

Palindromic Length in Linear Time

Kirill Borozdin¹, Dmitry Kosolobov², Mikhail Rubinchik³, and
Arseny M. Shur⁴

- 1 Ural Federal University, Ekaterinburg, Russia
borozdin.kirill@gmail.com
- 2 University of Helsinki, Helsinki, Finland
dkosolobov@mail.ru
- 3 Ural Federal University, Ekaterinburg, Russia
mikhail.rubinchik@gmail.com
- 4 Ural Federal University, Ekaterinburg, Russia
arseny.shur@urfu.ru

Abstract

Palindromic length of a string is the minimum number of palindromes whose concatenation is equal to this string. The problem of finding the palindromic length drew some attention, and a few $O(n \log n)$ time online algorithms were recently designed for it. In this paper we present the first linear time online algorithm for this problem.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.1 Combinatorics

Keywords and phrases palindrome, palindromic length, palindromic factorization, online

Digital Object Identifier 10.4230/LIPIcs.CPM.2017.23

1 Introduction

Algorithmic and combinatorial problems involving palindromes attracted the attention of researchers since the first days of stringology. Recall that a string $w = a_0a_1 \cdots a_{n-1}$ is a *palindrome* if it is equal to the string $\tilde{w} = a_{n-1} \cdots a_1a_0$. The early works [4, 6, 8, 11] considered palindromes as structures that might provide examples of (context-free) languages that are impossible to recognize in linear time, thus provably restricting the computational power of some models (RAM, in particular). Subsequently, it was shown that many of such languages are, in fact, linear recognizable. Recently it was proved [7] that the language \mathcal{P}^k , where \mathcal{P} is the set of all palindromes on a given alphabet, is recognizable online in $O(kn)$ time, where n is the length of the input string. Roughly at the same time, a closely related notion of *palindromic length* of a string was introduced: this is the minimal number k such that the string belongs to \mathcal{P}^k . In 2014–2015 three different algorithms that compute the palindromic length of a string of length n in $O(n \log n)$ time were presented in [3, 5, 10] (however, they all are based on similar principles). In this paper we present the first linear algorithm computing the palindromic length. Moreover, our algorithm is *online*, i.e., it reads the input string sequentially from left to right and computes the palindromic length for each prefix after reading the rightmost letter of that prefix. Thus, we prove the following theorem.

► **Theorem 1.** *Palindromic length of a string is computable online in linear time.*

The implementation of our algorithm and tests for it can be found in [9]. Due to a large constant under the big-O, it is slower in practice (for 32/64 bit machine words) than the existing $O(n \log n)$ solutions; the fastest algorithm is the one of [10].



© Kirill Borozdin, Dmitry Kosolobov, Mikhail Rubinchik, and Arseny M. Shur;
licensed under Creative Commons License CC-BY

28th Annual Symposium on Combinatorial Pattern Matching (CPM 2017).

Editors: Juha Kärkkäinen, Jakub Radoszewski, and Wojciech Rytter; Article No. 23; pp. 23:1–23:12

Leibniz International Proceedings in Informatics



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

The paper is organized as follows. Section 2 contains a high-level description of the algorithm: it starts with a naive $O(n^2)$ algorithm, then improves the time to $O(n \log n)$, and, finally, describes on a high level a modified $O(n)$ -time version of the $O(n \log n)$ algorithm. In Section 3 we discuss the main components of the linear algorithm in details.

1.1 Preliminaries

Let w be a string of length $n = |w|$. We write $w[i]$ for the i th letter of w ($i = 0, \dots, n-1$) and $w[i..j]$ for $w[i]w[i+1] \cdots w[j]$. A string u is a *substring* of w if $u = w[i..j]$ for some i, j . Such pair (i, j) is not necessarily unique; i specifies an *occurrence* of u at *position* i . A substring $w[0..j]$ (resp., $w[i..n-1]$) is a *prefix* (resp. *suffix*) of w . The *empty string* is denoted by ε . For any i, j , $[i..j]$ denotes the set $\{k \in \mathbb{Z} : i \leq k \leq j\}$; let $(i..j) = [i..j] \setminus \{i\}$, $[i..j) = [i..j] \setminus \{j\}$, $(i..j) = [i..j) \cap (i..j]$. Our notation for arrays is the same as for strings.

A substring (resp. suffix, prefix) that is a palindrome is called a *subpalindrome* (resp. *suffix-palindrome*, *prefix-palindrome*). If $w[i..j]$ is a subpalindrome of w , then the number $(j+i)/2$ is the *center* of $w[i..j]$ and the number $\lfloor (j-i+1)/2 \rfloor$ is the *radius* of $w[i..j]$. The following remarkable property of palindromic lengths is crucial for our algorithm.

► **Lemma 2** (see [10, Lemma 11]). *Denote by $\ell_0, \ell_1, \dots, \ell_{n-1}$, resp., the palindromic lengths of the prefixes $w[0..0], w[0..1], \dots, w[0..n-1]$ of a string w . Then, for any $i \in (0..n)$, $|\ell_i - \ell_{i-1}| \leq 1$.*

An integer p is a *period* of w if $w[i] = w[i+p]$ for any $i \in [0..n-p)$. As the previous results [3, 5, 10], our approach relies on a number of periodic properties of palindromes.

► **Lemma 3** (see [7, Lemmas 2, 3]). *For any palindrome w and any $p \in (0..|w|]$, the following conditions are equivalent: (1) p is a period of w , (2) there are palindromes u, v such that $|uv| = p$ and $w = (uv)^k u$ for some $k \geq 1$, (3) $w[p..|w|-1]$ ($w[0..|w|-p-1]$) is a palindrome.*

► **Lemma 4** (see [7, Lemma 7]). *Suppose that $w = (uv)^k u$ for $k \geq 1$ and for palindromes u and v such that $|uv|$ is the minimal period of w ; then, the center of any subpalindrome x of w such that $|x| \geq |uv|-1$ coincides with the center of some u or v from the decomposition.*

Henceforth, let s denote the input string of length n . We assume that the algorithm works in the unit-cost word-RAM model with $\Theta(\log n)$ -bit machine words (an assumption justified in, e.g., [2]) and standard operations like in the C programming language.

