

On the bit-complexity of Lempel-Ziv compression*

Paolo Ferragina[†]

Igor Nitto[†]

Rossano Venturini[†]

Abstract

One of the most famous and investigated lossless data-compression schemes is the one introduced by Lempel and Ziv about 30 years ago [37]. This compression scheme is known as “dictionary-based compressor” and consists of squeezing an input string by replacing some of its substrings with (shorter) codewords which are actually pointers to a dictionary of phrases built as the string is processed. Surprisingly enough, although many fundamental results are nowadays known about the speed and effectiveness of this compression process (see e.g. [23, 28] and references therein), “*we are not aware of any parsing scheme that achieves optimality when the LZ77-dictionary is in use under any constraint on the codewords other than being of equal length*” [28, pag. 159]. Here optimality means to achieve the *minimum* number of bits in compressing *each individual* input string, *without* any assumption on its generating source. In this paper we investigate three issues pertaining to the bit-complexity of LZ-based compressors, and we design algorithms which achieve bit-optimality in the compressed output size by taking efficient/optimal time and optimal space. These theoretical results will be sustained by some experiments that will compare our novel LZ-based compressors against the most popular compression tools (like `gzip`, `bzip2`) and state-of-the-art compressors (like the *booster* of [14, 13]).

1 Introduction

The problem of *lossless* data compression consists of compactly representing data in a format that can be faithfully recovered from the compressed file. Lossless compression is achieved by taking advantage of the *redundancy* present often in the data generated by either humans or machines. One of the most famous lossless data-compression schemes is the one introduced by Lempel and Ziv in the late 70s, and indeed many (non-)commercial programs are currently based on it—

like `gzip`, `zip`, `pkzip`, `arj`, `rar`, just to cite a few. This compression scheme is known as *dictionary-based compressor*, and consists of squeezing an input string $S[1, n]$ by replacing some of its substrings with (shorter) *codewords* which are actually pointers to a dictionary of phrases. The dictionary can be either *static* (in that it has been constructed before the compression starts) or *dynamic* (in that it is built as the input string is compressed). The well-known LZ77 and LZ78 compressors, proposed by Lempel and Ziv in [37, 38], and all their variants [29], are interesting examples of *dynamic* dictionary-based compressors. In LZ77, and its variants, the dictionary consists of all substrings starting in the last M scanned positions of the text, where M is called the *window size* and possibly depends on the text length. Each codeword consists of a triple $\langle d, \ell, c \rangle$ where d is the relative *offset* of the copied phrase ($d \leq M$), ℓ is the length of this phrase and c is the single (new) character following it. In LZ78, the dictionary is built upon phrases extracted from the previously scanned prefix of the input string, and each codeword consists of a pair $\langle \text{id}, c \rangle$ where *id* is the *identifier* of the dictionary phrase and c is the character following that phrase in the string.

Many theoretical and experimental results have been dedicated to LZ-compressors in these thirty years (see e.g. [29] and references therein); and, although today there are alternative solutions to the lossless data-compression problem (e.g., Burrows-Wheeler compression and Prediction by Partial Matching [35]), dictionary-based compression is still widely used for its unique combination of compression power and compression/decompression speed. Over the years dictionary-based compression has also gained importance as a general algorithmic tool, being employed in the design of compressed text indexes [27], in *universal* clustering [5] or classification tools [36], in designing *optimal* prefetching mechanisms [33], and in streaming or on-the-fly compression applications [8, 17].

In this paper we address some key issues which arise when dealing with the output-size in bits of so called *LZ-parsing schemes*, namely schemes which parse the input string in phrases extracted from the dynamic dictionaries built according to the LZ77 or LZ78 rules (detailed above). Classically, these parsers adopt a *greedy*

*It has been partially supported by Yahoo! Research, Italian MIUR Italy-Israel FIRB Project and PRIN MAINSTREAM. The authors' address is Dipartimento di Informatica, L.go B. Pontecorvo 3, 56127 Pisa, Italy. email: {ferragina,nitto,rossano}@di.unipi.it

[†]Dipartimento di Informatica, University of Pisa, Italy

rule, namely one that at each step takes the *longest* dictionary phrase which is a prefix of the currently unparsed suffix of the input string. This greedy parsing can be computed in $O(n \log \sigma)$ time and $O(M)$ space [16].¹ The greedy parsing is also optimal with respect to the *number of phrases* in which S can be parsed by any suffix-complete dictionary (like the LZ77-dictionary) or, a small variation of it (called *flexible*-parsing [25]), is optimal for prefix-complete dictionaries (like the LZ78-dictionary). Of course, the number of parsed phrases influences the compression ratio and, indeed, various authors [37, 23] proved that greedy parsing achieves asymptotically the (*empirical*) *entropy* of the source generating the input string S . However, these fundamental results have *not yet closed* the problem of optimally compressing S because the optimality in the number of parsed phrases *is not necessarily equal* to the optimality in the number of bits output by the final compressor on *each individual* input string S . In fact, if the phrases are compressed via an *equal-length* encoder, like in [23, 29, 37], then the produced output is *bit optimal*. But if one aims for higher compression, *variable-length encoders* should be taken into account (see e.g. [35, 12], and the software `gzip` [19]), and in this situation the greedy-parsing scheme is *no longer optimal* in terms of the number of bits output by the final compressor. Starting from these premises we address in this paper four main problems, both on the theoretical and the experimental side, which pertain with the bit-optimal compression of the input string S via parsers that deploy the LZ77-dictionary built on an unbounded window (namely it is $M = n$). These parsers will be hereafter called “LZ77-parsers”. Our results extend to the LZ78-dictionary, to windows of arbitrary size M , and to some other LZ-variants (see Section 6).

Problem 1. Let us consider the greedy LZ77-parser, and assume that we encode every parsed phrase w_i with a variable-length encoder. The value of $\ell_i = |w_i|$ is in some sense fixed by the greedy choice, being the length of the longest phrase occurring in the current LZ77-dictionary. Conversely, the value of d_i depends on the position of the copy of w_i in S . In order to minimize the number of bits output by the final compressor, the greedy parser should obviously select the *closest* copy of each phrase w_i in S , and thus the smallest possible d_i . Surprisingly enough, known implementations of greedy parsers are time optimal but not bit-optimal, because they select an arbitrary or the leftmost occurrence of the longest copied phrase (see

[9] and references therein), or they select the closest copy but take $O(n \log n)$ suboptimal time [1, 24]. In Section 3 we provide an elegant, yet simple, algorithm which computes at each parsing step the closest copy of the longest dictionary phrase in $O(n \frac{\log \sigma}{\log \log n})$ overall time and $O(n)$ space (Section 3, Lemma 3.1). This is optimal in terms of time/space performance when the alphabet has size *polylog*(n) (i.e. almost all texts of practical interest).

