Swap Matching in Time $O(n \log m \log |\Sigma|)$

Amihood Amir* Richard Cole† Ramesh Hariharan‡
Bar-Ilan University & Georgia Tech New York University Indian Institute of Science

Moshe Lewenstein§ Ely Porat¶
Bar-Ilan University Bar-Ilan University & Weizmann Institute

Abstract

Let a text string T of n symbols and a pattern string P of m symbols from alphabet Σ be given. A *swapped version* P' of P is a length m string derived from P by a series of *local swaps*, (i.e. $p'_\ell \leftarrow p_{\ell+1}$ and $p'_{\ell+1} \leftarrow p_\ell$) where each element can participate in *no more than one swap*.

The *Pattern Matching with Swaps* problem is that of finding all locations i of T for which there exists a swapped version P' of P with an exact matching of P' in location i of T. The best known deterministic upper bound for this problem was $O(nm^{1/3} \log m \log \sigma)$ for a general alphabet Σ, where $\sigma = min(m, |\Sigma|)$.

In this paper we present a deterministic algorithm that solves the pattern matching with swaps problem in time $O(n \log m \log \sigma)$.

Key Words: Design and analysis of algorithms, combinatorial algorithms on words, pattern matching, pattern matching with swaps, non-standard pattern matching.

1 Introduction

The *Pattern Matching with Swaps* problem (the *Swap Matching* problem, for short) requires finding all occurrences of a pattern of length m in a text of length n. The pattern is said to match the text at a given location i if adjacent pattern characters can be swapped, if necessary, so as to make the pattern identical to the substring of the text starting at location i. All the swaps are constrained to be disjoint, i.e., each character is involved in at most one swap.

*Department of Mathematics and Computer Science, Bar-Ilan University, 52900 Ramat-Gan, Israel, (972-3)531-8770; amir@cs.biu.ac.il. Partially supported by NSF grant CCR-96-10170, BSF grant 96-00509, and a BIU internal research grant.

†Courant Institute, New York University, cole@cs.nyu.edu. Supported in part by NSF grants CCR-95-03309 and CCR-98-00085.

‡Indian Institute of Science, Bangalore, ramesh@csa.iisc.ernet.in. Work done in part while visiting NYU and supported by NSF grants CCR-95-03309 and CCR-98-00085.

§Department of Mathematics and Computer Science, Bar-Ilan University, 52900 Ramat-Gan, Israel, (972-3)531-8407; moshe@cs.biu.ac.il. Partially supported by the Israel Ministry of Science Eshkol Fellowship 061-1-97.

¶Department of Mathematics and Computer Science, Bar-Ilan University, 52900 Ramat-Gan, Israel, (972-3)531-8407; porately@cs.biu.ac.il.

The importance of the swap matching problem lies in recent efforts to understand the complexity of various *generalized pattern matching* problems. The textbook problem of *exact string matching* that was first shown to be solvable in linear time by Knuth, Morris and Pratt [7] does not answer the growing requirements stemming from advances in Multimedia, Digital libraries and Computational Biology. To this end, pattern matching has to adapt itself to increasingly broader definitions of "matching" [14, 13]. In computational biology one may be interested in finding a "close" mutation, in communications one may want to adjust for transmission noise, in texts it may be desirable to allow common typing errors. In multimedia one may want to adjust for lossy compressions, occlusions, scaling, affine transformations or dimension loss.

All above applications motivated research of a new type – *Generalized Pattern Matching*. In generalized matching the input is still a text and pattern but the "matching" relation is defined differently. The output is all locations in the text where the pattern "matches" under the new definition of match. The different applications define the matching relation. An early generalized matching was the *string matching with don't cares* problem defined by Fischer and Paterson [6]. They allow a special "don't care" character that serves as a wildcard. It matches every other alphabet symbol. A text location where the pattern matches under these conditions is an exact match. They show that this problem can be solved in time $O(n \log m)$ for binary alphabets, and $O(n \log m \log \sigma)$ for general alphabets, where $\sigma = min(m, |\Sigma|)$. Another example of a generalized matching problem is the *less-than matching* [3] problem defined by Amir and Farach. In this problem both text and pattern are numbers. One seeks all text locations where every pattern number is less than its corresponding text number. Amir and Farach showed that the less-than-matching problem can be solved in time $O(n\sqrt{m \log m})$.

Muthukrishnan and Ramesh [11] prove that practically all general matching relations, where the generalization is in the definition of single symbol matches, are equivalent to the boolean convolutions, i.e. it is unlikely that they could be solved in time faster than $O(n \log m)$, where n is the text length and m is the pattern length. As we have seen, some examples have significantly worse upper bound than this. Muthukrishnan and Palem [10] further proved that in the *boolean convolutions model* there are non trivial lower bounds on the number of necessary convolutions for several non-standard string matching problems.

The swap matching problem is also a generalized matching problem. It arises from one of the edit operations considered by Lowrance and Wagner [8, 15] to define a distance metric between strings.

The fundamental question of which generalized matches are inherently hard computationally and what types are faster to compute has motivated much of the pattern matching research in the last twenty years. For the swap matching problem there were no known upper bounds better than the naive $O(nm)$ algorithm. Indeed, this problem was described in 1995 as one of the open problems in non-standard string matching [9].