2 High-Level Description of the Algorithm

Our aim is to maintain an array $\text{ans}[0..n-1]$ in which each element $\text{ans}[i]$ is the palindromic length of $s[0..i]$. We always assume n to be the length of the string s processed so far (i.e., $s = s[0..n-1]$). Processing the next letter $s[n]$, we compute $\text{ans}[n]$ and then increment n .

2.1 Naive approach

An easy quadratic-time approach is to maintain the list of all non-empty suffix-palindromes u_1, \dots, u_k of the string s and calculate $\text{ans}[n] = 1 + \min_{i \in [1..k]} \text{ans}[n - |u_i|]$. The list can be updated in linear time: the suffix-palindromes of wa have the form aua , where u is a suffix palindrome of w , plus the palindrome a and, optionally, aa . As a first speedup to this basic approach, we utilize the (palindromic) *iterator*, introduced in [7]; this data structure contains a string s and supports the following operations:

1. $\text{add}(a)$ appends the letter a to the end of s ;
2. maxPal returns the center of the longest suffix-palindrome of s ;
3. $\text{rad}(x)$ returns the radius of the longest subpalindrome of s with the center x ;
4. $\text{nextPal}(x)$ returns the center of the longest proper suffix-palindrome of the suffix-palindrome of s with the center x ; undefined if x is not the center of a suffix-palindrome.

The iterator can be implemented so that all its operations work in $O(1)$ time (amortized, for add) [7, Prop. 1]. The same time bound applies to computing length of the longest subpalindrome centered at x : $\text{len}(x) = 2 \cdot \text{rad}(x) + \lfloor x \rfloor - \lfloor x - \frac{1}{2} \rfloor$. Still, the iterator alone cannot lower the asymptotic time of the naive algorithm; its improved version looks as follows:

- 1: $\text{add}(s[n]); \text{ans}[n] \leftarrow +\infty$
- 2: **for** ($x \leftarrow \text{maxPal}; x \neq n + \frac{1}{2}; x \leftarrow \text{nextPal}(x)$) **do** $\text{ans}[n] \leftarrow \min\{\text{ans}[n], 1 + \text{ans}[n - \text{len}(x)]\}$

2.2 Algorithm working in $O(n \log n)$ time

All subquadratic algorithms for palindromic length heavily use grouping of suffix-palindromes into *series*. Let u_1, \dots, u_k be all non-empty suffix-palindromes of a string s in the order of decreasing length. Since u_j is a suffix of u_i for any $i < j$, any period of u_i is a period of u_j ; hence the sequence of minimal periods of u_1, \dots, u_k is non-increasing. The groups of suffix-palindromes with the same minimal period are *series of palindromes* (of s):

$$\underbrace{u_1, \dots, u_{i_1}}_{p_1}, \underbrace{u_{i_1+1}, \dots, u_{i_2}}_{p_2}, \dots, \underbrace{u_{i_{t-1}+1}, \dots, u_k}_{p_t}.$$

We refer to the longest and the shortest palindrome in a series as its *head* and *baby* respectively (they coincide in the case of a 1-element series); we enumerate the elements of a series from the head to the baby. Given an integer p , the p -series is the series with period p . A very useful observation [3, 5, 7] is that the length of a head is multiplicatively smaller than the length of the baby from the previous series, and thus every string of length n has $O(\log n)$ series. (As it was shown in [3], strings with $\Omega(\log n)$ series for $\Omega(n)$ prefixes do exist.)

The idea of the $O(n \log n)$ solution is to use the dynamic programming rule $\text{ans}[n] = 1 + \min_U \min_{u \in U} \text{ans}[n - |u|]$, where U runs through the series of s , and compute the internal minimum in $O(1)$ time using precalculations based on the structure of series. The structure of any series is described in the following lemma, which is easily implied by Lemmas 3, 4.

► **Lemma 5.** *For a string s and $p \geq 1$, let U be a p -series of palindromes. There exist $k \geq 1$ and unique palindromes u, v with $|uv| = p$, $v \neq \varepsilon$ such that one of three conditions hold:*

- $U = \{(uv)^{k+1}u, (uv)^k u, \dots, (uv)^2 u\}$ and the next series begins with uvu ,
- $U = \{(uv)^k u, (uv)^{k-1}u, \dots, uvu\}$ and the next series begins with u ,
- $U = \{v^k, v^{k-1}, \dots, v\}$, $p = 1$, $|v| = 1$, $u = \varepsilon$, and U is the last series for s .

Let U be a p -series for $s[0..n]$ with $k > 1$ palindromes (w.l.o.g., $U = \{(uv)^k u, \dots, uvu\}$). Updating $\text{ans}[n]$ using this series, we compute $m = \min\{\text{ans}[n - kp - |u|], \dots, \text{ans}[n - p - |u|]\}$. Now note that $s[0..n]$ ends with $(uv)^k u$ but not with $(uv)^{k+1}u$: otherwise, the latter string would belong to U . Then $s[0..n-p]$ ends with $(uv)^{k-1}u$ but not with $(uv)^k u$ and thus has the p -series $U' = \{(uv)^{k-1}u, \dots, uvu\}$. Thus, at that iteration we computed $m' = \min\{\text{ans}[n - kp - |u|], \dots, \text{ans}[n - 2p - |u|]\}$ for updating $\text{ans}[n - p]$. If we save m' into an auxiliary array, then $m = \min\{m', \text{ans}[n - p - |u|]\}$ is computable in constant time, as required. Let us implement this construction using the iterator.

We start an iteration calling $\text{add}(s[n])$. Let x be the center of a suffix-palindrome u . By Lemma 3, the minimal period p of u equals $\text{len}(x) - \text{len}(\text{nextPal}(x))$. Let $\text{cntr}(d)$ denote

the center of the length d suffix-palindrome of $s[0..n]$ (i.e., $\text{cntr}(d) = n - (d - 1)/2$). Let $x' = \text{cntr}(p + (\text{len}(x) \bmod p))$. By Lemma 5, x' is the center either of the baby of the p -series or of the head of the next series, depending on the period $\text{len}(x') - \text{len}(\text{nextPal}(x'))$ of this suffix-palindrome. All these computations take $O(1)$ amortized time using the iterator.

