positions. Now, for each *k-block*  $z$  and each  $y$  starting in its prefix of length  $\log n$ , with  $|y| = 2^{k-1} \log n$ , we check whether there are occurrences of  $y^R$  ending before  $z$  (one of them should be the occurrence of  $y^R$  in  $u^R$ ) that correspond to exactly  $2^{k-1}$  consecutive blocks of  $w'$ . Note that the occurrence of  $y$  in  $z$  may not necessarily correspond to a group of  $2^{k-1}$  consecutive blocks, but the one of  $y^R$  from  $u^R$  do. As  $u^R v u$  is  $\alpha$ -gapped and  $|u^R v| \leq \alpha 2^{k+2} \log n$ , then the occurrence of  $y^R$  from  $u^R$  starts at most  $(4\alpha + 1) \cdot 2^k \log n$  symbols before  $z$ . So, the block-encoding of  $y^R$  should occur in a factor of  $(4\alpha + 1) \cdot 2^k$  blocks of  $w'$ , to the left of the blocks corresponding to  $z$ .

For the considered  $z$  and some  $y$  as above, by binary searching the suffix array of  $w'$  (using *LCP*-queries on the word  $w w^R$  to compare the factors of  $\log n$  symbols of  $y^R$  and the blocks of  $w'$ , at each step of the search), we check whether there exists a factor  $y'$  of  $w'$  whose blocks correspond to  $y^R$ . If not, we try another possible  $y$ . If yes, we continue. Using Lemma 1 for  $w'$ , we obtain a representation of the occurrences of  $y'$  in the range of  $(4\alpha + 1) \cdot 2^k$  blocks of  $w'$ , occurring before the blocks that correspond to  $z$ ; note that this range corresponds to a range of length  $(4\alpha + 1) \cdot 2^k \log n$  of  $w$ .

If  $y'$  is aperiodic then there are  $\mathcal{O}(\alpha)$  such occurrences. Each factor corresponding to these occurrences might be the occurrence of  $y^R$  from  $u^R$ , so we try to extend it and the  $y$  from  $z$  in a similar way to see whether we obtain the longest  $\alpha$ -gapped palindrome.

If  $y'$  is periodic (so,  $y$  is periodic as well), the representation of its occurrences consists of  $\mathcal{O}(\alpha)$  separate occurrences and  $\mathcal{O}(\alpha)$  runs in which  $y'$  occurs. The separate occurrences are treated as above. Each run  $r'$





**Fig. 6:** Proof of Theorem 8: Segment of  $w$ , split into blocks of length  $\log n$ . In this segment,  $z$  is a  $k$ -block of length  $2^k \log n$ . For each factor  $y$ , of length  $2^{k-1} \log n$ , occurring in the first  $\log n$  symbols of  $z$  (not necessarily a sequence of blocks), we find the occurrences of  $y^R$  that correspond to sequences of  $2^{k-1}$  blocks, and start at most  $(8\alpha + 2)|y| = (4\alpha + 1) \cdot 2^k \log n$  symbols (or, alternatively,  $(4\alpha + 1) \cdot 2^k$  blocks) to the left of the considered  $z$ . These  $y^R$  factors may appear as runs or as separate occurrences. Some of them can be extended to form an  $\alpha$ -gapped palindrome  $u^R v u$  such that the respective occurrence of  $y^R$  is the mirror image of the initial  $y$  in  $u_2$ .

of  $w'$  where  $y'$  occurs is treated differently, depending on whether its corresponding run  $r$  from  $w$  (made of the blocks corresponding to  $r'$ ) supposedly starts inside  $u^R$ , ends inside  $u^R$ , or both starts and ends inside  $u^R$ . We can check all these three cases separately, each time trying to establish a correspondence between  $r$  and the run containing the occurrence of  $y$  from  $z$ , which should also end, start, or both start or end inside  $u$ , respectively. Then we define  $u^R$  and  $u$  as the longest mirrored words containing these matching runs on mirrored positions. In this way, for each separate occurrence of  $y'$  or run of such occurrences, we found a  $\alpha$ -gapped palindrome in  $w$ ; we just store the longest. This whole process takes  $\mathcal{O}(1)$  time for each run.