Problem 2. How good is the greedy LZ77-parsing of S whenever the compression cost is measured in terms of number of bits produced in output? Not surprisingly, we show that the greedy selection of the longest dictionary phrase at each parsing step is not optimal! But surprisingly, we show that the loss in using the greedy parser with respect to the parser that achieves bit-optimality in the compressed output size is *not* negligible, and diverges by a multiplicative factor $\Omega(\log n / \log \log n)$, which is unbounded asymptotically (Section 4). Additionally, we show that this lower-bound is tight up to a factor $\Theta(\log \log n)$, and we support these theoretical figures with some experimental results, see Table 1 below, which stress the practical importance of finding the bit-optimal parsing of S .

Problem 3. How much efficiently (in time and space) we can compute the bit-optimal LZ77-parsing of S ? Several solutions are indeed known for this problem but they are either inefficient [30, 10], in that they take $\Theta(n^2)$ worst-case time and space, or they are approximate [20], or they rely on heuristics [22, 31, 3, 6, 10] which do not provide any guarantee on the time/space performance of the compression process. This is the reason why Rajpoot and Sahinalp stated in [28, pag. 159] that “*We are not aware of any on-line or off-line parsing scheme that achieves optimality when the LZ77-dictionary is in use under any constraint on the codewords other than being of equal length*”. In this paper we investigate this question by considering a general class of variable-length codeword encodings which are typically used in data compression (e.g. `gzip`) and in the design of search engines and compressed indexes [27, 29, 35]. Our final result is a time efficient and space optimal solution for the problem above (Theorem 5.2). Due to space limitations, we will detail our results only for the LZ77-parser, and discuss the case of the LZ78-dictionary in the last Section 6.

Technically speaking, we follow [30] and model the search for a bit-optimal LZ77-parsing of an input string S as a *single-source shortest path* problem (shortly, SSSP) on a *weighted DAG* $\mathcal{G}(S)$ consisting of n nodes, one per character of S , and e edges, one per possible LZ77-parsing step. Every edge is weighted according to the length in bits of the codeword adopted to compress the

¹Recently, [9] showed how to achieve the optimal $O(n)$ time and space when the alphabet has size $O(n)$ and the window is unbounded, i.e. $M = n$.

corresponding phrase. Since these codewords are tuples of integers (see above), we consider a natural class of codeword encoders which satisfy the so called *increasing cost property*: the greater is the integer to be encoded, the longer is the codeword. This class encompasses most of the encoders used in the literature to design data compressors (see [12] and `gzip` [19]), compressed full-text indexes [27] and search engines [35]. We prove new combinatorial properties for this SSSP-problem and show that the computation of the SSSP in $\mathcal{G}(S)$ can be restricted onto a subgraph $\tilde{\mathcal{G}}(S)$ whose structure depends on the integer-encoding functions adopted to compress the LZ77-phrases, and whose size is *provably smaller* than the complete graph generated by [30] (see Theorem 5.1). Additionally, we design an algorithm that solves the SSSP on the subgraph $\tilde{\mathcal{G}}(S)$ without materializing it *all at once*, but it creates and explores its edges *on-the-fly* in optimal $O(1)$ amortized time per edge and $O(n)$ optimal space overall. As a result, our novel LZ77-compressor achieves bit-optimality in $O(n)$ optimal working space and in time proportional to $|\tilde{\mathcal{G}}(S)|$ (hence, it is optimal in its size). The latter is $O(n \log n)$ for a large class of integer encoders, like Elias and Fibonacci codes [35, 12], and it is the optimal $O(n)$ for (most of) the encodings used by `gzip` [19]. This is the first result providing a *positive answer* to Rajpoot-Sahinalp’s question above!

Problem 4. How much efficient are *in practice* our LZ77-parsers compared to known compressors? To establish this, we have taken several freely available text collections, and compared our LZ77-based compressors against the classic `gzip` and `bzip2`, as well as against the state-of-the-art *boosting* compressor of [14, 13]. Table 1 reports some experimental figures. Let us first consider algorithm `Fixed-LZ77`, which uses an unbounded window and equal-length encoders for the distance of the copied phrases: its compression performance shows that an unbounded window may introduce a significant compression gain wrt to a bounded one, as used by `gzip` and `bzip2` (see e.g. HTML), thus witnessing the presence in current (Web/text) collections of surprisingly many long repetitions at large distances. Then we consider `Rightmost-LZ77` (Problem 1), and notice that it improves `Fixed-LZ77`: as expected, variable-length encoders achieve higher compression ratios. This was the starting point of our theoretical investigation! Finally, we test `BitOptimal-LZ77` (Problem 3) finding that it improves `Rightmost-LZ77`, as theoretically predicted in Problem 2. Surprisingly, `BitOptimal-LZ77` significantly improves `bzip2` (which uses a bounded window) and comes close to the booster (which uses an unbounded window). Additionally, since `BitOptimal-LZ77` adopts the same decompression al-

gorithm of `gzip`, it retains its *fast decompression speed* which is at least one order of magnitude faster than decompressing `bzip2`’s or booster’s compressed files (see table below). This is a nice combination which makes `BitOptimal-LZ77` practically relevant for a wide range of applications in which the paradigm is “compress once & decompress many times” (like in Web search engines and IR systems), or where the decompression system is less powerful than the compressor one (like a server that distributes data to clients, possibly mobile phones).

These results *resort* LZ-based approaches to compress large textual collections, and thus pave the way to future algorithmic engineering research devoted to design proper variable-length encoders for `BitOptimal-LZ77`, which possibly depend onto the structural features of the input data to be compressed (see Section 6).

Compressor	english [15]	C/C++/Java src [15]
<code>gzip -9</code>	37.52%	23.29%
<code>bzip2 -9</code>	28.40%	19.78%
<code>boosterOpt</code>	20.62%	17.36%
<code>Fixed-LZ77</code>	26.19%	24.63%
<code>Rightmost-LZ77</code>	23.81%	20.14%
<code>BitOptimal-LZ77</code>	21.62%	17.62%
Compressor	HTML [4]	Avg Dec. time (sec)
<code>gzip -9</code>	20.09%	0.7
<code>bzip2 -9</code>	10.63%	6.3
<code>boosterOpt</code>	3.89%	20.2
<code>Fixed-LZ77</code>	4.98%	0.8
<code>Rightmost-LZ77</code>	4.27%	0.9
<code>BitOptimal-LZ77</code>	3.87%	0.9

Table 1: Each text collection consists of 50 Mbytes of data. All the experiments were executed on a 2.6 GHz Pentium 4, with 1.5 GB of main memory, and running Fedora Linux.

2 Notation and terminology

Let $S[1, n]$ be a string drawn from an alphabet Σ of size σ . We use $S[i]$ to denote the i th symbol of S ; $S[i : j]$ to denote the substring (also called the *phrase*) extending from the i th to the j th symbol in S (extremes included); and $S_i = S[i : n]$ to denote the i -th suffix of S .

