

EFFICIENT VARIANTS OF THE BACKWARD-ORACLE-MATCHING ALGORITHM

SIMONE FARO

*Dipartimento di Matematica e Informatica, Università di Catania
Viale A.Doria n.6, 95125 Catania, Italy
faro@dmf.unict.it*

THIERRY LECROQ

*Université de Rouen, LITIS EA 4108
76821 Mont-Saint-Aignan Cedex, France
thierry.lecroq@univ-rouen.fr*

Received 28 November 2008

Accepted 1 March 2009

Communicated by J. Holub

In this article we present two efficient variants of the BOM string matching algorithm which are more efficient and flexible than the original algorithm. We also present bit-parallel versions of them obtaining an efficient variant of the BNDM algorithm. Then we compare the newly presented algorithms with some of the most recent and effective string matching algorithms. It turns out that the new proposed variants are very flexible and achieve very good results, especially in the case of large alphabets.

Keywords: String matching; experimental algorithms; text processing; automaton.

1. Introduction

Given a text t of length n and a pattern p of length m over some alphabet Σ of size σ , the *string matching problem* consists in finding *all* occurrences of the pattern p in the text t . It is an extensively studied problem in computer science, mainly due to its direct applications to such diverse areas as text, image and signal processing, speech analysis and recognition, information retrieval, computational biology and chemistry, etc.

Many string matching algorithms have been proposed over the years (see [8]). The Boyer-Moore algorithm [5] deserves a special mention, since it has been particularly successful and has inspired much work.

Automata play a very important role in the design of efficient pattern matching algorithms. For instance the well known Knuth-Morris-Pratt algorithm [16] uses a deterministic automaton that searches a pattern in a text by performing its transitions on the text characters. The main result relative to the Knuth-Morris-

Pratt algorithm is that its automaton can be constructed in $\mathcal{O}(m)$ -time and -space, whereas pattern search takes $\mathcal{O}(n)$ -time.

Automata based solutions have been also developed to design algorithms which have optimal sublinear performance on average. This is done by using factor automata [4, 9, 3, 1], data structures which identify all factors of a word. Among the algorithms which make use of a factor automaton the BOM (Backward Oracle Matching) algorithm [1] is the most efficient, especially for long patterns. Another algorithm based on the bit-parallel simulation [2] of the nondeterministic factor automaton, and called BNDM (Backward Nondeterministic Dawg Match) algorithm [19], is very efficient for short patterns.

In this article we present two efficient variations of the BOM string matching algorithm which turn out to be more efficient and flexible than the original BOM algorithm. We also present a bit-parallel version of the previous solution which efficiently extends the BNDM algorithm.

The article is organized as follows. In Section 2 we introduce basic definitions and the terminology used along the paper. In Section 3 we survey some of the most effective string matching algorithms. Next, in Section 4, we introduce two new variations of the BOM algorithm. Experimental data obtained by running under various conditions all the algorithms reviewed are presented and compared in Section 5. Finally, we draw our conclusions in Section 6.

2. Basic Definitions and Terminology

A string p of length m is represented as a finite array $p[0..m-1]$, with $m \geq 0$. In particular, for $m = 0$ we obtain the empty string, also denoted by ε . By $p[i]$ we denote the $(i+1)$ -st character of p , for $0 \leq i < m$. Likewise, by $p[i..j]$ we denote the substring of p contained between the $(i+1)$ -st and the $(j+1)$ -st characters of p , for $0 \leq i \leq j < m$. Moreover, for any $i, j \in \mathbb{Z}$, we put $p[i..j] = \varepsilon$ if $i > j$ and $p[i..j] = p[\max(i, 0), \min(j, m-1)]$ if $i \leq j$. A substring of the form $p[0..i]$ is called a *prefix* of p and a substring of the form $p[i..m-1]$ is called a *suffix* of p for $0 \leq i \leq m-1$. For any two strings u and w , we write $w \sqsupseteq u$ to indicate that w is a suffix of u . Similarly, we write $w \sqsubseteq u$ to indicate that w is a prefix of u . The *reverse* of a string $p[0..m-1]$ is the string built by the concatenation of its letters from the last to the first: $p[m-1]p[m-2] \cdots p[1]p[0]$.

A Finite State Automaton is a tuple $A = \{Q, q_0, F, \Sigma, \delta\}$, where Q is the set of states of the automaton, $q_0 \in Q$ is the initial state, $F \subseteq Q$ is the set of accepting states, Σ is the alphabet of characters labeling transitions and $\delta() : (Q \times \Sigma) \rightarrow Q$ is the transition function. If $\delta(q, c)$ is not defined for a state $q \in Q$ and a character $c \in \Sigma$ we say that $\delta(q, c)$ is an undefined transition and write $\delta(q, c) = \perp$.

Let t be a text of length n and let p be a pattern of length m . When the character $p[0]$ is aligned with the character $t[s]$ of the text, so that the character $p[i]$ is aligned with the character $t[s+i]$, for $i = 0, \dots, m-1$, we say that the pattern p has *shift* s in t . In this case the substring $t[s..s+m-1]$ is called the *current window* of the text. If $t[s..s+m-1] = p$, we say that the shift s is *valid*.

Most string matching algorithms have the following general structure. First, during a *preprocessing phase*, they calculate useful mappings, in the form of tables, which later are accessed to determine nontrivial shift advancements. Next, starting with shift $s = 0$, they look for all valid shifts, by executing a *matching phase*, which determines whether the shift s is valid and computes a *positive* shift increment.

For instance, in the case of the naive string matching algorithm, there is no preprocessing phase and the matching phase always returns a unitary shift increment, i.e. all possible shifts are actually processed.

In contrast the Boyer-Moore algorithm [5] checks whether s is a valid shift, by scanning the pattern p from right to left and, at the end of the matching phase, it computes the shift increment as the maximum value suggested by two heuristics: the *good-suffix heuristic* and the *bad-character heuristic*, provided that both of them are applicable (see [8]).

3. Very Fast String Matching Algorithms

In this section we briefly review the BOM algorithm and other efficient algorithms for exact string matching that have been recently proposed. In particular, we present algorithms in the Fast-Search family [7], algorithms in the q -Hash family [18] and some among the most efficient algorithms based on factor automata.

3.1. Fast-search and forward-fast-search algorithms

The Fast-Search algorithm [6] is a very simple, yet efficient, variant of the Boyer-Moore algorithm. Let p be a pattern of length m and let t be a text of length n over a finite alphabet Σ . The Fast-Search algorithm computes its shift increments by applying the bad-character rule if and only if a mismatch occurs during the first character comparison, namely, while comparing characters $p[m - 1]$ and $t[s + m - 1]$, where s is the current shift. Otherwise it uses the good-suffix rule.