If  $\alpha > \log n$ , we run this algorithm for all  $k \geq 1$  and find the longest  $\alpha$ -gapped palindrome  $u^R v u$ , with  $4 \log n \leq |u|$ , in  $\mathcal{O}(\alpha n)$  time.

If  $\alpha \leq \log n$ , we run this algorithm for all  $k \geq \log \log n$  and find the longest  $\alpha$ -gapped palindrome  $u^R v u$ , with  $2^{\log \log n + 1} \log n \leq |u|$ , in  $\mathcal{O}(\alpha n)$  time. If our algorithm did not find such a palindrome, we should look for  $\alpha$ -gapped palindrome with shorter arm. The length of this arm,  $|u|$ , is now upper bounded by  $2^{\log \log n + 1} \log n = 2(\log n)^2$ , so  $|u^R v u| \leq \ell_0$ , for  $\ell_0 = \alpha \cdot 2(\log n)^2 + 2(\log n)^2 = (2\alpha + 2)(\log n)^2$ . Such an  $\alpha$ -gapped palindrome  $u^R v u$  is, thus, contained in (at least) one factor of length  $2\ell_0$  of  $w$ , starting on a position of the form  $1 + m\ell_0$  for  $m \geq 0$ . So, we take the factors  $w[1 + m\ell_0..(m+2)\ell_0]$  of  $w$ , for  $m \geq 0$ , and apply for each such factor, separately, the same strategy as above. As noted in the previous proof, these words can be re-encoded in linear time as words over an alphabet of size  $\mathcal{O}(\log n)$ . The total time needed to do that is  $\mathcal{O}\left(\alpha \ell_0 \frac{n}{\ell_0}\right) = \mathcal{O}(\alpha n)$ . Hence, we found the longest  $\alpha$ -gapped palindromes  $u^R v u$ , with  $2^{\log \log(2\ell_0) + 1} \log(2\ell_0) \leq |u|$ . If our search was still fruitless, we search  $\alpha$ -gapped palindromes with  $|u| \leq 2^{\log \log(2\ell_0) + 1} \log(2\ell_0) \leq 16 \log n$  (a rough estimation, based on the fact that  $\alpha \leq \log \log n$ ).

Now in both cases (when  $\alpha > \log n$  or  $\alpha \leq \log n$ ) it is enough to find the  $\alpha$ -gapped palindromes with  $|u| \leq 16 \log n$ . The right arm  $u$  of such a repeat is contained in a factor  $w[m \log n + 1..(m+17) \log n]$  of  $w$ , while  $u^R$  surely occurs in a factor  $x = w[m \log n - 16\alpha \log n + 1..(m+17) \log n]$  (or, if  $m \log n - 16\alpha \log n + 1 \leq 0$ , then in a factor  $x = w[1..(m+17) \log n]$ ). In total, there are  $\mathcal{O}(n/\log n)$  such  $x$  factors, and after a linear time preprocessing we can ensure that they are all over integer alphabets with respect to their length. In each of them, we look for  $\alpha$ -gapped palindromes  $u^R v u$  with  $2^{k+1} \leq |u| \leq 2^{k+2}$ , where  $0 \leq k \leq \log \log n + 1$  (the case  $|u| < 2$  is trivial), and  $u_2$  occurs in the suffix of length  $9 \log n$

of this factor. Moreover,  $u$  contains a factor  $y$  of the form  $x[j2^k + 1..(j+1)2^k]$ . Using Lemma 3 and Remark 2, for each such possible  $y$  occurring in the suffix of length  $17 \log n$  of  $x$ , we assume it is the one contained in  $u$  and we produce in  $\mathcal{O}(\alpha)$  time a representation of the  $\mathcal{O}(\alpha)$  occurrences of  $y^R$  in the factor of length  $(4\alpha + 1)|y|$  preceding  $y$ . One of these should be the occurrence of  $y^R$  from  $u^R$ . Similarly to the previous cases, we check in  $\mathcal{O}(\alpha)$  time which is the longest  $\alpha$ -gapped palindrome obtained by pairing one of these occurrences to  $y$ , and extending them similarly to the left and right. The time needed for this is  $\mathcal{O}(\alpha \log n)$  per each of the  $\mathcal{O}(\frac{n}{\log n})$  factors  $x$  defined above. This adds up to an overall complexity of  $\mathcal{O}(\alpha n)$ , again.

