

A General Practical Approach to Pattern Matching over Ziv-Lempel Compressed Text

Gonzalo Navarro*

Mathieu Raffinot[†]

Abstract

We address in this paper the problem of string matching on Lempel-Ziv compressed text. The goal is to search a pattern in a text without uncompressing. This is a highly relevant issue, since it is essential to have compressed text databases where efficient searching is still possible. We develop a general technique for string matching when the text comes as a sequence of blocks. This abstracts the essential features of Lempel-Ziv compression. We then apply the scheme to each particular type of compression. We present the *first* algorithm to find all the matches of a pattern in a text compressed using LZ77. When we apply our scheme to LZ78, we obtain a much more efficient search algorithm, which is faster than uncompressing the text and then searching on it. Finally, we propose a new hybrid compression scheme which is between LZ77 and LZ78, being in practice as good to compress as LZ77 and as fast to search in as LZ78. We show also how to search some extended patterns on Lempel-Ziv compressed text, such as classes of characters and approximate string matching.

1 Introduction

String matching is one of the most pervasive problems in computer science, with applications in virtually every area. It is also one of the oldest and richest area of development. The *string matching problem* is: given a pattern $P = p_1 \dots p_m$ and a text $T = t_1 \dots t_u$, both sequences of symbols over a finite alphabet Σ of size σ , find all the occurrences of P in T . There are many algorithms to solve this problem, from classical to very recent [18, 8, 4, 14, 26, 9, 25]. The complexity of this problem is $O(u)$ in the worst case and $O(u \log(m)/m)$ on average, where $u = |T|$ and $m = |P|$, and there exist variants of [8, 9] which achieve this complexity. In practice, however, [26, 25] are the fastest algorithms in most cases.

Another old and rich area in computer science is text compression. Its aim is to exploit the redundancies of the text to reduce its space usage. There are many different compression schemes, among which the Ziv-Lempel family [31, 32] is one of the best in practice because of their good compression ratios combined with efficient compression and decompression times. Other compression schemes are Huffman coding [15] and arithmetic coding [28], among others.

Today's textual databases are an excellent example of applications where both problems are crucial: the texts should be kept compressed to save space and I/O time, and they should be efficiently searched. Surprisingly, these two combined requirements are not easy to achieve together, as the only solution before the 90's was to process queries by uncompressing the texts and then searching into them.

*Dept. of Computer Science, University of Chile. Blanco Encalada 2120, Santiago, Chile. gnavarro@dcc.uchile.cl

[†]Institut Gaspard Monge, Cité Descartes, Champs-sur-Marne, 77454 Marne-la-Vallée Cedex 2, France. raffinot@monge.univ-mlv.fr

The *compressed matching problem* was first defined in the work of Amir and Benson [1] as the task of performing string matching in a compressed text without decompressing it. Given a text T , a corresponding compressed string $Z = z_1 \dots z_n$, and a pattern P , the compressed matching problem consists in finding all occurrences of P in T , using only P and Z . A naive algorithm, which first decompresses the string Z and then performs standard string matching, takes time $O(u + m)$. An optimal algorithm takes worst-case time $O(n + m)$, where $n = |Z|$. In [2], a new criterion, called *extra space*, for evaluating compressed matching algorithms, was introduced. According to the extra space criterion, algorithms should use at most $O(n)$ extra space, optimally $O(m)$ in addition to the n -length compressed file.

We define now a variation, where we are required to report all the text positions which match the pattern. That is, given P and Z , report all the $|x|$ such that $T = xPy$. The optimal algorithm for this problem takes $O(n + R)$ time, where R is the number of matches (we disregard pattern preprocessing costs).

Two different approaches have emerged in the last years to combine compression and searching in textual databases. A first one is strongly oriented to natural language texts, which are assumed to be composed of words which follow some statistical rules. The basic idea is to compress the text using Huffman, where the words instead of the characters are taken as the symbols [7, 21]. As Huffman assigns a fixed code to each symbol, searching a given string is a matter of compressing it and searching it in the compressed text using a classical string matching algorithm with minor modifications [23, 22]. Despite its simplicity, this approach is very effective on natural language text, with better compression ratios than those of the Ziv-Lempel family, and search times which are between 2 and 8 times faster than the fastest algorithms for standard string matching over the uncompressed text. They are also able to search for complex patterns (such as regular expressions) and allow errors in the search, provided that words are matched against words. A weakness of this scheme is that it does not work well on small texts (say, less than 10 Mb), since in that case the vocabulary is almost as big as the text itself.

Another practical approach is an ad-hoc technique [19], which however is not so fast, obtains compression ratios of near 70% (against 30% to 40% of Ziv-Lempel algorithms) and relies on the ASCII encoding.

The second line of research considers Ziv-Lempel compression, which is based on finding repetitions in the text and replacing them with references to similar strings previously appeared. LZ77 is able to reference any substring of the text already processed, while LZ78 references only a single previous reference plus a new letter that is added. In both cases, the referenced text is to be found is normally limited by a *window* which precedes the current text position.

String matching in Ziv-Lempel compressed texts is much more complex, since the pattern can appear in different forms across the compressed text. In [2] a compressed matching algorithm for LZ78 is presented, which works in time $O(n + m^2)$ and space $O(n + m^2)$. For LZ77, the only result is [11], which is a randomized algorithm to determine in time $O(n \log^2(u/n) + m)$ whether a pattern is present or not in an LZ77-compressed text, but they do not find all the pattern occurrences. Other algorithms for different specific search problems have been presented in [13, 17]. This second branch is rather theoretical and, to the best of our knowledge, no actual implementations have been developed.

In this paper we aim at efficient algorithms for string matching on Ziv-Lempel compressed texts.

We present new theoretical developments but also give practical implementations and experiments on our algorithms. The main results of this work are

- We develop a general technique for string matching on a text which is given as a sequence of blocks. This abstracts the essential features of Ziv-Lempel compressed texts and is the basis for the algorithms which run over specific members of the family.
- We apply our technique for LZ77-compressed texts. The result is the *first* algorithm to search under this compression scheme (recall that [11] cannot find all the occurrences of the pattern). The algorithm, however, is $O(u)$ time at best. In practice, the algorithm is slower than uncompressing the text and searching it with a classical algorithm.
- We apply the technique to the LZ78 compression scheme. The result is an algorithm which turns out to be a practical implementation of the theoretical proposal of [2]. This algorithm is $O(n+R)$ time in the worst and average case, and is in practice twice as fast as decompressing and searching.
- We propose a hybrid compression scheme which is between LZ77 and LZ78, which keeps some of the good features of LZ77 and which can be searched in $O(\min(u, n \log m) + R)$ time on average (and $O(\min(u, mn) + R)$ in the worst case). In practice, the compression efficiency is similar to LZ77 and the search time is similar to LZ78.
- We show how to search some extended patterns in a sequence of blocks, such as how to allow classes of characters or approximate string matching, the last one being an open problem advocated in [2].

Our approach is oriented to practical text searching and relies on *bit-parallelism*. Bit-parallelism is a general technique to take advantage of the fact that the computer operates in parallel over all the bits of the machine word, so that if a process is so simple that it can be expressed with bit operations we can perform many of those steps in a single operation of the processor. If we call w the length in bits of the machine word (typically 32 or 64), then the possible speedups are up to $O(w)$. The complexity results presented assume that $m = O(w)$, otherwise we have to multiply the u and n of our complexities by m/w .

2 String Matching on Blocks

We describe in this section a general technique to perform exact and approximate string matching when the text is presented as a sequence of atomic strings (here called “blocks”) instead of a sequence of characters. This technique is the basis for all the different searching algorithms on Lempel-Ziv compressed text, which are described in the next sections.

Our general assumption is that the blocks either have just one letter (and we can know which is that letter) or are formed by a concatenation of previously seen blocks. We describe an online algorithm where we process the text block by block. At any moment of the search we denote T' the text already processed (and its length in characters is $|T'|$). When we finish the search, $T' = T$, i.e. the original text.