The Forward-Fast-Search algorithm [7] maintains the same structure of the Fast-Search algorithm, but it is based upon a modified version of the good-suffix rule, called *forward good-suffix* rule, which uses a look-ahead character to determine larger shift advancements. Thus, if the first mismatch occurs at position $i < m - 1$ of the pattern p , the forward good-suffix rule suggests to align the substring $t[s + i + 1 .. s + m]$ with its rightmost occurrence in p preceded by a character different from $p[i]$. If such an occurrence does not exist, the forward good-suffix rule proposes a shift increment which allows to match the longest suffix of $t[s + i + 1 .. s + m]$ with a prefix of p . This corresponds to advance the shift s by $\vec{gs}_P(i + 1, t[s + m])$ positions, where

$$\begin{aligned} \vec{gs}_P(j, c) =_{\text{def}} \min(\{0 < k \leq m \mid p[j - k .. m - k - 1] \sqsupseteq p \\ \text{and } (k \leq j - 1 \rightarrow p[j - 1] \neq p[j - 1 - k]) \\ \text{and } p[m - k] = c\} \cup \{m + 1\}) , \end{aligned}$$

for $j = 0, 1, \dots, m$ and $c \in \Sigma$.

The good-suffix rule and the forward good-suffix rule require tables of size m and $m \cdot |\Sigma|$, respectively. These can be constructed in time $\mathcal{O}(m)$ and $\mathcal{O}(m \cdot \max(m, |\Sigma|))$, respectively.

More effective implementations of the Fast-Search and Forward-Fast-Search algorithm are obtained along the same lines of the Tuned-Boyer-Moore algorithm [15] by making use of a fast-loop, using a technique described in Section 4.1 and shown in Figure 3(A). Then subsequent matching phase can start with the $(m - 2)$ -th character of the pattern. At the end of the matching phase the algorithms uses the good-suffix rule for shifting.

3.2. The q -Hash algorithms

Algorithms in the q -Hash family have been introduced in [18] where the author presented an adaptation of the Wu and Manber multiple string matching algorithm [23] to single string matching problem.

The idea of the q -Hash algorithm is to consider factors of the pattern of length q . Each substring w of such a length q is hashed using a function h into integer values within 0 and 255. Then the algorithm computes in a preprocessing phase a function $shift() : \{0, 1, \dots, 255\} \rightarrow \{0, 1, \dots, m - q\}$. Formally for each $0 \leq c \leq 255$ the value $shift(c)$ is defined by

$$shift(c) = \min \left(\{0 \leq k < m - q \mid h(p[m - k - q .. m - k - 1]) = c\} \cup \{m - q\} \right) .$$

The searching phase of the algorithm consists in reading, for each shift s of the pattern in the text, the substring $w = t[s + m - q .. s + m - 1]$ of length q . If $shift[h(w)] > 0$ then a shift of length $shift[h(w)]$ is applied. Otherwise, when $shift[h(w)] = 0$ the pattern x is naively checked in the text. In this case a shift of length sh is applied where $sh = m - 1 - i$ with

$$i = \max\{0 \leq j \leq m - q \mid h(x[j .. j + q - 1]) = h(x[m - q + 1 .. m - 1])\}.$$

3.3. The backward-automaton-matching algorithms

Algorithms based on the Boyer-Moore strategy generally try to match suffixes of the pattern but it is possible to match some prefixes or some factors of the pattern by scanning the current window of the text from right to left in order to improve the length of the shifts. This can be done by the use of factor automata and factor oracles.

The factor automaton [4, 9, 3] of a pattern p , $Aut(p)$, is also called the factor DAWG of p (for Directed Acyclic Word Graph). Such an automaton recognizes all the factors of p . Formally the language recognized by $Aut(p)$ is defined as follows

$$\mathcal{L}(Aut(p)) = \{u \in \Sigma^* : \text{exists } v, w \in \Sigma^* \text{ such that } p = vuw\}.$$

The factor oracle of a pattern p , $Oracle(p)$, is a very compact automaton which recognizes at least all the factors of p and slightly more other words. Formally $Oracle(p)$ is an automaton $\{Q, m, Q, \Sigma, \delta\}$ such that




Fig. 1. The factor automaton (A) and the factor oracle (B) of the reverse of pattern $p = baabbba$. The factor automaton recognizes all, and only, the factors of the reverse pattern. On the other hand note that the word aba is recognized by the factor oracle whereas it is not a factor.

- (1) Q contains exactly $m + 1$ states, say $Q = \{0, 1, 2, 3, \dots, m\}$
- (2) m is the initial state
- (3) all states are final
- (4) the language accepted by $Oracle(p)$ is such that $\mathcal{L}(Aut(p)) \subseteq \mathcal{L}(Oracle(p))$.

Despite the fact that the factor oracle is able to recognize words that are not factors of the pattern, it can be used to search for a pattern in a text since the only factor of p of length greater or equal to m which is recognized by the oracle is the pattern itself. The computation of the oracle is linear in time and space in the length of the pattern.

In Figure 1 are shown the factor automaton and the factor oracle of the reverse of pattern $p = baabbba$.

The data structures factor automaton and factor oracle are used respectively in [10, 11] and in [1] to get optimal pattern matching algorithms on the average. The algorithm which makes use of the factor automaton of the reverse pattern is called BDM (for Backward Dawg Matching) while the algorithm using the factor oracle is called BOM (for Backward Oracle Matching). Such algorithms move a window of size m on the text. For each new position of the window, the automaton of the reverse of p is used to search for a factor of p from the right to the left of the window. The basic idea of the BDM and BOM algorithms is that if its backward search failed on a letter c after the reading of a word u then cu is not a factor of p and moving the beginning of the window just after c is secure. If a factor of length m is recognized then we have found an occurrence of the pattern.

The BDM and BOM algorithms have a quadratic worst case time complexity but are optimal in average since they perform $\mathcal{O}(n(\log_{\sigma} m)/m)$ inspections of text characters reaching the best bound shown by Yao [24] in 1979.

3.4. The BNBM algorithm

The BNBM algorithm [19] (for Backward Nondeterministic Dawg Match) is a bit-parallel simulation [2] of the BDM algorithm. It uses a nondeterministic automaton instead of the deterministic one in the BDM algorithm.




Fig. 2. The nondeterministic factor automaton of the string *abbaab*.

Figure 2 shows the nondeterministic version of the factor automaton for the reverse of pattern $p = baabbba$. For each character $c \in \Sigma$, a bit vector $B[c]$ is initialized in the preprocessing phase. The i -th bit is 1 in this vector if c appears in the reversed pattern in position i . Otherwise the i -th bit is set to 0. The state vector D is initialized to 1^m . The same kind of right to left scan in a window of size m is performed as in the BDM algorithm while the state vector is updated in a similar fashion as in the Shift-And algorithm [2]. If the m -th bit is 1 after this update operation, we have found a prefix starting at position j where j is the number of updates done in this window. Thus if j is the first position in the window, a match has been found.

A simplified version of the BNDM, called SBNDM, has been presented in [13]. This algorithm differs from the original one in the main loop which starts each iteration with a test of two consecutive text characters. Moreover it implements a fast-loop to obtain better results on average. Experimental results show that this simplified variant is always more efficient than the original one.

4. New Variations of the BOM Algorithm

In this section we present two variations of the BOM algorithm which perform better in most cases. The first idea consists in extending the BOM algorithm with a fast-loop over oracle transitions, along the same lines of the Tuned-Boyer-Moore algorithm [15]. Thus we are able to perform factor searching if and only if a portion of the pattern has already matched against a portion of the current window of the text. We present this idea in Section 4.1.

