








Maximal Closed Substrings

Golnaz Badkobeh¹, Alessandro De Luca², Gabriele Fici³✉,
and Simon J. Puglisi⁴

¹ Department of Computing, Goldsmiths University of London, London, UK
`g.badkobeh@gold.ac.uk`

² DIETI, Università di Napoli Federico II, Naples, Italy
`alessandro.deluca@unina.it`

³ Dipartimento di Matematica e Informatica, Università di Palermo, Palermo, Italy
`gabriele.fici@unipa.it`

⁴ Department of Computer Science, University of Helsinki, Helsinki, Finland
`simon.puglisi@helsinki.fi`

Abstract. A string is closed if it has length 1 or has a nonempty border without internal occurrences. In this paper we introduce the definition of a *maximal closed substring* (MCS), which is an occurrence of a closed substring that cannot be extended to the left nor to the right into a longer closed substring. MCSs with exponent at least 2 are commonly called *runs*; those with exponent smaller than 2, instead, are particular cases of *maximal gapped repeats*. We show that a string of length n contains $\mathcal{O}(n^{1.5})$ MCSs. We also provide an output-sensitive algorithm that, given a string of length n over a constant-size alphabet, locates all m MCSs the string contains in $\mathcal{O}(n \log n + m)$ time.

Keywords: Closed word · Maximal closed substring · Run

1 Introduction

The distinction between open and closed strings was introduced by the third author in [8] in the context of Sturmian words.

A string is *closed* (or *periodic-like* [6]) if it has length 1 or it has a border that does not have internal occurrences (i.e., it occurs only as a prefix and as a suffix). Otherwise the string is *open*. For example, the strings a , $abaab$ and $ababa$ are closed, while ab and $ababaab$ are open. In particular, every string whose exponent — the ratio between the length and the minimal period — is at least 2, is closed [1].

In this paper, we consider occurrences of closed substrings in a string with the property that the substring cannot be extended to the left nor to the right into another closed substring. These are called the *maximal closed substrings* (MCS)

Gabriele Fici is partly supported by MIUR project PRIN 2017 ADASCOML – 2017K7XPAN. Simon J. Puglisi is partly supported by the Academy of Finland, through grant 339070.

of the string. For example, if $S = abaabab$, then the set of pairs of starting and ending positions of the MCSs of S is

$$\{(1, 1), (1, 3), (1, 6), (2, 2), (3, 4), (4, 8), (5, 5), (6, 6), (7, 7), (8, 8)\}$$

This notion encompasses that of a *run* (maximal repetition) which is a MCS with exponent 2 or larger. It has been conjectured by Kolpakov and Kucherov [12] and then finally proved, after a long series of papers, by Bannai et al. [2], that a string of length n contains less than n runs.

On the other hand, maximal closed substrings with exponent smaller than 2 are particular cases of *maximal gapped repeats* [11]. An α -gapped repeat ($\alpha \geq 1$) in a string S is a substring uvu of S such that $|uv| \leq \alpha|u|$. It is maximal if the two occurrences of u in it cannot be extended simultaneously with the same letter to the right nor to the left. Gawrychowski et al. [10] proved that there are words that have $\Theta(\alpha n)$ maximal α -gapped repeats.

In this paper, we address the following problems:

1. How many MCSs can a string of length n contain?
2. What is the running time of an algorithm that, given a string S of length n , returns all the occurrences of MCSs in S ?

We show that:

1. A string of length n contains $\mathcal{O}(n^{1.5})$ MCSs.
2. There is an algorithm that, given a string of length n over a constant-size alphabet, locates all m MCSs the string contains in $\mathcal{O}(n \log n + m)$ time.

2 Preliminaries

Let $S = S[1..n] = S[1]S[2] \cdots S[n]$ be a string of n letters drawn from an alphabet Σ of constant size. The length n of a string S is denoted by $|S|$. The *empty string* has length 0. A *prefix* (resp. a *suffix*) of S is any string of the form $S[1..i]$ (resp. $S[i..n]$) for some $1 \leq i \leq n$. A *substring* of S is any string of the form $S[i..j]$ for some $1 \leq i \leq j \leq n$. It is also commonly assumed that the empty string is a prefix, a suffix and a substring of any string.

An integer $p \geq 1$ is a *period* of S if $S[i] = S[j]$ whenever $i \equiv j \pmod{p}$. For example, the periods of $S = aabaaba$ are 3, 6 and every $n \geq 7 = |S|$.