In the rest of the paper we will concentrate on LZ77-compression with an unbounded window size; thus, we consider a dictionary-based compressor that parses the input text by using phrases extracted from the LZ77-dictionary, which therefore consists of all substrings starting in the text prefix already scanned. Our results are easily generalizable to arbitrary window sizes (possibly constant, like in `gzip`), and other LZ-compressors (like LZ78, see Section 6). The LZ77-compressor, as any dictionary-based compressor, works in two intermingled phases: *parsing* and *encoding*. Let w_1, w_2, \dots, w_{i-1} be

the phrases in which a prefix of S has been already parsed. At this step, the LZ77-dictionary consists of all substrings of S starting in the last M positions of $w_1 w_2 \cdots w_{i-1}$, where M is called the *window size* (hereafter assumed unbounded, for simplicity). The classic parser adopted by the LZ77-compressor selects the next phrase according to the so called *longest match heuristic*: that is, this phrase is taken as the *longest* phrase in the current dictionary which prefixes the remaining suffix of S . This will be hereafter called *greedy parsing*. After such a phrase is selected, the parser adds one further symbol to it and thus forms the next phrase w_i of S 's parsing. In the rest of the paper, and for simplicity of exposition, we will restrict to the LZ77-variant which avoids the additional symbol per phrase. This means that w_i is represented by the integer pair $\langle d_i, \ell_i \rangle$, where d_i is the relative *offset* of the copied phrase w_i within the prefix $w_1 \cdots w_{i-1}$ and ℓ_i is its length $|w_i|$. Every first occurrence of a new symbol c is encoded as $\langle 0, c \rangle$.

Once phrases are identified and represented via pairs of integers, their components are compressed via *variable-length integer encoders* which eventually produce the compressed output of S as a sequence of bits. In order to study and design bit-optimal parsing schemes, we therefore need to deal with such integer encoders. Let f be an integer-encoding function that maps any integer $x \in [n]$ into a (bit-)codeword $f(x)$ whose length is denoted by $|f(x)|$ bits. In this paper we consider variable-length encodings which use longer codewords for greater integers:

PROPERTY 2.1. (INCREASING COST PROPERTY) *For any $x, y \in [n]$ it is $x \leq y$ iff $|f(x)| \leq |f(y)|$.*

This property is satisfied by most of known integer encoders—like equal-length codewords, Elias codes [35], Fibonacci's codes [12]—which are used to design data compressors [29], compressed full-text indexes [27] and search engines [35].

3 An efficient and bit-optimal greedy LZ77-parsing

Let f and g be two integer encoders which satisfy the Increasing Cost Property (possibly $f = g$). We denote by $\text{LZ}_{f,g}(S)$ the compressed output produced by the (classic) greedy-parsing strategy in which we have used f to compress the distance d_i , and g to compress the length ℓ_i of any parsed phrase w_i . Thus, in $\text{LZ}_{f,g}(S)$ any phrase w_i is encoded in $|f(d_i)| + |g(\ell_i)|$ bits. Given that the parsing is the greedy one, ℓ_i is in some sense fixed (to be the length of the longest copy), so we minimize $|\text{LZ}_{f,g}(S)|$ by minimizing the *distance* d_i of w_i 's copy in S . If p_i is the starting position of w_i in S (namely $S[p_i, p_i + \ell_i - 1] = w_i$), many copies of the phrase w_i

could be present in $S[1, p_i - 1]$. To minimize $|\text{LZ}_{f,g}(S)|$ we should choose the copy which is the closest one to p_i , and thus requires the minimum number of bits to encode its distance d_i (recall the assumption $M = n$).

In this section we propose an elegant, yet simple, algorithm that selects the rightmost copy of each phrase w_i in $O(n \log \sigma / \log \log n)$ time. This algorithm is the fastest known in the literature [9, 24], and results optimal for alphabets with a *polylog*(n) size (i.e., almost all texts in practice). It requires the suffix tree \mathcal{ST} of S , preprocessed to support constant-time *lca*-queries, and the LZ77-parsing of S which consists of, say, $k \leq n$ phrases. We say that a node u of \mathcal{ST} is *marked* iff the string spelled out by the root-to- u path in \mathcal{ST} is equal to some phrase w_i . In this case we use the notation u_{p_i} to denote the node marked by phrase w_i which starts at position p_i of S . Since the same node may be marked by different phrases, but any phrase marks just one node, the total number of marked nodes is bounded by the number of phrases, hence k . Furthermore, if a node is assigned with many phrases, since the greedy LZ77-parsing takes the longest one, it must be the case that every such occurrences of w_i is followed by a distinct character. So the number of phrases assigned to the same marked node is bounded by σ .

All marked nodes can be computed in $O(k)$ time by executing k *lca*-queries on \mathcal{ST} . Let us now define $\mathcal{ST}_{\mathcal{C}}$ as the contracted version of \mathcal{ST} , namely a tree whose internal nodes are the marked nodes of \mathcal{ST} and whose leaves are the leaves of \mathcal{ST} . The parent of any node in $\mathcal{ST}_{\mathcal{C}}$ is its lowest marked ancestor in \mathcal{ST} . It is simple to build $\mathcal{ST}_{\mathcal{C}}$ in linear time via a top-down visit of \mathcal{ST} . $\mathcal{ST}_{\mathcal{C}}$ consists of $O(k)$ internal nodes and n leaves.

Given the properties of suffix trees, we can now rephrase our problem as follows: for each position p_i , we need to compute the largest position x_i which is smaller than p_i and whose leaf in $\mathcal{ST}_{\mathcal{C}}$ lies within the subtree rooted at u_{p_i} . Our algorithm processes the input string S from left to right and, at each position j , it maintains the following invariant: the parent v of any leaf in $\mathcal{ST}_{\mathcal{C}}$ stores the maximum position $h < j$ such that the leaf labeled h is attached to v . Maintaining this invariant is trivial: after that position j is processed, j is assigned to the leaf parent of the leaf labeled j in $\mathcal{ST}_{\mathcal{C}}$. The key point is how to compute the position x_i of the rightmost-copy of w_i whenever we discover that j is the starting position of a phrase (i.e. $j = p_i$ for some i). In this case, the algorithm visits the subtree of $\mathcal{ST}_{\mathcal{C}}$ rooted at u_j and computes the maximum position stored in its marked nodes. By the invariant, this position is the rightmost copy of the phrase w_i . This process takes $O(n + \sigma \sum_{i=1}^k \#(u_{p_i}))$ time, where $\#(u_{p_i})$ is the number of marked nodes in the subtree rooted at u_{p_i} in

\mathcal{ST}_C . In fact, by construction, there can be at most σ repetitions of the same phrase in the LZ77-parsing of S , and for each of them the algorithm performs a visit of the corresponding subtree.

As a final step we prove that $\sum_{i=1}^k \#(u_{p_i}) = O(n)$. By construction of suffix trees, the depth of u_{p_i} is smaller than $\ell_i = |w_i|$, and each (marked) node of \mathcal{ST}_C is visited as many times as the number of its (marked) ancestors in \mathcal{ST}_C (with their multiplicities). For each (marked) node u_{p_i} , this number can be bounded by $\ell_i = O(\ell_i)$. Summing up on all nodes, we get $\sum_{i=1}^k O(\ell_i) = O(n)$. Thus, the above algorithm requires $O(\sigma \times n)$ time, which is trivially optimal whenever $\sigma = O(1)$. Actually, we are able to further reduce the time complexity to $O(n \log \sigma / \log \log n)$ by properly combining a slightly modified variant of the tree covering procedure of [18] with a dynamic Range Maximum Query data structure [26, 34] applied on properly composed arrays of integers. Notice that this improvement leads our algorithm to require the optimal $O(n)$ time, for alphabets whose size is poly-logarithmic in n . Due to the lack of space the description of this solution is deferred to full paper.

