

BAT-LZ Out of Hell

Zsuzsanna Lipták ✉ 

Dipartimento di Informatica, University of Verona, Italy

Francesco Masillo ✉ 

Dipartimento di Informatica, University of Verona, Italy

Gonzalo Navarro ✉ 

Center for Biotechnology and Bioengineering (CeBiB)

Department of Computer Science, University of Chile, Chile

Abstract

Despite consistently yielding the best compression on repetitive text collections, the Lempel-Ziv parsing has resisted all attempts at offering relevant guarantees on the cost to access an arbitrary symbol. This makes it less attractive for use on compressed self-indexes and other compressed data structures. In this paper we introduce a variant we call BAT-LZ (for Bounded Access Time Lempel-Ziv) where the access cost is bounded by a parameter given at compression time. We design and implement a linear-space algorithm that, in time $O(n \log^3 n)$, obtains a BAT-LZ parse of a text of length n by greedily maximizing each next phrase length. The algorithm builds on a new linear-space data structure that solves 5-sided orthogonal range queries in rank space, allowing updates to the coordinate where the one-sided queries are supported, in $O(\log^3 n)$ time for both queries and updates. This time can be reduced to $O(\log^2 n)$ if $O(n \log n)$ space is used.

We design a second algorithm that chooses the sources for the phrases in a clever way, using an enhanced suffix tree, albeit no longer guaranteeing longest possible phrases. This algorithm is much slower in theory, but in practice it is comparable to the greedy parser, while achieving significantly superior compression. We then combine the two algorithms, resulting in a parser that always chooses the longest possible phrases, and the best sources for those. Our experimentation shows that, on most repetitive texts, our algorithms reach an access cost close to $\log_2 n$ on texts of length n , while incurring almost no loss in the compression ratio when compared with classical LZ-compression. Several open challenges are discussed at the end of the paper.

2012 ACM Subject Classification Theory of computation → Data structures design and analysis

Keywords and phrases Lempel-Ziv parsing, data compression, compressed data structures, repetitive text collections

Digital Object Identifier 10.4230/LIPIcs.CPM.2024.

Supplementary Material The code is available at <https://github.com/fmasillo/BAT-LZ>

Funding *Zsuzsanna Lipták*: Partially funded by the MUR PRIN project Nr. 2022YRB97K 'PINC' (Pangenome INformatiCs. From Theory to Applications) and by the INdAM-GNCS Project CUP_E53C23001670001 (Compressione, indicizzazione, analisi e confronto di dati biologici).

Gonzalo Navarro: Funded by Basal Funds FB0001, Mideplan, Chile, and Fondecyt Grant 1-230755, Chile.



© Zsuzsanna Lipták, Francesco Masillo, Gonzalo Navarro;
licensed under Creative Commons License CC-BY 4.0

35th Annual Symposium on Combinatorial Pattern Matching (CPM 2024).



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The sharply growing sizes of text collections, particularly repetitive ones, has raised the interest in compressed data structures that can maintain the texts all the time in compressed form [43, 42, 41]. For archival purposes, the original Lempel-Ziv (LZ) compression format [36] is preferred because it yields the least space among the methods that support compression and decompression in polynomial time—actually, Lempel-Ziv compresses and decompresses a text $T[1..n]$ in $O(n)$ time [48]. For using a compression format as a compressed data structure, however—in particular, to build a compressed text self-index on it [34]—, we need that arbitrary text snippets $T[i..i + \ell]$ can be extracted efficiently, without the need of decompressing the whole text up to the desired snippet. Grammar compression formats [31] allow extracting such text snippets in time $O(\ell + \log n)$ [5, 22], which is nearly optimal [51]. So, although the compression they achieve is always lower-bounded by the size of the LZ parse [49, 8], grammar compression algorithms are preferred over LZ compression in the design of text indexes [42, 12], and of compressed data structures in general.

The LZ compression algorithm parses the text T into a sequence of so-called phrases, where each phrase points backwards to a previous occurrence of it in T and stores the next symbol in explicit form. While this yields a simple linear-time left-to-right decompression algorithm, consider the problem of accessing a particular symbol $T[i]$. Unless it is the final explicit symbol of a phrase, we must determine the text position $j < i$ where $T[i] = T[j]$ was copied from. We must then determine $T[j]$, which again may be—with low chance—the end of a phrase, or it may—most likely—refer to an earlier symbol $T[j] = T[k]$, with $k < j$. The process continues until we hit an explicit symbol. The cost of extracting $T[i]$ is then proportional to the length of that *referencing chain* $i \rightarrow j \rightarrow k \rightarrow \dots$. Despite considerable interest in algorithms to access arbitrary text positions from the LZ compression format, and apart from some remarkable results on restricted versions of LZ [30], there has been no progress on the original LZ parse (which yields the strongest compression).

In this paper we introduce and study an LZ variant we call *Bounded Access Time Lempel-Ziv (BAT-LZ)*, which takes a compression parameter c and produces a parse where no symbol has a referencing chain longer than c , thereby guaranteeing $O(c)$ access time.¹ As opposed to classical LZ, BAT-LZ parses allow very fast access to the text, indeed, like a bat out of hell.

We design a *Greedy BAT-LZ parser*, which at each step of the compression chooses the longest possible phrase. Finding such a phrase boils down to solving a 4-sided orthogonal range query in a 3-dimensional grid (in rank space), where one of the coordinates undergoes updates as the parsing proceeds. We design such a data structure, which turns out to handle 5-sided queries and support updates on the coordinate where the query is one-sided. Our data structure handles queries and updates in time $O(\log^3 n)$, yielding a greedy BAT-LZ parsing in time $O(n \log^3 n)$ and space $O(n)$. We then design another BAT-LZ parser, referred to as *Minmax*, which runs on an enhanced suffix tree. It looks for the “best” possible sources of the chosen phrases, that is, with symbols having shorter referencing chains, while not necessarily choosing the longest possible phrase. Finally, we combine the two ideas, resulting in our *Greedier parser*, which runs again on an enhanced suffix tree. These last two algorithms, while their running time is upper bounded by $O(n^3 \log n)$, both run in decent time in practice.

We implemented and tested our three BAT-LZ parsers on various repetitive texts of different sorts, comparing them with the original LZ parse and with two simple baselines that

¹ A parsing like BAT-LZ was described as a baseline in the experimental results in previous work [33] of one of the authors, but without a parsing algorithm, see Sec. 3 for more details.

82 ensure BAT-LZ parses without any optimization. The results show that all three algorithms
 83 run in a few seconds per megabyte and produce much better parses than the baselines. For
 84 values of $c = O(\log n)$ with a small constant, they produce just a small fraction of extra
 85 phrases on top of LZ. In particular, Greedier increases the size of the LZ parse by less than
 86 1% with c values that are about $\log_2 n$ (i.e., 20–30 in our texts).