Amir et al [2] obtained the first non-trivial results on this problem. They showed that the case when the size of the alphabet set Σ exceeds 2 can be reduced to the case when it is exactly 2 with a time overhead of $O(\log^2 \sigma)$. (The reduction overhead was reduced to $O(\log \sigma)$ in the journal version [1].) They then showed how to solve the problem for alphabet sets of size 2 in time $O(nm^{1/3} \log m)$, which is the best deterministic time bound known to date. Amir et al. [4] also give certain special cases for which $O(m\text{polylog}(m))$ time can be obtained. However, these cases are rather restrictive. Cole and Hariharan [5] give a randomized algorithm that solves the swap matching problem over a binary alphabet in time $O(n \log n)$.

In this paper we present a deterministic algorithm that solves the swap matching problem over binary alphabets in time $O(n \log m)$. Coupled with the alphabet reduction of [1, 5] it gives an algorithm for swap matching over general alphabet whose running time is $O(n \log m \log \sigma)$.

There are three main contributions in this paper.

1. The surprising time complexity of $O(n \log m)$ for solving the swap matching problem over binary alphabets. Until recently it was open whether the problem had a $o(nm)$ solution.
2. The key lemma 2 which is the precise formulation needed to understand the swap matching problem. With this formulation the solution becomes apparent.
3. Some sophisticated uses of polynomial multiplication to eke out properties in cases where there is an inherent symmetry that seems to defy convolutions as the means to a solution.

Paper organization. This paper is organized in the following way. In section 2 we give basic definitions. The key lemma on which our algorithm is based is proved in section 3. In section 4 we give the algorithm overview. The algorithm requires separately solving a number of cases. These can be grouped into three main categories. The solutions to the three categories are shown in sections 5, 6 and 7.

2 Problem Definition and Preliminaries

Definition: Let $S = s_1 \dots s_n$ be a string over alphabet Σ. A *swap permutation* for S is a permutation $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ such that

1. if $\pi(i) = j$ then $\pi(j) = i$ (characters are swapped).
2. for all i, $\pi(i) \in \{i-1, i, i+1\}$ (only adjacent characters are swapped).
3. if $\pi(i) \neq i$ then $s_{\pi(i)} \neq s_i$ (identical characters are not swapped).

For a given string $S = s_1 \dots s_n$ and swap permutation π for S we denote $\pi(S) = s_{\pi(1)} s_{\pi(2)} \dots s_{\pi(n)}$. We call $\pi(S)$ a *swapped version* of S.

For pattern $P = p_1 \dots p_m$ and text $T = t_1 \dots t_n$, we say that P *swap matches at location* i if there exists a swapped version P' of P that matches T starting at location i, i.e. $p'_j = t_{i+j-1}$ for $j = 1, \dots, m$.

The *Swap Matching Problem* is the following:
INPUT: Pattern $P = p_1 \dots p_m$ and text $T = t_1 \dots t_n$ over alphabet Σ.
OUTPUT: All locations i where P swap matches T.

2.1 Definition Equivalence

This section is added for the sake of comparison with previous results. Earlier work on swap matching used a slightly different definition that ours. This section porves the equivalence of the definitions. The reader who is happy with our definition may skip this section with no ill effect.

Definition: *P swap text matches at location i* if there exists a swapped version T' of T such that P matches T' starting at location i, i.e. $p_j = t'_{i+j-1}$ for $j = 1, \ldots, m$. *P swap text matches with internal symbols at location i* if there exists a swap permutation π for T such that P matches $\pi(T)$ at location i and additionally (1) $\pi(i) \geq i$ and (2) $\pi(i+m-1) \leq i+m-1$.

The *Swap Text Matching Problem* is the following:
INPUT: Pattern $P = p_1 \ldots p_m$ and text $T = t_1 \ldots t_n$ over alphabet Σ.
OUTPUT: All locations i where P swap text matches T.

The *Swap Text Matching with Internal Symbols Problem* is the following:
INPUT: Pattern $P = p_1 \ldots p_m$ and text $T = t_1 \ldots t_n$ over alphabet Σ.
OUTPUT: All locations i where P swap text matches T with internal symbols.

In [2] the problem of Pattern Matching with Swaps was introduced and the definition used was that of Swap Text Matching. It was shown in Lemma 1, Section 2 in [2] that Swap Text Matching can be solved as fast as Swap Text Matching with Internal Symbols. In the following lemma we show that our definition of Swap Matching is equivalent to that of Swap Text Matching with Internal Symbols. This, combined with the lemma from [2], allows us to use the definition of Swap Matching without loss of generality.

Lemma 1 *P swap matches at location i of T iff P swap text matches with internal symbols at location i of T.*

Proof. Assume P swap text matches with internal symbols at location i of T with swap permutation τ on T. Define π on P by setting $\pi(k) = \tau(k+i-1) - i + 1$.

By the internal symbol restrictions and rule 2 of swap permutations it follows that τ restricted on $\{i, ..., i+m-1\}$ is a permutation. Therefore, π is a permutation. To show that π is a swap permutation on P the three rules must hold.

Rule 1: $\pi(k) = j \iff \tau(k+i-1) - i + 1 = j \iff \tau(k+i-1) = j+i-1 \iff$ (by Rule 1 for τ) $k+i-1 = \tau(j+i-1) \iff k = \tau(j+i-1) - i + 1 \iff k = \pi(j)$.
Rule 2: Since $\tau(k+i-1) \in \{k+i-2, k+i-1, k+i\}$, $\pi(k) = \tau(k+i-1) - i + 1 \in \{k-1, k, k+1\}$.
Rule 3: If $\pi(k) \neq k$, i.e. $\tau(k+i-1) - i + 1 \neq k$ it must be that $\tau(k+i-1) \neq k+i-1$ which implies that $t_{\tau(k+i-1)} \neq t_{k+i-1}$ and therefore $p_{\pi(k)} \neq p_k$.