Another efficient variation of the BOM algorithm can be obtained by applying the idea suggested by Sunday in the Quick-Search algorithm and then implemented in the Forward-Fast-Search algorithm. This consists in taking into account, while shifting, the character which follows the current window of the text, since it is always involved in the next alignment. Such a variation is presented in Section 4.2.

4.1. Extending the BOM algorithm with a fast-loop

In this section we present an extension of the BOM algorithm by introducing a fast-loop with the aim of obtaining better results on the average. We discuss the application of different variations of the fast-loop, listed in Figure 3, and present experimental results in order to identify the best choice.

The idea of a fast loop has been proposed in [5]. The fast-loop we are using here has been first introduced in the Tuned-Boyer-Moore algorithm [15] and later largely used in almost all variations of the Boyer-Moore algorithm. Generally a fast-loop is implemented by iterating the bad character heuristic in a checkless cycle, in order to quickly locate an occurrence of the rightmost character of the pattern. Suppose $bc() : \Sigma \rightarrow \{0, 1, \dots, m\}$ is the function which implements the bad-character heuristic defined, for all $c \in \Sigma$, by

$$bc(c) = \min(\{0 \leq k < m \mid p[m - 1 - k] = c\} \cup \{m\}) .$$

If we suppose that $t[j]$ is the rightmost character of the current window of the text for a shift s , i.e. $j = s + m - 1$, then the original fast-loop can be implemented in a form similar to that presented in Figure 3(A).

In order to avoid testing the end of the text we could append the pattern at the end of the text, i.e. set $t[n..n + m - 1]$ to p . Thus we exit the algorithm only when an occurrence of p is found. If this is not possible (because memory space is occupied) it is always possible to store $t[n - m..n - 1]$ in z then set $t[n - m..n - 1]$ to p and check z at the end of the algorithm without slowing it.

However algorithms based on the bad character heuristic obtain good results only in the case of large alphabets and short patterns. It turns out moreover from experimental results [18] that the strategy of using an automaton to match prefixes or factors is much better when the length of the pattern increases.

<p>(A)</p> <pre> k = bc(t_j) while (k ≠ 0) do j = j + k k = bc(t_j) </pre>	<p>(B)</p> <pre> q = δ(m, t_j) while (q == ⊥) do j = j + m q = δ(m, t_j) </pre>	<p>(C)</p> <pre> q = δ(m, t_j) if q ≠ ⊥ then p = δ(q, t_{j-1}) while (p == ⊥) do j = j + m - 1 q = δ(m, t_j) if q ≠ ⊥ then p = δ(q, t_{j-1}) </pre>	<p>(D)</p> <pre> q = λ(t_j, t_{j-1}) while (q == ⊥) do j = j + m - 1 q = λ(t_j, t_{j-1}) </pre>
--	---	---	---

Fig. 3. Different variations of the fast-loop where $t_j = t_{s+m-1}$ is the rightmost character of the current window of the text. (A) The original fast-loop based on the bad character rule. (B) A modified version of the fast-loop based on automaton transitions. (C) A fast-loop based on two subsequent automaton transitions. (D) An efficient implementation of the previous fast loop which encapsulate two subsequent transitions in a single λ table.

Experimental results with $\sigma = 8$						Experimental results with $\sigma = 16$					
m	BOM	(A)	(B)	(C)	(D)	m	BOM	(A)	(B)	(C)	(D)
4	157.6	95.9	135.9	109.0	55.3	4	103.2	66.8	86.2	93.5	40.6
8	85.4	58.6	78.7	58.6	34.1	8	71.5	38.7	60.2	44.0	21.7
16	43.0	43.3	43.0	37.1	26.8	16	39.6	26.5	35.7	23.6	14.9
32	26.6	35.0	28.2	25.9	21.2	32	18.6	21.8	18.8	15.7	12.7
64	17.3	28.0	17.1	17.0	14.4	64	12.6	20.0	12.5	12.7	12.4
128	15.2	23.6	15.7	15.8	12.8	128	14.2	19.3	14.1	12.3	10.3
256	10.7	19.8	9.6	9.7	8.5	256	8.8	19.0	8.1	7.8	6.8
512	6.1	14.2	6.1	5.7	4.7	512	4.6	17.7	4.6	4.5	3.6
1024	3.2	8.2	3.4	3.3	2.6	1024	2.3	11.4	2.6	2.8	2.6

Experimental results with $\sigma = 32$						Experimental results with $\sigma = 64$					
m	BOM	(A)	(B)	(C)	(D)	m	BOM	(A)	(B)	(C)	(D)
4	78.7	55.2	57.7	88.5	37.4	4	64.8	50.9	42.3	88.5	37.0
8	51.6	30.3	42.0	39.8	18.5	8	39.3	27.4	29.2	38.8	17.9
16	35.4	19.9	30.1	20.3	12.2	16	26.0	17.1	22.0	20.0	11.5
32	20.6	16.1	19.3	12.2	11.5	32	19.4	14.0	17.1	11.8	10.7
64	12.1	14.8	11.5	10.6	11.1	64	13.1	13.5	12.2	10.3	10.7
128	12.6	15.6	11.2	10.0	7.4	128	13.1	17.6	10.8	9.7	6.3
256	7.5	16.7	6.3	5.9	3.7	256	6.2	18.0	5.7	5.5	3.6
512	4.2	17.9	3.7	3.8	3.2	512	2.9	18.0	3.1	5.3	1.9
1024	2.8	14.1	2.6	2.7	2.0	1024	2.7	16.8	2.5	2.4	1.5

Fig. 4. Experimental results obtained by comparing the original BOM algorithm (in the first column) against variations implemented using the four fast-loop presented in Figure 3. The results have been obtained by searching 200 random patterns in a 40Mb text buffer with a uniform distribution over an alphabet of dimension σ . Running times are expressed in hundredths of seconds.

This behavior is due to the fact that for large patterns an occurrence of the rightmost character of the window, i.e. $t[j]$, can be found in the pattern and the probability that the rightmost occurrence is near the rightmost position increases for longer patterns and smaller alphabets. In this latter case an iteration of the fast loop leads to a short shift. In contrast, when using an oracle for matching, it is common that after a small number of characters we are not able to perform other transitions. So generally this strategy looks for a number of characters greater than 1, for each iteration, but leads to shift of length m .

As a first step we can translate the idea of the fast-loop over to automaton transitions. This consists in shifting the pattern along the text with no more check until a non-undefined transition is found with the rightmost character of the current window of the text. This can be translated in the fast-loop presented in Figure 3(B).

It turns out from experimental results presented in Figure 4 that the variation of the BOM algorithm which uses the fast-loop on transitions (col.B) performs better than the original algorithm (first column), especially for large alphabets. However it is not flexible since its performances decrease when the length of the pattern increases or when the dimension of the alphabet is small. This is the fast-loop finds only a small number of undefined transitions for small alphabets or long patterns.