This was the last case we needed to consider. In conclusion, we can find the longest  $\alpha$ -gapped palindrome  $u^R v u$  in  $\mathcal{O}(\alpha n)$  time.  $\square$

The algorithms presented in the previous two proofs were non-trivially extended by Gawrychowski et al. (2016) to algorithms that construct the sets of all maximal  $\alpha$ -gapped repeats and maximal  $\alpha$ -gapped palindromes (which have a non-empty gap) in  $\mathcal{O}(\alpha n)$  time. Essentially, instead of looking for the longest  $\alpha$ -gapped repeat (or palindrome) that contains a certain basic factor (as we did in this proof), we look for all the maximal  $\alpha$ -gapped repeats (respectively, palindromes) that contain the respective basic factor. Using a series of deep combinatorial observations on the structure of these maximal gapped repeats or palindromes, one can output all of them in  $\mathcal{O}(1)$  time per repeat or palindrome. Using the crucial fact that the number of both  $\alpha$ -gapped repeats with non-empty gap as well as  $\alpha$ -gapped palindromes with non-empty gap is  $\mathcal{O}(\alpha n)$  (in fact, the main result of Gawrychowski et al. (2016)), we get that they can all be identified and output in  $\mathcal{O}(\alpha n)$  time.

Accordingly, we further show that given the set  $S$  of all factors of a word which are maximal  $\alpha$ -gapped palindromes (respectively, repeats) we can compute the array  $LPal_\alpha$  (respectively,  $LRep_\alpha$ ) for that word in  $\mathcal{O}(n + |S|)$  time. As a consequence, for constant  $\alpha$ , these problems can be solved in linear time.

Note that, in this case, our strategy is fundamentally different from the ones we used in the cases of Problems 1 and 2. There we were able to construct the desired data structures without constructing first the set of all maximal gapped palindromes and maximal gapped repeats whose gap fulfilled the required restrictions. Here we first find all maximal  $\alpha$ -gapped repeats and  $\alpha$ -gapped palindromes using the algorithms of Gawrychowski et al. (2016), and then compute directly the desired data structures.


**Theorem 9** *Problem 3(a) can be solved in  $\mathcal{O}(\alpha n)$  time.*

**Proof:** We assume that we are given an input word  $w$ , for which the set  $S$  of all maximal  $\alpha$ -gapped palindromes is computed, using the algorithm of Gawrychowski et al. (2016).

Let us consider a maximal  $\alpha$ -gapped palindrome  $w[i..i+\ell-1]vw[j..j+\ell-1]$ , with  $w[i..i+\ell-1]^R = w[j..j+\ell-1]$ . For simplicity, let us denote by  $\delta = |v| = j - i - \ell$ , the length of the gap; here,  $i$  and  $j + \ell - 1$  will be called the *outer ends* of this palindrome, while  $j$  and  $i + \ell - 1$  are the *inner ends*.

It is not hard to see that from a maximal  $\alpha$ -gapped palindrome one can get a family of  $\alpha$ -gapped palindromes whose arms cannot be extended by appending letters simultaneously to their outer ends. We now show how this family of  $\alpha$ -gapped palindromes can be computed. Intuitively, we extend simultaneously the gap in both directions, decreasing in this way the length of the arms of the palindrome, until the gap becomes long enough to violate the  $\alpha$ -gapped restriction. The longest possible such extension of the gap can be easily computed.