LEMMA 3.1. *Given a string $S[1, n]$ drawn from an alphabet of size σ , we can design an algorithm that computes the greedy LZ77-parsing of S and reports the right-most copy of each phrase in $O(n \frac{\log \sigma}{\log \log n})$ time and $O(n)$ space.*

4 On the bit-efficiency of the greedy LZ77-parsing

We have already noticed that $\text{LZ}_{f,g}(S)$ is not necessarily bit optimal, so we will hereafter use $\text{OPT}_{f,g}(S)$ to denote the *Bit-Optimal LZ77-parsing* of S relative to f and g , namely a parsing of S which uses phrases extracted from the LZ77-dictionary and *minimizes* the total number of bits produced by using the encoding functions f and g . Of course $|\text{LZ}_{f,g}(S)| \geq |\text{OPT}_{f,g}(S)|$, but this does not provide us with any estimate of how much worse the greedy parsing can be with respect to the bit-optimal one. In what follows we identify an infinite family of strings S for which $\frac{|\text{LZ}_{f,g}(S)|}{|\text{OPT}_{f,g}(S)|} = \Omega(\frac{\log n}{\log \log n})$, so the gap may be asymptotically unbounded thus stressing the need for an (f, g) -optimal parser, as requested by [28].

Our argument holds for any choice of f and g from the family of encoding functions that represent an integer x with a bit string of size $\Theta(\log x)$ bits (thus the well-known Elias' and Fibonacci's coders belong to this family). Taking inspiration from the proof of Lemma 4.2 in [23], we consider the infinite family of strings $S_l = ba^l c^{2^l} ba^2 ba^3 \dots ba^l$, parameterized in the positive value l . The greedy LZ77-parser partitions S_l

as²: $(b) (a) (a^{l-1}) (c) (c^{2^l-1}) (ba) (ba^2) (ba^3) \dots (ba^l)$, where the symbols forming a parsed phrase have been delimited within a pair of brackets. Thus it copies the latest l phrases from the beginning of S_l and takes at least $l |f(2^l)| = \Theta(l^2)$ bits.

A more parsimonious parser selects the copy of ba^{i-1} (with $i > 1$) from its immediately previous occurrence thus parsing S_l as: $(b) (a) (a^{l-1}) (c) (c^{2^l-1}) (b) (a) (ba) (a) (ba^2) (a) \dots (ba^{l-1}) (a)$. Hence the encoding of this parsing, called $\text{rOPT}(S_l)$, takes $|g(2^l)| + |g(l)| + \sum_{i=2}^l [|f(i)| + |g(i)| + f(0)] + O(l) = O(l \log l)$ bits.

LEMMA 4.1. *There exists an infinite family of strings such that, for any of its elements S , it is $|\text{LZ}_{f,g}(S)| \geq \Theta(\log |S| / \log \log |S|) |\text{OPT}_{f,g}(S)|$.*

Proof. Since $|\text{OPT}(S_l)| \leq |\text{rOPT}(S_l)|$, we can conclude that: $\frac{|\text{LZ}_{f,g}(S_l)|}{|\text{OPT}_{f,g}(S_l)|} \geq \frac{|\text{LZ}_{f,g}(S_l)|}{|\text{rOPT}(S_l)|} \geq \Theta\left(\frac{l}{\log l}\right)$. Since $|S_l| = 2^l + l^2 - O(l)$, we have that $l = \Theta(\log |S_l|)$ for sufficiently long strings. ■

The experimental results reported in Table 1 show that this gap is not negligible in practice too, just look at the entries **Fixed-LZ77** and **BitOptimal-LZ77**. Additionally we can prove that this lower bound is tight up to a $\log \log |S|$ multiplicative factor, by easily extending to Property 2.1 and to the LZ77-dictionary (which is dynamic), a result proved in [21] for static dictionaries. Precisely (details in the full paper), we can show that $\frac{|\text{LZ}_{f,g}(S)|}{|\text{OPT}_{f,g}(S)|} \leq \frac{|f(n)| + |g(n)|}{|f(0)| + |g(0)|}$, which is upper bounded by $O(\log n)$ because $|S| = n$, $|f(n)| = |g(n)| = \Theta(\log n)$ and $|f(0)| = |g(0)| = O(1)$.

5 Bit-Optimal Parsing and SSSP-problem

Following [30], we model the design of a bit-optimal LZ77-parsing strategy for a string S as a Single-Source Shortest Path problem (shortly, **SSSP-problem**) on a weighted DAG $\mathcal{G}(S)$ defined as follows. Graph $\mathcal{G}(S) = (V, E)$ has one vertex per symbol of S plus a dummy vertex v_{n+1} , and its edge set E is defined so that $(v_i, v_j) \in E$ iff (1) $j = i+1$ or (2) the substring $S[i : j-1]$ occurs in S starting from a (previous) position $p < i$. Clearly $i < j$ and thus $\mathcal{G}(S)$ is a DAG. Every edge (v_i, v_j) is labeled with the pair $\langle d_{i,j}, \ell_{i,j} \rangle$ which is set to $\langle 0, S[i] \rangle$ in case (1), or it is set to $\langle p - i, j - i \rangle$ in case (2). The second case corresponds to copying a phrase longer than one single character.

²Recall the variant of LZ77 we are considering in this paper, which uses just a pair of integers per phrase, and thus drops the char following that phrase in S . See section 2.

It is easy to see that the edges outgoing from v_i denote all possible parsing steps that can be taken by any parsing strategy which uses a LZ77-dictionary. Hence, there exists a correspondence between paths from v_1 to v_{n+1} in $\mathcal{G}(S)$ and LZ77-parsings of the whole string S . If we weight every edge $(v_i, v_j) \in E$ with an integer $c(v_i, v_j) = |f(d_{i,j})| + |g(\ell_{i,j})|$, which accounts for the cost of encoding its label (phrase) via the encoding functions f and g , then the length in bits of the encoded parsing is equal to the cost of the corresponding weighted path in $\mathcal{G}(S)$. The problem of determining $\text{OPT}_{f,g}(S)$ is thus reduced to computing the shortest path from v_1 to v_{n+1} in $\mathcal{G}(S)$.

Given that $\mathcal{G}(S)$ is a DAG, its shortest path from v_1 to v_{n+1} can be computed in $O(|E|)$ time and space. However, this is $\Theta(n^2)$ in the worst case (take e.g. $S = a^n$ [30, 20]) thus resulting inefficient and actually unusable in practice even for strings of few Megabytes. In what follows we show that the computation of the SSSP can be restricted to a subgraph of $\mathcal{G}(S)$ whose size depends on the choice of f and g satisfying Property 2.1, and is $O(n \log n)$ for most known integer-encoding functions. Then we will design efficient algorithms and data structures that will allow us to generate this subgraph *on-the-fly* by taking $O(1)$ amortized time per edge and $O(n)$ space overall. These algorithms will be therefore time-and-space optimal for the subgraph in hand, and will provide the first positive answer to Rajpoot-Sahinalp's question we mentioned at the beginning of this paper (see [28, pag. 159]).