Moreover, we show that $\pi(P) = T$ based on our knowledge that $P = \tau(T)$. We need to show that for $1 \leq k \leq m$, $p_{\pi(k)} = t_{i+k-1}$. Following rule 2 there are two cases: (1) $\tau(i+k-1) = i+k-1$. This implies $\pi(k) = \tau(i+k-1) - i + 1 = i+k-1-i+1 = k$. Since $p_k = t_{\tau(i+k-1)}$ it follows that $p_{\pi(k)} = t_{i+k-1}$. (2) $\tau(i+k-1) - (i+k-1) = \rho = \pm 1$. So, $\pi(k) = \tau(k+i-1) - i + 1 = k + \rho$. By rule 1 and rule 2 for τ and the fact that $P = \tau(T)$ it follows that $t_{\tau(i+k-1)-\rho} = p_{k+\rho}$. However, this implies that $t_{i+k-1} = t_{\tau(i+k-1)-\rho} = p_{k+\rho} = p_{\pi(k)}$.

Conversely, assume that P swap matches at location i of T with swap permutation π on P. Define $\tau(k) = k$ for $k < i$ and for $k > i+m-1$ and define $\tau(k) = \pi(k-i+1) + i - 1$. In an almost identical fashion this can be shown to be a swap permutation for T such that P matches $\tau(T)$. □

2.2 Reducing Large Alphabets to Binary Alphabets

In [2] it was shown how to reduce the swap matching problem over unbounded alphabets to the problem over a two letter alphabet with an $O(\log^2 |\Sigma|)$ multiplicative overhead. Here we outline a reduction requiring only an $O(\log |\Sigma|)$ factor detailed in [1]. Another unpublished $O(\log |\Sigma|)$ reduction appears in [5].

Definition: A $(\Sigma, 3)$-*universal set* is a set $S = \{\chi_1, \ldots, \chi_k\}$ of characteristic functions, $\chi_j : \Sigma \rightarrow \{0, 1\}$ such that for every $a, b, c \in \Sigma$, and for each of the eight possible combinations of $0-1$s, there exists χ_j such that $\chi_j(a), \chi_j(b), \chi_j(c)$ equals this combination.

We extend the definition of the functions χ_j to strings in the usual manner, i.e. for $S = s_1 \ldots s_n$, $\chi_j(S) = \chi_j(s_1)\chi_j(s_2) \ldots \chi_j(s_n)$.

Theorem 1 *[2] Let P be a pattern, T a text, both over an arbitrary alphabet Σ, and let $S = \{\chi_1, \ldots, \chi_k\}$ be a $(\Sigma, 3)$-universal set. P swap matches T at location i iff for all j, $\chi_j(P)$ swap matches $\chi_j(T)$ at location i.*

In [12] it was shown how to construct $(\Sigma, 3)$-universal set of cardinality $k = O(\log |\Sigma|)$ yielding the following.

Corollary 1 *A solution of the Swap Matching problem over alphabet $\{a, b\}$ of time $O(f(n, m))$ implies a solution of time $O(\log |\Sigma| f(n, m))$ over a general alphabet Σ.*

3 The Key Property

We assume that the text and the pattern both have only a's and b's. An alternating segment of a string $S \in \{a, b\}^*$ is a substring alternating between as and bs. A maximal alternating segment, or *segment* for short, is an alternating segment such that the character to the left of the leftmost character x in the alternating segment, if any, is identical to x, and similarly, the character to the right of the rightmost character y, if any, is identical to y.

We now show the key property necessary for describing matches of the pattern in the text. To this end we partition the text and pattern into segments.

Lemma 2 *The pattern does not match in a particular alignment if and only if there exists a segment A in the text and a segment B in the pattern with odd-length overlap such that the characters of A and B misalign in the overlap.*

Proof. First, we show the *if* part. For any overlap of segments that is of even length the overlapping portions of A and B can be made identical by swapping, if necessary. Note that these swaps stay within the overlapping portion so other segments are not affected. If the overlap is odd and the overlapping portions match exactly then nothing needs to be done. So clearly, if the pattern does not match it must be the case that there exists a segment A in the text and a segment B in the pattern that overlap with odd-length overlap and with misaligning characters in the overlap.

Consider the *only if* part now. Suppose A and B overlap such that the overlap is of odd length and, without loss of generality, the overlap has $u = (ab)^*a$ in the text and $v = (ba)^*b$ in the pattern. Then, for a match to occur, either the leftmost b in v must be swapped with the character to the left or the rightmost b in v must be swapped with the character to the right.

Clearly, the portion of B overlapping A cannot be a prefix of B for the above swap to be effective (the character preceding B, if any, is a b). It follows that the portion of B overlapping A must be a prefix of A. Once again the above swap is ineffective (the character preceding A is an a). Thus it follows that the pattern does not match at this alignment. □

4 Algorithm Overview

We now restrict ourselves to the situation where both the text and pattern are strings over binary alphabet $\Sigma = \{a, b\}$. Lemma 2 motivates us to identify all text locations where there exists some overlapping segments of the text and pattern where the length of the overlap is odd and where for every location of the overlap the text symbol is not equal to the pattern symbol.

This property of swaps allows us to consider pairs of segments *separately*. As soon as a pair is found where the property holds, it immediately leads to the conclusion that there is no swap match in that location. Conversely, if all appropriate pairs of segments from the text and pattern are considered and none of them has the property of lemma 2 we can conclude that there is a swap match at that location.