The variation of the algorithm we propose tries two subsequent transitions for each iteration of the fast-loop with the aim to find with higher probability an undefined transition. This can be translated in the fast-loop presented in Figure 3(C). From experimental results it turns out that such a variation (Figure 4, col.C) obtains better results than the previous one only for long pattern and large alphabets. This is for each iteration of the fast-loop the algorithm performs two subsequent transitions affecting the overall performance.

To avoid this problem we could encapsulate the two first transitions of the oracle in a function $\lambda() : (\Sigma \times \Sigma) \rightarrow Q$ defined, for each $a, b \in \Sigma$, by

$$\lambda(a, b) = \begin{cases} \perp & \text{if } \delta(m, a) = \perp \\ \delta(\delta(m, a), b) & \text{otherwise.} \end{cases}$$

Thus the fast loop can be implemented as presented in Figure 3(D). At the end of the fast-loop the algorithm could start standard transitions with the Oracle from state $q = \lambda(t[j], t[j - 1])$ and character $t[j - 2]$. The function λ can be implemented with a two dimensional table in $\mathcal{O}(\sigma^2)$ time and space.

The resulting algorithm, here named Extended-BOM algorithm, is very fast and flexible. Its pseudocode is presented in Figure 6(A). From experimental results in Figure 4 it turns out that the Extended-BOM algorithm (col.D) is the best choice in most cases and, differently from the original algorithm, it has very good performance also for short patterns.

4.2. Looking for the forward character

The idea of looking for the forward character for shifting has been originally introduced by Sunday in the Quick-Search algorithm [21] and then efficiently implemented in the Forward-Fast-Search algorithm [7]. Specifically, it is based on the following observation: when a mismatch character is encountered while comparing the pattern with the current window of the text $t[s .. s + m - 1]$, the pattern is always shifted to the right by at least one character, but never by more than m characters. Thus, the character $t[s + m]$ is always involved in testing for the next alignment.

In order to take into account the forward character of the current window of the text without skip safe alignment we construct the *forward factor oracle* of the reverse pattern. The forward factor oracle of a word p , $FOracle(p)$, is an automaton which recognizes at least all the factors of p , eventually preceded by a word $x \in \Sigma \cup \{\varepsilon\}$. More formally the language recognized by $FOracle(p)$ is defined by

$$\mathcal{L}(FOracle(p)) = \{xw \mid x \in \Sigma \cup \{\varepsilon\} \text{ and } w \in \mathcal{L}(Oracle(p))\}$$

Observe that in the previous definition the prefix x could be the empty string. Thus if w is a word recognized by the factor oracle of p then the word cw is recognized by the forward factor oracle, for all $c \in \Sigma \cup \{\varepsilon\}$.

The forward factor oracle of a word p can be constructed, in time $\mathcal{O}(m + \Sigma)$, by simply extending the factor oracle of p with a new initial state which allows




Fig. 5. (A) The forward factor oracle of the reverse pattern $p = baabbba$ (B) The nondeterministic version of the forward factor automaton of the reverse pattern $p = baabbba$

to perform transitions starting at the text character of position $s + m$ of the text, avoiding to skip valid shift alignments.

Suppose $Oracle(p) = \{Q, m, Q, \Sigma, \delta\}$, for a pattern p of length m . We construct $FOracle(p)$ by adding a new initial state $(m + 1)$ and introducing transitions from state $(m + 1)$. More formally, given a pattern p of length m , $FOracle(p)$ is an automaton $\{Q', (m + 1), Q, \Sigma, \delta'\}$, where

- (1) $Q' = Q \cup \{(m + 1)\}$
- (2) $(m + 1)$ is the initial state
- (3) all states are final
- (4) $\delta'(q, c) = \delta(q, c)$ for all $c \in \Sigma$, if $q \neq (m + 1)$
- (5) $\delta'(m + 1, c) = \{m, \delta(m, c)\}$ for all $c \in \Sigma$

Figure 5(A) shows the forward factor oracle of the reverse pattern $p = baabbba$. The dashed transitions are those outgoing from the new initial state. A transition labeled with all characters of the alphabet has been introduced from state $(m + 1)$ to state m . Note that, according to rule n.5, the forward factor oracle of the reverse pattern p is a non-deterministic automaton. For example, starting from the initial state 8 in Figure 5(A), after reading the couple of characters aa , both states 6 and 1 are active.

Observe moreover that we have to read at least two consecutive characters to find an undefined transition. This is state m is always active after reading any character of the alphabet.

Suppose we start transitions from the initial state of $FOracle(p)$. Then after reading a word $w = au$, with $a \in \Sigma$ and $u \in \Sigma^+$, at most two different states could be active, i.e., state $x = \delta^*(w)$ and state $y = \delta^*(u)$. Where we recall that δ is the transition function of $Oracle(p)$ and where $\delta^* : \Sigma^* \rightarrow Q$ is the final-state function induced by δ and defined recursively by

$$\delta^*(w) = \delta(\delta^*(w'), c), \text{ for each } w = w'c, \text{ with } w' \in \Sigma^*, c \in \Sigma.$$

The idea consists in simulating the behavior of the nondeterministic forward factor oracle by following transition for only one of the two active states. More

(A)	(B)
EXTENDED-BOM(p, m, t, n) 1. $\delta \leftarrow \text{precompute-factor-oracle}(p)$ 2. for $a \in \Sigma$ do 3. $q \leftarrow \delta(m, a)$ 4. for $b \in \Sigma$ do 5. if $q = \perp$ then $\lambda(a, b) \leftarrow \perp$ 6. else $\lambda(a, b) \leftarrow \delta(q, b)$ 7. $t[n..n+m-1] \leftarrow p$ 8. $j \leftarrow m-1$ 9. while $j < n$ do 10. $q \leftarrow \lambda(t[j], t[j-1])$ 11. while $q = \perp$ do 12. $j \leftarrow j+m-1$ 13. $q \leftarrow \lambda(t[j], t[j-1])$ 14. $i \leftarrow j-2$ 15. while $q \neq \perp$ do 16. $q \leftarrow \delta(q, t[i])$ 17. $i \leftarrow i-1$ 18. if $i < j-m+1$ then 19. output(j) 20. $i \leftarrow i+1$ 21. $j \leftarrow j+i+m$	FORWARD-BOM(p, m, t, n) 1. $\delta \leftarrow \text{precompute-factor-oracle}(p)$ 2. for $a \in \Sigma$ do 3. $q \leftarrow \delta(m, a)$ 4. for $b \in \Sigma$ do 5. if $q = \perp$ then $\lambda(a, b) \leftarrow \perp$ 6. else $\lambda(a, b) \leftarrow \delta(q, b)$ 7. $q \leftarrow \delta(m, p[m-1])$ 8. for $a \in \Sigma$ do $\lambda(a, p[m-1]) \leftarrow q$ 9. $t[n..n+m-1] \leftarrow p$ 10. $j \leftarrow m-1$ 11. while $j < n$ do 12. $q \leftarrow \lambda(t[j+1], t[j])$ 13. while $q = \perp$ do 14. $j \leftarrow j+m$ 15. $q \leftarrow \lambda(t[j+1], t[j])$ 16. $i \leftarrow j-1$ 17. while $q \neq \perp$ do 18. $q \leftarrow \delta(q, t[i])$ 19. $i \leftarrow i-1$ 20. if $i < j-m+1$ then 21. output(j) 22. $i \leftarrow i+1$ 23. $j \leftarrow j+i+m$