The method works as follows. We process the blocks one by one. For each new block B , we compute a *description* for B which has all the information of the block which is relevant for the search. This description is denoted $D(B) = (L, O, S, P, M)$, where

- $L = |B|$, that is, the length of B in characters
- $O = \text{Offs}(B)$ = the length in characters of the text we had processed when B appeared
- $S = \text{Suff}(B)$ = all the pattern positions¹ which either start a complete occurrence of B inside the pattern, or start a proper pattern suffix which matches with a prefix of B . Formally,

$$\text{Suff}(B) = \{|x|, P = xBy\} \cup \{|x|, |x| > 0 \wedge |z| > 0 \wedge P = xz \wedge B = zy\}$$

- $P = \text{Pref}(B)$ = all the pattern positions which either follow a complete occurrence of B inside the pattern, or follow a proper pattern prefix which matches with a suffix of B . Formally,

$$\text{Pref}(B) = \{|xB|, P = xBy \wedge |y| > 0\} \cup \{|z|, |z| > 0 \wedge |y| > 0 \wedge P = zy \wedge B = xz\}$$

- $M = \text{Matches}(B)$ = all the block positions where the pattern occurs (\emptyset if $|B| < |P|$). Formally,

$$\text{Matches}(B) = \{|x|, B = xPy\}$$

Figure 1 illustrates these concepts.




Figure 1: Prefixes (P) and suffixes (S) for a long and a short block. The pattern has the diagonal tiling and the possible blocks have a bar tiling. The suffixes (dotted lines) and prefixes (dashed lines) are pattern positions. Prefixes are marked after the position where they finish, suffixes are marked at the position they start.

The description $D(B)$ of a new block B is obtained in two forms: (a) the block is an explicit letter and then we obtain the description directly, or (b) the block is a concatenation of other blocks previously known, and we obtain its description by operating on the descriptions of the previous blocks.

Once the description of the new block is computed, we use that description to update the state of the search. This concludes the processing of the block and we move to the next one. The state of the search contains the matches that have already occurred and the potential matches in progress, that is,

¹To simplify the notation, we number pattern positions starting at zero.

- $\text{Res}(T') =$ the text positions that matched up to now, formally

$$\text{Res}(T') = \{|x|, T' = xPy\}$$

- $\text{Active}(T') =$ the set of positions following the pattern prefixes which match a suffix of the current text. Formally,

$$\text{Active}(T') = \{|x|, |x| > 0 \wedge |y| > 0 \wedge P = xy \wedge T' = zx\}$$

Hence, when we complete the text processing and T' is not a text prefix anymore but the whole text, $\text{Res}(T)$ is our answer. The initial state of the search is $\text{Res}(\epsilon) = \text{Active}(\epsilon) = \emptyset$, and $T' = \epsilon$.

We have defined already the information we keep, and consider now how to compute that information. For the formulas that follow, we define some auxiliary functions, namely

- Left_i , which receives a set of $\text{Suff}()$ positions not smaller than i , subtracts i to all them and then adds new pattern positions filling the hole left by the shift. Formally,

$$\text{Left}_i(X) = \{x - i, x \in X\} \cup \{m - i, m - i + 1, \dots, m - 1\}$$

- Right_i , which does the same for $\text{Pref}()$ positions, in the other direction. Formally,

$$\text{Right}_i(X) = \{x + i, x \in X\} \cup \{1, 2, \dots, i\}$$

- $\text{Add}_i(X) = \{i + x, x \in X\}$, which adds i to all the elements of the set.
- $\text{Subtr}_i(X) = \{i - x, x \in X\}$, which subtracts all the elements of the set from i .

2.1 Description of a Letter

The base case of our scheme is to obtain the description of a block which is a letter a . We have

- $|B| = 1$
- $\text{Offs}(B) = |T'|$
- $\text{Suff}(B) = \{|x|, P = xay\}$
- $\text{Pref}(B) = \{|xa|, P = xay \wedge |y| > 0\}$
- $\text{Matches}(B) = \text{if } P = a \text{ then } \{0\} \text{ else } \emptyset$

which is a direct application of the general formulas.

2.2 Concatenating Two Blocks

Assume that our block B is defined as the concatenation of one or more previous blocks. If only one previous block B' is referenced, we just copy its definition. We show now how to concatenate two blocks, since the case of more than two blocks is a simple iteration over this procedure. We are given two blocks B_1 and B_2 , and we have to obtain the description for their concatenation $D(B) = D(B_1B_2) = D(B_1) \cdot D(B_2)$ (where we define \cdot as the concatenation of block descriptions). The formulas are as follows

- $|B| = |B_1| + |B_2|$
- $\text{Offs}(B) = |T'|$
- $\text{Suff}(B) = \text{Suff}(B_1) \cap \text{Left}_{|B_1|}(\text{Suff}(B_2))$
- $\text{Pref}(B) = \text{Pref}(B_2) \cap \text{Right}_{|B_2|}(\text{Pref}(B_1))$
- $\text{Matches}(B) = \text{Matches}(B_1) \cup \text{Add}_{|B_1|}(\text{Matches}(B_2))$
 $\cup (\text{Subtr}_{|B_1|}(\text{Pref}(B_1) \cap \text{Suff}(B_2)) \cap \{0, 1, 2, \dots, |B| - m\})$

We explain now the rationale for the formulas (see Figure 2). The first two are immediate. We explain $\text{Suff}(B)$. If the number i belongs to $\text{Suff}(B_1B_2)$ then either

- $i \geq m - |B_1|$, that is, a prefix of B_1B_2 is a suffix of P . Notice that in this case also a prefix of B_1 is a suffix of P . Since $\text{Left}_{|B_1|}$ will add all these positions, they will appear in the result if and only if they are present in $\text{Suff}(B_1)$, which is correct.
- $i < m - |B_1|$, that is, B_1 appears inside P and is immediately followed by an occurrence of B_2 (which can be a complete occurrence or share a prefix with the pattern suffix). If we subtract $|B_1|$ to the elements in $\text{Suff}(B_2)$, then we are interested in the positions which also appear in $\text{Suff}(B_1)$ (which since $i < m - |B_1|$ can only correspond to complete occurrences of B_1).




Figure 2: Suffixes of the concatenation of two blocks. It is possible that the result involves only B_1 (rightmost pair) or that it involves both. In this case B_1 is completely inside the pattern and B_2 may or may not be totally inside (leftmost and middle pairs, respectively).

The rationale for $\text{Pref}()$ is analogous to $\text{Suff}()$. For $\text{Matches}(B)$, there are three parts. The first one is the matches which are inside B_1 , and the second one is the same for B_2 (displaced since

now B_2 comes after B_1 in B). The third one accounts for matches that appear only when B_1 and B_2 are concatenated. If a prefix of the pattern is at the end of B_1 , and the corresponding suffix is at the beginning of B_2 , then we have the pattern in B_1B_2 . The *Subtr* converts pattern to block positions and the final set which is intersected with the results ensures that we have really prefixes and suffixes instead of substrings of the blocks.

2.3 Updating the Search State

We want now to update the state of our search by processing a new block B whose description has just been computed. The formulas to obtain the new $\text{Res}(T'B)$ and $\text{Active}(T'B)$ values from the old $\text{Res}(T')$ and $\text{Active}(T')$ ones are

- $\text{Active}(T'B) = \text{Right}_{|B|}(\text{Active}(T')) \cap \text{Pref}(B)$
- $\text{Res}(T'B) = \text{Res}(T') \cup \text{Add}_{|T'|}(\text{Matches}(B))$
 $\cup \text{Subtr}_{|T'|}(\text{Active}(T') \cap \text{Suff}(B) \cap \{m - |B|, m - |B| + 1, \dots, m - 1\})$

The new $\text{Active}(T'B)$ value considers that, since a new block B has been added to T' , the pattern prefixes that are suffixes of $T'B$ are those that are already suffixes of B (i.e. $\text{Pref}(B)$), or those which are suffixes of T' and are followed by B in the pattern. As before, *Right* does the trick of considering both cases in a single formula.