Indeed, let  $r = \lfloor \frac{(\alpha-1)\ell-\delta}{\alpha+1} \rfloor$ . It is not hard to check that for  $r' \leq r$  we have that  $w[i..i+\ell-r'-1]v'w[j+r'..j+\ell-1]$ , with  $v' = w[i+\ell-r'-1..i+\ell-1]vw[j..j+r'-1]$ , is an  $\alpha$ -gapped palindrome



**Fig. 7:** Proof of Theorem 9:  $p = uvu$  is a maximal  $\alpha$ -gapped palindrome (here  $\alpha = 2$ ) with  $|v| = \delta$  and  $|u| = \ell$ . We define  $r = \lfloor \frac{(\alpha-1)\ell-\delta}{\alpha+1} \rfloor = \lfloor \frac{\ell-\delta}{3} \rfloor$ . This allows us to define the interval  $I_p$ , where the right arm of an  $\alpha$ -gapped palindrome obtained from  $p$  may start. Further,  $y^R v'' y$  is an example of such an  $\alpha$ -gapped palindrome;  $t^R v' t$  is a gapped palindrome obtained from  $p$ , which is not  $\alpha$ -gapped: its right arm does not start in  $I_p$ . The interval  $I_p$  has weight  $j + \ell - 1$ .

whose left arm cannot be extended by appending letters to their outer ends. For  $r' > r$  we have that  $w[i..i + \ell - r' - 1]v'w[j + r'..j + \ell - 1]$ , with  $v' = w[i + \ell - r' - 1..i + \ell - 1]vw[j..j + r' - 1]$ , is still a gapped palindrome, but it is not  $\alpha$ -gapped anymore. So, for a maximal  $\alpha$ -gapped palindrome  $p = w[i..i + \ell - 1]vw[j..j + \ell - 1]$ , we associate the interval  $I_p = [j, j + r]$ , and associate to it a weight  $g(I_p) = j + \ell - 1$ . Intuitively, we know that at each position  $j' \in I_p$  there exists a factor  $u$ , ending at position  $j + \ell - 1$ , such that  $u^R v$  is a suffix of  $w[1..j' - 1]$  for some  $v$ .

On the other hand, if  $u$  is the longest factor starting at some position  $j' \leq n$  such that  $u^R v$  is a suffix of  $w[1..j - 1]$ , then the factor  $w[i..j] = u^R v u$  is, in fact, a maximal  $\alpha$ -gapped palindrome  $x^R y x$  (i.e.,  $u^R$  is a prefix of  $x^R$  and  $u$  is a suffix of  $x$ ). In other words,  $u$  and  $u^R$  could be extended simultaneously inside the gap, but not at the outer ends.

Consequently, to compute  $LPal_\alpha[i]$  for some  $i \leq n$  we have to find the  $\alpha$ -gapped palindromes  $p \in S$  for which the interval  $I_p$  contains  $i$ . Then, we identify which of these intervals has the greatest weight. Say, for instance, that the interval  $I_p$  which contains  $i$ , is the one that weight maximal weight  $k$  from all the intervals containing  $i$ . Then  $LPal_\alpha[j] = k - j + 1$ . Indeed, from all the factors  $u$  starting at position  $j$ , such that  $u^R v$  is a suffix of  $w[1..j' - 1]$  for some  $v$ , there is one that ends at position  $k$ , while all the other end before  $k$  (otherwise, the intervals associated, respectively, to the maximal  $\alpha$ -gapped palindromes containing each of these factors  $u^R v u$  would have a greater weight). So, the palindrome ending at position  $k$  is the longest of them all.

This allows us to design the following algorithm for the computation of  $LPal_\alpha[j]$ . We first use the algorithm of Kolpakov and Kucherov (2009a) to compute the set  $S$  of all maximal  $\alpha$ -gapped palindromes of  $w$ . For each maximal  $\alpha$ -gapped palindrome  $p = w[i..i + \ell - 1]w[i + \ell..j - 1]w[j..j + \ell - 1]$ , we associate the interval  $I_p = [j, j + r]$ , where  $r = \lfloor \frac{(\alpha-1)\ell-\delta}{\alpha+1} \rfloor$  and  $\delta = j - i - \ell$ , and associate to it the weight  $g(I_p) = j + \ell - 1$ . We process these  $|S|$  intervals, with weights and bounds in  $[1, n]$ , in  $\mathcal{O}(n + |S|)$  time as in Lemma 4, to compute for each  $j \leq n$  the maximal weight  $H[j]$  of an interval containing  $j$ . Then we set  $LPal_\alpha[j] = H[j] - j + 1$ .