Our algorithm maintains an array $\text{left}[1..n]$: for $p \in [1..n]$, if there is a p -series, then $s[\text{left}[p]+1..n]$ is the longest suffix (which is not necessarily a palindrome) of $s[0..n]$ with period p ; otherwise, $\text{left}[p]$ is undefined. E.g., if $s[0..n] = \cdots aaabaaba$ and $p = 3$, then the mentioned suffix is $s[n-6..n] = aabaaba$ and $\text{left}[3] = n - 7$. Computing $\text{left}[p]$ in $O(1)$ time is done as follows. Let $w = (uv)^k u$ be the head of the p -series (see Lemma 5), x be the center of w , and $z(uv)^k u$ be the longest suffix of $s[0..n]$ with period p (in our example, $u = \varepsilon$, $v = aba$, $x = n - 5/2$, $z = a$). Then z is a proper suffix of uv . Hence $\text{len}(x_1) = 2|z| + |u|$, where x_1 is the center of the prefix-palindrome u of w (in the example, $x_1 = n - 11/2$, $\text{len}(x_1) = |aa|$). Note that $|u| = \text{len}(x) \bmod p$ and $x_1 = 2x - x_2$, where $x_2 = \text{cntr}(\text{len}(x) \bmod p)$ is the center of the suffix u of w . Thus, $|z|$ and $\text{left}[p] = n - \text{len}(x) - |z|$ are computed in $O(1)$ time.

All precalculated minimums are stored in an array $\text{pre}[1..n]$, where each $\text{pre}[p]$ is, in turn, an array $\text{pre}[p][0..p-1]$ (we discuss in the next subsection why only $O(n)$ of possible $O(n^2)$ elements of pre are actually stored). For each j such that $n - j > \text{left}[p]$, the string $s[0..n-j]$ usually has a suffix-palindrome with period p and thus can have a p -series; the array $\text{pre}[p][0..p-1]$ contains the precalculations made for all these series. Formally,

$$\begin{aligned} \text{pre}[p][i] &= \min\{\text{ans}[t] : \\ &\quad (t - \text{left}[p]) \bmod p = i \text{ and } s[t+1..n] \text{ has a prefix-palindrome of minimal period } p\}; \end{aligned}$$


$\text{pre}[p][i]$ is undefined if there is no such t (i.e., no p -series for the corresponding string). So if u_1, \dots, u_k is a p -series for $s[0..n]$, then $\text{pre}[p][n - |u_1| - \text{left}[p]] = \min\{\text{ans}[n - |u_i|] : i \in [1..k]\}$. Hence, given a new letter $s[n]$, we compute $\text{ans}[n]$ as follows:

- 1: $\text{add}(s[n]); \text{ans}[n] \leftarrow +\infty;$
- 2: **for** ($x \leftarrow \text{maxPal}; x \neq n + \frac{1}{2}; x \leftarrow \text{nextPal}(\text{cntr}(d))$) **do** \triangleright goes to next head each time
- 3: $p \leftarrow \text{len}(x) - \text{len}(\text{nextPal}(x));$ \triangleright min. period of the suf.-pal. centered at x
- 4: $d \leftarrow p + (\text{len}(x) \bmod p);$ \triangleright length of the baby in the p -series
- 5: **if** $\text{len}(\text{cntr}(d)) - \text{len}(\text{nextPal}(\text{cntr}(d))) \neq p$ **then** $d \leftarrow d + p;$ \triangleright corrected length
- 6: **compute** $\text{left}[p];$ \triangleright in $O(1)$ time, see above
- 7: **if** $\text{len}(x) = d$ **then** $\text{pre}[p][n - \text{len}(x) - \text{left}[p]] \leftarrow \text{ans}[n - d];$
- 8: $\text{pre}[p][n - \text{len}(x) - \text{left}[p]] \leftarrow \min\{\text{pre}[p][n - \text{len}(x) - \text{left}[p]], \text{ans}[n - d]\};$
- 9: $\text{ans}[n] \leftarrow \min\{\text{ans}[n], 1 + \text{pre}[p][n - \text{len}(x) - \text{left}[p]]\};$

Let u_1, \dots, u_k be a p -series, $i = n - |u_1| - \text{left}[p]$. If $k = 1$, there was no p -series p iterations ago, so we set the undefined value $\text{pre}[p][i]$ to $\text{ans}[n - |u_k|]$ in line 7. Otherwise, by the definition of pre , we have $\text{pre}[p][i] = \min\{\text{ans}[n - |u_1|], \dots, \text{ans}[n - |u_{k-1}|]\}$. We update this value using $\text{ans}[n - |u_k|]$ in line 8. So pre is correctly maintained, and the above algorithm computes the array ans in $O(n \log n)$ time due to logarithmic number of series.

2.3 Sketch of the linear algorithm

The idea of the linear solution is to perform the above log-time processing of all series of the current string not n times, but only $O(\frac{n}{\log n})$ times during the run of the algorithm. (However, we are able to make $\Theta(n)$ calls to the iterator.) To achieve this, during the processing of a series we replace each computation of the minimum $\text{ans}[n] \leftarrow \min\{\text{ans}[n], 1 + z\}$, for a precomputed value z from pre , with the simultaneous computation (“prediction”) for the range of values $\text{ans}[n..n+b]$, where $b = \lfloor \frac{\log n}{8} \rfloor$: we compute in advance $\text{ans}[j] \leftarrow \min\{\text{ans}[j], 1 + z_j\}$



■ **Figure 1** Predictable extensions.

for all $j \in [n..n+b]$ and corresponding precomputed z_j from `pre`. It is proved below that the arrays `ans` and `pre` can be organized so that, after a linear time preprocessing, such range operations on $O(b)$ elements of `ans` will take $O(1)$ time (this type of bit compression techniques is referred to as the four Russians' trick [1]).