The main idea of the algorithm is to separate the segments of the text and pattern into a small number of groups. In each of these groups it will be possible to check for the lemma 2 property in time $O(n \log m)$ using polynomial multiplications (which can be done in time $O(n \log m)$ using FFT in a model with word length m bits). In the following sections we handle the different cases. Some of these cases necessitate new and creative uses of convolutions.

4.1 Grouping the Segments

4.1.1 Odd a and Even a Text Segments

The first two categories of text segments we consider are the *Even a segments*, where all the a's fall on even text locations, and the *Odd a segments*, where all the a's fall on odd text locations. We construct two new text strings, T_{odd-a} and T_{even-a}, each of them having length n. T_{odd-a} is a text whose segments are all the odd a segments of T in their exact locations. The text locations where there are even a segments are all assumed to contain the symbol ϕ (*don't care*) and never contribute an error. (In the polynomial multiplication there will always be a 0 in these text locations.) Similarly, T_{even-a} is the text whose segments are all the even a segments of T in their exact locations, with all locations where T has odd a segments having the don't care symbol ϕ.

Example: Let $T = ababbbaaababababbaababbaaababababa$.
The segments are: *abab b ba a ababab ba abab ba a abababa*
The odd a segments are the first, third, fifth, seventh, and ninth. Thus the texts we will consider are $T_{odd-a} = abab\phi ba\phi abababab\phi\phi abab\phi\phi a\phi\phi\phi\phi\phi\phi\phi$ and $T_{even-a} = \phi\phi\phi\phi b\phi\phi a\phi\phi\phi\phi\phi\phi ba\phi\phi\phi\phi ba\phi abababa$.

4.1.2 Parity of Starting and Ending Location of Text Segment

Consider T_{odd-a}. It will be important for us to know whether the segment we are dealing with starts at an odd or even text location. We also would like to know whether it ends at an odd or even text location. Consequently, we define new texts in a manner similar to the one of subsection 4.1.1 where each text has *exactly* the subsegments of a given start and end parity, with all other text elements defined as ϕ.

Definition: T^{oo}_{odd-a} is a string of length n where every location i has the same symbol as T if t_i is in an odd a segment whose first element is in an odd location and whose last element is in an odd location. In all other locations j, $T^{oo}_{odd-a}[j] = \phi$.

In a similar fashion define $T^{oe}_{odd-a}, T^{eo}_{odd-a}, T^{ee}_{odd-a}, T^{oo}_{even-a}, T^{oe}_{even-a}, T^{eo}_{even-a}$ and T^{ee}_{even-a} .

Example: In the above example, the first segment in an odd a segment that starts in text location 1 and ends at text location 4, thus it is in T^{oe}_{odd-a}.
In fact, $T^{oe}_{odd-a} = abab\phi\phi\phi\phi ababab\phi\phi ababab\phi\phi\phi\phi\phi\phi\phi\phi\phi\phi$.
For another string in this example note that $T^{ee}_{even-a} = \phi\phi\phi\phi\phi\phi\phi a\phi\phi\phi\phi\phi\phi\phi\phi\phi\phi\phi\phi\phi\phi\phi\phi ababaaba$.

Note that the segments of T^{oo}_{odd-a} are exactly the odd a segments of T that start and end at an odd location. Thus, it is clear that in every location where *none* of the above eight text strings raise the property of lemma 2 there is a swap match.

4.1.3 Grouping the Pattern Segments

The pattern segments are defined in exactly the same way as the text segments and we would like to group them in the same way. However, there is a difficulty here in "nailing down" the parity of a location, since the pattern is shifted and compared to every text location. However, since the only property we have used in our grouping was the parity of the text location, it is clear that in all the pattern alignments that start in odd locations, the parity of the a occurrences and the parity of the start and end segment locations will be the same. Similarly, these parities will be the same for all pattern alignments that start in even text locations.

Thus we limit ourselves to two cases. The *Odd Result Case* and the *Even Result Case*. The *Odd Result* case is simply the pattern alignmentss starting in odd text locations. The *Even Result* case is pattern alignments starting in even text locations.

When we fix an alignment result case, it is possible to define the odd a and even a segments of the pattern, as well as the parity of the starting and ending locations of each pattern segment. Thus we can get, for the odd result case, patterns $PO^{oo}_{odd-a}, PO^{oe}_{odd-a}, PO^{eo}_{odd-a}, PO^{ee}_{odd-a}, PO^{oo}_{even-a}, PO^{oe}_{even-a}$, PO^{eo}_{even-a} and PO^{ee}_{even-a}. Similarly, for the even result case we consider patterns $PE^{oo}_{odd-a}, PE^{oe}_{odd-a}, PE^{eo}_{odd-a}$, $PE^{ee}_{odd-a}, PE^{oo}_{even-a}, PE^{oe}_{even-a}, PE^{eo}_{even-a}$ and PE^{ee}_{even-a}.

We are now ready for the algorithm.

begin Algorithm

for $X = O, E$ do:
  for $k = odd - a, even - a$ do
    for $ti = o, e$ do
      for $tj = o, e$ do

for $pi = o, e$ do
  for $pi = o, e$ do
    if $k = odd - a$ then $\ell = even - a$, else $\ell = odd - a$
    check lemma 2 property for $T_k^{ti,tj}$ and $PX_\ell^{pi,pj}$
    { Note that when $X = O$ we only consider the results in the odd text locations and when $X = E$ we only consider the results in the even text locations. }

end Algorithm

It should be noted that we only check the cases of odd a text segments versus even a pattern segments, and even a text segments versus odd a pattern segments. The reason is that the overlap of any two odd a (or even a) segments is identical, thus can never generate a swap match problem.