Fig. 6. (A) The Extended-BOM algorithm which extend the original BOM algorithm by using an efficient fast-loop. (B) The Forward-BOM algorithm which performs a look ahead for character of position $t[j+1]$ in text to obtain larger shift advancements.

precisely we are interested only in transitions from state q where

$$q = \begin{cases} y = \delta^*(u) & \text{if } u[0] = p[m-1] \\ x = \delta^*(w) & \text{otherwise} \end{cases}$$

To prove the correctness of our strategy, suppose first we have read a word $w = au$, as defined above, and $u[0] \neq p[m-1]$. If $Oracle(p)$ recognizes a word u (i.e. $\delta^*(u) \neq \perp$) then by definition $FOracle(p)$ recognize the word au , since $a \in \Sigma \cup \{\varepsilon\}$.

Suppose now that $u[0] = p[m-1]$. If $Oracle(p)$ recognizes a word w then it recognizes also word u which is a suffix of w . Thus by definition $FOracle(p)$ recognizes the word xu , with $x = \varepsilon$.

The simulation of the forward factor oracle can be done by simply changing the computation of the λ table in the following way

$$\lambda(a, b) = \begin{cases} \delta(m, b) & \text{if } \delta(m, a) = \perp \vee b = p[m-1] \\ \delta(\delta(m, a), b) & \text{otherwise} \end{cases}$$

Figure 6(B) shows the code of the Forward-BOM algorithm. Here the fast loop has been modified to take into account also the forward character of position $t[s+m]$. However if there is no transition for the first two characters, $t[s+m]$ and $t[s+m-1]$, the algorithm can shift the pattern of m position to the right. Line 1 of the preprocessing phase can be performed in $\mathcal{O}(m)$ -time, lines 2 to 6 in $\mathcal{O}(\sigma^2)$ and line 8 in $\mathcal{O}(\sigma)$. Thus the preprocessing phase can be performed in $\mathcal{O}(m + \sigma^2)$ time and space.

```

FORWARD-SBNDM( $p, m, t, n$ )
1.   for all  $c \in \Sigma$  do  $B[c] \leftarrow 1$ 
2.   for  $i = 0$  to  $m - 1$  do  $B[p[i]] \leftarrow B[p[i]] \mid (1 \ll (m - i))$ 
3.    $j \leftarrow m - 1$ 
4.   while  $j < n$  do
5.      $D \leftarrow (B[t[j + 1]] \ll 1) \& B[t[j]]$ 
6.     if  $D \neq 0$  then
7.        $pos \leftarrow j$ 
8.       while  $D \leftarrow (D + D) \& B[t[j - 1]]$  do  $j \leftarrow j - 1$ 
9.        $j \leftarrow j + m - 1$ 
10.      if  $j = pos$  then
11.        output( $j$ )
12.         $j \leftarrow j + 1$ 
13.      else  $j \leftarrow j + m$ 

```

Fig. 7. The Forward SBNDM algorithm which simulates using bit-parallelism the non deterministic forward automaton of the reverse pattern.

This idea can be applied also to the SBNDM algorithm based on bit-parallelism. In this latter case we have to add a new first state and change the preprocessing in order to perform correct transitions from the first state. Moreover we need $m + 1$ bits for representing the NFA, thus we are able to search only for patterns with $1 \leq m < w$, if w is the dimension of a machine word. Figure 7 shows the code of the Forward-SBNDM algorithm.

5. Experimental Results

We present next experimental data which allow to compare in terms of running time the following string matching algorithms under various conditions:

- Algorithms in the Boyer-Moore family
 - The Fast-Search algorithm [6] (FS)
 - The Forward-Fast-Search algorithm [7] (FFS)
 - The Two Sliding Window algorithm [14] (TSW)
 - The SSABS algorithm [20] (SSABS).
 - The TVSBS algorithm [22] (TVSBS)
- Three algorithms in the q -Hash family [18] (q -HASH with $q = 3, 5, 8$)
- Algorithms in the BOM family
 - The Backward Oracle Matching algorithm [1] (BOM)
 - The Extended-BOM algorithm (EBOM)
 - The Forward-BOM algorithm (FBOM)
- Algorithms based on bit-parallelism
 - The Simplified BNDM algorithm [13] (SBNDM)
 - The Forward-SBNDM algorithm (FSBNDM)
 - The Bit-parallel Length Invariant Matcher [17] (BLIM)
 - The Fast Average Optimal Shift Or algorithm [12] (FAOSO)

All algorithms have been implemented in the **C** programming language and were used to search for the same strings in large fixed text buffers on a PC with Intel Core2 processor of 1.66GHz. In particular, the algorithms have been tested on seven $\text{Rand}\sigma$ problems, for $\sigma = 2, 4, 8, 16, 32, 64, 128$, on a genome, on a protein sequence and on a natural language text buffer. Searching have been performed for patterns of length $m = 2, 4, 8, 16, 32, 64, 128, 256$ and 512.

For the sake of completeness we report also that algorithms in the BOM family have been implemented by constructing factor oracles in quadratic space with a transition matrix.

Moreover we remember that algorithms which make use of bit-parallelism, with the exception of the BLIM algorithm, can handle only pattern of length $m \leq w$. In our machine such a bound is equal to 32. In order to handle patterns with a length $m > 32$ algorithms in such a family search for all occurrences of the prefix p' of the input pattern p , where $|p'| = m' = \min(m, 32)$. If $m' < m$ and an occurrence of p' is found at a given position of the text then the algorithms naively check for the whole occurrence of the pattern p .

Finally we report that the FAOSO algorithm consists in considering sparse q -grams of the pattern and unrolling u shifts, thus $q(\lceil m/q \rceil + u) \leq w$ should hold, where w is again the number of bits of a machine word. In our test we have set $u = 4$. Then, for each pattern, we have run the FAOSO algorithm with different values of the input parameter q and have reported the best running time. In particular for a pattern of length m the value of q ranges between 2 and $\lfloor m/2 \rfloor$.

In the following tables, running times are expressed in hundredths of seconds. The best results have been boldfaced.

5.1. Running times for random problems

For the case of random texts the algorithms have been tested on seven $\text{Rand}\sigma$ problems. Each $\text{Rand}\sigma$ problem consists of searching a set of 400 random patterns of a given length in a 20Mb random text over a common alphabet of size σ , with a uniform distribution of characters.