We recall the following classical result:

Lemma 1 (Periodicity Lemma (weak version) [9]). *If a string S has periods p and q such that $p + q \leq |S|$, then $\gcd(p, q)$ is also a period of S .*

Given a string S , we say that a string $\beta \neq S$ is a *border* of S if β is both a prefix and a suffix of S (we exclude the case $\beta = S$ but we do consider the case $|\beta| = 0$). Note that if β is a border of S , then $|S| - |\beta|$ is a period of S ; conversely, if $p \leq |S|$ is a period of S , then S has a border of length $|S| - p$.

The following well-known property of borders holds:

Property 1. If a string has two borders β and β' , with $|\beta| < |\beta'|$, then β is a border of β' .

The *border array* $B_S[1..n]$ of string $S = S[1..n]$ is the integer array where $B_S[i]$ is the length of the longest border of $S[1..i]$. When the string S is clear from the context, we will simply write B instead of B_S .

For any $1 \leq i \leq n$, let $B^1[i] = B[i]$ and $B^j[i] = B[B^{j-1}[i]]$ for $j \geq 2$. We set

$$B^+[i] = \{|\beta| \mid \beta \text{ is a border of } S[1..i]\}.$$

By Property 1, we have $B^+[i] = \{B^j[i] \mid j \geq 1\}$.

For example, in the string $S = aabaaaabaaba$, we have $B^+[6] = \{0, 1, 2\}$. Indeed, $B[6] = 2$, and $B^2[6] = B[2] = 1$, while $B^j[6] = 0$ for $j > 2$.

The *OC array* [5] $OC_S[1..n]$ of string S is a binary array where $OC_S[i] = 1$ if $S[1..i]$ is closed and $OC_S[i] = 0$ otherwise.

We also define the array P_S where $P_S[i]$ is the length of the longest repeated prefix of $S[1..i]$, that is, the longest prefix of $S[1..i]$ that has at least two occurrences in $S[1..i]$. Again, if S is clear from the context, we omit the subscripts.

Let S be a string of length n . Since for every $1 \leq i \leq n$, the longest repeated prefix v_i of $S[1..i]$ is the longest border of $S[1..j]$, where $j \leq i$ is the ending position of the second occurrence of v_i , we have that

$$P[i] = \max_{1 \leq j \leq i} B[j]. \quad (1)$$

Lemma 2 ([7]). *Let S be a string of length n . For every $1 \leq i \leq n$, one has*

$$P[i] = \sum_{j=1}^i OC[j] - 1, \quad (2)$$

that is, $P[i]$ is the rank of 1's in $OC[1..i]$ minus one.

Proof. For every repeated prefix v of S , the second occurrence of v in S determines a closed prefix of S ; conversely, every closed prefix of S of length greater than 1 ends where the second occurrence of a repeated prefix of S ends. Indeed, the length of the longest repeated prefix increases precisely in those positions in which we have a closed prefix. That is, $P[i] = P[i-1] + OC[i]$, for any $1 < i \leq n$, which, together with $P[1] = 0 = OC[1] - 1$, yields (2). \square

As a consequence of (1) and (2), if two strings have the same border array, then they have the same OC array, but the converse is not true in general (take for example $aaba$ and $aabb$).

The OC array of a string can be obtained from its P array by taking the differences of consecutive values, putting 1 in the first position (cf. [8]). Since the border array can be easily computed in linear time [13], it is possible to compute the OC array in linear time.

Example 1. The OC, B, and P arrays for $S = aabaaaabaaba$ are shown in the following table:

i	1	2	3	4	5	6	7	8	9	10	11	12
S	a	a	b	a	a	a	a	b	a	a	b	a
OC	1	1	0	0	1	0	0	1	1	1	0	0
B	0	1	0	1	2	2	2	3	4	5	3	4
P	0	1	1	1	2	2	2	3	4	5	5	5

3 A Bound on the Number of MCS

The goal of this section is to prove our bound $\mathcal{O}(n^{1.5})$ in the number of MCSs in a string of length n . This will be derived from a bound on the number of runs (in the sense of maximal blocks of identical symbols, not to be confused with maximal repetitions mentioned in the introduction) in the OC array.

In the next lemmas, we gather some structural results on the OC array.