87 We note that, unlike the original LZ parse, a greedy parsing does not guarantee obtaining
 88 the minimal BAT-LZ parse. Indeed, finding the optimal BAT-LZ parse has recently been
 89 shown to be NP-hard for all constant c , and also hard to approximate for any constant
 90 approximation ratio [10]. Our results show that, on repetitive texts, a polylog-linear time
 91 greedy algorithm can nonetheless achieve good compression while guaranteeing fast access to
 92 text snippets. The other two algorithms are still polynomial time and offer fast access with
 93 almost no loss in compression compared to the classical LZ-compression. In our scenarios of
 94 interest (i.e., accessing the compressed text at random) the data is compressed only once
 95 and accessed many times, so slower compression algorithms can be afforded in exchange for
 96 faster access. We discuss at the end this and some other problems our work opens.

97 2 Basic Data Structures

98 A string (or text) T is a finite sequence of characters from an alphabet Σ . We write
 99 $T = T[1..n]$ for a string T of length n , and assume that the final character is a unique
 100 end-of-string marker $\$$. We index strings from 1 and write $T[i..j]$ for the substring $T[i]..T[j]$,
 101 $T[i..]$ for the suffix starting in position i , and $T[..j]$ for the prefix ending in position j .

102 **Bitvectors and Wavelet Matrices.** A bitvector $B[1..n]$ can be stored using n bits, or
 103 actually $\lceil n/w \rceil$ words on a w -bit word machine, while providing access and updates to
 104 arbitrary bits in constant time. If the bitvector is static (i.e., does not undergo updates) then
 105 it can be preprocessed to answer *rank* queries in $O(1)$ time using $o(n)$ further bits [11, 39]:
 106 $rank_b(B, i)$, where $b \in \{0, 1\}$ and $0 \leq i \leq n$, is the number of times bit b occurs in $B[1..i]$.

107 A wavelet matrix [13] is a data structure that can be used, in particular, to represent a
 108 discrete $[1, n] \times [1, n]$ grid, with exactly one point per column, using $n \log_2 n + o(n \log_2 n)$
 109 bits. Let $S[1..n]$ be such that $S[i]$ is the row of the point at column i . The first wavelet
 110 matrix level contains a bitvector $B_1[1..n]$ with the highest (i.e., $\lceil \log_2 n \rceil$ th) bit of every
 111 value in S . For the second level, the sequence values are stably sorted by their highest bit,
 112 and the wavelet matrix stores a bitvector $B_2[1..n]$ with the second highest bits in that order.
 113 To build the third level, the values are stably sorted by their second highest bit, and so on.
 114 Every level i also stores the number $z_i = rank_0(B_i, n)$ of zeros in its bitvector.

115 The value $S[i]$ can be retrieved from the wavelet matrix in $O(\log n)$ time. Its highest bit
 116 is $b_1 = B_1[i_1]$, with $i_1 = i$. The second highest bit is $b_2 = B_2[i_2]$, with $i_2 = rank_0(B_1, i_1)$ if
 117 $b_1 = 0$ and $i_2 = z_1 + rank_1(B_1, i_1)$ if $b_1 = 1$. The other bits are obtained analogously.

118 The wavelet matrix can also obtain the grid points that fall within a rectangle $[x_1, x_2] \times$
 119 $[y_1, y_2]$ (i.e., the values $(i, S[i])$ such that $x_1 \leq i \leq x_2$ and $y_1 \leq S[i] \leq y_2$) in time $O(\log n)$,
 120 plus $O(\log n)$ per point reported. We start at the first level, in the range $B_1[sp_1, ep_1] =$
 121 $B_1[x_1, x_2]$. We then map the range into two ranges of the second level: the positions i where
 122 $B_1[i] = 0$ are all mapped to the range $B_2[sp_2, ep_2] = B_2[rank_0(B_1, sp_1 - 1) + 1, rank_0(B_1, ep_1)]$,
 123 and those where $B_1[i] = 1$ are mapped to $B_2[sp'_2, ep'_2] = B_2[z_1 + rank_1(B_1, sp_1 - 1) + 1, z_1 +$
 124 $rank_1(B_1, ep_1)]$. The recursive process stops when the range becomes empty; when the
 125 sequence of highest bits makes the possible set of values either disjoint with $[y_1, y_2]$ or
 126 included in $[y_1, y_2]$; or when we reach the last level. It can be shown that the recursion ends

127 in $O(\log n)$ ranges, at most two per level, so that every value in those ranges is an answer.
 128 The corresponding y values can be obtained by tracking them downwards as explained.

129 These data structures, and our results, hold in the RAM model with computer word size
 130 $w = \Theta(\log n)$. The wavelet matrix is then said to use $O(n)$ space—i.e., linear space—, which
 131 is counted in w -bit words. The wavelet matrix is easily built in $O(n \log n)$ time, and less [40].

132 Another relevant functionality that can be offered within $2n + o(n)$ bits is the so-called
 133 *range maximum query (RMQ)*: given a static array $A[1..n]$, we preprocess it in $O(n)$ time
 134 so that we can answer RMQs in $O(1)$ time [19]: $rmq(A, i, j)$ is a position p , $i \leq p \leq j$, such
 135 that $A[p] = \max\{A[k], i \leq k \leq j\}$. The data structure does not need to maintain A . In this
 136 paper we will use RMQs where A can undergo updates, see Sec. 5.

137 **Suffix Arrays and Trees.** The suffix tree [52] is a classic data structure on texts which
 138 is able to answer efficiently many different kinds of string processing queries [24, 1], which
 139 uses linear space and can be built in linear time [52, 38, 17, 50]. We give a brief recap; see
 140 Gusfield [24] for more details.

141 The suffix tree $ST(T)$ of a text T is the compact trie of the suffixes of T ; it is a rooted
 142 tree whose edges are labeled by substrings of T (stored as two pointers into T), and whose
 143 inner nodes are branching. The *label* $L(v)$ of a node v is the concatenation of the labels
 144 of the edges on the root-to- v path. There is a one-to-one correspondence between leaves
 145 and suffixes of T ; $leaf_i$ is then the unique leaf whose label equals the i th suffix $T[i..]$. The
 146 *stringdepth* $sd(v)$ of a node v is the length of its label, and we assume $sd(v)$ is stored in v .

147 The suffix array SA of T is a permutation of the index set $\{1, \dots, n\}$ such that $SA[i] = j$
 148 if the j th suffix of T is the i th in lexicographic order among all suffixes. The suffix array can
 149 be computed from the suffix tree, or directly from the text, in linear time and space [47, 45].
 150 The inverse suffix array, denoted ISA , can be computed in linear time using $ISA[SA[i]] = i$.