The new value $\text{Res}(T'B)$ adds to $\text{Res}(T')$ not only the matches which are completely inside B , but also those which appear when T' is concatenated to B . For this sake, we consider pattern prefixes which are suffixes of T' (i.e. $\text{Active}(T')$), and which are followed by the corresponding pattern suffix in B . The final intersection ensures that the complete pattern has appeared. Figure 3 illustrates.




Figure 3: Updating the state of the search. In the first case we illustrate the updating of $\text{Active}(T')$ (a short block is added). In the second case we show how the matches are updated (when a long block is added). In general both updates are necessary.

2.4 Extended Patterns

We show now how to handle some extended patterns in this paradigm. A first alternative, which is easily implemented using bit-parallelism (see next section) is to allow *classes of characters*, i.e. the pattern at position i matches not just one letter but a set of letters. The pattern can then be seen as a sequence of sets of characters $P = c_1 \dots c_m$, $c_i \subseteq \Sigma$. This is easily implemented by modifying our equations for a single letter a , so that instead of $P = xay$ we require $P = xay \wedge a \in A$.

Approximate searching can also be performed in this scenario. In this case we allow that the pattern does not match exactly but up to k errors. An alternative definition is that we want all the text segments t such that $dist(t, P) \leq k$, where $dist(a, b)$ gives the minimum number of operations (errors) to perform on P or t to transform one into the other. There are many choices to define what is an error, but the most commonly used are: allow substitutions (Hamming distance) or allow substitutions, insertions, and deletions (Levenshtein distance). If we want to search all the pattern occurrences with up to $0 < k < m$ substitution errors, we keep for each block and each $i \in 0..k$ a description $D_i(B)$. It represents all the block information when the matches are allowed to contain up to i errors. There is a different search state for each i (i.e. $Active_i(T')$), representing that a pattern prefix matches a text suffix with up to i errors. $Res(T')$ is the same for any i , and keeps track of the matches allowing k errors.

We write $Pref_i(B)$ and $Suff_i(B)$ when we refer to $D_i(B)$, while the other components are independent on i . $Matches(B)$ refers to the matches allowing up to k errors which occurred completely inside the block.

Given $D(B) = (L, O, S, P, M)$ and $D(B') = (L, O, S', P', M)$ we define $D(B) \cup D(B') = (L, O, S \cup S', P \cup P', M)$ as their union. With this notation we can express the concatenation of two blocks $D_i(B) = D_i(B_1 B_2)$, allowing i errors:

$$D_i(B) = \bigcup_{j=0}^i D_j(B_1) \cdot D_{i-j}(B_2)$$

(where we recall that \cdot represents the concatenation of descriptions). The reason for this formula is as follows: imagine that we search with $k = 2$ errors. Then, we can pair a prefix that matched with zero errors with a suffix that matched with two errors, or a prefix that matched with one error with a suffix that matched with one error, or a prefix that matched with two errors with a suffix that matched with zero errors. In general, the sum of errors between prefix and suffix must be k . This is easily generalized if we are interested in $i \leq k$ errors.

To update the $Active_i(T')$ values we use a similar idea, i.e.

$$Active_i(T'B) = \bigcup_{j=0}^i Right_{|B|}(Active_j(T')) \cap Pref_{i-j}(B)$$

where the rationale is the same as before: we match with i errors if we already matched a pattern prefix with j errors and the block starts with the corresponding pattern suffix matched with $i - j$ errors. Figure 4 illustrates.




Figure 4: Updating $Active(T')$ when 2 errors are allowed.

The $\text{Res}(T')$ value is interested only in k errors:

$$\begin{aligned} \text{Res}(T'B) &= \text{Res}(T') \cup \text{Add}_{|T'|}(\text{Matches}(B)) \\ &\quad \cup \text{Subtr}_{|T'|}(\text{Active}_k(T') \cap \text{Suff}_k(B) \cap \{m - |B|, m - |B| + 1, \dots, m - 1\}) \end{aligned}$$

while however the other $\text{Active}_i(T')$ values are necessary to maintain $\text{Active}_k(T')$.

The values for an individual letter a is also modified:

- $\text{Suff}_i(B) = \{0, 1, 2, \dots, m - 1\}$
- $\text{Pref}_i(B) = \{1, 2, 3, \dots, m - 1\}$
- $\text{Matches}(B) = \text{if } (m = k + 1 \wedge P = xay) \text{ then } \{0\} \text{ else } \emptyset$

Notice that if $k > 0$ (i.e. our case of interest), then a single letter matches at any pattern position. On the other hand, the pattern matches inside the letter only if we can delete all its letters and leave a single one which is equal to a (the case of deleting all the letters is not considered because it implies $m = k$, which is a trivial problem).

The case of the Levenshtein distance is more complicated, because the matches may not have the same length. In this case we need to store, instead of initial and final positions of matches ($\text{Pref}_i(B)$ and $\text{Suff}_i(B)$), the segments of the pattern that match with i errors or less. We define

$$\text{Segm}_i(B) = \{(i, j), \text{dist}(P_{i..j-1}, B) \leq k\}$$

instead of $\text{Pref}_i(B)$ and $\text{Suff}_i(B)$. For block concatenation, instead of intersecting prefixes of B_1 with suffixes of B_2 , we consider the segments of B_1 immediately followed by segments of B_2 and take their concatenation. We do not work out the details because, as explained in the next section, we do not have an efficient implementation (the naive implementation is $O(m^2k^2n)$ if we perform $O(n)$ block concatenations).

3 A Bit-Parallel Implementation

Until now, we have defined our algorithms in terms of sets of pattern positions. We present now a very well-suited implementation paradigm which allows to convert the previous algorithms into efficient implementations.

We use the technique called *bit-parallelism* [3]. This technique takes advantage of the fact that the processor works in parallel on all the bits of the computer word. We call w the number of bits of the computer word, which is 32 or 64 in current architectures. If one is able to map the elements of a set on bits, and to express the operations to perform on them by using only the operators provided by the processor (which are rather limited, i.e. bit shifts, masking, etc.), then one can effectively parallelize the work on the set, obtaining speedups of up to $O(w)$ over the original algorithm.

This paradigm was invented in 1989 by Baeza-Yates and Gonnet [4] for a text searching algorithm called Shift-Or. If we consider $m \leq w$, then we keep the state of the search in a computer word D , whose i -th bit tells whether the prefix of length i of the pattern matches the current text

suffix. All the bits start with value zero, and a match is reported whenever the m -th bit of D signals a match. The update formula upon reading a new text character is

$$D' = (D \ll 1) | T[a]$$

where $T[a]$ is a mask whose i -th bit tells whether $P_i = a$, we are assuming that 0 represents a match and a 1 a mismatch, “|” is the bitwise-or of the computer word, and “ $\ll \ell$ ” is a bit shift operation which assigns the i -th bit to the $(i + \ell)$ -th, setting the first ℓ to zero. Other operations allowed in most architectures are bitwise-and (&), shift to the other direction (\gg), and, which is more sophisticated, arithmetic operations such as addition and subtraction which operate on the bit mask as if it was a number.

The Shift-Or algorithm is $O(n)$ provided $m \leq w$. If the computer word is too short to hold one bit per pattern position, then $\lceil m/w \rceil$ computer words are used for the simulation, and the search takes in the worst case $O(mn/w)$ time. It is not hard to show that on average it takes $O(n)$, since $O(1)$ computer words have active states on average.