Running times for a $\text{Rand}2$ problem.

m	4	8	16	32	64	128	256	512
FS	155.08	118.12	83.48	62.58	48.67	39.26	33.75	28.32
FFS	123.63	89.88	60.16	41.72	30.86	23.94	19.81	16.56
TSW	211.76	190.66	178.55	174.26	176.21	180.35	173.20	176.29
SSABS	173.55	199.57	206.72	203.78	204.80	208.75	206.29	204.61
TVSBS	162.46	168.24	172.58	165.70	170.39	167.47	170.16	170.16
3-HASH	162.62	77.97	51.76	46.17	45.00	44.30	44.88	45.83
5-HASH	–	67.97	25.55	14.72	11.41	10.54	10.35	10.39
8-HASH	–	–	34.45	13.36	7.19	8.32	7.11	6.17
BOM	210.63	134.38	72.42	38.20	20.90	12.11	8.63	6.49
EBOM	168.87	104.18	59.26	31.10	17.19	11.72	8.40	5.28
FBOM	177.97	115.39	63.32	33.98	18.91	15.74	9.88	5.54
SBNDM	141.88	83.86	43.55	22.18	22.23	22.27	22.18	22.19
FSBNDM	148.75	82.62	42.03	21.84	21.80	21.72	21.76	21.88
BLIM	159.61	98.28	56.99	41.91	30.43	31.68	34.02	38.75
FAOSO	193.32	98.02	43.25	52.09	51.95	52.78	51.28	51.64

Running times for a Rand4 problem.

<i>m</i>	4	8	16	32	64	128	256	512
FS	83.91	61.17	49.26	42.43	37.15	32.50	28.87	25.63
FFS	75.19	52.03	38.79	29.49	23.00	18.01	15.23	13.60
TSW	95.11	61.80	41.02	29.92	26.52	26.14	26.64	26.48
SSABS	99.65	83.36	80.70	79.65	78.05	80.20	79.92	78.16
TVSBS	89.34	60.08	42.34	33.05	30.70	30.51	30.23	30.16
3-HASH	118.12	43.16	22.93	16.05	13.47	12.58	12.34	12.34
5-HASH	–	56.64	20.27	10.47	7.15	6.83	6.88	7.07
8-HASH	–	–	33.83	13.09	7.34	8.20	7.07	6.25
BOM	111.45	61.29	35.86	20.08	11.53	8.12	6.29	3.94
EBOM	58.55	43.83	29.02	16.80	9.96	7.66	5.55	3.44
FBOM	84.80	51.76	31.37	18.09	10.66	9.80	6.21	3.56
SBNDM	59.53	39.69	23.63	12.89	12.54	12.58	12.58	12.66
FSBNDM	75.47	41.56	22.42	12.46	12.42	12.54	12.54	12.50
BLIM	86.56	47.23	32.46	23.79	18.95	20.16	22.62	27.35
FAOSO	81.29	41.91	27.70	18.49	19.04	18.50	18.83	18.83

Running times for a Rand8 problem.

<i>m</i>	4	8	16	32	64	128	256	512
FS	46.65	29.73	22.81	20.59	19.34	18.05	17.07	16.64
FFS	44.73	28.39	21.02	17.15	14.37	12.27	11.29	11.05
TSW	58.40	36.17	21.36	13.28	9.26	9.26	9.26	9.18
SSABS	54.41	38.12	31.09	28.75	28.21	29.22	28.83	28.48
TVSBS	52.11	32.66	19.96	12.93	9.50	8.13	8.04	8.20
3-HASH	106.14	37.46	18.12	11.29	8.63	7.70	7.70	7.85
5-HASH	–	54.80	19.22	9.42	6.52	7.58	6.92	6.37
8-HASH	–	–	33.98	13.17	7.23	8.20	6.95	6.17
BOM	79.53	43.05	22.03	13.47	8.60	6.91	5.43	3.04
EBOM	27.07	17.11	13.12	10.39	7.30	6.83	4.22	2.62
FBOM	50.08	30.08	19.11	12.62	7.97	7.38	4.25	2.50
SBNDM	33.79	18.83	12.81	9.41	9.42	9.45	9.49	9.49
FSBNDM	39.72	23.05	14.61	9.49	9.65	9.61	9.61	9.65
BLIM	58.32	34.49	22.62	15.94	14.53	15.67	18.09	22.73
FAOSO	48.79	25.18	15.46	14.70	14.86	14.96	14.82	14.72

Running times for a Rand16 problem.

<i>m</i>	4	8	16	32	64	128	256	512
FS	33.16	19.14	13.17	11.29	10.67	10.35	10.35	10.39
FFS	32.66	18.87	13.17	10.90	9.80	9.18	9.18	9.61
TSW	45.19	27.82	16.32	9.96	7.19	8.98	7.58	6.33
SSABS	36.64	23.40	16.80	14.06	13.05	13.01	12.89	13.05
TVSBS	37.15	22.85	13.48	8.39	6.60	8.32	6.95	6.13
3-HASH	102.38	35.31	16.30	9.38	6.84	6.80	6.80	6.88
5-HASH	–	54.18	18.83	9.14	6.29	8.20	6.96	6.09
8-HASH	–	–	34.02	13.20	7.23	8.28	7.11	6.21
BOM	52.93	35.82	19.77	9.57	6.37	7.35	4.61	2.62
EBOM	20.39	10.82	7.27	6.21	6.05	5.27	3.24	1.99
FBOM	38.98	22.34	13.16	8.67	7.19	10.15	4.45	2.00
SBNDM	28.44	13.87	8.28	6.21	6.29	6.29	6.33	6.37
FSBNDM	28.04	15.31	9.18	6.68	6.64	6.64	6.56	6.60
BLIM	47.93	27.23	16.52	11.37	10.47	11.76	14.06	18.71
FAOSO	40.59	20.82	14.50	14.04	13.67	13.54	13.63	13.79

Running times for a Rand32 problem.

m	4	8	16	32	64	128	256	512
FS	27.77	15.39	9.65	8.13	7.62	7.62	7.73	7.73
FFS	27.58	15.28	9.72	8.09	7.38	7.74	8.09	8.75
TSW	45.47	27.93	16.44	10.04	7.42	7.62	5.20	3.71
SSABS	29.77	18.01	11.72	9.34	8.55	8.20	8.28	8.44
TVSBS	31.13	19.26	11.45	7.23	6.25	9.26	6.13	4.37
3-HASH	100.86	34.45	15.59	8.44	6.33	7.54	6.76	6.41
5-HASH	–	54.18	18.71	9.06	6.21	8.12	6.95	5.90
8-HASH	–	–	34.02	13.08	7.23	8.32	7.11	6.21
BOM	36.41	24.76	17.42	10.23	5.82	6.21	3.56	2.19
EBOM	18.63	9.30	6.02	5.31	5.24	3.71	2.27	1.60
FBOM	35.78	20.23	11.60	7.23	6.60	10.59	6.05	3.16
SBNDM	26.99	12.58	7.11	5.59	5.51	5.54	5.27	5.55
FSBNDM	23.91	12.81	7.58	5.78	5.62	5.46	5.59	5.54
BLIM	40.16	23.24	15.82	11.72	11.33	12.15	14.49	19.03
FAOSO	38.36	19.72	12.28	12.27	12.73	12.61	12.03	12.44

Running times for a Rand64 problem.