151 **3 The Lempel-Ziv (LZ) Parsing and its Bounded Version (BAT-LZ)**

152 The Lempel-Ziv (LZ) parsing of a text $T[1..n]$ [36] produces a sequence of z “phrases”,
 153 which are substrings of T whose concatenation is T . Each phrase is formed by the longest
 154 substring that has an occurrence starting earlier in T , plus the character that follows it.

155 **► Definition 1.** A leftward parse of $T[1..n]$ is a sequence of substrings $T[i..i+\ell]$ (called
 156 phrases) whose concatenation is T and such that there is an occurrence of each $T[i..i+\ell-1]$
 157 starting before i in T (the occurrence is called the source of the phrase). The LZ parse of T
 158 is the leftward parse of T that, in a left-to-right process, chooses the longest possible phrases.

159 The algorithm moves a pointer i along T , from $i = 1$ to $i = n$. At each step, the algorithm
 160 has already processed $T[1..i-1]$, and it must form the next phrase. As said, the phrase is
 161 formed by (1) the longest prefix $T[i..i+\ell-1]$ of $T[i..]$ that has an occurrence in T starting
 162 before position i , and (2) the next symbol $T[i+\ell]$. If $\ell > 0$, then the occurrence of (1),
 163 $T[s..s+\ell-1] = T[i..i+\ell-1]$ with $s < i$, is called the source of $T[i..i+\ell-1]$. Once suitable
 164 s and ℓ have been determined, the next phrase is $T[i..i+\ell]$ and the algorithm proceeds
 165 from $i \leftarrow i + \ell + 1$ onwards. The phrase $T[i..i+\ell]$ is encoded as the triple $(s, \ell, T[i+\ell])$,
 166 and if $\ell = 0$ we can encode just the character $(T[i+\ell])$.

167 This greedy parsing, which maximizes the phrase length at each step, turns out to be
 168 optimal [36], that is, it produces the least number z of phrases among all the leftward parses
 169 of T . Further, it can be computed in $O(n)$ time [48, 9, 46, 25, 26, 23, 20, 32, 3, 27, 21].

170 Note that phrases can overlap their sources, as sources must start—but not necessarily
 171 end—before i . For example, the LZ parse of $T = \mathbf{a}^{n-1}\$$ is $(\mathbf{a}) (0, n-1, \$)$. For illustrative

XX:4 BAT-LZ Out of Hell

172 purposes, we describe the parsings by writing bars, “|”, between the formed phrases. The
 173 parsing of the example is then written as $a|a^{n-1}\$$. To illustrate the access problem, consider
 174 the LZ parsing of the text `alabaralalabarda$` (disregard for now the numbers below):

$$175 \quad \begin{array}{c} a \mid 1 \mid a \ b \mid a \ r \mid a \ 1 \ a \ 1 \mid a \ b \ a \ r \ d \mid a \ \$ \mid \\ 0 \mid 0 \mid 1 \ 0 \mid 1 \ 0 \mid 1 \ 1 \ 2 \ 0 \mid 2 \ 1 \ 2 \ 1 \ 0 \mid 1 \ 0 \end{array}$$

176 Assume we want to extract $T[11] = a$. The position is the first of the 6th phrase, `abard`,
 177 and it is copied from the third phrase, `ab`. In turn, the first position of that phrase is copied
 178 from the first phrase, where `a` is stored in explicit form. We need then to follow a *chain*
 179 of length two in order to extract $T[11]$, so the length of that chain is the access cost. The
 180 numbers we wrote below the symbols in the parse are the lengths of their chains.

181 **Bounded Access Time Lempel-Ziv (BAT-LZ).** We define a leftward parse we call Bounded
 182 Access Time Lempel-Ziv (BAT-LZ), which takes as a parameter the maximum length c any
 183 chain can have. A BAT-LZ parse is a leftward parse where no chain is longer than c . Note
 184 that we do not require a BAT-LZ parse to be of minimal size. For example, a BAT-LZ parse
 185 for the above text with $c = 1$ is as follows:

$$186 \quad \begin{array}{c} a \mid 1 \mid a \ b \mid a \ r \mid a \ 1 \ a \mid 1 \ a \mid b \ a \mid r \ d \mid a \ \$ \mid \\ 0 \mid 0 \mid 1 \ 0 \mid 1 \ 0 \mid 1 \ 1 \ 0 \mid 1 \ 0 \mid 1 \ 0 \mid 1 \ 0 \mid 1 \ 0 \end{array}$$

187 When the LZ parse produces the phrase $T[i..i+\ell]$ from the source $T[s..s+\ell-1]$ and
 188 the extra symbol $T[i+\ell]$, the character $T[i+\ell]$ is stored in explicit form, and thus its chain
 189 is of length zero. The chain length of every other phrase symbol, $T[i+l]$ for $0 \leq l < \ell$, is one
 190 more than the chain length of its source symbol, $T[s+l]$.

191 A special case occurs when sources and targets overlap. If we want to extract $T[n-1]$
 192 from $T = a^{n-1}\$$, we could note that it is copied from $T[n-2]$, which is in turn copied
 193 from $T[n-3]$, and so on, implying a chain of length $n-1$. Instead, we can note that our
 194 phrase $T[2..n]$ overlaps its source $T[1..n-2]$. In general, when the phrase $T[i..i+\ell-1]$
 195 overlaps its source $T[s..s+\ell-1]$ by $0 < b = i-s$ characters, this implies that the word
 196 $S = T[s..s+\ell-1] = T[i..i+\ell-1]$ has a *border* (a prefix which is also a suffix) of length b .
 197 It is well known that if S has a border of length b , then S has a period $p = |S| - b$, see [37,
 198 Ch. 8]. Therefore, S can be written in the form $S = U^{\lfloor |S|/p \rfloor} V$, where U is the p -length
 199 prefix of S and V a proper prefix of U , and thus, for all $l > p$, $S[l] = S[l \bmod p]$.

200 **► Definition 2 (Chain length).** Let $T[i..i+\ell]$ be a phrase in a leftward parse of $T[1..n]$,
 201 whose source is $T[s..s+\ell-1]$. The chain length of the explicit character is $C[i+\ell] = 0$. If
 202 $\ell \leq i-s$ (i.e., there is no overlap between the source and the phrase), then for all $0 \leq l < \ell$,
 203 $C[i+l] = C[s+l] + 1$. Otherwise, for $0 \leq l < i-s$, the chain length is $C[i+l] = C[s+l] + 1$,
 204 and for $i-s \leq l < \ell$, the chain length is $C[i+l] = C[i+(l \bmod (i-s))]$.

205 We remark that a parsing like BAT-LZ is described as a baseline in the experimental
 206 results of one of the current authors’ previous work [33], under the name LZ-Cost, but as no
 207 efficient parsing algorithm was devised for it, it could be tested only on the tiny texts of the
 208 Canterbury Corpus (<https://corpus.canterbury.ac.nz>). It also did not handle overlaps
 209 between sources and targets, so it did not perform well on the text $T = a^n$. For testing the
 210 BAT-LZ parsing on large repetitive text collections we need an efficient parsing algorithm.