In the subsequent sections we will show how to implement the check for the lemma 2 property for each of the cases in time $O(n \log m)$. The intrepid reader will realize that there are 64 cases to consider. Don't panic! Because of symmetry reasons, there are only three cases we need to describe. All others are similar.

The cases we describe are the following. First fix the result parity (say to the Odd Result case) and fix the text segment type to the odd a type. It is clear that all other cases are similar. Since we don't need the X and k parameters in the algorithm, we will henceforth ignore them for the sake of a simpler notation. We are now down to the combinations $T^{ti,tj}$ and $P^{pi,pj}$. This gives us 16 cases. We will handle separately only the following three types of cases:

1. $T^{ti,tj}$ and $P^{pi,pj}$ where either $ti = pi$ or $tj = pj$. (This type covers 12 cases.) These situations are handled in section 5.
2. $T^{ti,tj}$ and $P^{pi,pj}$ where $ti, tj = oe$ and $pi, pj = eo$; or where $ti, tj = eo$ and $pi, pj = oe$. These cases are handled in section 6.
3. $T^{ti,tj}$ and $P^{pi,pj}$ where $ti, tj = oo$ and $pi, pj = ee$; or where $ti, tj = ee$ and $pi, pj = oo$. These cases are handled in section 7.

5 Segments with Equal Parity Start

Consider the case $T^{ti,tj}$ and $P^{pi,pj}$ where $ti = pi$.

Observation 1 *: For every two segments, S_t in $T^{ti,tj}$, starting at location x and S_p in $P^{pi,pj}$, starting at location y, $|x - y|$ is always even. See all possibilities in figure 1 below.*

We are interested in the length of the segment overlaps (the shaded areas in figure 1). A segment overlap has odd length iff the property of lemma 2 holds. We now show a convolution for which the resulting value at location i is 0 iff there does not exists an odd-length segment overlap.

The Convolution: Construct a pattern $P' = p'_1 \cdots p'_m$ where

$$p'_i = \begin{cases} 0, & \text{if } P^{pi,pj}[i] = \phi; \\ 1, & \text{otherwise.} \end{cases}.$$

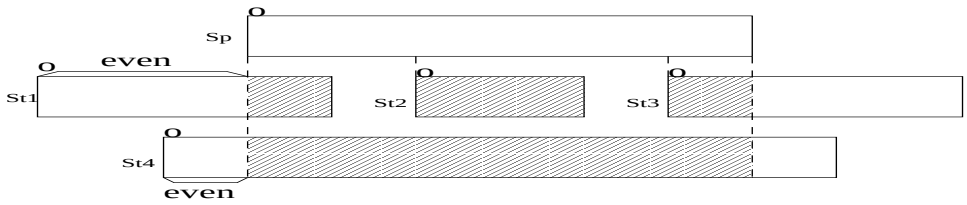

Figure 1: The cases of both text and pattern segments starting in locations witht he same parity.

Construct a text $T' = t'_1 \cdots t'_n$ where every ϕ in $T^{ti,tj}$ is replaced by 0; and every segment in $T^{ti,tj}$ is replaced by an alternating segment of 1 and -1, starting at 1.

It is clear that for all cases where the starting location of a pattern segment is smaller than the starting location of a text segment the result of the convolution will be 1 if the length of the overlap is odd and 0 if it is even (since every text segment starts with a 1 and then alternates between -1 and 1). Because of observation 1, even when the text segment starts at a smaller location than the pattern segment, the difference between the starting locations has even length. Therefore in the area of the overlap, the text starts with a 1 and alternates between -1 and 1. Thus the convolution gives us the desired result.

This solves all eight cases of $T^{ti,tj}$ and $P^{pi,pj}$ where $ti = pi$. For the additional four cases where $tj = pj$ simply reverse the text and pattern and achieve the case considered above.

6 The Odd-Even Even-Odd Segments

Consider the case T^{oe} and P^{eo} (the case of T^{eo} and P^{oe} is symmetric).

Terminology: Let S_t be a text segment whose starting location is s_1 and whose ending location is f_1. Let S_p be a pattern segment being compared to the text at starting position s_2 and ending position f_2. If $s_1 < s_2 < f_2 < f_1$ then we say that S_t *contains* S_p. If $s_2 < s_1 < f_1 < f_2$ then we say that S_p *contains* S_t. If $s_1 < s_2 < f_1 < f_2$ then we say that S_t has a *left overlap* with S_p. If $s_2 < s_1 < f_2 < f_1$ then we say that S_t has a *right overlap* with S_p. We will sometimes refer to a left or right overlap as a *side overlap.*

Observation 2 *: For every two segments, S_t in T^{oe} and S_p in P^{eo} if either S_p is contained in S_t or S_t is contained in S_p then the overlap is of even length. If the overlap is a left overlap or right overlap then it is of odd length. See all possibilities in figure 2 below.*

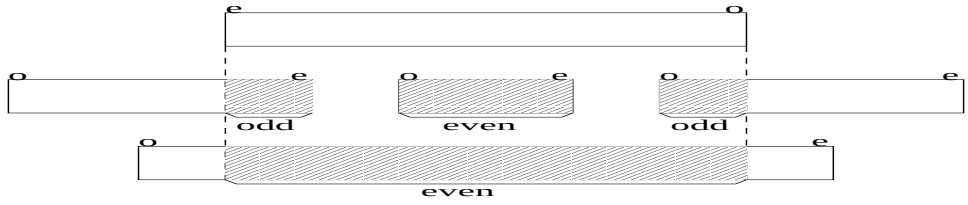

Figure 2: The cases where the text segment starts at an odd location and ends at an even location; the pattern segment does the opposite.