m	4	8	16	32	64	128	256	512
FS	25.47	13.87	8.32	6.80	6.64	8.48	8.55	8.59
FFS	25.12	13.75	8.32	6.84	6.68	8.63	8.91	8.95
TSW	45.74	28.09	17.54	10.94	8.05	7.19	4.50	2.97
SSABS	26.60	15.74	9.69	7.31	6.88	8.36	8.48	8.48
TVSBS	28.83	17.85	10.90	7.11	6.25	9.65	5.66	3.40
3-HASH	99.61	33.95	15.20	8.13	6.17	8.25	6.80	6.02
5-HASH	–	53.90	18.67	9.14	6.21	8.28	6.96	6.05
8-HASH	–	–	33.59	13.05	7.19	8.28	7.15	6.13
BOM	29.65	18.28	12.73	9.53	6.25	6.09	3.13	1.99
EBOM	18.44	8.99	5.71	5.00	5.31	3.36	1.91	1.25
FBOM	34.22	19.61	11.37	7.50	6.57	10.35	5.71	3.36
SBNDM	26.84	12.27	6.75	5.35	5.27	5.35	5.35	5.28
FSBNDM	22.11	11.72	6.95	5.35	5.24	5.27	5.24	5.23
BLIM	35.59	19.92	12.85	9.22	8.87	10.31	12.66	17.50
FAOSO	37.97	19.59	9.83	8.03	7.88	7.54	8.14	9.19

Running times for a Rand128 problem.

m	4	8	16	32	64	128	256	512
FS	24.22	13.01	7.70	6.41	6.37	10.00	9.38	8.90
FFS	24.14	12.85	7.74	6.52	6.52	10.20	9.69	9.22
TSW	46.05	28.20	18.09	11.72	8.36	7.26	4.22	2.66
SSABS	25.16	14.64	8.83	6.60	6.37	9.49	9.30	8.79
TVSBS	41.80	25.82	15.27	9.34	6.41	9.92	5.66	3.28
3-HASH	99.69	33.83	15.23	8.05	6.09	8.28	6.80	5.94
5-HASH	–	54.41	18.83	9.14	6.13	8.12	6.88	6.09
8-HASH	–	–	33.94	13.09	7.11	8.32	7.11	6.21
BOM2	26.53	15.23	9.84	7.50	6.21	6.05	3.36	1.95
EBOM	23.40	11.21	6.68	5.36	5.36	3.48	2.07	1.45
FBOM	50.63	28.79	16.29	9.72	6.84	10.62	5.82	3.52
SBNDM	26.64	12.19	6.68	5.31	5.24	5.24	5.23	5.20
FSBNDM	21.21	10.98	6.52	5.27	5.08	5.12	5.12	5.12
BLIM	29.96	16.29	10.04	7.77	8.04	10.51	12.70	17.19
FAOSO	37.81	19.47	8.70	6.99	6.19	5.48	5.44	5.29

Experimental results show that the Extended-BOM and the Forward-BOM algorithms obtain the best run-time performance in most cases. In particular for small alphabets and short patterns the presented variations are second to algorithms in the q -Hash family. Moreover for large alphabets and short patterns algorithms based on bit-parallelism are the best choice. Note however that for alphabets of

medium dimension, when the pattern is short, the performance of the Extended-BOM algorithm outperform those of bit-parallel algorithms that until now have been considered the best choice for short patterns.

5.2. Running times for real world problems

The tests on real world problems have been performed on a genome sequence and on a natural language text buffer. A genome is a DNA sequence composed of the four nucleotides, also called base pairs or bases: Adenine, Cytosine, Guanine and Thymine. The genome we used for these tests is a sequence of 4,638,690 base pairs of *Escherichia coli*. We used the file E.coli of the Large Canterbury Corpus (<http://www.data-compression.info/Corpora/CanterburyCorpus/>).

The tests on the protein sequence has been performed using a 2.4Mb file containing a protein sequence from the human genome with 22 different characters.

For the experiments on the natural language text buffer we used the file world192.txt (The CIA World Fact Book) of the Large Canterbury Corpus. The alphabet is composed of 94 different characters. The text is composed of 2,473,400 characters.

Running times for a genome sequence ($\sigma = 4$).

m	4	8	16	32	64	128	256	512
FS	18.52	13.32	11.06	9.49	8.09	7.23	6.33	5.67
FFS	16.68	11.80	8.75	6.56	4.96	4.06	3.60	3.36
TSW	21.29	13.98	9.34	6.88	5.98	5.82	5.94	5.78
SSABS	22.46	18.56	17.70	17.39	17.38	17.62	17.58	17.54
TVSBS	19.49	13.40	9.26	7.30	6.68	6.52	6.68	6.52
3-HASH	23.59	8.32	3.91	2.38	1.79	1.68	1.60	1.56
5-HASH	–	12.03	4.22	1.96	1.49	1.60	1.56	1.25
8-HASH	–	–	7.54	3.00	1.64	1.79	1.60	1.45
BOM	23.64	13.67	8.04	4.57	2.62	1.92	1.60	1.33
EBOM	13.32	10.12	6.53	3.87	2.34	1.91	1.41	1.09
FBOM	19.02	11.49	6.95	4.07	2.50	2.11	1.60	1.10
SBNDM	13.32	9.03	5.35	2.97	3.05	2.93	3.04	3.04
FSBNDM	17.46	9.50	5.20	2.97	2.89	2.97	2.97	2.85
BLIM	19.06	10.62	7.34	5.90	5.08	6.21	8.48	13.01
FAOSO	17.93	9.48	5.62	4.48	3.51	3.03	3.56	3.80

Running times for a protein sequence ($\sigma = 22$).

m	4	8	16	32	64	128	256	512
FS	4.53	2.54	1.64	1.48	1.29	1.29	1.29	1.25
FFS	4.45	2.58	1.68	1.41	1.33	1.29	1.49	1.72
TSW	6.49	4.06	2.58	1.76	1.45	1.56	1.33	0.90
SSABS	4.92	3.12	2.23	1.76	1.60	1.44	1.53	1.52
TVSBS	5.12	3.08	1.92	1.25	0.90	1.29	0.86	0.70
3-HASH	13.95	4.84	2.15	1.29	0.86	0.98	0.90	0.86
5-HASH	–	7.54	2.69	1.25	0.90	1.21	0.94	0.78
8-HASH	–	–	4.69	1.83	1.02	1.21	1.02	0.82
BOM	6.83	4.65	2.73	1.37	0.97	1.05	0.74	0.55
EBOM	2.85	1.45	1.09	0.98	0.98	0.82	0.58	0.55
FBOM	5.51	3.09	1.88	1.33	1.09	1.60	0.67	0.47
SBNDM	3.95	1.95	1.09	0.90	0.86	0.90	0.90	0.82
FSBNDM	3.98	2.15	1.25	0.94	0.82	0.90	0.78	0.90
BLIM	6.25	3.87	2.46	2.00	2.46	3.55	5.86	10.43
FAOSO	4.58	2.60	1.90	1.40	1.22	1.13	1.17	1.06

Running times for a natural language text buffer ($\sigma = 93$).