4 A Greedy Parsing Algorithm for BAT-LZ

In this section we describe an algorithm that, using $O(n)$ space and $O(n \log^3 n)$ time, produces a BAT-LZ parse of a text $T[1..n]$ by maximizing the next phrase length at each step. We then show how to reduce the time to $O(n \log^2 n)$ at the price of increasing the space to $O(n \log n)$. Of course, unlike in LZ, this greedy algorithm does not in general produce an optimal BAT-LZ parse, since the problem is NP-hard.

► **Definition 3.** A BAT-LZ parse of $T[1..n]$ with maximum chain length c is a leftward parse of T where the chain length of no position exceeds c . A greedy BAT-LZ parse is a BAT-LZ parse where each phrase, processed left to right, is as long as possible.

Let $T[1..i-1]$ be already processed. We call a prefix $T[i..i+\ell-1]$ of $T[i..]$ valid if $C[j] \leq c$ for all $j = i, \dots, i+\ell-1$. A leftward parse of T is therefore a BAT-LZ parse if and only if all phrases are valid. Our Greedy BAT-LZ parser proceeds then analogously to the original LZ parser. At each step, it has already processed $T[1..i-1]$, and it must find the next phrase, which is formed by (1) the longest valid prefix $T[i..i+\ell-1]$ of $T[i..]$ that has an occurrence $T[s..s+\ell-1]$ with $s < i$, and (2) the next symbol $T[i+\ell]$. In other words, the algorithm enforces that every symbol in $T[s..s+\ell-1]$ must have a chain length less than c , the maximum chain length allowed. The phrase $T[i..i+\ell]$ is encoded just as in the standard LZ, as a triple $(s, \ell, T[i+\ell])$.

To efficiently find s and ℓ , our BAT-LZ parsing algorithm stores the following structures:

1. The suffix array $\text{SA}[1..n]$ of T , represented as a wavelet matrix [13].
2. The inverse suffix array $\text{ISA}[1..n]$ of T , represented in plain form.
3. An array $C[1..n]$, where $C[i]$ is the chain length of i . Note that $C[i]$ is defined only for the already parsed positions of T .
4. An array $D[1..n]$, where $D[s]$ is the minimum $d \geq 0$ such that $C[s+d] = c$. If no such d exists (in particular, because $C[i]$ is defined only for the parsed prefix), then $D[s] = \infty$ (which holds initially for all s).
5. For each level of the wavelet matrix of SA , a special *dynamic RMQ* data structure to track the text positions that can be used. This structure is related to the values of D and therefore it changes along the parsing.

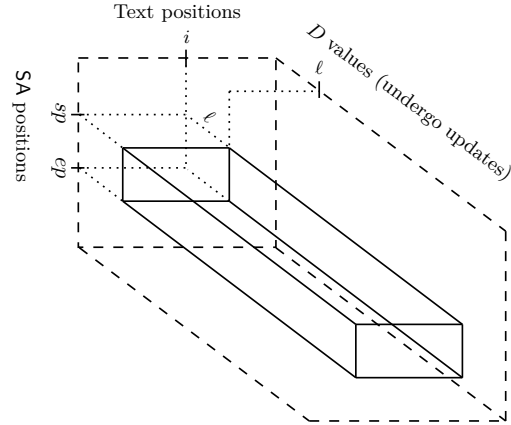
Note that the definition of BAT-LZ implies that, if the source of $T[i..i+\ell-1]$ is $T[s..s+\ell-1]$, then it must be that $\ell \leq D[s]$. This motivates the following observation:

► **Observation 4.** Let $T[1..i-1]$ be already processed. A prefix $T[i..i+\ell-1]$ of $T[i..]$ is valid if and only if there exists a source $T[s..s+\ell-1]$ such that

- (i) its lexicographic position satisfies $\text{ISA}[s] \in [sp..ep]$, where $[sp..ep]$ is the suffix array range of $T[i..i+\ell-1]$ (i.e., $T[s..s+\ell-1] = T[i..i+\ell-1]$);
- (ii) its starting position in T is $s < i$; and
- (iii) it does not use forbidden text positions, that is, $\ell \leq D[s]$.

The parsing then must find the longest valid prefix $T[i..i+\ell-1]$ of $T[i..]$. We do so by testing the consecutive values $\ell = 1, 2, \dots$. Note that, once we have determined the next phrase $T[i..i+\ell]$, we must update C and D as follows: (1) $C[i+\ell] \leftarrow C[s+\ell] + 1$ for all $0 \leq l < \ell$, and $C[i+\ell] \leftarrow 0$ ², and (2) Every time we obtain $C[t] = c$ in the previous point,

² Recall that a special case occurs if $T[i..i+\ell-1]$ overlaps $T[s..s+\ell-1]$: we start copying from $k = s$ and increasing k and, whenever $k = s + l = i$, we restart copying from $k = s$.



■ **Figure 1** General scheme of our translation of queries onto a 3-dimensional data structure.

252 we set $D[k] \leftarrow t - k$ for all $k' < k \leq t$, where k' is the last position where $D[k'] < \infty$ (so
 253 $k' = 0$ in the beginning and we reset $k' \leftarrow t$ after this process).


254 Note that points (i) and (ii) above correspond to the classic LZ parsing problem. In
 255 particular, they correspond to determining whether there are points in the range $[sp, ep] \times$
 256 $[1, i - 1]$ of the grid represented by our wavelet matrix, which represents the points $(j, SA[j])$.
 257 As the wavelet matrix answers this query in time $O(\log n)$, this yields an $O(n \log n)$ LZ parsing
 258 algorithm. Point (iii), however, is exclusive to BAT-LZ. It can be handled by converting
 259 the grid into a three-dimensional mesh, where we store the values $(j, SA[j], D[SA[j]])$ and
 260 look for the existence of points in the range $[sp, ep] \times [1, i - 1] \times [\ell, n]$. Note that we need to
 261 determine whether the range is empty and, if it is not, retrieve a point from it (whose second
 262 coordinate is the desired s). In addition, as the array D is modified along the parsing, we
 263 need a dynamic 3-dimensional data structure: every time we modify D in point 2 above, our
 264 data structure changes (this occurs up to n times). See Fig. 1.