The correctness of the observation is immediate. segments of these types have even length thus if one contains the other the overlap is necessarily of even length. Conversely, in case of a left or right overlap then the overlap starting and ending locations have the same parity, making the length of the overlap odd.

We will now show a convolution where segments that are contained in each other contribute a 0, and side overlaps contribute positive numbers. This convolution will be more complex than the previous one. The reason is that there is some inherent relation between text segments that *start* within a pattern segment. Likewise, there is a commonality between text segments that *end* within a pattern segment. Our problem is that we need to differentiate between cases that are naturally easier handled together.

Consider the following case. Let S_t be a segment of T^{oe} starting at text location s_1 and having length ℓ_t. Assume that the pattern P^{eo} of length ℓ_p is aligned to start at the text location such that the first symbol of segment S_p occurs at location s_2. Further assume that one of the segments is contained in the other. If we replace S_t by $s_1 1^{\ell_t - 2} - s_1$ and S_p by $s_2 1^{\ell_p - 2} - s_2$ (and ϕ by 0) then multiplication of these two segments will yield $k-2$, where k is the size of the overlap. See figure 3.

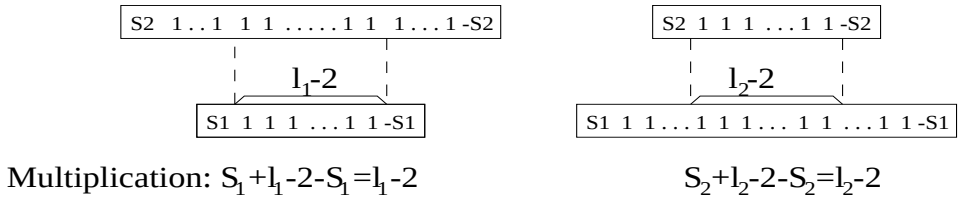

Figure 3: Containment cases for text T^{oe} and pattern P^{oe}.

Conversely, if there is a side overlap of S_t with S_p then the multiplication will yield $\max(s_1, s_2) - \min(s_1, s_2) + k - 2$, where k is the size of the overlap. See figure 4.

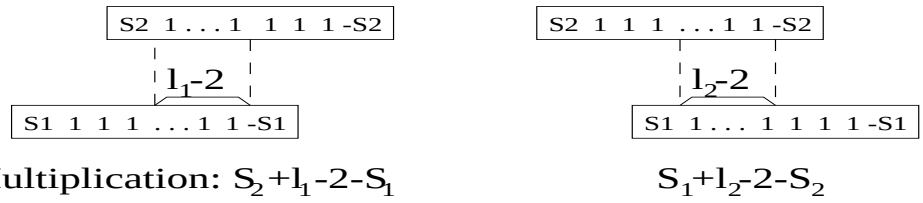

Figure 4: Edge intersection cases for text T^{oe} and pattern P^{eo}.

If we could solve two problems, we are done.

1. Remove the size of the overlap -2 from the result. (This will make the multiplication result be 0 when all overlaps are containments.)
2. Be able to actually replace every segment by its starting point, 1's and the negation of its starting point. (This will make the multiplication result positive when there exist side overlaps.) Being able to insert the starting point of the segment for the sake of multiplication is not at all simple to do since the same segment has $n - m$ different starting points!

However, if the above two problems could be solved by a convolution, then the result would be 0 iff all overlaps are containments iff all overlaps have even length (by observation 2).

Solution 1. The first problem can be easily solved. The convolution below provides, for every text location i, the sum $\sum_j (k_j - 2)$, where k_j are all the overlap lengths. All that is necessary, then, is to subtract the result of this convolution from the value obtained for each text location.

The Overlap Length Convolution: Replace every segment of length ℓ in both text and pattern by $01^{\ell-2}0$. Replace all ϕ's by 0's. Multiply.

Solution 2. The starting locations of all text segments are known in advance and thus the substitution of the segment by $s_1 1^{\ell_t - 2} - s_1$ can be done in linear time for the entire text. The pattern segment is problematic since it has $n - m$ different alignments. Therefore, we treat the pattern segments as if they all start relative to location 1.

The Zero Containment Convolution: Replace every segment of length ℓ_t starting in location s_1 of the text by $s_1 1^{\ell_t - 2} - s_1$. Replace every segment of length ℓ_p starting at location sp_1 of the *pattern* by $sp_1 1^{\ell_p - 2} - sp_1$. Replace all ϕ's by 0's. Multiply.

The problem is that now the result of the multiplication of side overlaps (after subtraction of the overlap length convolution) may sometimes be positive and sometimes be negative. This may cause the zero containment convolution result to be a 0 even for overlaps that are not contained, i.e. of odd length. We will then not be able to distinguish between them and the cases of all overlaps being containments.

We now show how to adjust the zero containment convolution to give the desired result. The zero containment convolution (after subtracting the result of the overlap length convolution) causes all contained segments to give the result 0. Thus in location i_0 we get the value

$$\sum_{\ell \in A} sp_\ell - \sum_{k \in B} sp_k + \sum_{i \in C} s_i - \sum_{j \in D} s_j,$$

where A is the the set of pattern segments that have text segments overlapping them from the left, B is the the set of pattern segments that have text segments overlapping them from the right, C is the set of text segments that have right overlap with pattern segments and D is the set of text segments that have left overlap with pattern segments.