Let us extend $s[0..n-1]$ with $s[n] = a$. We say that a suffix-palindrome u of $s[0..n-1]$, centered at x , *survives* if $s[0..n]$ has the suffix aua (i.e. x remains the center of a suffix-palindrome), and *dies* otherwise. We say that an extension of $s[0..n-1]$ by $s[n]$ is *predictable* if it retains `maxPal`, i.e., if the longest suffix-palindrome survives. From `maxPal` it can be calculated which of the other suffix-palindromes survive. If a suffix-palindrome of s centered at x survives $d \geq 0$ consecutive predictable extensions but dies after the $(d+1)$ th such extension (or the $(d+1)$ th predictable extension is not possible), we write $\text{live}(x) = d$. We have $\text{live}(\text{maxPal}) = n - \text{len}(\text{maxPal})$ and $\text{live}(x) = \text{rad}(\text{refl}(x)) - \text{rad}(x)$ for $x \neq \text{maxPal}$; here $\text{refl}(x) = 2 \cdot \text{maxPal} - x$ is the position symmetric to x w.r.t. `maxPal`. (See Fig. 1 for clarification; e.g., in Fig. 1 $\text{live}(x) = 2$ and $\text{live}(\text{maxPal}) = 6$.)

Suppose that $\text{ans}[n+j] = +\infty$ for $j \in [0..b]$. Having performed `add(s[n])`, we get access to the suffix-palindromes of $s[0..n]$. If, for the center x of each such palindrome, we perform

$$\text{ans}[n+j] \leftarrow \min\{\text{ans}[n+j], 1 + \text{ans}[n - \text{len}(x) - j]\} \text{ for all } j \in [0.. \min\{b, \text{live}(x)\}], \quad (1)$$

then we accumulate all information we can obtain from these palindromes during the next b predictable extensions. Thus we get an approximation of $\text{ans}[n..n+b]$, which later will be updated using suffix-palindromes with the centers $x \geq n + \frac{1}{2}$. One phase of our algorithm is roughly as follows:

- append $s[n]$ to the iterator, update precalculations, and “predict” $\text{ans}[n..n+b]$ with the assignments (1), using operations on blocks of bits ($\text{ans}[n]$ is computed exactly);
- append subsequent letters, each time updating the predictions with either one or two new palindromes (after processing $s[n+j]$, $\text{ans}[n..n+j]$ contains correct values);
- stop after b iterations or at the moment when an unpredicted letter is encountered;
- discard unused predictions and start a new phase with the first unpredicted letter.

For arrays α, β and numbers $i, j, \ell \geq 0$, denote by $\alpha[i..i+\ell] \stackrel{\min}{\leftarrow} \beta[j..j+\ell]$ the sequence of assignments $\alpha[i+k] \leftarrow \min\{\alpha[i+k], \beta[j+k]\}$ for all $k \in [0..\ell]$. Let `increv`(i, j) be the function returning an array $a[0..j-i]$ such that $a[k] = 1 + \text{ans}[j-k]$ for $k \in [0..j-i]$ (“increment & reverse”). The predictions are made by the function `predict` that uses precalculations stored in `pre` to perform in a fast way the assignments $\text{ans}[n..n+c] \stackrel{\min}{\leftarrow} \text{increv}(n - \text{len}(x) - c, n - \text{len}(x))$, where $c = \min\{b, \text{live}(x)\}$, for all centers x of suffix-palindromes. (Hence `predict` computes the value $\text{ans}[n]$ correctly even if $c = 0$ for some x .) Let `precalc` be a function that updates (possibly once in several iterations) the array `pre` to the actual state. The implementations of `predict` and `precalc` are discussed in Section 3. Our algorithm is as follows:

- 1: **for** ($n \leftarrow 0, \text{end} \leftarrow 0$; **not**(`end_of_input`); $n \leftarrow n + 1$) **do**
- 2: **if** $n = \text{end}$ **or** $\text{len}(\text{maxPal}) = n$ **or** $s[n] \neq s[n - \text{len}(\text{maxPal}) - 1]$ **then** \triangleright new phase
- 3: `add(s[n]); precalc; predict; end` $\leftarrow n + b$
- 4: **else** `add(s[n])` \triangleright old phase continues, $s[n]$ is predictable
- 5: $c \leftarrow \min\{b, \text{live}(n)\}$; $\text{ans}[n..n+c] \stackrel{\min}{\leftarrow} \text{increv}(n-1-c, n-1)$
- 6: **if** $s[n] = s[n-1]$ **then** $c \leftarrow \min\{b, \text{live}(n-\frac{1}{2})\}$; $\text{ans}[n..n+c] \stackrel{\min}{\leftarrow} \text{increv}(n-2-c, n-2)$

This algorithm computes the same values $\text{ans}[n]$ as the $O(n \log n)$ algorithm above, because finally all suffix-palindromes of $s[0..n]$ are used. So, the algorithm is correct.

Let t be the number of series in the current string $s[0..n]$ and q is the time required to perform all the calls $\text{add}(s[n]), \text{add}(s[n-1]), \dots, \text{add}(s[n'+1])$, where $s[0..n']$ is the string for which precalc was called last time. Below we show that predict and precalc work in $O(t)$ and $O(t+q)$ time respectively, and the array ans can be organized so that the range operations in lines 5–6 can be performed in $O(1)$ time using the four Russians' trick. Let us estimate the running time of the algorithm under these assumptions.