265 Our 3-dimensional problem, then, (a) is essentially a range emptiness query (where we
 266 must return one point if there are any), (b) the search is 4-sided (though our solution handles
 267 5-sided queries), and (c) the updates in D occur only to convert some $D[k] = \infty$ into a
 268 smaller value, so each value $D[k]$ changes at most once along the parsing process (yet, our
 269 solution handles arbitrary updates along the coordinate where the query is one-sided). We
 270 have found no linear-space solutions to this problem in the literature; only solutions to less
 271 general ones or using super-linear space (indeed, more than $O(n \log n)$): (1) linear space for
 272 *two dimensions*, with $O(\log n)$ query time and $O(\log^{3+\epsilon} n)$ update time [44]; (2) linear space
 273 for three dimensions with *no updates*, with $O(\log n / \log \log n)$ query time [6]; (3) *super-linear*
 274 space (at least $O(n \log^{1.33} n)$ for three dimensions), with $O((\log n / \log \log n)^2)$ query time
 275 and $O(\log^{1.33+\epsilon} n)$ update time [7]. In the next section we describe our data structures for
 276 this problem: one uses linear space and $O(\log^3 n)$ query and update time; the other uses
 277 $O(n \log n)$ space and $O(\log^2 n)$ query time. This yields our first main result.

278 ► **Theorem 5.** *A Greedy BAT-LZ parse of a text $T[1..n]$ can be computed using $O(n)$ space*
 279 *and $O(n \log^3 n)$ time, or $O(n \log n)$ space and $O(n \log^2 n)$ time.*

280 5 A Geometric Data Structure

281 To solve the 3-dimensional search problem we associate, with each level of the wavelet matrix,
 282 a data structure that represents the sequence of values $D[k]$ in the order the text positions k



■ **Figure 2** On the left, we reach a candidate area $[sp_3, ep_3]$ of the wavelet matrix and must obtain its maximum D value using the (dynamic) RMQ data structure for D_3 . The tree H_3 for this RMQ structure is shown on the right. Arrows point to the child holding the maximum value in D_3 . Blue diamonds are the roots v_3^1, \dots, v_3^4 of the subtrees that cover the query area $[sp_3, ep_3]$ and red circles are the candidates in the range. The left plot shows how we find the actual value of one of those circles by tracking it down in the wavelet matrix.

283 are listed in that level. Because in linear space we cannot store the actual values in every
 284 wavelet matrix level, we store only a dynamic RMQ data structure on the internal levels,
 285 and store the explicit values only in (the order corresponding to) the last level (in a wavelet
 286 matrix, that final level is not the text order, thus we need another array to map it to D).

287 Let D_l be the array D permuted in the way it corresponds to level l of the wavelet
 288 matrix. The dynamic RMQ structure for level l is then a heap-shaped perfectly balanced
 289 tree $H_l[1..n]$ whose leaves (implicitly) point to the entries of D_l . The nodes $H_l[p]$ store
 290 only one bit, 0 indicating that the maximum in the subtree is to the left and 1 indicating
 291 that it is to the right. By navigating H_l from the root p of any subtree, moving to $H_l[2p]$
 292 if $H_l[p] = 0$ and $H_l[2p + 1]$ if $H_l[p] = 1$, we arrive in $O(\log n)$ time at the position p where
 293 $D_l[p]$ is maximum below that subtree. The actual value $D_l[p]$ is obtained in other $O(\log n)$
 294 time by tracking position p downwards in the wavelet matrix, from level l until the last level,
 295 where the values of D are explicitly stored. See Fig. 2 (right); ignore the query for now.

296 **Updates.** When a value $D[k]$ decreases from ∞ , we obtain its position in the top-level of the
 297 wavelet matrix as $p = \text{ISA}[k]$; thus we must reflect in H_1 the decrease in the value of $D_1[p]$.
 298 By halving p successively we arrive at its ancestors, $H_1[p_h]$ for $p_h = \lfloor p/2^h \rfloor$, $h = 1, 2, \dots$
 299 We traverse the path upwards, recomputing the maximum value m below p_h and modifying
 300 accordingly the bits of $H_1[p_h]$. Initially, this new maximum is $m = D_1[p] = D[k]$. At any
 301 point in the traversal, if the parent $H_1[p_h]$ of the current node indicates that the maximum
 302 below p_h descends from the *other* child of p_h , then we can stop updating of H_1 , because
 303 decreasing $D_1[p]$ does not require further changes. Otherwise, we must obtain the maximum
 304 value m' below the other child of $H_1[p_h]$ and compare it with m . The value m' is obtained
 305 in $O(\log n)$ time as explained in the previous paragraph. We set $H_1[p_h]$ depending on which
 306 is larger between m and m' , update $m \leftarrow \max(m, m')$, and continue upwards. This process
 307 takes $O(\log^2 n)$ time as we traverse all the levels of H_1 . We then track position p downwards
 308 to the second level of the wavelet matrix, update H_2 in the same way, and continue updating
 309 H_l on all the wavelet matrix levels l , for a total update time of $O(\log^3 n)$.

310 **Searches.** The search for a range $[sp, ep] \times [1, i - 1] \times [\ell, n]$ first determines, as in the normal
 311 wavelet matrix search algorithm, the $O(\log n)$ maximal ranges that cover $[1, i - 1]$ along the
 312 wavelet matrix levels l (there is at most one range per level because the range $[1, i - 1]$ is
 313 one-sided; otherwise there could be two), and maps $[sp, ep]$ to $[sp_l, ep_l]$ on each such range
 314 (see Sec. 2), all in time $O(\log n)$. We then need to determine if there is some value $D_l[p] \geq \ell$
 315 below some of the ranges $[sp_l, ep_l]$ (see the top-left part of Fig. 2). Each such range is then,
 316 again, decomposed into $O(\log n)$ maximal nodes v_l^1, v_l^2, \dots of H_l (see the right of Fig. 2). We
 317 find, in $O(\log n)$ time, the maximum value of D_l below each node v_l^j , stopping as soon as we
 318 find some value $\geq \ell$. Note that we use $O(\log n)$ time to find the *position* of the maximum
 319 in D_l using H_l , and then $O(\log n)$ time to find the *value* of that maximum by tracking the
 320 position down in the wavelet matrix (see the bottom left of Fig. 2). Since we have $O(\log n)$
 321 ranges $[sp_l, ep_l]$, each yielding $O(\log n)$ candidates v_l^i , and the maximum of each candidate is
 322 computed in $O(\log n)$ time, the whole search process takes time $O(\log^3 n)$.

323 **Generalizations.** Though not necessary for our problem, we remark that our update process
 324 can be extended to arbitrary updates on the third coordinate, $D[k]$, not only to reductions
 325 in value. Further, our search could support five-sided ranges, not only four-sided, because we
 326 would still have $O(\log n)$ ranges $[sp_l, ep_l]$ if the range of the second coordinate was two-sided.
 327 Only the range of the third coordinate (the one supporting the updates) must be one-sided.