Our implementation can indeed be seen as a Shift-Or algorithm working on blocks instead of letters. The sets $\text{Pref}(B)$, $\text{Suff}(B)$, and $\text{Active}(T')$ are represented by bit masks. Hence, for blocks of one letter a we have $\text{Suff}(B) = T[a]$ and $\text{Pref}(B) = (T[a] \ll 1)$. The formulas to concatenate blocks are directly translated by noticing that Left_ℓ and Right_ℓ are converted into “ $\gg \ell$ ” and “ $\ll \ell$ ”, respectively (taking care of the borders which must get active bits), and union and intersection are converted into “|” and “&” respectively. Hence, all those operations on sets are performed in $O(1)$ time if $m \leq w$, and $O(m/w)$ time in general. In practical text searching we can assume $m = O(w)$.

On the other hand, the sets $\text{Res}(T')$ and $\text{Matches}(B)$ are explicitly stored in an array. However, it is not difficult to see that the total amount of work to handle them is $O(R)$, where R is the number of occurrences of the pattern in the text. The cost cannot be $o(R)$ if we report all the occurrences.

Hence, if $f(n)$ concatenations are performed along all the process, our total search cost is $O(f(n) + R)$. The value of $f(n)$ depends on the compression algorithm.

The bit-parallel paradigm allows to seamlessly expand the type of patterns we are able to search. Since the T table is the only connection between the pattern and the search, we can for instance allow having classes of characters, that is, each pattern position matches with a set of characters instead of just one character. To achieve this, just set the i -th bit of $T[a]$ to “match” for any $a \in P_i$. Other extended patterns considered in [29], such as regular expressions, are not easily adapted to this scheme. It is also possible to handle errors in the matches, such as replacement errors [4] (at $O(m \log(k)/w)$ cost per character) or insertions, deletions and replacements at $O(mk/w)$ [29, 5] or even $O(m/w)$ [24] cost per character. The implementation of our technique to handle mismatches is $O(k^2 f(n) + R)$ cost, while the extension for Levenshtein distance is not easily implemented.

In all cases, the space complexity of our algorithms is $O(n + R)$, since we need to store the descriptions of the blocks already seen and the matches found. Notice that this n refers in fact to the size of the compression window.

Finally, we consider the practical problem of uncompressing a neighborhood of the occurrences. In practice it is undesirable that we just give the text positions matching the pattern. It is much better to uncompress and show a neighborhood of the match. This neighborhood can be defined

as the line holding the occurrence, the record (delimited by some given pattern), a fixed number of characters, etc.

Assume that we know a pattern position and want to show a neighborhood. We just decompress the surrounding blocks forward and backward, until from the plain text obtained we determine that the neighborhood has been decompressed. To decompress a block we have two cases: (a) the block is a letter, in which case we deliver the letter, (b) the block is a concatenation of other blocks, in which case we decompress each of those blocks in turn. This process takes $O(N)$ time at most (where N is the size of the decompressed neighborhood), since at each step we either obtain one character of N or split the final text to be obtained, and it is not possible to split it more than $O(N)$ times. This shows that it is practical to show a part of a Ziv-Lempel compressed file without necessarily uncompressing the whole file.

4 LZ78 Compression

4.1 Compression Algorithm

The Ziv-Lempel compression algorithm of 1978 (usually named LZ78 [32]) is based on a dictionary of blocks, in which we add every new block computed. At the beginning of the compression, the dictionary contains a single block b_0 of length 0. The current step of the compression is as follows: if we assume that a prefix $t_1 \dots t_i$ of T has been already compressed in a sequence of blocks $Z = b_1 \dots b_c$, all them in the dictionary, then we look for the longest prefix of the rest of the text $t_{i+1} \dots t_u$ which is a block of the dictionary. Once we found this block, say b_k of length l_k , we construct a new block $b_{c+1} = (k, t_{i+l_k+1})$, we write the pair at the end of the compressed file Z , i.e. $Z = b_1 \dots b_c b_{c+1}$, and we add the block to the dictionary. It is easy to see that this dictionary is prefix-closed (i.e. any prefix of an element is also an element of the dictionary) and a natural way to represent it is a trie.

We give as an example the compression of the word *ananas* in Figure 5. The first block is $(0, a)$, and next $(0, n)$. When we read the next a , a is already the block 1 in the dictionary, but an is not in the dictionary. So we create a third block $(1, n)$. We then read the next a , a is already the block 1 in the dictionary, but as do not appear. So we create a new block $(1, s)$.




Figure 5: Compression of the word *ananas* with the algorithm LZ78.

The compression algorithm is $O(u)$ in the worst case and efficient in practice if the dictionary is stored as a trie, which allows rapid searching of the new text prefix (for each character of T we move once in the trie). The decompression needs to build the same dictionary (the pair that defines the block c is read at the c -th step of the algorithm), although this time it is not convenient to have a trie, and an array implementation is preferable. Compared to LZ77, the compression is rather fast but decompression is slow. LZ78 is used by Unix's *Compress* program.

Many variations on LZ78 exist, which deal basically with the best way to code the pairs in the compressed file, or with the best way to update the window. A particularly interesting variant is from Welch, called LZW [27]. In this case, the extra letter (second element of the pair) is not coded, but it is taken as the first letter of the next block (the dictionary is started with one block per letter). A variant over this is presented by Miller and Wegman [20] (which we call LZMW), where the new block is not the previous one plus the first letter of the new one, but simply the concatenation of the previous and the new one.

4.2 Pattern Matching in LZ78 Compressed Files

Our general algorithm for searching in a sequence of blocks $Z = b_1 \dots b_n$ can be directly applied if we consider the new letter added after each block created by the LZ78 compression algorithm as a separate block. That is, each new pair (k, a) read at step c is taken as a reference to a previous block (b_k) followed by a literal block (a) . Hence, we compute the description of the concatenation of b_k and a and add it as the new block b_c to our dictionary. At the same time, we update the state of the search using the description of b_c just computed. Of course, in practice we manage this one-letter block in a special way, to speed-up the block concatenation. We keep all the descriptions of the blocks b_k in an array which is directly accessed.

The algorithm we obtain is quite the same as in [2]. The main differences are that we obtain this algorithm as a particular case of a general string search algorithm for text that comes in blocks, that their algorithm is originally designed for LZW compression, and that we search all the occurrences of the pattern, not only the first one. Moreover, we present a practical implementation based on bit-parallelism, while [2] is a theoretical work that has not been implemented. To our knowledge ours is the first real implementation of this algorithm. It is quite easy to adapt our algorithm to work on other variants of LZ78, such as LZW or LZMW. In particular we can easily adapt to different window management policies. The simplest one is that when the compressor memory is full, the dictionary is deleted and compression is restarted. Others try to remove the least interesting blocks from the dictionary, e.g. [12]. Our searcher can follow the same steps of the compressor along the search, using the same amount of memory.

4.3 Analysis

The theoretical complexity of the pattern matching algorithm is $O(n + R)$ (recall that, as we use bit-parallelism, we have $O(mn/w + R)$ time for long patterns). If $n = o(u)$, this is faster than searching in the uncompressed text. In practical terms, the algorithm is rather efficient since no extra work apart from one block concatenation and one update of the search is performed per element of the compressed file.

Our experimental results, however (Section 7), show that the algorithm takes in practice twice

the time of a Shift-Or run on the uncompressed text. This is because Shift-Or is very simple, and although we process many characters of the uncompressed text in one shot, in practice the cost of each step is big enough to amortize any possible gain due to compression. A specific problem is the locality of reference: the compressed matching algorithm reads random positions in the array of block definitions, while the uncompressed algorithm works basically in-place. The caching mechanism of the computer largely favors this last approach.

However, there is a positive result. Searching the compressed file with this algorithm is twice as fast as decompressing it and then searching the uncompressed file. For this comparison we are assuming that the file is compressed with LZ77 (which is much faster than LZ78 to decompress) and consider the times of *gunzip*, which is an optimized decompression software. Hence, if the text collection is kept compressed (which is definitely of interest) then it is much faster to search directly the compressed files.