Lemma 3 ([7, Remark 8]). *If $\text{OC}[i] = 1$, then $\text{B}[i] = \text{P}[i]$, and $\text{B}[i-1] = \text{P}[i-1]$ (provided $i > 1$).*

Lemma 4. *For all i and k such that $\text{OC}[i+1..i+k+1] = 0^k 1$, if $\text{P}[i] \geq k$ then $\text{P}[i] - k \in \text{B}^+[i]$.*

Proof. By Lemma 2 and Lemma 3, $\text{P}[i+k+1] = \text{P}[i] + 1$ is the length of the longest border of S at position $i+k+1$. The assertion is then a consequence of the following simple observation: Let u, v and x be strings; if ux is a border of vx , then u is a border of v . In fact, letting $v = S[1..i]$, and $x = S[i+1..i+k+1]$, as $\text{B}^+[i+k+1] > k$, the longest border of vx can be written as ux for some u of length $\text{P}[i] + 1 - k - 1 = \text{P}[i] - k$. \square

Lemma 5. *For all i and k such that $\text{OC}[i..i+k+1] = 10^k 1$, if $\text{P}[i] \geq k$ then $\text{P}[i] - k \in \text{B}^+[\text{P}[i]]$.*

Proof. Immediate by Lemmas 3 and 4, as $\text{B}[i] = \text{P}[i]$ and $\text{P}[i] - k \in \text{B}^+[i]$. \square

Lemma 6. *If $\text{OC}[i..i+k_1+k_2+t+1] = 10^{k_1} 1^t 0^{k_2} 1$ and $k_1, k_2 > 0$, then $\text{P}[i] < k_1 + k_2$.*

Proof. By contradiction. Assume $\text{P}[i] \geq k_1 + k_2$. Then by Lemma 5 we have $\text{P}[i] - k_1 \in \text{B}^+[\text{P}[i]]$, which implies that k_1 is a period of $S[1..\text{P}[i]]$. Similarly, k_2 is a period of $S[1..\text{P}[i]+t]$ and then of $S[1..\text{P}[i]+1]$ and $S[1..\text{P}[i]]$, since $\text{P}[i] \geq k_2$. By the Periodicity Lemma 1 we know that $K = \text{gcd}(k_1, k_2)$ is also a period of $S[1..\text{P}[i]]$. Note that $k_1 - k_2$ is divisible by K .

Furthermore, $S[i+1] \neq S[i+1+k_1]$ because $\text{OC}[i+1]$ is not 1. By Lemma 4, we have $\text{P}[i] + 1 - k_1 \in \text{B}^+[i+1]$, which implies $S[i+1] = S[\text{P}[i] + 1 - k_1]$.

However, $S[i+1+k_1] = S[\text{P}[i]+1] = S[\text{P}[i]+1-k_2] = S[\text{P}[i]+1-k_2 - (k_1 - k_2)] = S[\text{P}[i]+1-k_1] = S[i+1]$, which is a contradiction. \square

Theorem 1. *Let S be a string of length n . Then the number of runs in its OC array is $\mathcal{O}(\sqrt{n})$.*

Proof. Let $\text{OC}_S = 1^{t_1}0^{k_1} \dots 1^{t_m}0^{k_m}$, where $k_m \geq 0$ and all other exponents are positive. By Lemma 6, we have for $1 < i < m$,

$$k_{i-1} + k_i \geq \sum_{r=1}^{i-1} t_r \geq i - 1.$$

This implies

$$n = \sum_{i=1}^m (t_i + k_i) \geq m + \sum_{j=1}^{\lfloor \frac{m-1}{2} \rfloor} (k_{2j-1} + k_{2j}) \geq m + \sum_{j=1}^{\lfloor \frac{m-1}{2} \rfloor} (2j - 1) = m + \left\lfloor \frac{m-1}{2} \right\rfloor^2$$

so that $n = \Omega(m^2)$ and then $m = \mathcal{O}(\sqrt{n})$. □

The bound in the previous proposition is tight. Indeed, there exists a binary string whose OC array is $\prod_{k>0} 10^k$. Actually, the string is uniquely determined by its OC array and can be defined by $u = a \prod_{k>0} \overline{u[k]} u[1..k] = abaaabbabababaa \dots$.

The following proposition is a direct consequence of the definition of MCS. Essentially, it says that we can check if $S[i..j]$ is a MCS by looking at the OC array of the suffixes starting at position i and $i - 1$.

Proposition 1. *Let S be a string of length n . If $S[i..j]$ is a MCS, then $\text{OC}_{S[i..n]}[j - i + 1] = 1$ and either $j - i + 1 = n$ or $\text{OC}_{S[i..n]}[j - i + 2] = 0$. Moreover, either $i = 1$ or $\text{OC}_{S[i-1..n]}[j - i + 2] = 0$.*

Example 2. Let $S = aabaaabaaba$. The OC arrays of the first few suffixes of S are displayed below.