328 **Faster and larger.** By storing the values of D_l in each node of H_l for each wavelet matrix
 329 level l , the space increases to $O(n \log n)$ but the time of updates and searches decreases to
 330 $O(\log^2 n)$, as we have now the maximum below any $H_l[p_n]$ readily available in $O(1)$ time.

331 6 The Minmax Parsing Algorithm

332 We note that our Greedy BAT-LZ algorithm does not necessarily produce the smallest greedy
 333 parse, because it may fail in choosing the best *source* for the longest phrase. Consider, say,
 334 the text $T = \text{alabaralalabarda\$}$ and $c = 2$. Our implementation parses it into 8 phrases as

335
$$\begin{array}{c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|} \mathbf{a} & \mathbf{1} & \mathbf{a} & \mathbf{b} & \mathbf{a} & \mathbf{r} & \mathbf{a} & \mathbf{1} & \mathbf{a} & \mathbf{1} & \mathbf{a} & \mathbf{b} & \mathbf{a} & \mathbf{r} & \mathbf{d} & \mathbf{a} & \mathbf{\$} \\ \hline 0 & 0 & 1 & 0 & 2 & 0 & 1 & 1 & 2 & 0 & 2 & 1 & 0 & 1 & 0 & 1 & 0 \end{array}$$

336 because it chooses $T[3]$ as the source for the 4th phrase, \mathbf{ar} , and then $T[5]$ has a chain of
 337 length two and cannot be used again. If, instead, we choose $T[1]$ as the source of the 4th
 338 phrase, the chain of $T[5]$ will be of length 1 and we could parse T into 7 phrases, just as the
 339 first parse shown in Sec. 3.

340 Our second algorithm, the Minmax parser, always chooses a source that minimizes the
 341 maximum chain length in the phrase, among all possible sources. It compromises however
 342 on the *length* of the phrase, by not always choosing the longest admissible phrase. As we will
 343 see, this is well worth it: Minmax always produces a much better compression than Greedy.

344 **High-level description of the Minmax parser.** Let $T[1 \dots i - 1]$ be already processed. We
 345 will call a prefix $T[i \dots i + \ell - 1]$ of $T[i \dots]$ *admissible* if it has a source $T[s \dots s + \ell - 1]$ with
 346 $\max C[s \dots s + \ell - 1] < c$. We would ideally like to find the longest admissible prefix of $T[i \dots]$,
 347 and then choose its best source if there is more than one. We will use an enhanced suffix
 348 tree of the text; this will allow us to store additional information in the nodes. Navigating in
 349 the suffix tree, we will then be able to choose the longest admissible prefix *which ends in*
 350 *some node* (i.e., not necessarily the longest), and then choose the best source of this prefix.

351 In order to do this, we will match the current suffix $T[i..]$ in the usual way in the suffix
 352 tree, using the desired information written in the nodes. As this information is dynamic,
 353 however, we will have to update it during the algorithm. The algorithm thus proceeds by (1)
 354 matching the suffix $T[i..]$ in the suffix tree and returning the next phrase and its source,
 355 and (2) updating the annotation.

356 **Annotation of the suffix tree.** On the suffix tree of T , we annotate each node v with
 357 three variables $\text{minmax}(v)$, $\text{txtpos}(v)$, and a Boolean $\text{real}(v)$, initializing $\text{minmax}(v)$ to $+\infty$,
 358 $\text{txtpos}(v)$ to -1 , and $\text{real}(v)$ to 0. Recall that $L(v)$ is the label of v and $sd(v)$ its length. The
 359 variables $\text{minmax}(v)$ and $\text{txtpos}(v)$ will point to the current best candidate of an occurrence
 360 of $L(v)$, with $\text{txtpos}(v)$ its starting position and $\text{minmax}(v)$ the maximum C -value within this
 361 occurrence. The Boolean $\text{real}(v)$ indicates whether this value is *realistic* ($\text{real}(v) = 1$), i.e., a
 362 full occurrence with this value has already been seen, or only *optimistic* ($\text{real}(v) = 0$), meaning
 363 that no full occurrence has yet been seen. More formally, let i be the current position, and let
 364 us first assume that $\text{real}(v) = 1$. Then $\text{minmax}(v) = x$ if $x = \min\{\max C[s..s + sd(v) - 1] :$
 365 $T[s..s + sd(v) - 1] = L(v)$ and $s + sd(v) - 1 < i\}$, and $\text{txtpos}(v) = s_0$ for one such s_0 , i.e., (i)
 366 $T[s_0..s_0 + sd(v) - 1] = L(v)$, (ii) $s_0 + sd(v) - 1 < i$, and (iii) $\max C[s_0..s_0 + sd(v) - 1] = x$.

367 Now let us look at the case $\text{real}(v) = 0$, we have yet to see an occurrence of $L(v)$. Initially,
 368 $\text{minmax}(v) = +\infty$; when we encounter a non-empty prefix of $L(v)$, of length $0 < d \leq sd(v)$,
 369 starting, say, in position s_0 , we update $\text{minmax}(v)$ to $\max C[s_0..s_0 + d - 1]$ and $\text{txtpos}(v)$
 370 to s_0 . Thus, we have seen an occurrence of a prefix of $L(v)$ but not yet a full occurrence of
 371 $L(v)$, and we are optimistic since we are hoping to find a full occurrence whose max does
 372 not exceed the current one. However, as soon as we find the first full occurrence (and set
 373 $\text{real}(v) = 1$), from that point on we only update $\text{minmax}(v)$ and $\text{txtpos}(v)$ if we see another
 374 full occurrence. Therefore, $\text{real}(v)$ is updated exactly once during the algorithm.

375 **Finding an admissible phrase and choosing its source.** Let us now assume that we have
 376 processed $T[1..i - 1]$ and want to find the next phrase and source. We match $T[i..]$ in the
 377 suffix tree, making sure during navigation that we only get admissible prefixes of $T[i..]$. In
 378 particular, if we are in node v and should go to child u of v next (because $T[i + sd(v)]$ is the
 379 first character of the edge label (v, u)), then we first check if $\text{minmax}(u) < c$. If so, then we
 380 can descend to u and continue from there, skipping over the next $sd(u) - sd(v)$ positions in
 381 T . Otherwise, $\text{minmax}(u) \geq c$ and we return the new phrase $(\text{txtpos}(v), sd(v), T[i + sd(v)])$.
 382 Moreover, the C -array for $j = i, \dots, i + \ell$ is set according to Def. 2.

383 **Updating the suffix tree annotation.** After the new phrase has been computed, we need
 384 to update the annotations in the suffix tree. For $j \leq i + \ell$, going backward in the string, we
 385 will update the nodes on the leaf-to-root path from leaf j . The idea is the following.