During predictable extensions, line 3 is reached iff $n = \text{end}$, i.e., at most $O(\frac{n}{b})$ times. Since add works in $O(1)$ amortized time (see [7, Prop. 1]), the sum of all q 's in the working time of precalc is $O(n)$. Since $O(t) = O(\log n)$, all predictable extensions take $O(n + \frac{n}{b} \log n) = O(n)$ overall time. To estimate the running time of unpredictable extensions, consider the value $\gamma_i = \text{live}[\text{maxPal}] = i - \text{len}(\text{maxPal})$ after processing $s[0..i]$. If $s[i+1]$ is predictable, one has $\gamma_{i+1} = (i+1) - (\text{len}(\text{maxPal}) + 2) = \gamma_i - 1$. If $s[i+1]$ is unpredictable, $\gamma_{i+1} \geq (i+1) - (\text{len}(\text{nextPal}(\text{maxPal})) + 2)$; by Lemma 5, $\gamma_{i+1} - \gamma_i \geq p - 1$, where p is the minimal period of the longest suffix-palindrome of $s[0..i]$. By Lemmas 4 and 5, the length of the longest suffix-palindrome whose minimal period differs from p is less than $2p$. Therefore, predict and precalc take $O(p+q)$ time during this unpredictable extension (actually, $O(\log p + q)$). Since $\gamma_n - \gamma_1 < n$, the sum of the working times of all calls to predict and precalc is $O(n)$.

2.4 Organization of the arrays ans and pre

Informally, the four Russians' trick allows us to compute any operation on structures of size $\leq \varepsilon \log n$ bits in $O(1)$ time using a precomputed table of size $O(n^\varepsilon \log^{O(1)} n)$ bits. For example, let a $\lfloor \frac{\log n}{2} \rfloor$ -bit integer x encode a sequence $x_1, \dots, x_{\lfloor \log n/4 \rfloor}$ so that $x_j = 1 - (\lfloor x/2^{2j-2} \rfloor \bmod 4)$, i.e., $(2j-1)$ th and $(2j-2)$ th bits of x encode x_j . We can compute, for $j \in [1.. \log n/4]$, the sum $x_1 + \dots + x_j$ in $O(1)$ time using a table $T[0..\lfloor \sqrt{n} \rfloor][1..\log n/4]$ such that $T[x][j] = x_1 + \dots + x_j$ for any $x \in [0..2^{\log n/2}] = [0..\sqrt{n}]$ and $j \in [1.. \log n/4]$. The table T can be precomputed in $O(\sqrt{n} \log^{O(1)} n)$ time.

In our case, we split ans into blocks of length b . By Lemma 2, adjacent elements of ans differ by at most one. This allows us to encode each block $\text{ans}[ib+1..(i+1)b]$ as the number $\text{ans}[ib+1]$ and the sequence x_1, x_2, \dots, x_b such that $x_j \in \{-1, 0, 1\}$ and $\text{ans}[ib+j] = \text{ans}[ib+1] + x_1 + \dots + x_j$ for any $j \in [1..b]$. This sequence x_1, x_2, \dots, x_b is encoded in a $2b$ -bit integer exactly as in the example above (note that $2b \leq \lfloor \frac{\log n}{4} \rfloor$). Using a precomputed table of size $O(\sqrt[4]{nb})$, we can extract any element $\text{ans}[j]$ in $O(1)$ time. It is shown in Sect. 3 that arrays in pre can be encoded in a similar way (with some additional complications).

Applying a similar trick, one can perform many other operations. Let $c[0..\ell]$ be an array of integers such that $\ell \in [0..b]$, $|c[i-1] - c[i]| \leq 1$, and c is encoded, like a block of ans , by $c[0]$ and a $2b$ -bit integer. Let us show how to perform in $O(1)$ time the operation $\text{ans}[n..n+\ell] \stackrel{\text{min}}{\leftarrow} c[0..\ell]$ as in lines 5–6 of the algorithm (similar operations are also performed in predict). We first check whether $c[0] > \text{ans}[n] + 2\ell$: if so, then ans remains unchanged. It is guaranteed by the algorithm that $c[0] \geq \text{ans}[n-1] - 1$. Then, we concatenate bit representations of all required components: the (at most) two blocks $\text{ans}[ib+1..(i+1)b]$ and $\text{ans}[(i+1)b+1..(i+2)b]$ encoding the subarray $\text{ans}[n..n+\ell]$ are stored as two $2b$ -bit sequences (encoding the differences $\text{ans}[i] - \text{ans}[i-1]$ for $i \in [ib+2..(i+2)b]$ as above), $c[0..\ell]$ is also stored as a $2b$ -bit sequence, the offset $(n-ib)$ and the difference $c[0] - \text{ans}[n-1]$ are stored as $O(\log b)$ -bit integers; $6b + O(\log b)$ bits in total. We precompute a table T that, for a given

combined bit representation, stores two $2b$ -bit sequences encoding two blocks that represent the resulting modified $\text{ans}[n..n+\ell]$ array. It should be noted that the information provided in the given representation suffices to compute the result and, since the resulting array ans satisfies Lemma 2, we may put $\text{ans}[n+\ell+j] = \min\{\text{ans}[n+\ell+j], \text{ans}[n+\ell] + j\}$ for $j \geq 1$ so that the structure of the last block is preserved. Since $6b \leq \frac{3}{4} \log n$, the size of the table T is $O(b \cdot 2^{6b+O(\log b)}) = O(n^{3/4} \log^k n)$ for some $k = O(1)$. Obviously, T can be precomputed in $O(n^{3/4} \log^{k+O(1)} n)$ time. Analogously, we precompute tables that allow us to calculate $\text{incrv}(i, j)$ in $O(1)$ time if $j - i \leq b$; the resulting array of incrv is encoded, like the array c , by the first element and a $2b$ -bit integer. Thus, all range operations in lines 5–6 of the algorithm can be performed in $O(1)$ time.

We use a number of different range operations on the arrays ans and pre in Section 3 but all of them are similar to the discussed ones, so we omit detailed descriptions.

3 Implementation of the Main Functions

Now it remains to describe the functions predict and precalc and prove their time complexity.

3.1 Function predict

At the beginning, the function predict sets $\text{ans}[n+j] \leftarrow \text{ans}[n-1]+j+1$ for $j \in [0..b]$. By Lemma 2, the assigned values are upper bounds for the elements of $\text{ans}[n..n+b]$. The assignments are performed in $O(1)$ time using range operations. Then predict processes each of the t series; let us describe precisely how we process a p -series u_1, \dots, u_k .