m	4	8	16	32	64	128	256	512
FS	3.79	2.23	1.36	1.09	0.94	0.98	0.94	0.67
FFS	3.75	2.11	1.33	1.02	0.94	1.17	1.21	1.21
TSW	5.98	3.75	2.34	1.56	1.29	1.37	0.98	0.66
SSABS	4.22	2.58	1.76	1.25	1.13	1.02	0.97	0.70
TVSBS	4.25	2.81	1.76	1.06	0.86	1.05	0.70	0.39
3-HASH	11.60	3.98	1.88	1.02	0.78	0.90	0.70	0.47
5-HASH	–	6.25	2.23	1.13	0.70	0.86	0.70	0.55
8-HASH	–	–	3.91	1.52	0.86	0.98	0.74	0.47
BOM	5.59	3.79	2.46	1.41	0.90	0.90	0.55	0.51
EBOM	2.66	1.49	1.05	0.94	0.94	0.78	0.47	0.51
FBOM	4.76	2.89	1.83	1.21	0.94	1.22	0.74	0.58
SBNDM	3.48	1.76	1.09	0.82	0.86	0.82	0.86	0.90
FSBNDM	3.48	1.91	1.13	0.90	0.86	0.86	0.90	0.86
BLIM	5.90	3.40	2.23	1.95	2.27	3.32	5.62	10.28
FAOSO	4.58	–	–	–	–	–	–	–

From experimental results it turns out that the Extended-BOM algorithm obtains in most cases the best results and sporadically is second to algorithms of the q -Hash family. Again better results are obtained for medium dimensions of the alphabet.

6. Conclusion

We presented two efficient variants of the Backward Oracle Matching algorithm which is considered one of the most effective algorithm for exact string matching. The first variation, called Extended-BOM, introduces an efficient fast-loop over transitions of the oracle by reading two consecutive characters for each iteration. The second variation, called Forward-BOM, extends the previous one by using a look-ahead character at the beginning of transitions in order to obtain larger shift advancements.

It turns out from experimental results that the new proposed variations are very fast in practice and obtain the best results in most cases, especially for long patterns and alphabets of medium dimension.

References

- [1] C. Allauzen, M. Crochemore, and M. Raffinot. Factor oracle: a new structure for pattern matching. In J. Pavelka, G. Tel, and M. Bartosek, editors, *SOFSEM'99, Theory and Practice of Informatics*, number 1725 in Lecture Notes in Computer Science, pages 291–306, Milovy, Czech Republic, 1999. Springer-Verlag, Berlin.
- [2] R. Baeza-Yates and G. H. Gonnet. A new approach to text searching. *Commun. ACM*, 35(10):74–82, 1992.
- [3] A. Blumer, J. Blumer, A. Ehrenfeucht, D. Haussler, M. T. Chen, and J. Seiferas. The smallest automaton recognizing the subwords of a text. *Theor. Comput. Sci.*, 40(1):31–55, 1985.
- [4] A. Blumer, J. Blumer, A. Ehrenfeucht, D. Haussler, and R. McConnel. Linear size finite automata for the set of all subwords of a word: an outline of results. *Bull. Eur. Assoc. Theor. Comput. Sci.*, 21:12–20, 1983.
- [5] R. S. Boyer and J. S. Moore. A fast string searching algorithm. *Commun. ACM*, 20(10):762–772, 1977.

- [6] D. Cantone and S. Faro. **Fast-Search**: a new efficient variant of the Boyer-Moore string matching algorithm. *WEA 2003*, LNCS 2647(4/5):247–267, 2003.
- [7] D. Cantone and S. Faro. Fast-Search Algorithms: New Efficient Variants of the Boyer-Moore Pattern-Matching Algorithm. *J. Autom. Lang. Comb.*, 10(5/6):589–608, 2005.
- [8] C. Charras and T. Lecroq. *Handbook of exact string matching algorithms*. King’s College Publications, 2004.
- [9] M. Crochemore. Optimal factor transducers. In A. Apostolico and Z. Galil, editors, *Combinatorial Algorithms on Words*, volume 12 of *NATO Advanced Science Institutes, Series F*, pages 31–44. Springer-Verlag, Berlin, 1985.
- [10] M. Crochemore, A. Czumaj, L. Gąsieniec, S. Jarominek, T. Lecroq, W. Plandowski, and W. Rytter. Speeding up two string matching algorithms. *Algorithmica*, 12(4/5):247–267, 1994.
- [11] M. Crochemore and W. Rytter. *Text algorithms*. Oxford University Press, 1994.
- [12] K. Fredriksson and S. Grabowski. Practical and optimal string matching. In M. P. Consens and G. Navarro, editors, *SPIRE*, volume 3772 of *Lecture Notes in Computer Science*, pages 376–387. Springer, 2005.
- [13] J. Holub and B. Durian. Talk: Fast variants of bit parallel approach to suffix automata. In *The Second Haifa Annual International Stringology Research Workshop of the Israeli Science Foundation*, <http://www.cri.haifa.ac.il/events/2005/string/presentations/Holub.pdf>, 2005.
- [14] A. Hudaib, R. Al-Khalid, D. Suleiman, M. Itriq, and A. Al-Anani. A fast pattern matching algorithm with two sliding windows (TSW). *Journal of Computer Science*, 4(5):393–401, 2008.
- [15] A. Hume and D. M. Sunday. Fast string searching. *Softw. Pract. Exp.*, 21(11):1221–1248, 1991.
- [16] D. E. Knuth, J. H. Morris, Jr, and V. R. Pratt. Fast pattern matching in strings. *SIAM J. Comput.*, 6(1):323–350, 1977.
- [17] M. Oguzhan Kulekci. A method to overcome computer word size limitation in bit-parallel pattern matching. In *Proceedings of the 19th International Symposium on Algorithms and Computation, ISAAC 2008*, pages 497–507, 2008.
- [18] T. Lecroq. Fast exact string matching algorithms. *Inf. Process. Lett.*, 102(6):229–235, 2007.
- [19] G. Navarro and M. Raffinot. A bit-parallel approach to suffix automata: Fast extended string matching. In M. Farach-Colton, editor, *Proceedings of the 9th Annual Symposium on Combinatorial Pattern Matching*, number 1448, pages 14–33, Piscataway, NJ, 1998. Springer-Verlag, Berlin.
- [20] S.S. Sheik, S.K. Aggarwal, A. Poddar, N. Balakrishnan, and K. Sekar. A fast pattern matching algorithm. *J. Chem. Inf. Comput.*, 44:1251–1256, 2004.
- [21] D. M. Sunday. A very fast substring search algorithm. *Commun. ACM*, 33(8):132–142, 1990.
- [22] R. Thathoo, A. Virmani, S. Sai Lakshmi, N. Balakrishnan, and K. Sekar. TVSBS: A fast exact pattern matching algorithm for biological sequences. *Journal of the Indian Academy of Science, Current Science*, 91(1):47–53, 2006.
- [23] S. Wu and U. Manber. A fast algorithm for multi-pattern searching. Report TR-94-17, Department of Computer Science, University of Arizona, Tucson, AZ, 1994.
- [24] A. C. Yao. The complexity of pattern matching for a random string. *SIAM J. Comput.*, 8(3):368–387, 1979.