5.1 A useful, small, subgraph of $\mathcal{G}(S)$. We use $FS(v)$ to denote the *forward star* of a vertex v , namely the set of vertices pointed to by v in $\mathcal{G}(S)$; and we use $BS(v)$ to denote the *backward star* of v , namely the set of vertices pointing to v in $\mathcal{G}(S)$. Since $\mathcal{G}(S)$ is a DAG, all of its edges (v_i, v_j) are oriented rightward and thus $i < j$. Actually the indices of the vertices in $FS(v)$ and $BS(v)$ form a *contiguous* range:

FACT 5.1. *Given a vertex v_i , it is*

- $FS(v_i) = \{v_{i+1}, \dots, v_{i+x-1}, v_{i+x}\}$ and
- $BS(v_i) = \{v_{i-y}, \dots, v_{i-2}, v_{i-1}\}$.

Furthermore, x, y are smaller than the length of the longest repeated substring in S .

Proof. By definition of (v_i, v_{i+x}) , string $S[i : i+x-1]$ occurs at some position $p < i$ in S . Any prefix $S[i : k-1]$ of $S[i : i+x-1]$ also occurs at that position p , thus $v_k \in FS(v_i)$. The bound on x derives from the definition of (v_i, v_{i+x}) . A similar argument holds for $BS(v_i)$. ■

This means that if an edge does exist in $\mathcal{G}(S)$, then exist also all edges which are *nested* within it and

are incident into one of its extremes. The following property relates the indices of the vertices $v_j \in FS(v_i)$ with the cost of their connecting edge (v_i, v_j) , and not surprisingly shows that the smaller is j (i.e. shorter edge), the smaller is the cost of encoding the phrase $S[i : j-1]$ ³

FACT 5.2. *Given a vertex v_i , for any pair of vertices $v_{j'}, v_{j''} \in FS(v_i)$ such that $j' < j''$, we have that $c(v_i, v_{j'}) \leq c(v_i, v_{j''})$. The same property holds for $v_{j'}, v_{j''} \in BS(v_i)$.*

Proof. We have that $d_{i,j'} \leq d_{i,j''}$ and $\ell_{i,j'} < \ell_{i,j''}$ because $S[i : j'-1]$ is a prefix of $S[i : j''-1]$ and thus the first substring occurs wherever the latter occurs. The property holds because f and g satisfy the Increasing Cost Property 2.1. ■

Given these monotonicity properties, we are ready to characterize a special subset of the vertices in $FS(v_i)$, and their connecting edges.

DEFINITION 1. *An edge $(v_i, v_j) \in E$ is called*

- (1) *d -maximal iff the next edge from v_i takes more bits to encode its distance: $|f(d_{i,j})| < |f(d_{i,j+1})|$.*
- (2) *ℓ -maximal iff the next edge from v_i takes more bits to encode its length: $|g(\ell_{i,j})| < |g(\ell_{i,j+1})|$.*

Edge (v_i, v_j) is called maximal if it is either d -maximal or ℓ -maximal: thus $c(v_i, v_j) < c(v_i, v_{j+1})$.

The number of maximal edges depends on the functions f and g (which satisfy Property 2.1). Let $Q(f, n)$ (resp. $Q(g, n)$) be the number of different codeword lengths generated by f (resp. g) when applied to integers in the range $[n]$. We can partition $[n]$ into contiguous sub-ranges $I_1, I_2, \dots, I_{Q(f,n)}$ such that the integers in I_i are mapped by f to codewords (strictly) shorter than the codewords for the integers in I_{i+1} . Similarly, g partitions the range $[n]$ in $Q(g, n)$ contiguous sub-ranges.

LEMMA 5.1. *There are at most $Q(f, n) + Q(g, n)$ maximal edges outgoing from any vertex v_i .*

Proof. By Fact 5.1, vertices in $FS(v_i)$ have indices in a range R , and by Fact 5.2, $c(v_i, v_j)$ is monotonically non-decreasing as j increases in R . Moreover we know that f (resp. g) cannot change more than $Q(f, n)$ (resp. $Q(g, n)$) times, so that the statement follows. ■

³Recall that $c(v_i, v_j) = |f(d_{i,j})| + |g(\ell_{i,j})|$, if the edge does exist, otherwise we set $c(v_i, v_j) = +\infty$.

To speed up the computation of a SSSP from v_1 to v_{n+1} , we construct a subgraph $\tilde{\mathcal{G}}(S)$ of $\mathcal{G}(S)$ which is formed by maximal edges only, it is smaller than $\mathcal{G}(S)$ and contains one of those SSSP.

THEOREM 5.1. *There exists a shortest path in $\mathcal{G}(S)$ from v_1 to v_{n+1} that traverses maximal edges only.*

Proof. By contradiction assume that every such shortest path contains at least one non-maximal edge. Let $\pi = v_{i_1}v_{i_2}\dots v_{i_k}$, with $i_1 = 1$ and $i_k = n + 1$, be one of these shortest paths, and let $\gamma = v_{i_1}\dots v_{i_r}$ be the longest initial subpath of π which traverses maximal edges only. Assume w.l.o.g. that π is the shortest path maximizing the value of $|\gamma|$. We know that $(v_{i_r}, v_{i_{r+1}})$ is a non-maximal edge, and thus we can take the maximal edge (v_{i_r}, v_j) that has the same cost. By definition of maximal edge, it is $j > i_{r+1}$; furthermore, we must have $j < n + 1$ because we assumed that no path is formed by maximal edges only. Now, since $\mathcal{G}(S)$ is a DAG and indices in π are increasing, it must exist an index $i_h \geq i_r$ such that the index of that maximal edge j lies in $[i_h, i_{h+1}]$. Since $(v_{i_h}, v_{i_{h+1}})$ is an edge of π , it does exist the edge $(v_j, v_{i_{h+1}})$ (by Fact 5.1), and by Fact 5.2 on $BS(v_{i_{h+1}})$ we can conclude that $c(v_j, v_{i_{h+1}}) \leq c(v_{i_h}, v_{i_{h+1}})$. Consequently, the path $v_{i_1}\dots v_{i_r}v_jv_{i_{h+1}}\dots v_{i_k}$ is also a shortest path but its longest initial subpath of maximal edges consists of $|\gamma| + 1$ vertices, which is a contradiction! ■

Theorem 5.1 implies that the distance between v_1 and v_{n+1} is the same in $\mathcal{G}(S)$ and $\tilde{\mathcal{G}}(S)$, with the advantage that computing SSSP in $\tilde{\mathcal{G}}(S)$ can be done faster and in reduced space, because $|FS(v)| \leq Q(f, n) + Q(g, n)$ (Lemma 5.1). Thus, subgraph $\tilde{\mathcal{G}}(S)$ consists of $n+1$ vertices and at most $n(Q(f, n) + Q(g, n))$ edges. For Elias' codes [11], Fibonacci's codes [12], and most practical integer encoders used for search engines and data compressors [29, 35], it is $Q(f, n) = Q(g, n) = O(\log n)$. Therefore $|\tilde{\mathcal{G}}(S)| = O(n \log n)$, so it is smaller than the complete graph built and used by previous papers [30, 20, 10]. For the encoders used in `gzip`, it is $Q(f, n) = Q(g, n) = O(1)$ and $|\tilde{\mathcal{G}}(S)| = O(n)$.