Let u, v be the palindromes described in Lemma 5, x_i be the center of u_i for $i \in [1..k]$. If $\text{len}(x_i) < n - \text{left}[p]$ (i.e., either $i > 1$ or $i = 1$ and u_1 is not the longest suffix of s with period p), then x_i will remain the center of a new suffix-palindrome after the appending of $s[n]$ iff $s[n] = s[n-p] = v[0]$. In this case, the period p “extends” and x_i remains the center of a suffix-palindrome with the minimal period p . In the remaining case $\text{len}(x_1) = n - \text{left}[p]$ (u_1 is the longest suffix with period p) x_1 will remain the center of a suffix-palindrome iff $s[n] = s[\text{left}[p]]$; the period p breaks and the palindrome $s[n]u_1s[n]$ will belong to a different series.

Suppose that d upcoming predictable extensions extend the period p of the suffix $s[\text{left}[p]+1..n-1]$ and the $(d+1)$ st predictable extension breaks this period. It follows from the previous paragraph that the only suffix-palindrome u_i that can survive the $(d+1)$ st extension (in other words, for which $\text{live}(x_i) > d$) must have length $n - \text{left}[p] - d$ (see Fig. 2). So if d is known, we check whether $x = \text{cntr}(n - \text{left}[p] - d)$ is the center of a suffix-palindrome (i.e., $\text{cntr}(\text{len}(x)) = x$) and, if so, we compute $\text{ans}[n..n+c] \stackrel{\min}{\leftarrow} \text{incrv}(n - \text{len}(x) - c, n - \text{len}(x))$, where $c = \min\{b, \text{live}(x)\}$, in $O(1)$ time using range operations.

Now it remains to find d and change $\text{ans}[n..n + \min\{b, d\}]$ taking u_1, \dots, u_k into account. Since predictable extensions append the letters $s[n - \text{len}(\text{maxPal})], s[n - \text{len}(\text{maxPal}) - 1], \dots$ to the right of s , we can approximately find d looking at the string $s[0..n - \text{len}(\text{maxPal}) - 1]$. Put $d' = \min\{\text{live}(\text{cntr}(|u|)), \text{live}(\text{cntr}(|uvu|))\}$ (see Fig. 2). Let us show that we can use d' instead of d . If $d' < n - \text{left}[p] - |uvu|$, then the longest suffix-palindrome is preceded by the reversed prefix of $(vu)^\infty$ of length d' . In turn, this prefix either is preceded by a letter that breaks the period p of this prefix (the letter e in Fig. 2) or is a prefix of the whole string. In either case, $d' = d$. If $d' \geq n - \text{left}[p] - |uvu|$, then the longest suffix-palindrome is also preceded by the reversed prefix of $(vu)^\infty$ of length d' but $d \geq d'$ in general. However, even in this case, we can use d' in the sequel since none of the suffix-palindromes from our series survives after $n - \text{left}[p] - |uvu|$ predictable extensions; therefore, also, the possible processing of a suffix-palindrome of length $n - \text{left}[p] - d'$ mentioned above is not required.

Proof. For $i \in [0..p)$ and $j \geq 0$, denote $i(j) = \text{left}[p] + i + jp$. Let k be an integer such that, for $i = p - 1$, we have $i(k) < n$ and $i(k) + p \geq n$. So, for $i \in [0..p)$, we obtain $\phi_i = j'_i - 1$, where j'_i is the minimal integer such that $j'_i \in [0..k]$ and the string $s[i(j'_i)+1..n]$ has no prefix-palindromes with the minimal period p . While i descends from $p-1$ to 0 with step one, some of the suffixes $s[i(j)+1..n]$ may acquire prefix-palindromes with the minimal period p and some may lose such prefix-palindromes thus changing the value of ϕ_i (see Fig. 3).

Let us choose $i \in [0..p)$ that maximizes the value of ϕ_i . Denote $j' = \phi_i$ for this i . If $j' \geq 0$, then $s[i(j')+1..n]$ has a prefix-palindrome w with the minimal period p ; by Lemma 3, there are palindromes u and v such that $|uv| = p$ and $w = (uv)^r u$ for $r \geq 1$. Thus, for any $j'' \in [0..j'-2]$, the suffix $s[i(j'')+1..n]$ has prefix-palindromes $(uv)^3 u$ and $(uv)^2 u$ both having the minimal period p . When i further decreases to 0, the prefix-palindrome $(uv)^2 u$ “grows” together with $s[i(j'')+1..n]$ and, when i increases, $(uv)^3 u$ “shrinks”; in both cases $s[i(j'')+1..n]$ retains a prefix-palindrome with the minimal period p while $i \in [0..p)$. Hence, only suffixes $s[i(j'-1)+1..n]$ and $s[i(j')+1..n]$ may lose or acquire a prefix-palindrome with the minimal period p while i changes from $p-1$ to 0, i.e., ϕ_i varies in the range $[j'-2..j']$.

Let us prove that any suffix $s[i(j)+1..n]$ can lose a prefix-palindrome with the minimal period p at most once during the descending of i from $p-1$ to 0. Then, the existence of the desired numbers k_0, k_1, \dots, k_7 follows from a simple analysis of possible cases.