The problem is that the result we really want is

$$\sum_{\ell \in A} s_\ell - \sum_{k \in B} s_k + \sum_{i \in C} s_i - \sum_{j \in D} s_j.$$

In other words, we want the starting positions of the aligned pattern segments, rather than their starting positions relative to location 1. However, note that for location i_0, $s_x = i_0 + sp_x$. This means that the value we want in location i_0 is

$$\sum_{\ell \in A} (sp_\ell + i_0) - \sum_{k \in B} (sp_k + i_0) + \sum_{i \in C} s_i - \sum_{j \in D} s_j$$

This equals

$$\sum_{\ell \in A} sp_\ell - \sum_{k \in B} sp_k + \sum_{i \in C} s_i - \sum_{j \in D} s_j + \sum_{\ell \in A} i_0 - \sum_{k \in B} i_0.$$

We conclude that the desired value will be obtained if we add to the result of the zero containment convolution at every location i_0 the value $\sum_{\ell \in A} i_0 - \sum_{k \in B} i_0$. These values can be obtained by the following convolution.

The Shifting Convolution: Replace every pattern segment of length ℓ_p by $10^{\ell_p-2}-1$ and every text segment of length ℓ_t by $01^{\ell_t-2}0$. Replace every ϕ by 0. Multiply.

It is easy to see that the result of this convolution in location i_0 is $\sum_{\ell\in A} i_0 - \sum_{k\in B} i_0$.

Definition: Let T be a text with segments $St_1,\ldots,St_x$ starting at locations $s_1,\ldots,s_x$, respectively. Let P be a pattern such that, when placed at location i_0, its segments $Sp_1,\ldots,Sp_y$ start at locations $v_1,\ldots,v_y$. Let $\{s_{i_1},\ldots,s_{i_z}\}$ be the starting locations of text segments that have a left overlap with pattern segments, and let $\{v_{j_1},\ldots,v_{j_z}\}$ be the starting location of the corresponding pattern segments. Similarly let $\{s_{k_1},\ldots,s_{k_w}\}$ be the starting locations of text segments that have a right overlap with pattern segments, and $\{v_{\ell_1},\ldots,v_{\ell_w}\}$ be the starting locations of the corresponding pattern segments.

The *external side length* of the pattern at location i_0 is

$$\sum_{r=1}^{y}(v_{i_r}-s_{j_r})+\sum_{r=1}^{z}(s_{k_r}-v_{\ell_r}).$$

Note that containments do not contribute to the external side length. See figure 5 for an illustration.

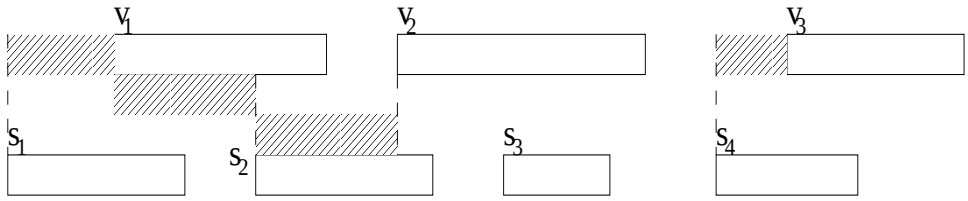

Figure 5: The sum of the shaded areas is the external side length:
$$(v_1-s_1)+(v_2-s_2)+(v_3-s_4)+(s_2-v_1).$$
Note that we are not counting external lengths of containments.

Lemma 3 *It is possible in time $O(n\log m)$ to provide, for every text location i, a value that is 0 if all overlaps of the pattern and text segments are containments, and positive otherwise.*

Proof: Add the results of the shifting convolution and the zero containment convolution and subtract from it the result of the overlap length convolution. We get the external side length which is 0 if all overlaps are containments and positive, otherwise. □

7 The Odd-Odd Even-Even Segments

Consider the case T^{oo} and P^{ee} (the case of T^{ee} and P^{oo} is symmetric).

Observation 3 *: For every two segments, S_t in T^{oo} and S_p in P^{ee} if either S_p is contained in S_t or S_t is contained in S_p then the overlap is of odd length. If the overlap is a left overlap or right overlap then it is of even length. See all possibilities in figure 6 below.*

The correctness of the observation is immediate. Segments of these types have odd lengths thus if one contains the other the overlap is necessarily of odd length. Conversely, in case of a left or right

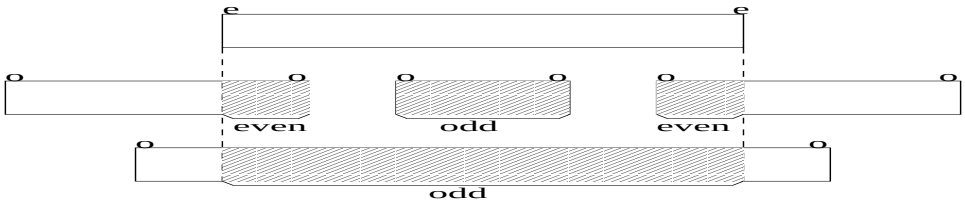

Figure 6: Every containment has odd length; every side overlap has even length.

overlap then the overlap starting and ending locations have the opposite parity, making the length of the overlap even.