S	a	a	b	a	a	a	a	b	a	a	b	a
$\text{OC}_{S[1..n]}$	1	1	0	0	1	0	0	1	1	1	0	0
$\text{OC}_{S[2..n]}$		1	0	1	0	0	0	1	1	1	0	0
$\text{OC}_{S[3..n]}$			1	0	0	0	0	1	1	1	0	0
$\text{OC}_{S[4..n]}$				1	1	1	1	0	0	0	0	0
$\text{OC}_{S[5..n]}$					1	1	1	0	0	0	0	0
$\text{OC}_{S[6..n]}$						1	1	0	0	1	1	1

One can check for instance that $S[4..7]$ is a MCS, because the $4 = (7 - 4 + 1)$ th entry of $\text{OC}_{S[4..n]}$ is a 1 which does not have another 1 on its right nor on top of it (i.e., in the OC array of the previous suffix). Similarly, $S[6..12]$ is a MCS because the last entry of $\text{OC}_{S[6..n]}$ is 1 with a 0 on top.

As a consequence of the previous proposition, the number of MCSs in S is bounded from above by the total number of runs of 1s in all the OC arrays of the suffixes of S .

From Theorem 1, we therefore have a bound of $\mathcal{O}(n\sqrt{n})$ on the number of MCSs in a string of length n .

4 An Algorithm for Locating All MCS

In the previous section, we saw that one can locate all MCSs of S by looking at the OC arrays of all suffixes of S . However, since the OC array of a string of length n requires $\Omega(n)$ time to be constructed, this yields an algorithm that needs $\Omega(n^2)$ time to locate all MCSs.

We now describe a more efficient algorithm for computing all the maximal closed substrings in a string S of length n . For simplicity of exposition we assume that S is on a binary alphabet $\{a, b\}$, however the algorithm is easily adapted for strings on any constant-sized alphabet. The running time is asymptotically bounded by $n \log n$ plus the total number of MCSs in S .

The inspiration for our approach is an algorithm for finding maximal pairs under gap constraints due to Brodal, Lyngsø, Pedersen, and Stoye [3]. The central data structure is the suffix tree of the input string, which we now define.

Definition 1 (Suffix tree). *The suffix tree $T(S)$ of the string S is the compressed trie of all suffixes of S . Each leaf in $T(S)$ represents a suffix $S[i..n]$ of S and is annotated with the index i . We refer to the set of indices stored at the leaves in the subtree rooted at node v as the leaf-list of v and denote it $LL(v)$. Each edge in $T(S)$ is labelled with a nonempty substring of S such that the path from the root to the leaf annotated with index i spells the suffix $S[i..n]$. We refer to the substring of S spelled by the path from the root to node v as the path-label of v and denote it $L(v)$.*

At a high level, our algorithm for finding MCSs processes the suffix tree (which is a binary tree, for binary strings) in a bottom-up traversal. At each node the leaf lists of the (two, for a binary string) children are intersected. For each element in the leaf list of the smaller child, the successor in the leaf list of the larger child is found. Note that because the element from the smaller child and its successor in the larger child come from different subtrees, they represent a pair of occurrences of substring $L(v)$ that are right-maximal. To ensure left maximality, we must take care to only output pairs that have different preceding characters. We explain how to achieve this below.

Essential to our algorithm are properties of AVL trees that allow their efficient merging, and the so-called “smaller-half trick” applicable to binary trees. These properties are captured in the following lemmas.

Lemma 7 (Brown and Tarjan [4]). *Two AVL trees of size at most n and m can be merged in time $\mathcal{O}(\log \binom{n+m}{n})$.*

Lemma 8 (Brodal et al. [3], Lemma 3.3). *Let T be an arbitrary binary tree with n leaves. The sum over all internal nodes v in T of terms that are $\mathcal{O}(\log \binom{n_1+n_2}{n_1})$, where n_1 and n_2 are the n_1 numbers of leaves in the subtrees rooted at the two children of v , is $\mathcal{O}(n \log n)$.*

As stated above, our algorithm traverses the suffix tree bottom up. At a generic step in the traversal, we are at an internal node v of the suffix tree. Let the two children of node v be v_ℓ and v_r (recall the tree is a binary suffix tree, so every internal node has two children). The leaf lists of each child of v are maintained in two AVL trees — note, there are *two AVL trees for each of the two children*, two for v_ℓ and two for v_r . For a given child, say v_r , one of the two AVL trees contains positions where $L(v_r)$ is preceded by an a symbol, and the other AVL tree contains positions where $L(v_r)$ is preceded by a b symbol in S . Call these the a -tree and b -tree, respectively.