We have tried to further improve our algorithm. For instance, we have created a variant called Mark-LZ78. In this compression algorithm, we mark with a bit flag for each block if the block is a leaf of the dictionary trie or not, to avoid storing the block description if this block is not used anymore. However, as we show in the experiments, the performance does not improve.

5 LZ77 Compression

5.1 Compression Algorithm

The Ziv-Lempel compression algorithm of 1977 (usually named LZ77 [31]) is, in some sense, simpler than LZ78, since the basic idea is just to recognize two repeated segments of the text and to mark the second as a reference (position in the text and length of the repeated part) to the first one. More formally, assume that a prefix $t_1 \dots t_i$ of T has been already compressed in a sequence of block $Z = b_1 \dots b_c$. We look for the longest prefix v of $t_{i+1} \dots t_u$ which appears already in $t_1 \dots t_i t_{i+1} \dots t_{i+|v|-1}$. Once we have it, say that we find it v starting at position $j \leq i$, we add a new block $(j, |v|)$ to the compressed file. A special case occurs if v is empty, in which case t_{i+1} is a new letter and we code it with a special block $(0, t_{i+1})$. With the same example *ananas*, we obtained: $(0, a)$ *nanas*; $(0, a)(0, n)$ *anas*; $(0, a)(0, n)(1, 3)$ *s*; $(0, a)(0, n)(1, 3)(0, s)$.

Notice that the above definition allows that the referenced block overlaps the one which is being compressed. Another variant avoids this for simplicity, i.e. v must be found in $t_1 \dots t_i$. In this case the compression of *ananas* becomes: $(0, a)$ *nanas*; $(0, a)(0, n)$ *anas*; $(0, a)(0, n)(1, 2)$ *as*; $(0, a)(0, n)(1, 2)(1, 1)$ *s*; $(0, a)(0, n)(1, 2)(1, 1)(0, s)$.

Yet another variant codes the repeated block and then the letter which follows it in the still uncompressed text. There are many other variants as well, mainly related to how to represent the pairs in the compressed file and how to compress fast. In general, the position j is coded as the difference $i + 1 - j$, since the last occurrence of the block is used and v is normally restricted to not appear too far away from t_i .

LZ77 compresses more than LZ78, both in theory and in practice. From a theoretical point of view, the variant which allows overlaps can obtain a compressed file of $O(1)$ blocks in the best case, while the one not allowing overlaps obtains at most $O(\log u)$. LZ78, on the other case, cannot obtain less than $O(\sqrt{u})$. This is easily seen by considering the best-case file $T = a^u$. In practice it is also true that LZ77 compresses more than LZ78. LZ77 is implemented in the Gnu *gzip* program.

Compression is rather slow with LZ77. It is expensive in time and space memory to find the longest prefix of the uncompressed part of the file that appears already in the compressed part. In theory, the compression is $O(u)$ in space and in time by the use of a suffix tree or a DAWG automaton [31, 30]. In practice, the search is done in a buffer window and an large hash table is normally used, as in *gzip*. An experimental comparison of different techniques to find the prefix can be found in [6]. The decompression algorithm, on the other hand, is very fast (faster than for LZ78) because to decompress a block is it just necessary to copy a part of the text and no dictionary has to be kept.

5.2 Pattern Matching in LZ77 Compressed Files

Our algorithm for LZ77 is an adaptation of the general algorithm on blocks, with a main difference. On LZ77 compressed files, when we want to process a new block, the situation shown in Figure 6 generally occurs: the new block references a sequence of r contiguous previously processed blocks, but it overlaps with the first and last one (u and v in the Figure). That is, the new block does not exactly correspond to previously processed blocks. Therefore, we do not have all the information on the blocks u and v that we need to concatenate the blocks.

We solve this by computing recursively the descriptions of the two blocks u and v with the same method. That is, we simulate that we are back in the text, where those blocks appeared, and compute their description (this may trigger more recursive invocations with the same purpose). When we finally obtain the descriptions of u and v , we concatenate all the referenced blocks to obtain the description of the new block. Another possibility is that the new block is completely inside another block already processed, in which case we have to recursively consider the blocks that define the referenced block.




Figure 6: Recursive computation of the description of a block in LZ77 compressed files.

We explain now a technique to concatenate the r blocks in low average time. Instead of computing $\text{Pref}(B)$ and $\text{Suff}(B)$ of the first block, then concatenating with the second, then to the third, until the r blocks are concatenated, we compute $\text{Suff}(B)$ from the first block to the r -th and $\text{Pref}(B)$ from the r -th block to the first one. We analyze this shortly.

5.3 Analysis and Improvements

We analyze now the many aspects of our algorithm and propose some improvements.

Block concatenation. If we use the proposed block concatenation technique, we have that in the worst case only the first m blocks can affect $\text{Suff}(B)$ and only the last m blocks can affect

$\text{Pref}(B)$, so the worst case time for concatenating the blocks becomes $O(\min(u, mn))$.

We show now that on average only $O(\log m)$ blocks are processed until $\text{Suff}(B)$ becomes stable. Each new block character we process will either extend the current suffixes of the set $\text{Suff}(B)$ or make them disappear from the set. Each suffix is removed from the set with probability $1 - 1/\sigma$ (i.e. if the new block character block cannot extend it). Before we read the block characters all the m pattern positions are in $\text{Suff}(B)$, and therefore on average no pattern positions remain in the set after $O(\log m)$ block characters are read (after the i -th character is read, the pattern positions $m - i$ to $m - 1$ cannot be removed from the set, but their situation cannot change anyway).

Even if we consider all blocks of length 1 (worst case), we work $O(n \log m)$ because of concatenations. The same reasoning holds for $\text{Pref}(B)$.

The only part of the block concatenation which cannot skip blocks is the computation of $\text{Matches}(B)$. However, this adds up $O(R)$ along all the search. Therefore, the total time for block concatenation is $O(\min(u, n \log m) + R)$ on average.

Finding the blocks. We consider now how to find the indices of the block that define a text position j . We keep an array with the blocks already seen. Binary searching the text position among these blocks adds $O(n \log n)$ to the cost. Instead, we keep a table of $O(n)$ entries where the element i points to the block where the text position $\lfloor iu/n \rfloor$ is defined. By accessing this table we directly arrive at the correct block with an average inaccuracy of $O(u/n)$, and a final binary search finds the correct position, for a total cost of $O(n \log(u/n))$ (in practice a linear search is faster for the final part). This gives good results in practice. Another alternative is that the compressor does not store the text position and length of the repeated part, but instead it gives the block numbers involved and the offsets inside u and v . Since a text position needs $O(\log u)$ bits and a block number plus an offset inside the block needs on average $O(\log n + \log(u/n)) = O(\log u)$ bits, the *order* of compression ratio should not worsen. We show in the experiments that this version of the algorithm (called Block-LZ77) is faster than the plain version, since no searching of the text position is necessary. However, compression ratios worsen significantly in practice due to round-offs.

Computing partial blocks. However, the really costly part of the algorithm is not here, but in the recursive computation of the partial blocks u and v . If we consider that each time we perform a recursive call we “split” the original block B in a new place, then it is clear that at most $|B|$ recursive calls can be done until we have split it in single characters and therefore we have found the definition of each one. This shows that the total cost of the recursive calls is $O(u)$ in the worst case. Our experiments suggest that this is also the average case, but we were not able to analyze it.

Consider now the cost of the recursive invocations in the case where the new block B is strictly inside its referencing block. For instance, a letter which repeats inside a large block could trigger a long chain of recursive invocations until its real definition is found. In the worst case, we could have a block of size s which references one of size $s - 1$, and this one references another of size $s - 2$, and so on. We would work $O(s)$, but the size of the text at that point is $O(s^2)$. Hence, at text position i we cannot work more than \sqrt{i} , which gives a total worst-case cost of $O(n\sqrt{u})$, which is too high. This problem does not disappear if the compressor always stores the first occurrence of the repeated block instead of the last one, because we may not point to the first occurrence when

we consider partial blocks.