The correctness of the above algorithm follows from the remarks at the beginning of this proof. Its complexity is clearly  $\mathcal{O}(\alpha n)$ .  $\square$

The solution of Problem 3(b) is very similar.

**Theorem 10** *Problem 3(b) can be solved in  $\mathcal{O}(\alpha n)$  time.*

**Proof:** We first use the algorithm of Gawrychowski et al. (2016) to compute the set  $S$  of all maximal  $\alpha$ -gapped repeats with non-empty gap of  $w$ . For each maximal  $\alpha$ -gapped repeat  $p = w[i..i + \ell - 1]w[i + \ell..j - 1]w[j..j + \ell - 1]$ , we associate the interval  $I_p = [j, j + r]$ , where  $r = \left\lfloor \frac{(\alpha-1)\ell-\delta}{\alpha} \right\rfloor$  and  $\delta = j - i - \ell$ , and associate to it the weight  $g(I_p) = j + \ell - 1$ . We process these  $|S|$  intervals, with weights and bounds in  $[1, n]$ , in  $\mathcal{O}(n + |S|)$  time as in Lemma 4, to compute for each  $j \leq n$  the maximal weight  $H[j]$  of an interval containing  $j$ . Now, we use Lemma 5 to compute the values  $SC[j]$  for each  $j \leq n$ . We set  $LRep_\alpha[j] = \max\{H[j] - j + 1, SC[j]\}$ .

The complexity of this algorithm is  $\mathcal{O}(\alpha n)$ , as  $|S| \in \mathcal{O}(\alpha n)$  (see Gawrychowski et al. (2016)).

The correctness of the algorithm follows from the following remark. For a maximal  $\alpha$ -gapped repeat  $p = w[i..i + \ell - 1]w[i + \ell..j - 1]w[j..j + \ell - 1]$  let  $r = \left\lfloor \frac{(\alpha-1)\ell-\delta}{\alpha} \right\rfloor$ , where  $\delta = j - i - \ell$ . Then the factors  $w[i + r'..i + \ell - 1]w[i + \ell..j + r' - 1]w[j + r'..j + \ell - 1]$  are  $\alpha$ -gapped repeats for all  $r' \leq r$ , whose right arm cannot be extended anymore to the right. Moreover, the factors  $w[i + r'..i + \ell - 1]w[i + \ell..j + r' - 1]w[j + r'..j + \ell - 1]$  are gapped repeats which are not  $\alpha$ -gapped for all  $r' > r$ . The rest of the arguments showing the soundness of our algorithm are similar to those of Theorem 9.  $\square$

## 6 Future Work

In this paper we proposed a series of algorithms that construct data structures giving detailed information on the longest gapped repeats and palindromes occurring in a given word. There are several directions in which the work presented here can be continued.

Firstly, it seems interesting to us whether Problem 1(b) (the construction of the array  $LPF_{g,G}[\cdot]$ ) can be solved faster. An intermediate problem could be to check whether we can find in linear time the longest gapped repeats with the length of the gap between a given lower bound and a given upper bound.

Secondly, although the algorithms we propose in Theorems 7 and 8 do not rely on computing and going through all the maximal  $\alpha$ -gapped repeats and palindromes when looking for the longest such structure, they have asymptotically the same complexity as the (optimal) algorithms finding all such structures. Thus, it seems natural and interesting to design algorithms finding the longest  $\alpha$ -gapped repeat or palindrome of a word that run faster than the algorithms we proposed here (and, in particular, than the algorithms finding all these structures). Also, it is interesting whether we can construct the data structures defined in Problem 3 without producing first the list of all  $\alpha$ -gapped repeats and palindromes.