Suppose that $s[i(j)+1..n]$ has a prefix-palindrome centered at x with the minimal period p . When i decreases, $s[i(j)+1..n]$ grows and the prefix-palindrome “grows” simultaneously. Then, before $s[i(j)+1..n]$ loses the prefix-palindrome, we have $|s[i(j)+1..n]| = \text{len}(x)$ for some $i \in [0..p)$. By Lemma 4, there are palindromes u' and v' such that $|u'v'| = p$ and $s[n - \text{len}(x) + 1..n] = (u'v')^{r'} u'$ for $r' \geq 1$. If, for some smaller $i \in [0..p)$, $s[i(j)+1..n]$ again acquires a prefix-palindrome with the minimal period p , then, by Lemma 4, the center x' of this prefix-palindrome must coincide with the center of u' or v' from the decomposition. Hence $x' \leq x - p/2$. Then, this prefix-palindrome can be lost only after p decrements of i once we have had $|s[i(j)+1..n]| = \text{len}(x)$. This proves the claim. ◀

We partition $\text{pre}[p][0..p-1]$ into subarrays $\text{pre}[p][k_0..k_1-1], \dots, \text{pre}[p][k_6..k_7-1]$ according to Lemma 6. Consider a segment $[a..b] \subset [0..p)$ such that $\phi_{i_1} = \phi_{i_2}$ and $\phi_{i_1} \neq -1$ whenever $i_1, i_2 \in [a..b]$. Since $\text{pre}[p][i] = \min\{\text{ans}[\text{left}[p] + i + jp] : j \in [0..\phi_i]\}$ and, by Lemma 2, $|\text{ans}[j] - \text{ans}[j-1]| \leq 1$ for any $j \in (0..n)$, we easily obtain $|\text{pre}[p][i] - \text{pre}[p][i-1]| \leq 1$ for any $i \in (a..b)$. Therefore, by Lemma 6, each of the subarrays of pre either contains only $+\infty$ or has a structure similar to the structure of ans described in Lemma 2. We do not store the subarrays that contain $+\infty$ and encode all other subarrays in a way described for ans in Sect. 2.4: we split them into blocks of length b and encode each block as its starting element and a $2b$ -bit integer encoding the differences between adjacent elements (the last block may contain less than b elements). The linear size of pre measured in machine words (but not in the number of elements) follows from the overall linear running time of the function precalc maintaining pre ; this analysis is given below.

To perform $\text{ans}[n+j] \stackrel{\text{min}}{\leftarrow} \text{pre}[p][r(j)]$ for all $j \in [0..\min\{b, d'\}]$, we concatenate $2b$ -bit integers from the blocks covering the subarray $\text{ans}[n..n + \min\{b, d'\}]$, $2b$ -bit integers from a constant number of blocks covering the subarrays of $\text{pre}[p][0..p-1]$ containing positions $r(j)$ for $j \in [0..\min\{b, d'\}]$, and some other lightweight auxiliary data similar to the data used in the operation $\stackrel{\text{min}}{\leftarrow}$ considered above; then we compute the resulting array $\text{ans}[n..n + \min\{b, d'\}]$ using the obtained bit string and a precomputed table of size $o(n)$. This might require to duplicate the content of $\text{pre}[p]$ if $p < \min\{b, d'\}$ (see the shaded region in Fig. 2); these duplications must be already precalculated in the tables. Note that thus defined changes of ans may affect the whole subarray $\text{ans}[n..n+b]$ and not only $\text{ans}[n..n + \min\{b, d'\}]$: e.g., if we

$\text{pre}[p][(r-m) \bmod p] \stackrel{\text{min}}{\leftarrow} \text{ans}[n'-|w|-m]$, where $r = n' - |w| - \text{left}[p]$, for all such m . Among these assignments there is the required $\text{pre}[p][i] \stackrel{\text{min}}{\leftarrow} \text{ans}[\text{left}[p]+i+jp-1]$ (see Fig. 4). Since $n - n' \leq b$, once x is known, we can perform all these assignments in $O(1)$ time using range operations and precomputed tables; the boundaries of subarrays of pre can be adjusted appropriately after these calculations. Now it remains to find all such centers x .

By Lemma 3, there are palindromes \tilde{u} and \tilde{v} such that $p = |\tilde{u}\tilde{v}|$ and $w = (\tilde{u}\tilde{v})^r \tilde{u}$ for $r \geq 1$ (see Fig. 4). If the minimal period of $(\tilde{u}\tilde{v})^{r-1} \tilde{u}$ is p , then all strings $s[h..n']$, for $h \in (\text{left}[p]..n'-|w|)$, have prefix-palindromes of the form $\alpha(\tilde{u}\tilde{v})^{r-1} \tilde{u} \tilde{\alpha}$, where α is a suffix of $\tilde{u}\tilde{v}$, with the minimal period p . But, by our assumption, $s[\text{left}[p]+i+jp..n']$ cannot have such a prefix-palindrome. Therefore, w is the baby in the p -series of the string $s[0..n']$, i.e., either $w = \tilde{u}\tilde{v}\tilde{u}$ or $w = \tilde{u}\tilde{v}\tilde{u}\tilde{v}\tilde{u}$. We find the baby in $O(1)$ time by the techniques described above using an instance of the iterator and the list of all series of suffix-palindromes for the string $s[0..n']$; these iterator and list are further discussed below.

Case 2. It remains to detect all x such that x is the center of a prefix-palindrome of $s[\text{left}[p]+i+jp..n]$ with the minimal period p , for some $i \in [0..p)$ and $j \geq 0$, and either $x > n'$ or the minimal period of any subpalindrome in $s[0..n']$ centered at x is not p . Hence, a subpalindrome with the minimal period p and the center x appeared after several extensions of $s[0..n']$ and, thus, was a suffix-palindrome at that moment. To catch the moments when growing suffix-palindromes acquire new minimal periods, we need a device tracking changes of periods in all suffix-palindromes after extensions. The iterator can serve as such a device.

Let w be a suffix-palindrome of $s[0..n']$ with the minimal period p' . By Lemma 3, we have $p' = |w| - |u|$, where u is the longest proper suffix-palindrome of w . Suppose that $s[0..n']$ is extended by the letter $a = s[n'+1]$ and awa is a suffix-palindrome of the new string. By Lemma 3, awa has period p' iff awa is a suffix-palindrome of $s[0..n'+1]$. Thus, to detect new suffix-palindromes with a given period p , we can track, during the extensions of s , changes in the list of the centers of all suffix-palindromes. The iterator maintains such list. The following lemma is a straightforward corollary of the proof of [7, Prop. 1].