386 Fix $j \leq i + \ell$. The prefix $T[j..i + \ell]$ of $T[j..]$ now has the C -array filled in, so its
 387 max-value $m = \max C[j..i + \ell]$ is known. This may or may not necessitate updates in the
 388 nodes on the path from leaf $leaf_j$ to the root. First, for the leaf j itself, if $i \leq j$, then the
 389 minmax is still $+\infty$, so we set $\text{minmax}(leaf_j) \leftarrow m$. Otherwise, we are seeing a longer prefix
 390 of $T[j..]$ than before, so we update $\text{minmax}(leaf_j) \leftarrow \max(\text{minmax}(leaf_j), m)$. Regarding the
 391 nodes v on the path from $leaf_j$ to the root: their labels are increasingly shorter prefixes of
 392 suffix $T[j..]$, so they need to be updated only as long as $j + sd(v) - 1 \geq i$, since otherwise,
 393 the prefix $L(v)$ does not overlap with the newly assigned subinterval $C[i..i + \ell]$.

394 So let $j + sd(v) \geq i$, there are two cases. First, if $j + sd(v) - 1 \leq i + \ell$, then m is a realistic
 395 value, since the entire corresponding C -array interval has been filled in. Therefore, we can then

396 compute m in a more clever way by using an RMQ on C , i.e., $m = \text{RMQ}(C, j, j + \text{sd}(v) - 1)$.
 397 So if $\text{real}(v) = 0$, then we update $\text{minmax}(v) \leftarrow m$ and $\text{real}(v) \leftarrow 1$. Otherwise (if $\text{real}(v) = 1$),
 398 an update is needed only if $\text{minmax}(v) > m$, in which case we set $\text{minmax}(v) \leftarrow m$ and
 399 $\text{txtpos}(v) \leftarrow j$; since $\text{real}(v) = 1$, we have seen the label $L(v)$ before and already had a
 400 realistic value for its minmax value. Second, if $j + \text{sd}(v) - 1 > i + \ell$, then m is an optimistic
 401 value only, and therefore, we update the annotation of v only if $\text{real}(v) = 0$; in that case, we
 402 set $\text{minmax}(v) \leftarrow m$ and $\text{txtpos}(v) \leftarrow j$.

403 Finally, we use the following criterion for how far back in the string we need to go with j .
 404 If no node in the path from leaf_j to the root can be effected by the new phrase, then we do
 405 not need to consider position j at all in the current iteration. This holds if the label of the
 406 parent node does not reach i , i.e., if $j + \text{sd}(\text{parent}(\text{leaf}_j)) - 1 < i$. We compute an auxiliary
 407 array $E[1..n]$ s.t. $E[j] = j + \text{sd}(\text{parent}(\text{leaf}_j)) - 1$. It is easy to see that $E[j] \leq E[j']$ if $j < j'$.
 408 This means that, moving back-to-front, we can stop at the first j for which $E[j] < i$.

409 A worst-case time complexity for a Minmax parse producing z' phrases is $O(z'n^2) \subseteq O(n^3)$,
 410 as in principle one can consider every $j \in [1..i-1]$ for every new phrase $T[i..i+\ell]$, traverse
 411 the $O(n)$ ancestors of leaf_j , and run an RMQ operation on each. While the RMQ structure
 412 we use on C is dynamic, it only undergoes appends to the right, in which case it is possible
 413 to support updates in $O(1)$ amortized time and queries in $O(1)$ time [18, p. 5]. We do not
 414 know if this cubic complexity is tight, however. In practice we expect z' to be much less
 415 than n on highly repetitive texts, and the height of the suffix tree to be logarithmic, yielding
 416 a time complexity of $O(z'n \log n)$, which thus becomes practical on repetitive data.

417 ► **Example 6.** In Fig. 3, we can see the suffix tree ST for $T = \text{alabaralalabarda}\$$ with
 418 some additional annotations in some nodes. In this example, the first three phrases, i.e.,
 419 **a**, **l**, and **ab**, have been already computed, with the corresponding chain lengths in C and
 420 annotations in ST. The annotations exhibit non-trivial updates using the colour red, namely
 421 updates that are different than changing the starting value for minmax , i.e., changing minmax
 422 from $+\infty$ to a finite value. Nodes whose annotation is not shown have not been updated
 423 yet, therefore, they have $\text{minmax} = +\infty$, $\text{txtpos} = -1$, and $\text{real} = 0$. The updates caused
 424 by the new phrase **ar** are highlighted in blue. First, to find the longest previous factor we
 425 have to descend to the child with label **a**, then we check whether the child with label **ar**
 426 has $\text{minmax} < c$. In this case, it was not less than c ($\text{minmax} = +\infty$), so we stopped the
 427 search and output the new phrase **ar**. Then all suffixes j with $1 \leq j \leq 6$ undergo an update
 428 starting from the corresponding leaf; e.g., leaves 5 and 6 and corresponding ancestors get
 429 updated to some non-initial value, whereas inner nodes with label **abar**, **alabar**, **bar** and
 430 **labar** change real to 1 because $j + \text{sd}(v) - 1 \leq i + \ell$.

431 **7 The Greedier Parser: Combining Greedy with Minmax**

432 We now combine the ideas of the Greedy and the Minmax parsers, using the enhanced suffix
 433 tree to consider only longest admissible phrases. Consider when the Minmax algorithm stops
 434 in a node v and returns $(\text{txtpos}(v), \text{sd}(v), T[i + \text{sd}(v)])$. It did not descend to the next child
 435 u because $\text{minmax}(u) = c$, i.e., every occurrence of $L(v)$ seen so far has a value c somewhere
 436 in the C -array. However, it is possible that in one of these occurrences, the position of this c
 437 is after $L(v)$; in other words, that we could have gone down the edge some way towards u .

438 To check this, we will use the D -array from Sec. 4, in addition to the C -array and
 439 the enhanced suffix tree. Let v and u be as before, i.e., v is parent of u , $\text{minmax}(v) < c$,
 440 $\text{minmax}(u) \geq c$, $L(v)$ is a prefix of $T[j..]$ and $T[j + \text{sd}(v)]$ is the first character of the label