Lastly, following the problems studied by Gawrychowski et al. (2013a, 2014), one could be interested in finding the longest gapped pseudo-repeats. More precisely, for a literal bijective anti-/morphism  $f$ , we want to find the longest word (or words)  $u$  such that a given word  $w$  contains a factor  $f(u)vu$  with  $|v|$  subject to different length-restrictions (e.g., between a lower and an upper bound, or shorter than  $|u|$  multiplied by a factor, like in the case of  $\alpha$ -gapped repeats and palindromes, etc.). Such repeats and palindromes are sometimes used to formalise repeats and palindromes occurring in DNA sequences. In this setting one works with the alphabet  $\{A, C, G, T\}$ . When we are interested in direct repeats we may take  $f$  work as a morphism and model the Watson-Crick complementarity:  $f(A) = T, f(C) = G, f(G) = C$ , and  $f(T) = A$ ; when we are interested in inverted repeats in the genetic sequence we may take  $f$  work as an antimorphism, still defined by the Watson-Crick complementarity. It is not hard to see

that our algorithms can also be adapted in a straightforward manner to work in the context of such gapped pseudo-repeats.

## Acknowledgements

The authors thank the anonymous referees of this paper, as well as those of the conference papers which we extend here, for their valuable remarks, suggestions, and comments, that improved the quality of this manuscript. The work of Florin Manea was supported by the DFG grant 596676.

## References

- D. Angluin. Finding patterns common to a set of strings. *J. Comput. Syst. Sci.*, 21(1):46–62, 1980.
- T. M. Apostol. *Introduction to analytic number theory*. Springer, 1976.
- H. Bannai, T. I. S. Inenaga, Y. Nakashima, M. Takeda, and K. Tsuruta. The Runs Theorem. *CoRR*, abs/1406.0263, 2014.
- H. Bannai, T. I. S. Inenaga, Y. Nakashima, M. Takeda, and K. Tsuruta. A new characterization of maximal repetitions by lyndon trees. In *Proc. SODA*, pages 562–571, 2015.
- M. A. Bender and M. Farach-Colton. The LCA problem revisited. In *Proc. LATIN*, volume 1776 of *LNCS*, pages 88–94, 2000.
- J. Berstel and L. Boasson. Partial words and a theorem of Fine and Wilf. *Theoret. Comput. Sci.*, 218: 135–141, 1999.
- G. S. Brodal, R. B. Lyngsø, C. N. S. Pedersen, and J. Stoye. Finding maximal pairs with bounded gap. In *Proc. CPM*, volume 1645 of *LNCS*, pages 134–149. Springer, 1999.
- E. Chiniforooshan, L. Kari, and Z. Xu. Pseudopower avoidance. *Fundamenta Informaticae*, 114(1): 55–72, 2012.
- C. Choffrut and J. Karhumäki. Combinatorics of words. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 1, pages 329–438. Springer-Verlag, 1997.
- S. Constantinescu and L. Ilie. Generalised Fine and Wilf’s theorem for arbitrary number of periods. *Theoret. Comput. Sci.*, 339(1):49–60, 2005.
- S. Constantinescu and L. Ilie. Fine and Wilf’s theorem for abelian periods. *Bulletin of EATCS*, 89: 167–170, 2006.
- T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms (3. ed.)*. MIT Press, 2009.
- M. Crochemore. An optimal algorithm for computing the repetitions in a word. *Inf. Process. Lett.*, 12(5): 244–250, 1981.