5.2 An efficient bit-optimal parser. From a high level, our solution is a variant of a classic linear-time algorithm for SSSP over a DAG (see [7, Section 24.2]), here applied to work on the subgraph $\tilde{\mathcal{G}}(S)$. Therefore its correctness follows directly from Theorem 24.5 of [7] and our Theorem 5.1. However, the key difficulty in implementing this approach consists of how to *generate on-the-fly and efficiently (in time and space) the maximal edges outgoing from vertex v_i* . We will refer to this problem as the *forward-star generation* problem, and

use FSG for brevity. In what follows we consider the case $\sigma \leq n$, and show that FSG takes $O(1)$ amortized time per edge and $O(n)$ space in total. In case of a larger alphabet, we need to add $T_{\text{sort}}(n, \sigma)$ time because of the sorting/remapping of S 's symbols into $[n]$. Since we have n vertices, with no more than $Q(f, n) + Q(g, n)$ maximal edges each (Lemma 5.1), we will obtain the following:

THEOREM 5.2. *Given a string $S[1, n]$ drawn from an alphabet of size σ , and two integer-encoding functions f and g that satisfy Property 2.1, there exists a compressor that computes the (f, g) -optimal LZ77-parsing of S in $O(n(Q(f, n) + Q(g, n)) + T_{\text{sort}}(n, \sigma))$ time and $O(n)$ space in the worst case.*

We know that the edges outgoing from v_i can be partitioned into no more than $Q(f, n)$ groups according to the distance from $S[i]$ of the copied string they represent (proof of Lemma 5.1). Let $I_1, I_2, \dots, I_{Q(f, n)}$ be the intervals of distances such that all distances in I_k are encoded with the same number of bits by f . Take now the d -maximal edge (v_i, v_{h_k}) for the interval I_k . We can infer that substring $S[i : h_k - 1]$ is the *longest* substring having a copy at distance within I_k because, by Definition 1 and Fact 5.2, any edge following (v_i, v_{h_k}) denotes a longer substring which must lie in a subsequent interval (by d -maximality of (v_i, v_{h_k})), and thus must have longer distance from $S[i]$. Once d -maximal edges are known, the computation of the ℓ -maximal edges is then easy because it suffices to further decompose the edges between successive d -maximal edges, say between $(v_i, v_{h_{k-1}+1})$ and (v_i, v_{h_k}) , according to the distinct values assumed by the encoding function g on the lengths in the range $[h_{k-1}, \dots, h_k - 1]$. This takes $O(1)$ time per ℓ -maximal edge, because it needs some algebraic calculations, and the corresponding copied substring can then be inferred as a prefix of $S[i : h_k - 1]$ (details in the full paper).

So, let us concentrate on the computation of d -maximal edges outgoing from vertex v_i . We remark that we could use the solution proposed in [16] on each of the $Q(f, n)$ ranges of distances in which a phrase copy can be found. Unfortunately, this approach would pay another multiplicative factor $\log \sigma$ per symbol and its space complexity would be super-linear in n . Conversely, our solution overcomes these drawbacks by deploying two key ideas:

(1) The first idea aims at minimizing the space usage by achieving the optimal $O(n)$ working-space bound. It consists of proceeding in $Q(f, n)$ passes, one per interval I_k of possible d -costs for the edges in $\tilde{\mathcal{G}}(S)$. During the k th pass, we logically partition the vertices of $\tilde{\mathcal{G}}(S)$ in blocks of $|I_k|$ contiguous vertices, say

$v_{i_k}, v_{i_k+1}, \dots, v_{i_k+|I_k|-1}$, and compute all d -maximal edges which spread out from that block and have copy-distance within I_k (thus they all have the same d -cost, say $c(I_k)$). These edges are kept in memory until they are used by `Optimal-Parser`, and discarded as soon as the first vertex of the next block, i.e. $v_{i_k+|I_k|}$, needs to be processed. The next block of $|I_k|$ vertices is then fetched and the process repeats. Actually, all passes are executed in *parallel* to guarantee that all d -maximal edges of v_i are available when processing this vertex. There are $n/|I_k|$ distinct blocks at each pass, and each d -maximal edge of a vertex is considered in some pass (because it has d -cost in some I_k). The space is $\sum_{k=1}^{Q(f,n)} |I_k| = O(n)$ because we keep one d -maximal edge per vertex at any pass.

(2) The second key idea aims at computing the d -maximal edges for that block of $|I_k|$ contiguous vertices in $O(|I_k|)$ time and space. This is what we address below, being the most sophisticated technicality of our solution. As a result, we show that the time complexity of `FSG` is $\sum_{k=1}^{Q(f,n)} (n/|I_k|)O(|I_k|) = O(n Q(f, n))$, i.e., $O(1)$ amortized time per d -maximal edge. Combining this fact with the previous observation on the computation of the ℓ -maximal edges, we get Theorem 5.2 above.




Figure 1: Interval $B = [i, j]$ with $j = i + |I_k| - 1$, window \mathcal{W}_B and its two halves $\mathcal{W}'_B, \mathcal{W}''_B$.

Let us assume that the alphabet size is $\sigma \leq n$, and consider the k th pass of `FSG` in which we assume that $I_k = [l, r]$. Recall that all distances in I_k can be f -encoded in the same number of, say, $c(I_k)$ bits. Let $B = [i, i + |I_k| - 1]$ be the block of (indices of) vertices for which we wish to compute *on-the-fly* the d -maximal edges of cost $c(I_k)$. This means that the d -maximal edge from vertex v_h , $h \in B$, represents a phrase that starts at $S[h]$ and has a copy starting in the *window* $W_h = [h - r, h - l]$. Thus the distance of that copy can be f -encoded in $c(I_k)$ bits, and so we will say that the edge has d -cost $c(I_k)$. Since this computation must be done for all vertices in B , it is useful to consider the window $\mathcal{W}_B = W_i \cup W_{i+|I_k|-1}$ which merges the first and last window and thus spans all positions that can be the (copy-)reference of any d -maximal edge outgoing from B . Note that $|\mathcal{W}_B| = 2|I_k|$ (see Figure 1). The

next fact is crucial to fast compute all these d -maximal edges via indexing data structures built over S :

FACT 5.3. *If there exists a d -maximal edge outgoing from v_h and having d -cost $c(I_k)$, then this edge can be found by determining a position $s \in W_h$ whose suffix S_s shares the maximum longest common prefix (shortly, `lcp`) with S_h .*