► **Lemma 7.** *The iterator maintains a linked list of the centers of all suffix-palindromes of $s[0..n]$. The function $\text{add}(a)$ removes a number of centers from the list, adds the centers $n+\frac{1}{2}$ (if $a = s[n]$) and $n+1$ to the end of the list, and thus obtains a new list for the string $s[0..n]a$; all in $\Omega(1+c)$ time, where c is the number of removed centers.*

We maintain an instance of the iterator for the previously processed string $s[0..n']$ and store the list of the centers of all suffix-palindromes of $s[0..n']$ since the last call of precalc . The function precalc performs $\text{add}(s[n'+1]), \dots, \text{add}(s[n])$ and thus consecutively obtains the lists of the centers of all suffix-palindromes of $s[0..n'+1], \dots, s[0..n]$.

Consider, for $n'' \in (n'..n]$, such list x_1, \dots, x_k for $s[0..n''-1]$ so that $x_1 < \dots < x_k$. By Lemma 7, the call to $\text{add}(s[n''])$ gives us a sublist x_{i_1}, \dots, x_{i_c} of the centers removed from x_1, \dots, x_k . By Lemma 3, for $x_i \notin \{x_{i_1}, \dots, x_{i_c}\}$ the minimal period of the suffix-palindrome with the center x_i has changed iff $x_{i+1} \in \{x_{i_1}, \dots, x_{i_c}\}$. We easily find all such x_i parsing the list x_{i_1}, \dots, x_{i_c} from left to right and compute the new period as $p = \text{len}(x_i) - \text{len}(\text{nextPal}(x_i))$. Denote by ℓ the number that is equal to $\text{len}(x_i)$ for $s[0..n'']$. By the definition of pre , we must perform $\text{pre}[p][r] \stackrel{\text{min}}{\leftarrow} \text{ans}[n''-\ell]$, where $r = (n'' - \ell - \text{left}[p]) \bmod p$, if the string $s[0..n]$ has a p -series. In this case, we must also perform $\text{pre}[p][(r-m) \bmod p] \stackrel{\text{min}}{\leftarrow} \text{ans}[n''-\ell-m]$ for all $m \in [0.. \min\{n - n'', n'' - \ell - \text{left}[p]\}]$ because the strings $s[n''-\ell-m..n]$ have prefix-palindromes of length $\ell+2m$ centered at x_i with the minimal period p ; after $n'' - \ell - \text{left}[p] + 1$ extensions, such palindrome dies since it reaches $\text{left}[p]$ by its leftmost position and thus its period breaks. Since $n - n'' \leq b$, these assignments, for all such m , can be performed by

range operations on `pre` and `ans` in $O(1)$ time using precomputed tables like those described in Sect. 2.4 (subarrays of `pre[p]` can be also adjusted appropriately).

Thus, `precalc` works in $O(t + q)$ time as required, where t is the number of series in $s[0..n]$ and q is the time required to perform the sequence of calls `add(s[n'+1]), ..., add(s[n])`. This finishes the proof of the linear time complexity of main algorithm.

References

- 1 Vladimir Arlazarov, Efim Dinic, Mikhail Kronrod, and Igor Faradzev. On economical construction of the transitive closure of a directed graph. *Dokl. Akad. Nauk SSSR*, 194(11):1209–1210, 1970.
- 2 Stefan Burkhardt and Juha Kärkkäinen. Fast lightweight suffix array construction and checking. In Ricardo A. Baeza-Yates, Edgar Chávez, and Maxime Crochemore, editors, *Proc. of the 14th Annual Symposium on Combinatorial Pattern Matching (CPM 2003)*, volume 2676 of *LNCS*, pages 55–69. Springer, 2003. doi:10.1007/3-540-44888-8_5.
- 3 Gabriele Fici, Travis Gagie, Juha Kärkkäinen, and Dominik Kempa. A subquadratic algorithm for minimum palindromic factorization. *J. Discrete Algorithms*, 28:41–48, 2014. doi:10.1016/j.jda.2014.08.001.
- 4 Zvi Galil and Joel I. Seiferas. A linear-time on-line recognition algorithm for “palstar”. *J. ACM*, 25(1):102–111, 1978. doi:10.1145/322047.322056.
- 5 Tomohiro I, Shiho Sugimoto, Shunsuke Inenaga, Hideo Bannai, and Masayuki Takeda. Computing palindromic factorizations and palindromic covers on-line. In Alexander S. Kulikov, Sergei O. Kuznetsov, and Pavel A. Pevzner, editors, *Proceedings of the 25th Annual Symposium on Combinatorial Pattern Matching (CPM 2014)*, volume 8486 of *LNCS*, pages 150–161. Springer, 2014. doi:10.1007/978-3-319-07566-2_16.
- 6 Donald E. Knuth, James H. Morris Jr., and Vaughan R. Pratt. Fast pattern matching in strings. *SIAM J. Comput.*, 6(2):323–350, 1977. doi:10.1137/0206024.
- 7 Dmitry Kosolobov, Mikhail Rubinchik, and Arseny M. Shur. Pal^k is linear recognizable online. In Giuseppe F. Italiano, Tiziana Margaria-Steffen, Jaroslav Pokorný, Jean-Jacques Quisquater, and Roger Wattenhofer, editors, *Proceedings of the 41st International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM 2015)*, volume 8939 of *LNCS*, pages 289–301. Springer, 2015. doi:10.1007/978-3-662-46078-8_24.
- 8 Glenn 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. doi:10.1145/321892.321896.
- 9 Palindromic Length. Sources and tests for linear palindromic length problem, 2017. URL: <https://github.com/kborozdin/palindromic-length>.
- 10 Mikhail Rubinchik and Arseny M. Shur. EERTREE: an efficient data structure for processing palindromes in strings. In Zsuzsanna Lipták and William F. Smyth, editors, *Proc. of the 26th International Workshop on Combinatorial Algorithms (IWOCA 2015)*, volume 9538 of *LNCS*, pages 321–333. Springer, 2015. doi:10.1007/978-3-319-29516-9_27.
- 11 Anatol O. Slisenko. A simplified proof of the real-time recognizability of palindromes on turing machines. *J. Soviet. Math.*, 15(1):68–77, 1977. doi:10.1007/BF01404109.