- M. Crochemore and L. Ilie. Computing longest previous factor in linear time and applications. *Inf. Process. Lett.*, 106(2):75–80, Apr. 2008.
- M. Crochemore and W. Rytter. Usefulness of the Karp-Miller-Rosenberg algorithm in parallel computations on strings and arrays. *Theoret. Comput. Sci.*, 88(1):59–82, 1991.
- M. Crochemore and W. Rytter. Squares, cubes, and time-space efficient string searching. *Algorithmica*, 13(5):405–425, 1995.
- M. Crochemore and W. Rytter. *Jewels of Stringology: Text algorithms*. World Scientific, Singapore, 2002.
- M. Crochemore and G. Tischler. Computing longest previous non-overlapping factors. *Inf. Process. Lett.*, 111(6):291–295, Feb. 2011.
- M. Crochemore, C. Hancart, and T. Lecroq. *Algorithms on strings*. Cambridge University Press, 2007.
- M. Crochemore, C. S. Iliopoulos, M. Kubica, W. Rytter, and T. Walen. Efficient algorithms for two extensions of LPF table: The power of suffix arrays. In *Proc. SOFSEM 2010*, volume 5901 of *LNCS*, pages 296–307, 2010.
- M. Crochemore, L. Ilie, C. S. Iliopoulos, M. Kubica, W. Rytter, and T. Walen. Computing the longest previous factor. *Eur. J. Comb.*, 34(1):15–26, 2013.
- M. Crochemore, R. Kolpakov, and G. Kucherov. Optimal bounds for computing  $\alpha$ -gapped repeats. In *Proc. LATA 2016*, volume 9618 of *LNCS*, pages 245–255, 2016.
- E. Czeizler, L. Kari, and S. Seki. On a special class of primitive words. *Theoret. Comput. Sci.*, 411:617–630, 2010.
- M. Dumitran and F. Manea. Longest gapped repeats and palindromes. In *Proc. MFCS 2015*, volume 9234 of *LNCS*, pages 205–217. Springer, 2015.
- J. Duval, R. Kolpakov, G. Kucherov, T. Lecroq, and A. Lefebvre. Linear-time computation of local periods. In *Proc. MFCS*, volume 2747 of *LNCS*, pages 388–397, 2003.
- A. Ehrenfeucht and G. Rozenberg. Finding a Homomorphism Between Two Words is NP-Complete. *Inf. Process. Lett.*, 9(2):86–88, 1979.
- M. Farach. Optimal suffix tree construction with large alphabets. In *Proc. FOCS 1997*, pages 137–143, 1997.
- N. J. Fine and H. S. Wilf. Uniqueness theorem for periodic functions. *Proceedings of the American Mathematical Society*, 16:109–114, 1965.
- 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.
- H. N. Gabow and R. E. Tarjan. A linear-time algorithm for a special case of disjoint set union. In *Proc. STOC*, pages 246–251, 1983.

- P. Gawrychowski. Pattern matching in Lempel-Ziv compressed strings: Fast, simple, and deterministic. In *Proc. ESA*, volume 6942 of *LNCS*, pages 421–432, 2011.
- P. Gawrychowski and F. Manea. Longest  $\alpha$ -gapped repeat and palindrome. In *Proc. FCT 2015*, volume 9210 of *LNCS*, pages 27–40. Springer, 2015.
- P. Gawrychowski, F. Manea, R. Mercas, D. Nowotka, and C. Tiseanu. Finding Pseudo-repetitions. In *Proc. STACS 2013*, volume 20 of *LIPIcs*, pages 257–268, 2013a.
- P. Gawrychowski, F. Manea, and D. Nowotka. Discovering hidden repetitions in words. In *Proc. CiE*, page to appear, 2013b.
- P. Gawrychowski, F. Manea, and D. Nowotka. Testing generalised freeness of words. In *Proc. STACS 2014*, volume 25 of *LIPIcs*, pages 337–349, 2014.
- P. Gawrychowski, T. I. S. Inenaga, D. Köppl, and F. Manea. Efficiently finding all maximal alpha-gapped repeats. In *Proc. STACS 2016*, volume 47 of *LIPIcs*, pages 39:1–39:14, 2016.
- T. H. Grönwall. Some asymptotic expressions in the theory of numbers. *Trans. Amer. Math. Soc.*, 14: 113–122, 1913.
- D. Gusfield. *Algorithms on strings, trees, and sequences: computer science and computational biology*. Cambridge University Press, New York, NY, USA, 1997.
- J. Justin. On a paper by Castelli, Mignosi, Restivo. *ITA*, 34(5):373–377, 2000.
- L. Kari and S. Seki. An improved bound for an extension of Fine and Wilf’s theorem, and its optimality. *Fundamenta Informaticae*, 191:215–236, 2010.
- J. Kärkkäinen, P. Sanders, and S. Burkhardt. Linear work suffix array construction. *J. ACM*, 53:918–936, 2006.
- D. E. Knuth. *The Art of Computer Programming*, volume 1: Fundamental Algorithms. Addison-Wesley, 1968.
- D. E. Knuth, J. H. M. Jr., and V. R. Pratt. Fast pattern matching in strings. *SIAM J. Comput.*, 6(2): 323–350, 1977.
- T. Kociumaka, J. Radoszewski, W. Rytter, and T. Walen. Efficient data structures for the factor periodicity problem. In *Proc. SPIRE*, volume 7608 of *LNCS*, pages 284–294, 2012.
- R. Kolpakov and G. Kucherov. Finding maximal repetitions in a word in linear time. In *Proc. FOCS*, pages 596–604, 1999.
- R. Kolpakov and G. Kucherov. Searching for gapped palindromes. *Theor. Comput. Sci.*, 410(51):5365–5373, 2009a.
- R. Kolpakov and G. Kucherov. Searching for gapped palindromes. *Theoretical Computer Science*, 410(51):5365 – 5373, 2009b. ISSN 0304-3975. Combinatorial Pattern Matching.