Proof. Among all positions s in W_h take one whose suffix S_s shares the maximum `lcp` with S_h , and let q be the length of this `lcp`. Of course, there may exist many such positions, we take just one of them. The edge (v_h, v_{h+q+1}) has d -cost $c(I_k)$, because $s \in W_h$, and is d -maximal because any other position $s' \in W_h$ induces an edge $(v_h, v_{h+q'+1})$ whose length $q' \leq q$, by maximality of q . ■

Hereafter we call the position s of Fact 5.3 *maximal position* for vertex v_h , and note that it does exist only if v_h has a d -maximal edge of cost $c(I_k)$. Our algorithm will compute the maximal positions of every vertex v_h in B whenever they do exist, otherwise it will assign an *arbitrary* position to v_h . The net result is that we will generate a *supergraph* of $\tilde{\mathcal{G}}(S)$ which is still guaranteed to have the size stated in Lemma 5.1 and can be created efficiently in $O(|I_k|)$ time and space, as we required above.

Fact 5.3 has related the computation of maximal positions for the vertices in B to `lcp`-computations between suffixes in B and suffixes in \mathcal{W}_B . Therefore it is natural to resort some indexing data structure, like the compact trie \mathcal{T}_B , built over the suffixes of S which start in the range of positions $B \cup \mathcal{W}_B$. Trie \mathcal{T}_B takes $O(|B| + |\mathcal{W}_B|) = O(|I_k|)$ space, and this bound is within our required space complexity. It is not easy to build \mathcal{T}_B in $O(|I_k|)$ time and space, because this time complexity is *independent* of the length of the indexed suffixes and the alphabet size. The proof of this result, deferred to full paper, deploy the fact that the algorithm we detail below does not make any assumption on the edge-ordering of \mathcal{T}_B , because it just computes (sort of) `lca`-queries on its structure.

Given the trie \mathcal{T}_B , we notice that the maximal position s for a vertex v_h in B having d -cost $c(I_k)$ can be computed by *finding the leaf of \mathcal{T}_B which is labeled with an index $s \in W_h$ and has the deepest lowest common ancestor (shortly, `lca`) with the leaf labeled h* . We need to answer this query in $O(1)$ amortized time per vertex v_h , since we aim at achieving an $O(|I_k|)$ time complexity over all vertices in B . This is not easy because this is *not* the classic `lca`-query since we do not know s , which is actually the position we are searching for! Furthermore, since the leaf s is the closest one to h in \mathcal{T}_B

among the leaves with index in W_h , one could think to use proper predecessor/successor queries on a suitable dynamic set of suffixes in W_h . Unfortunately, this would take $\omega(1)$ time because of well-known lower bounds [2]. Therefore, in order to answer this query in constant (amortized) time per vertex of B , we deploy proper structural properties of the trie \mathcal{T}_B and the problem at hand.

Let u be the lca of the leaves labeled h and s in \mathcal{T}_B . For simplicity, we assume that the window W_h strictly precedes B and that s is the unique maximal position for v_h (our algorithm deals with these cases too, details deferred to full paper) We observe that h must be the smallest index that lies in B and labels a leaf descending from u in \mathcal{T}_B . In fact assume, by contradiction, that a smaller index $h' < h$ does exist. By definition $h' \in B$ and thus v_h would not have a d -maximal edge of d -cost $c(I_k)$ because it could copy from the closer h' a possibly longer phrase, instead of copying from the farther set of positions in W_h . This observation implies that we have to search only for one maximal position per node u of \mathcal{T}_B , and this position refers to the vertex $v_{a(u)}$ whose index $a(u)$ is the smallest one that lies in B and labels a leaf descending from u . Computing a -values clearly takes $O(|\mathcal{T}_B|) = O(|I_k|)$ time and space via a traversal of the trie \mathcal{T}_B .

Now we need to compute the maximal position for $v_{a(u)}$, for each node $u \in \mathcal{T}_B$. We cannot traverse the subtree of u searching for the maximal position for $v_{a(u)}$, because this would take quadratic time complexity overall. Conversely, we define \mathcal{W}'_B and \mathcal{W}''_B to be the first and the second half of \mathcal{W}_B , respectively, and observe that any window W_h has its left extreme in \mathcal{W}'_B and its right extreme in \mathcal{W}''_B (see Figure 1). Therefore the window $W_{a(u)}$ containing the maximal position s for $v_{a(u)}$ overlaps both \mathcal{W}'_B and \mathcal{W}''_B . If s does exist for $v_{a(u)}$, then s belongs to either \mathcal{W}'_B or to \mathcal{W}''_B , and the leaf labeled s descends from u . Hence the maximum (resp. minimum) among the elements in \mathcal{W}'_B (resp. \mathcal{W}''_B) that label leaves descending from u must belong to $W_{a(u)}$.

This suggests to compute for each node u the rightmost position in \mathcal{W}'_B and the leftmost position in \mathcal{W}''_B that label a leaf descending from u , denoted respectively by $\max(u)$ and $\min(u)$. This takes $O(|I_k|)$ time with a post-order visit of \mathcal{T}_B . We can now efficiently compute $\text{mp}[h]$ as the maximal position for v_h , if it exists, or otherwise set $\text{mp}[h]$ arbitrarily. We initially set all mp 's entries to nil ; then we visit \mathcal{T}_B in post-order and perform, at each node u , the following two checks whenever $\text{mp}[a(u)] = \text{nil}$: If $\min(u) \in W_{a(u)}$, we set $\text{mp}[a(u)] = \min(u)$; if $\max(u) \in W_{a(u)}$, we set $\text{mp}[a(u)] = \max(u)$. At the end of the visit, if $\text{mp}[a(u)]$ is still nil we set $\text{mp}[a(u)] = a(\text{parent}(u))$ whenever

$a(u) \neq a(\text{parent}(u))$. This last check is needed to manage the case in which $S[a(u)]$ can copy the phrase starting at its position from position $a(\text{parent}(u))$ and, additionally, we have that B overlaps \mathcal{W}_B (which may occur depending on f). Since \mathcal{T}_B has size $O(|I_k|)$, the overall algorithm requires $O(|I_k|)$ time and space in the worst case, and hence Theorem 5.2 follows. Correctness follows from lemma below (proof in the full paper).

LEMMA 5.2. *For each position $h \in B$, if there exists a d -maximal edge outgoing from v_h and having d -cost $c(I_k)$, then $\text{mp}[h]$ is equal to its maximal position.*

6 Conclusions

Our bit-optimal parsing scheme can be extended to variants of LZ77 which deploy parsers that refer to a bounded compression-window (the typical scenario of `gzip` and its derivatives [29]). In this case, LZ77 selects the next phrase by looking only at the most recent M input symbols. Since M is usually a constant of few Kbs [29], the running time of our algorithm is $O(n Q(g, n))$, given that $Q(f, M)$ turns out to be a constant. This complexity could be further refined as $O(n Q(g, \ell))$ by considering the length ℓ of the longest repeated substring in S . Furthermore, if S is generated by an ergodic source [32] and g is taken to be the classic Elias' code, we have $Q(g, \ell) = O(\log \log n)$ so that the complexity of our algorithm results $O(n \log \log n)$ time and $O(n)$ space for this class of strings.