Definition: Let T be a text with segments $St_1, \ldots, St_x$ starting at locations $s_1, \ldots, s_x$, respectively. Let P be a pattern such that, when placed at location i_0, its segments $Sp_1, \ldots, Sp_y$ start at locations $v_1, \ldots, v_y$. Let $\{s_{i_1}, \ldots, s_{i_z}\}$ be the starting locations of text segments that have a left overlap with pattern segments or that contain a pattern segments, and let $\{v_{j_1}, \ldots, v_{j_z}\}$ be the starting location of the corresponding pattern segments. Similarly let $\{s_{k_1}, \ldots, s_{k_w}\}$ be the starting locations of text segments that have a right overlap with pattern segments or that are contained in pattern segments, and $\{v_{\ell_1}, \ldots, v_{\ell_w}\}$ be the starting locations of the corresponding pattern segments.

The *external overlap length* of the pattern at location i_0 is

$$\sum_{r=1}^{y}(v_{i_r} - s_{j_r}) + \sum_{r=1}^{z}(s_{k_r} - v_{\ell_r}).$$

Note that unlike the external side length, containments also contribute to the external overlap length. See figure 7 for an illustration.

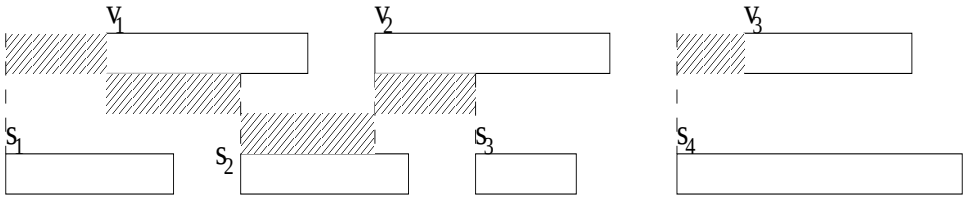

Figure 7: The sum of the shaded areas is the external overlap length:
$(v_1 - s_1) + (v_2 - s_2) + (v_3 - s_4) + (s_2 - v_1) + (s_3 - v_2)$.
Note that $(s_3 - v_2) + (v_3 - s_4)$ would not be counted in the external side length.

We need a method to indicate whether there exist any overlaps. We will show that the external overlap length can be calculated by two convolutions. Since we know from section 6 how to compute the external side length, it suffices to see the difference between these two values. If it is 0 then there are no containments (and therefore no odd-length overlaps), otherwise there are containments.

The sum of the following two convolutions gives the external overlap length.

The External Contained and Right Length Convolution: Replace every segment of length ℓ in the pattern by $0123\cdots\ell-1$, and every segment of length ℓ in the text by $10^{\ell-1}$. Replace every ϕ by 0. Multiply.

The resulting value at every location i_0 is precisely the external length of the right overlaps and of the text segments contained in pattern segment overlaps. Left overlaps and text containing pattern

give a 0. The next convolution computes precisely the remaining external lengths.

The External Containing and Left Length Convolution: Replace every segment of length ℓ in the text by $0123\cdots\ell-1$, and every segment of length ℓ in the pattern by $10^{\ell-1}$. Replace every ϕ by 0. Multiply.

Adding the results of the above two convolutions gives the external overlap length.

Lemma 4 *It is possible in time $O(n \log m)$ to provide, for every text location i, a value that is 0 if there are no containments either of text in pattern segments or of pattern in text segments, and positive otherwise.*

Proof: Subtract the external side length from the external overlap length. If there is a containment, we will get a positive number, otherwise we get 0. □

References

[1] A. Amir, Y. Aumann, G. Landau, M. Lewenstein, and N. Lewenstein. Pattern matching with swaps. Submitted for publication.

[2] A. Amir, Y. Aumann, G. Landau, M. Lewenstein, and N. Lewenstein. Pattern matching with swaps. *Proc. 38th IEEE FOCS*, pages 144–153, 1997.

[3] A. Amir and M. Farach. Efficient 2-dimensional approximate matching of half-rectangular figures. *Information and Computation*, 118(1):1–11, April 1995.

[4] A. Amir, G.M. Landau, M. Lewenstein, and N. Lewenstein. Efficient special cases of pattern matching with swaps. *Information Processing Letters*, 68(3):125–132, 1998.

[5] R. Cole and R. Harihan. Randomized swap matching in $o(m \log m \log |\sigma|)$ time. Technical Report TR1999-789, New York University, Courant Institute, September 1999.

[6] M.J. Fischer and M.S. Paterson. String matching and other products. *Complexity of Computation, R.M. Karp (editor), SIAM-AMS Proceedings*, 7:113–125, 1974.

[7] D.E. Knuth, J.H. Morris, and V.R. Pratt. Fast pattern matching in strings. *SIAM J. Computing*, 6:323–350, 1977.

[8] R. Lowrance and R. A. Wagner. An extension of the string-to-string correction problem. *J. of the ACM*, pages 177–183, 1975.

[9] S. Muthukrishnan. New results and open problems related to non-standard stringology. In *Proc. 6th Combinatorial Pattern Matching Conference*, pages 298–317. Lecture Notes in Computer Science 937, Springer-Verlag, 1995.

[10] S. Muthukrishnan and K. Palem. Non-standard stringology: Algorithms and complexity. In *Proc. 26th Annual Symposium on the Theory of Computing*, pages 770–779, 1994.

[11] S. Muthukrishnan and H. Ramesh. String matching under a general matching relation. *Information and Computation*, 122(1):140–148, 1995.

[12] Joseph (Seffi) Naor and Moni Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM Journal on Computing*, pages 838–856, 1993.

[13] M. V. Olson. A time to sequence. *Science*, 270:394–396, 1995.

[14] A. Pentland. Invited talk. NSF Institutional Infrastructure Workshop, 1992.

[15] R. A. Wagner. On the complexity of the extended string-to-string correction problem. In *Proc. 7th ACM STOC*, pages 218–223, 1975.