Without loss of generality, let v_r be the smaller of v 's children. We want to search for the successor and predecessor of each of the elements of v_r 's a -tree amongst the elements v_ℓ 's b -tree, and, similarly the elements of v_r 's b -tree with the elements from v_ℓ 's a -tree. Observe that the resulting pairs of elements represent a pair of occurrences of $L(v)$ that are both right and left maximal: they have different preceding characters and so will be left maximal, and they are siblings in the suffix tree and so will be right maximal. These are candidate MCSs. What remains is to discard pairs that are not consecutive occurrences of $L(v)$, to arrive at the MCSs. Discarding is easy if we process the elements of $LL(v_r)$ in order (which is in turn easy, because they are stored in two AVL trees). To see this, consider two consecutive candidates that have the same right border position (a successor found in $LL(v_\ell)$; discarding for left borders is similar). The first of these candidates can clearly be discarded because there is an occurrence of $L(v)$ (from $LL(v_r)$) in between the two borders, preventing the pair of occurrences from forming an MCS. Because we only compute a successor/predecessor for each of the elements of the smaller of v 's children, by Lemma 8 the total time for all successor/predecessor searches will be $\mathcal{O}(n \log n)$ (discarding also takes time proportional to the smaller subtree, and so does not increase this complexity). After this, the a -tree and b -tree of the smaller child are merged with their counterparts from the larger child.

Thus, by Lemmas 7 and 8, the overall processing is bounded by $\mathcal{O}(n \log n)$ in addition to the number of MCSs that are found.

The above approach is easily generalized from strings on binary alphabets to those on any alphabet of constant size by replacing nodes of the suffix tree having degree $d > 2$ with binary trees of height $\log d$. This does not increase the height of the suffix tree asymptotically and so preserves the runtime stated above. It would be interesting to design algorithms for general alphabets, and we leave this as an open problem.

References

1. Badkobeh, G., Fici, G., Lipták, Z.: On the number of closed factors in a word. In: Dediu, A.-H., Formenti, E., Martín-Vide, C., Truthe, B. (eds.) LATA 2015. LNCS, vol. 8977, pp. 381–390. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-15579-1_29
2. Bannai, H.I.T., Inenaga, S., Nakashima, Y., Takeda, M., Tsuruta, K.: The “runs” theorem. *SIAM J. Comput.* **46**(5), 1501–1514 (2017)
3. Brodal, G.S., Lyngsø, R.B., Pedersen, C.N.S., Stoye, J.: Finding maximal pairs with bounded gap. In: Crochemore, M., Paterson, M. (eds.) CPM 1999. LNCS, vol. 1645, pp. 134–149. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48452-3_11
4. Brown, M.R., Tarjan, R.E.: A fast merging algorithm. *J. ACM* **26**(2), 211–226 (1979)
5. Bucci, M., De Luca, A., Fici, G.: Enumeration and structure of trapezoidal words. *Theor. Comput. Sci.* **468**, 12–22 (2013)
6. Carpi, A., de Luca, A.: Periodic-like words, periodicity, and boxes. *Acta Informatica* **37**(8), 597–618 (2001). <https://doi.org/10.1007/PL00013314>
7. De Luca, A., Fici, G., Zamboni, L.Q.: The sequence of open and closed prefixes of a Sturmian word. *Adv. Appl. Math.* **90**, 27–45 (2017)
8. Fici, G.: Open and Closed Words. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS* **123**, 140–149 (2017)
9. Fine, N.J., Wilf, H.S.: Uniqueness theorems for periodic functions. *P. Am. Math. Soc.* **16**(1), 109–114 (1965)
10. Gawrychowski, P., I, T., Inenaga, S., Köppl, D., Manea, F.: Tighter bounds and optimal algorithms for all maximal α -gapped repeats and palindromes. *Theor. Comput. Syst.* **62**(1), 162–191 (2017). <https://doi.org/10.1007/s00224-017-9794-5>
11. Kolpakov, R., Podolskiy, M., Posypkin, M., Khrapov, N.: Searching of gapped repeats and subrepetitions in a word. *J. Discrete Algorithms* **46–47**, 1–15 (2017). https://doi.org/10.1007/978-3-319-07566-2_22
12. Kolpakov, R.M., Kucherov, G.: Finding maximal repetitions in a word in linear time. In: 40th Annual Symposium on Foundations of Computer Science, FOCS 1999, 17–18 October 1999, New York, NY, USA, pp. 596–604. IEEE Computer Society (1999)
13. Morris, J.H., Pratt, V.R.: A linear pattern-matching algorithm. Technical Report 40, University of California, Berkeley (1970)