Hence the total amount of work is $\omega(u)$ in the worst case whenever $n = \omega(\sqrt{u})$, and we conjecture that this is also the average case. See the left plot of Figure 7, where we have experimented with the English text described in Section 7. Least squares fitting shows that a good model for the number of recursive invocations per text character is $0.177 + 0.1 \ln u$ (with less than 0.5% error in the approximation). The experiment suggests that the algorithm is $O(u \log u)$ on average. This is, unfortunately, worse than uncompressing and searching. We present now some techniques to improve this situation.

Improvements. A first improvement we tried consisted in storing more information than simply one description per block. For instance, when we compute the description for the partial blocks u and v (which are not part of the original sequence of blocks), we could store instead of discarding them. If later another block needs the description of u and v , we have already computed them. Figure 7 (right plot) shows that the total amount of recursive calls is reduced using this technique, and we conjecture that in this case we work $O(u)$ (least squares fitting yields a complexity of $O(u^{0.99927})$). These blocks, however, cannot be easily stored in the array of blocks since they do not belong to the sequence. A hashing implementation gave bad results in practice, that is, the cost to add the new blocks outweighed the gains of having them already computed. This could change for longer texts, if the orders of the two algorithms are different.




Figure 7: Number of recursive invocations (thick line) and block concatenations (thin line) per text character, for natural language text. The left plot shows the basic algorithm and the right plot shows the improvement of adding the computed blocks.

Another improvement, and which gave good practical results, was to try to compute *less* (instead of more) information. Our aim was to avoid the recursive computation of u and v . Hence, instead of computing their descriptions recursively, we just assume that they match all the pattern positions. If they are short enough we will not have a match even assuming this, and we could process them without actually obtaining their descriptions. Only when we find a (possible) match we backtrack to the point where it could have been started and compute correctly the involved blocks. For each block, we store whether it has been correctly or pessimistically computed. As we show in the

experiments, this improves search times for patterns of length 15 or more in practice. However, the method is limited since we cannot skip more than m characters of T without having at least one character correctly computed, hence in the very best case we pay $O(u/m)$ with this speedup. We call this algorithm Skip-LZ77 (and combined with Block-LZ77 it yields Skip-Block-LZ77).

Final remarks. Even with all these improvements, the experiments show that this algorithm is much slower than decompressing (with *gunzip*) and searching (with Shift-Or). Although ours is the first algorithm to search in LZ77-compressed text, we believe that it is not possible in practice to beat a decompress-then-search approach. The root of this limitation lies in the need to recursively compute u and v . Another consequence of the existence of partial blocks is that, even if the compressor uses a window of fixed size to select the strings to repeat, we need to keep in memory all the previous blocks, since even if they are not directly referenced anymore, we may need to resort to them in case of partial blocks. We propose in the next section a slightly different compression scheme which gets rid of all the aspects of LZ77 compression that degrade the searching performance.

We finish this section with a couple of comments. First, as it is clear from the algorithm, we do not handle the case of overlapping compression, i.e. when the referenced block can overlap with the new block B . Although we could handle it, the result is the same in cost as if the compressor avoided such overlapping (i.e. performing many steps, where a step ends when an overlap occurs). Second, other variants of LZ77 are easily accommodated. Finally, we notice that a neighborhood of size N around the occurrences can be obtained using the general mechanism at $O(N\sqrt{u})$ cost (or, according to the empirical results, $O(N \log u)$ cost). This is because of the cost to find the definitions of the incomplete blocks.

6 A New Hybrid Compression Algorithm

It became clear in the previous section that the worst part of the cost of the algorithm was due to the cost of recursively computing partial blocks, and of finding the block corresponding to a text position. We design a new compression algorithm between the LZ78 and LZ77, to have multiple-block compression (not just one block like in LZ78), but also to avoid the recursive situation which appears in searching LZ77-compressed files (Figure 6).

We propose the following algorithm. Assume that a prefix $t_1 \dots t_i$ of T has been already compressed in a sequence of block $Z = b_1 \dots b_c$. We look now for the longest prefix v of $t_{i+1} \dots t_u$ which is equal to a sequence $b_r \dots b_{r+h}$ already present in the compressed file. If there are many alternative choices for the same v , we take the one with the minimum of blocks (to reduce the cost of concatenations). And if still several possibilities occur, we take the first occurrence (the minimum in the number of the first block). We code this new block by (r, h) . As in LZ77, if v is empty (i.e the letter t_{i+1} is new), we code a special block $(0, t_{i+1})$. With the same example *ananas*, we obtain: $(0, a)$ *nanas*; $(0, a)(0, n)$ *anas*; $(0, a)(0, n)(1, 1)$ *as*; $(0, a)(0, n)(1, 1)(1, 0)$ *s*; $(0, a)(0, n)(1, 1)(1, 0)(0, s)$.

The main advantage of this compression scheme is that it avoids the recursive case in the LZ77 pattern matching (Figure 6), because we know already that the new block corresponds directly to a concatenation of already processed blocks. Moreover, we do not need to search the text position in the blocks, since we can directly access the relevant blocks.

The compression can still be performed in $O(u)$ time by using a sparse suffix tree [16] where only the block beginnings are inserted and when we fall out of the trie we consider the last node visited which corresponds to a block ending. Decompression is slower than for LZ77, since we need to keep track of the blocks already seen to be able to retrieve the appropriate text. Finally, the compression ratio is in principle worse than for LZ77 since we are limited in the text segments that we can use. On the other hand, our numbers are smaller since we code block positions in $O(\log n)$ bits instead of text positions in $O(\log u)$ bits. Moreover, if we use a simple trick, the compression is in general better than for LZ78 since we are not limited to using just one block. The trick is: we mark with one bit in the r of a block (r, h) whether h is greater than 0 or not, obtaining r' . If h is greater than 0, we replace the block (r, h) with (r', h) , and if not with (r') .

The searching algorithm is like that of LZ77 except because we do not need to search for the blocks and we do not have to recursively find the partial blocks u and v (they simply do not exist now). From the analysis of the LZ77 pattern matching algorithm we have that we work $O(\min(u, n \log m) + R)$ on average and $O(\min(u, mn) + R)$ in the worst case (thanks to the improved algorithm to concatenate blocks). In practice, this algorithm performance is very close to LZ78 pattern matching. We also tried a marked version (called Mark-Hybrid) where for each block a bit is stored which tells whether or not the block will be used again, but as for LZ78, the search times do not improve in practice.

Unlike LZ77, we can use less memory if the compressor restricts the references to a window of the text. Since there are no recursive references, those blocks which are far away in the past need not be stored since they will not be referenced anymore. Hence, as in LZ78, we need the same memory as the compressor. A window of size N can be displayed in $O(N)$ time.

7 Experimental Results

We show in this section our empirical results on the behavior of our search and compression schemes. We first study the compression techniques and later the searching performance.

We use two files for the experiments. One is an English literary text (from B. Franklin) of 1.29 Mb, and the other is the DNA chain of “h.influenzae”, of 1.36 Mb.

7.1 Compression Performance

It is interesting to study the compression performance of the algorithms for two reasons: first, we propose a hybrid compression scheme which we have to evaluate in terms of compression ratios. Second, our search algorithms use a technique to code the pairs which speeds up search time but which is suboptimal: the numbers are stored in as many bytes as needed (using the highest bit to denote if there are more bytes or not).

We first compare the number of bits needed to code a file with our hybrid compression scheme against the same number for LZ77 and LZ78. We call this approach “bit-coding”. This is aimed to give an idea of the expected compression performance when the file is compressed with a real technique (such as Elias codes [10]). Many other tricks are possible, for instance *gzip* uses Huffman coding on the pairs, since they are not uniformly distributed. A deeper study of the best technique to code the pairs in our hybrid compressor is deferred for future study.