We finally notice that, although we have mainly dealt with the LZ77-dictionary, the techniques presented in this paper can be extended to design efficient bit-optimal compressors for other on-line dictionary construction schemes, like LZ78 (details in the full paper). The main open question we leave with this paper is to extend our results to statistical encoding functions like Huffman or Arithmetic coders applied on the integral range $1 \dots n$ [35]. They do not necessarily satisfy Property 2.1 because it might be the case that $|f(x)| > |f(y)|$, whenever the integer y occurs more frequently than the integer x in the parsing of S . We argue that it is not trivial to design a bit-optimal compressor for these encoding functions because their codeword lengths change as it changes the set of distances and lengths used in the parsing process.

Acknowledgments

Part of this work was done while the third author was visiting the Max-Planck-Institut für Informatik in Saarbrücken.

References

- [1] A. Amir, G. M. Landau, and E. Ukkonen. On-line timestamped text indexing. *Inf. Process. Lett.*, 82(5):253–259, 2002.
- [2] P. Beame and F. E. Fich. Optimal bounds for the predecessor problem. In *Procs ACM STOC*, pages 295–304, 1999.
- [3] J. Bekesi, G. Galambos, U. Pferschy, and G.J. Woeginger. Greedy algorithms for on-line data compression. *Journal of Algorithms*, 25(2):274–289, 1997.
- [4] P. Boldi, B. Codenotti, M. Santini, and S. Vigna. Ubcrawler: A scalable fully distributed web crawler. *Software: Practice & Experience*, 34(8):711–726, 2004.
- [5] R. Cilibrasi and P. M. B. Vitányi. Clustering by compression. *IEEE Transactions on Information Theory*, 51(4):1523–1545, 2005.
- [6] M. Cohn and R. Khazan. Parsing with prefix and suffix dictionaries. In *Procs IEEE DCC*, pages 180–189, 1996.
- [7] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Second Edition*. The MIT Press and McGraw-Hill Book Company, 2001.
- [8] G. Cormode and S. Muthukrishnan. Substring compression problems. In *Procs ACM-SIAM SODA*, pages 321–330, 2005.
- [9] M. Crochemore, L. Ilie, and W. F. Smyth. A simple algorithm for computing the Lempel-Ziv factorization. In *Procs IEEE DCC*, pages 482–488, 2008.
- [10] G. Della Penna, F. Mignosi, A. Langiu, and A. Ulisse. On dictionary-symbolwise data compression. Personal Communication, 2008.
- [11] P. Elias. Universal codeword sets and representations of the integers. *IEEE Transactions on Information Theory*, 21(2):194–203, 1975.
- [12] P. Fenwick. Universal codes. In *Lossless Compression Handbook*, pages 55–78. Academic Press, 2003.
- [13] P. Ferragina, R. Giancarlo, and G. Manzini. The engineering of a compression boosting library: Theory vs practice in bwt compression. In *Procs ESA*, LNCS 4168, Springer, pages 756–767, 2006.
- [14] P. Ferragina, R. Giancarlo, G. Manzini, and M. Sciortino. Boosting textual compression in optimal linear time. *Journal of the ACM*, 52:688–713, 2005.
- [15] P. Ferragina and G. Navarro. Site pizza&chili (<http://pizzachili.di.unipi.it> or <http://pizzachili.dcc.uchile.cl>).
- [16] E. R. Fiala and D. H. Greene. Data compression with finite windows. *Communications of the ACM*, 32(4):490–505, 1989.
- [17] T. Gagie and G. Manzini. Space-conscious compression. In *Procs MFCS*, LNCS 4708, Springer, pages 206–217, 2007.
- [18] R. F. Geary, R. Raman, and V. Raman. Succinct ordinal trees with level-ancestor queries. *ACM Transactions on Algorithms*, 2(4):510–534, 2006.
- [19] Gzip home page. <http://www.gzip.org>.
- [20] J. Katajainen and T. Raita. An approximation algorithm for space-optimal encoding of a text. *Computer Journal*, 32(3):228–237, 1989.
- [21] J. Katajainen and T. Raita. An analysis of the longest match and the greedy heuristics in text encoding. *Journal of the ACM*, 39(2):281–294, 1992.
- [22] S. T. Klein. Efficient optimal recompression. *Computer Journal*, 40(2/3):117–126, 1997.
- [23] R. Kosaraju and G. Manzini. Compression of low entropy strings with Lempel-Ziv algorithms. *SIAM Journal on Computing*, 29(3):893–911, 1999.
- [24] M. Lewenstein, V., Mäkinen, and S. J. Puglisi. Personal communications.
- [25] Y. Matias and S.C. Şahinalp. On the optimality of parsing in dynamic dictionary based data compression. In *Procs ACM-SIAM SODA*, pages 943–944, 1999.
- [26] E. M. McCreight. Priority search trees. *SIAM Journal on Computing*, 14(2):257–276, 1985.
- [27] G. Navarro and V. Mäkinen. Compressed full-text indexes. *ACM Computing Surveys*, 39(1), 2007.
- [28] N. Rajpoot and C. Sahinalp. *Handbook of Lossless Data Compression*, chapter Dictionary-based data compression, pages 153–167. Academic Press, 2002.
- [29] D. Salomon. *Data Compression: the Complete Reference, 4th Edition*. Springer Verlag, 2006.
- [30] E. J. Schuegraf and H. S. Heaps. A comparison of algorithms for data base compression by use of fragments as language elements. *Information Storage and Retrieval*, 10(9-10):309319, 1974.
- [31] M. E. Gonzalez Smith and J. A. Storer. Parallel algorithms for data compression. *Journal of the ACM*, 32(2):344–373, 1985.
- [32] W. Szpankowski. Asymptotic properties of data compression and suffix trees. *IEEE Transactions on Information Theory*, 39(5):1647–1659, 1993.
- [33] J. S. Vitter and P. Krishnan. Optimal prefetching via data compression. *Journal of the ACM*, 43(5):771–793, 1996.
- [34] D. E. Willard. Examining computational geometry, van emde boas trees, and hashing from the perspective of the fusion tree. *SIAM Journal on Computing*, 29(3), 2000.
- [35] I. H. Witten, A. Moffat, and T. C. Bell. *Managing Gigabytes: Compressing and Indexing Documents and Images*. Morgan Kaufmann Publishers, Los Altos, CA 94022, USA, second edition, 1999.
- [36] J. Ziv. Classification with finite memory revisited. *IEEE Transactions on Information Theory*, 53(12):4413–4421, 2007.
- [37] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Transaction on Information Theory*, 23:337–343, 1977.
- [38] J. Ziv and A. Lempel. Compression of individual sequences via variable-rate coding. *IEEE Transactions on Information Theory*, 24(5):530–536, 1978.