- R. Kolpakov, M. Podolskiy, M. Posypkin, and N. Khrapov. Searching of gapped repeats and subrepetitions in a word. In *Proc. CPM*, volume 8486 of *LNCS*, pages 212–221, 2014.
- R. M. Kolpakov and G. Kucherov. Finding repeats with fixed gap. In *Proc. SPIRE*, pages 162–168, 2000.
- S. R. Kosaraju. Computation of squares in a string (preliminary version). In *Proc. CPM*, volume 807 of *LNCS*, pages 146–150, 1994.
- D. Kosolobov. Lempel-Ziv factorization may be harder than computing all runs. In *Proc. STACS*, volume 30 of *LIPICs*, pages 582–593, 2015.
- M. Lothaire. *Combinatorics on Words*. Cambridge University Press, 1997.
- M. G. Main. Detecting leftmost maximal periodicities. *Discrete Applied Mathematics*, 25(1-2):145–153, 1989.
- G. K. Manacher. A new linear-time on-line algorithm for finding the smallest initial palindrome of a string. *J. ACM*, 22(3):346–351, 1975.
- F. Manea, R. Mercas, and D. Nowotka. Fine and Wilf’s theorem and pseudo-repetitions. In *Proc. MFCS*, volume 7464 of *LNCS*, pages 668–680. Springer, 2012a.
- F. Manea, M. Müller, and D. Nowotka. The avoidability of cubes under permutations. In *Proc. DLT*, volume 7410 of *LNCS*, pages 416–427. Springer, 2012b.
- G. K. Maxime Crochemore, Roman Kolpakov. Optimal searching of gapped repeats in a word. *ArXiv e-prints*, 2015.
- J. Shallit. <http://www.cs.uwaterloo.ca/~shallit/Talks/wilf3.pdf>.
- E. B. S.R. Chan. Telomeres and telomerase. volume 359 of *Philos. Trans. R. Soc. Lond. B. Biol. Sci.*, pages 109–121, 2004.
- Y. Tanimura, Y. Fujishige, T. I. S. Inenaga, H. Bannai, and M. Takeda. A faster algorithm for computing maximal  $\alpha$ -gapped repeats in a string. In *Proc. SPIRE 2015*, volume 9309 of *LNCS*, pages 124–136. Springer, 2015.
- P. van Emde Boas. Preserving order in a forest in less than logarithmic time. In *Proc. FOCS*, pages 75–84, 1975.
- Z. Xu. A minimal periods algorithm with applications. In *Proc. CPM*, volume 6129 of *LNCS*, pages 51–62, 2010.