Table 1 shows the results. The “Ideal” row counts exactly the bits used by each number stored in the compressed file, while both “Elias” rows count the number of bits needed to represent the numbers using these codes² [10]. The letters, however, are always represented with 6 bits (English lowercase text) and 2 bits (DNA text). The percentages in parenthesis represent different variants of the compressors: for LZ77 it refers to the version where blocks and offsets are referenced instead of text positions, while for LZ78 and Hybrid it represents the versions where the blocks are marked to know if they will be used again or not. With our Hybrid compression method, we obtain estimated compression ratios comparable to LZ77. The Hybrid and LZ77 compression is better than LZ78 except for DNA text, where only two bits are necessary to code a letter. The block version of the LZ77 compresses quite badly, and its *Elias- γ* code gives it a size bigger than the original file on natural language text.

Method	File	Bit-LZ77 (blocks)	Bit-LZ78 (marked)	Bit-Hybrid (marked)
Ideal	English	29.67% (52.45%)	39.01% (40.87%)	29.28% (31.24%)
	DNA	28.03% (47.21%)	24.98% (26.45%)	29.08% (31.15%)
Elias- γ	English	59.34% (104.90%)	66.87% (85.17%)	58.57% (62.48%)
	DNA	56.06% (94.43%)	47.01% (67.30%)	58.18% (62.30%)
Elias- δ	English	48.96% (74.25%)	54.90% (57.57%)	46.17% (48.75%)
	DNA	45.77% (73.14%)	37.39% (39.59%)	46.40% (49.03%)

Table 1: Estimated compression ratios with three different methods. For each number in the compressed file, if we note n the bits needed to code it, then *Ideal* counts only n , *Elias- γ* counts $2n$ and *Elias- δ* counts $n + 2\lceil \log_2 n \rceil$.

We now perform a practical comparison using our byte-coding techniques against good LZ77 and LZ78 compressors, namely *gzip* and *Compress* respectively. This is to show how much compression are we losing in order to ease the searching process.

Table 2 shows the compression ratios achieved. The percentages in parenthesis have the same meaning as before. Interestingly, *Compress* is better than *gzip* on DNA, which rarely happens on natural language texts. Our compression ratios are more or less 10% worse than the one of *gzip* for our LZ77 and our hybrid compression and 15% worse for our LZ78. Our byte compression method is very simple, and these results show in which proportion our compression ratios could be improved by engineering techniques, keeping in mind that complicating the encoding of the numbers risks to slow down the pattern matching process.

File	<i>gzip</i>	<i>Compress</i>	Byte-LZ77 (blocks)	Byte-LZ78 (marked)	Byte-Hybrid (marked)
English	35.58%	38.90%	44.49% (79.32%)	54.41% (56.20%)	43.29% (45.24%)
DNA	30.44%	27.96%	41.07% (75.24%)	43.17% (44.90%)	42.23% (44.22%)

Table 2: Compression ratios for classical compressors and our byte versions.

²Recall that Elias- γ precedes the number x by its length in unary, while Elias- δ uses Elias- γ to code that length that precedes the number.

7.2 Searching Algorithms

We compare now the search times for our algorithms against the decompressing and searching approach. The experiments were run on a Sun UltraSparc-1 of 167 MHz, with 64 Mb of RAM, running Solaris 2.5.1. We consider user times, which are within 2% of accuracy with 95% confidence. The times are expressed in seconds.

In general, searching a compressed text has the additional advantage over the uncompressed text that it performs less I/O. However, this is relevant if we compare compressed searching versus uncompressed searching. This is not what we compare here: we consider that the text is always compressed. Hence, we measure the cost of searching it without decompressing versus the cost of decompressing it and then searching. Clearly the last task can be done using an intermediate buffer in main memory, and therefore the I/O is the same in both cases.

Figure 8 compares the marked and unmarked versions of LZ78 and the Hybrid compressor. As it can be seen, there is no advantage in practice by the use of marking. Therefore, we do not further consider the marked versions. Another conclusion we take from the figure is that the searcher for Hybrid compression is slightly faster than for LZ78 on English text but slower for DNA. This may be related to the good performance of the LZ78 compressor on DNA.

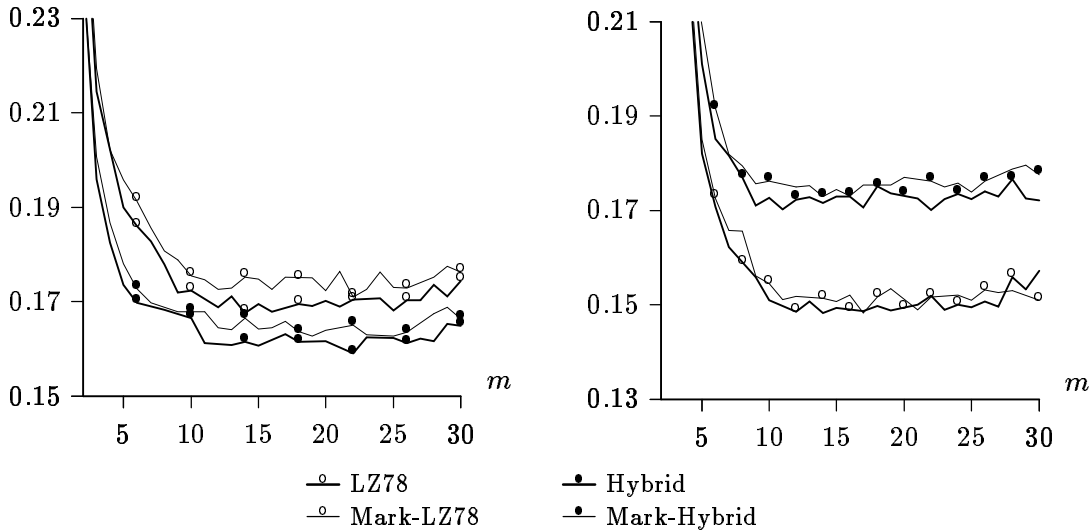


Figure 8: Comparison of the marked and unmarked versions of LZ78 and Hybrid compressors. The left plot is for English text and the right one for DNA.

Figure 9 compares all the search algorithms together, as well as decompression (with *gunzip*) plus search times (with Shift-Or and BNDM [25], a bit-parallel searcher which is the fastest in practice together with [26]). It can be seen that Block-LZ77 improves significantly over LZ77, and that the Skip-LZ77 versions improve as the pattern length grows. However, all the LZ77 search algorithms are not competitive against decompressing and searching, especially on DNA. On the other hand, both the Hybrid and LZ78 search algorithms are twice as fast as decompressing and searching.

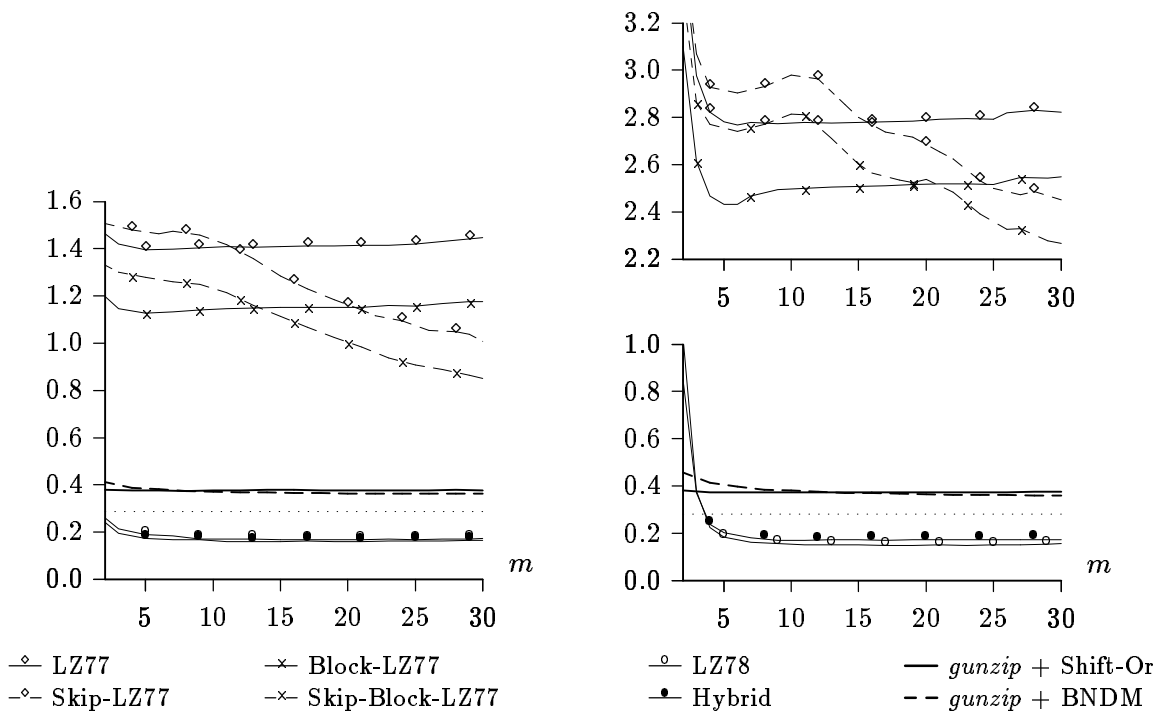


Figure 9: Comparison of the search algorithms. The dotted line is the time taken by *gunzip* alone. The left plot is for English text and the right one for DNA.

8 Conclusions

We have focused in the problem of string matching on Ziv-Lempel compressed text. This is an important practical problem, as it is of interest to be able to keep the texts compressed and at the same time being able to efficiently search on it.

We presented a general paradigm to search in a text that is expressed as a sequence of blocks, which abstracts the main features of Ziv-Lempel compression. Then, we applied the technique to the different variants, i.e. LZ77 and LZ78. For LZ78, we are able to search in half the time of uncompressing and searching, while for LZ77 our algorithm, although much slower, is the first one proposed to search on LZ77 compressed text. This motivated us to present a new hybrid compression technique which allows to search as fast as in LZ78 but which keeps many of the features of LZ77 compression, being in practice similar in compression ratios.

Therefore, we are able to search in a compressed text faster than uncompressing and then searching. In general, on the other hand, searching on compressed text at the same speed of on uncompressed text seems difficult to achieve in practice because of a basic problem of locality of reference.

Future work involves studying better the performance of our hybrid compression, both in theory and in practice (especially on finding better methods to encode the numbers). We also plan to work more in understanding the behavior of the algorithm to search in LZ77 compressed text. This is a field where important theoretical and practical development is necessary, and we have presented new results in both aspects. We hope that more improvements are to come.

References

- [1] A. Amir and G. Benson. Efficient two-dimensional compressed matching. In *Proc. Second IEEE Data Compression Conference*, pages 279–288, March 1992.
- [2] A. Amir, G. Benson, and M. Farach. Let sleeping files lie: Pattern matching in Z-compressed files. *Journal of Computer and System Sciences*, 52(2):299–307, 1996.
- [3] R. Baeza-Yates. Text retrieval: Theory and practice. In *12th IFIP World Computer Congress*, volume I, pages 465–476. Elsevier Science, September 1992.
- [4] R. Baeza-Yates and G. Gonnet. A new approach to text searching. *Communications of the ACM*, 35(10):74–82, October 1992.
- [5] R. Baeza-Yates and G. Navarro. A faster algorithm for approximate string matching. In *Proc. CPM'96*, number 1075 in LNCS, pages 1–23, 1996.
- [6] T. Bell and D. Kulp. Longest-match string searching for Ziv-Lempel compression. *Software-Practice and Experience*, 23(7):757–771, July 1993.
- [7] J. Bentley, D. Sleator, R. Tarjan, and V. Wei. A locally adaptive data compression scheme. *Communications of the ACM*, 29:320–330, 1986.

- [8] R. S. Boyer and J. S. Moore. A fast string searching algorithm. *Communications of the ACM*, 20(10):762–772, 1977.
- [9] A. Czumaj, Maxime Crochemore, L. Gasieniec, S. Jarominek, Thierry Lecroq, W. Plandowski, and W. Rytter. Speeding up two string-matching algorithms. *Algorithmica*, 12:247–267, 1994.
- [10] P. Elias. Universal codeword sets and representations of the integers. *IEEE Transactions on Information Theory*, 21:194–203, 1975.
- [11] M. Farach and M. Thorup. String matching in Lempel-Ziv compressed strings. In *27th ACM Annual Symposium on the Theory of Computing*, pages 703–712, 1995.
- [12] E. Fiala and D. Greene. Data compression with finite windows. *Communications of the ACM*, 32(4):490–505, 4 1989.
- [13] L. Gasieniec, M.Karpinski, W.Plandowski, and W. Rytter. Efficient algorithms for Lempel-Ziv encodings. In *Proc. of the Scandinavian Workshop on Algorithm Theory*, 1996.
- [14] R. N. Horspool. Practical fast searching in strings. *Software Practice and Experience*, 10:501–506, 1980.
- [15] D. Huffman. A method for the construction of minimum-redundancy codes. *Proc. of the I.R.E.*, 40(9):1090–1101, 1952.
- [16] J. Kärkkäinen and E. Ukkonen. Sparse suffix trees. In *COCOON'96*, pages 219–230, 1996. LNCS v. 1090.
- [17] M. Karpinski, A. Shinohara, and W. Rytter. Pattern matching problem for strings with short descriptions. *Nordic Journal of Computing*, 4(2):172–186, 1997.
- [18] D. E. Knuth, J. H. Morris, Jr, and V. R. Pratt. Fast pattern matching in strings. *SIAM Journal on Computing*, 6(1):323–350, 1977.
- [19] U. Manber. A text compression scheme that allows fast searching directly in the compressed file. *ACM Transactions on Information Systems*, 15(2):124–136, 1997.
- [20] V. Miller and M. Wegman. Variations on a theme by Ziv and Lempel. In *Combinatorial Algorithms on Words*, volume 12 of *NATO ASI Series F*, pages 131–140. Springer-Verlag, 1985.
- [21] A. Moffat. Word-based text compression. *Software Practice and Experience*, 19(2):185–198, 1989.
- [22] E. Moura, G. Navarro, N. Ziviani, and R. Baeza-Yates. Direct pattern matching on compressed text. In *Proc. SPIRE'98*, pages 90–95. IEEE CS Press, 1998.
- [23] E. Moura, G. Navarro, N. Ziviani, and R. Baeza-Yates. Fast searching on compressed text allowing errors. In *Proc. SIGIR'98*, pages 298–306. York Press, 1998.

- [24] G. Myers. A fast bit-vector algorithm for approximate pattern matching based on dynamic programming. In *Proc. CPM'98*, LNCS v. 1448, pages 1–13, 1998.
- [25] G. Navarro and M. Raffinot. A bit-parallel approach to suffix automata: Fast extended string matching. In *Proc. CPM'98*, LNCS v. 1448, pages 14–33, 1998.
- [26] D. Sunday. A very fast substring search algorithm. *Communications of the ACM*, 33(8):132–142, August 1990.
- [27] T. A. Welch. A technique for high performance data compression. *IEEE Computer Magazine*, 17(6):8–19, June 1984.
- [28] I. Witten, R. Neal, and J. Cleary. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–541, 1987.
- [29] S. Wu and U. Manber. Fast text searching allowing errors. *Communications of the ACM*, 35(10):83–91, October 1992.
- [30] M. Zipstein. Data compression with factor automata. *Theor. Comput. Sci.*, 92(1):213–221, 1992.
- [31] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Trans. Inf. Theory*, 23:337–343, 1977.
- [32] J. Ziv and A. Lempel. Compression of individual sequences via variable length coding. *IEEE Trans. Inf. Theory*, 24:530–536, 1978.