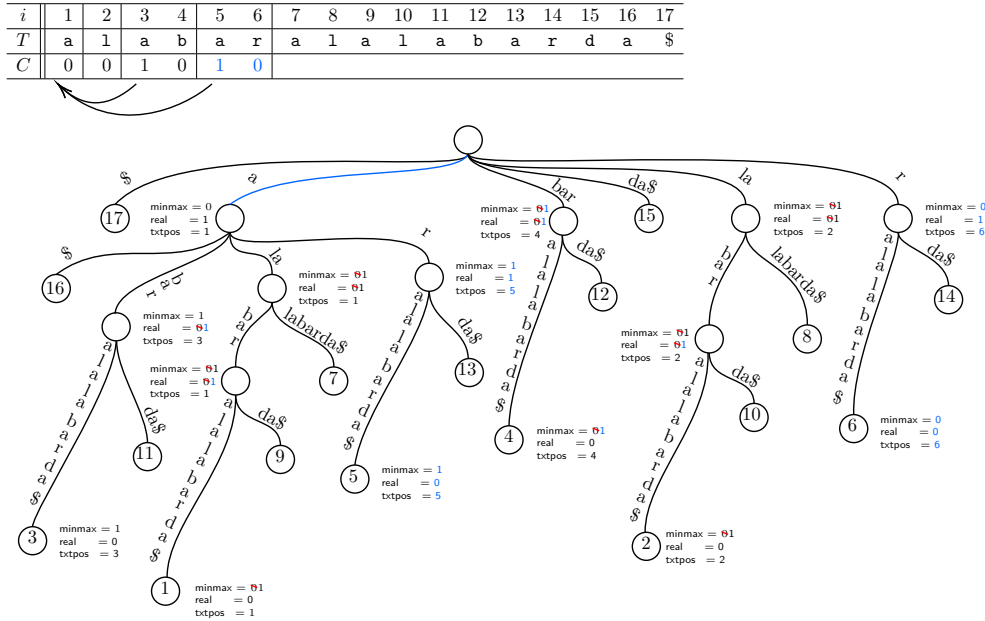


Figure 3 Example of the Minmax algorithm using the suffix tree of $T = alabaralalabarda\$$. The vertical bars are for delimiting already parsed phrases. See Example 6 for more details.

441 of (v, u) . Let d be the maximum value of $D[k]$ for some occurrence of $L(u)$ that we have
 442 already processed, so d is the largest distance from the start of an occurrence of $L(u)$ to the
 443 next c in the C -array. We return $(\text{txtpos}(v), sd(v), T[i + sd(v)])$, as before, if $d \leq sd(v)$, and
 444 $(k, D[k], T[i + D[k]])$, where k is a leaf in u 's subtree with $D[k] = d$, otherwise. As for updating
 445 the annotations, if node v has $\text{minmax}(v) = c$ and some $\text{txtpos}(v) = x$, then, when performing
 446 the traversal from leaf_j up to the root, we want to change $\text{txtpos}(v) \leftarrow j$ if $D[j] > D[\text{txtpos}(v)]$.
 447 It is easy to see that the Greedier algorithm now returns the longest admissible phrases;
 448 otherwise, it works similarly to the Minmax algorithm. The time complexity increases to
 449 $O(z'n^2 \log n) \subseteq O(n^3 \log n)$, because we need dynamic RMQs on array D as well, which
 450 undergoes updates at arbitrary positions.

451 8 Experiments

452 We implemented the BAT-LZ parsing algorithms in C, and ran our experiments on an AMD
 453 EPYC 7343, with 32 cores at 1.5 GHz, with a 32 MB cache and 1 TB of RAM. We used the
 454 repetitive files from Pizza&Chili (<http://pizzachili.dcc.uchile.cl>) and compared the
 455 number of phrases produced by BAT-LZ using different maximum values c for the chains,
 456 with the number of phrases produced by LZ (i.e., with no limit c). We used a classic LZ
 457 implementation [26] where the source of each phrase is its lexicographically closest suffix.³

458 As a reference point, we also implemented two simple baselines that obtain a BAT-LZ
 459 parse. The first, called BAT-LZ1, runs the classic LZ parse and then cuts the phrases at
 460 the points where the chain lengths reach $c + 1$. Since the symbol becomes explicit, its chain

³ It is likely that using the variant called “rightmost LZ parse”, which chooses the rightmost source, gives better results because it tends to distribute the uses of the sources more uniformly. Such a parse seems to be nontrivial to compute [4, 15], however, and we are not aware of practical implementations.

File	σ	n	z	n/z	max c	g	h
<code>coreutils</code>	236	205,281,779	1,286,070	160	66	2,409,429	28
<code>kernel</code>	160	257,961,616	705,791	365	70	1,374,651	32
<code>einstein</code>	139	467,626,545	75,779	6,171	1,736	212,902	47
<code>leaders</code>	89	46,968,181	155,937	301	60	399,667	27
<code>para</code>	5	429,265,758	1,879,635	228	38	5,344,477	26
<code>influenza</code>	15	154,808,555	557,349	278	63	1,957,370	26

■ **Table 1** Our repetitive text collections and some statistics: alphabet size σ , length n , number z of phrases in the LZ parse, average phrase length n/z , maximum chain length in our LZ parse, size g of a balanced grammar, and height h of that grammar.

461 length becomes zero and the chain lengths of the symbols referencing it decrease by $c + 1$.
 462 We should then find the new positions that reach $c + 1$, and so on. It is not hard to see that
 463 this laborious postprocessing can be simulated by just adding, to the original z value of LZ,
 464 the number of positions i where $C[i] \bmod (c + 1) = 0$.

465 The second baseline, BAT-LZ2, is slightly stronger: when it detects that it has produced
 466 a text position exceeding the maximum c , it cuts the phrase there (making the symbol
 467 explicit), and restarts the LZ parse from the next position. This gives the chance of choosing
 468 a better phrase starting after the cut, unlike BAT-LZ1, which maintains the original source.

469 Despite some optimizations, our Greedy BAT-LZ parser consistently reaches the $\Theta(\log^3 n)$
 470 time complexity per text symbol, making it run at about 3 MB per minute. The Greedier
 471 and the Minmax parsers, despite their cubic worst-case time complexity, run at a similar
 472 pace: 1.9–4.7 MB per minute: our upper bound is utterly pessimistic, and perhaps not tight.

473 Table 1 shows the main characteristics of the collections chosen. We included two
 474 versioned software repositories (`coreutils` and `kernel`, where the versioning has a tree
 475 structure), two versioned documents (`einstein` and `leaders`, where the versioning has a
 476 linear structure), and two biological sequence collections (`para` and `influenza`, where all
 477 the sequences are pairwise similar). The average phrase length is in the range 160–365 and
 478 the maximum chain length of a symbol is in the range 38–70. The exception is `einstein`,
 479 which is extremely compressible and also has a very large c value.

480 As a point of comparison, the table also includes the grammar size and height obtained
 481 with a balanced version of RePair [35].⁴ We modified the RePair grammar so as to remove
 482 the nonterminals that are referenced only once, inserting their right-hand side in that unique
 483 referencing place. The maximum grammar height is comparable with c as a measure of access
 484 cost in the grammar-compressed text. We can see that the height is considerably smaller
 485 than c , for the price of a weaker compression method.

486 Fig. 4 shows how the quotient between the number of phrases generated by the BAT-LZ
 487 parsers and by the optimal number of LZ phrases evolves as we allow longer chains. It
 488 can be seen that our Greedy BAT-LZ parser sharply outperforms the baselines in terms of
 489 compression performance. Our Greedy parser is, in turn, outperformed by Minmax, and
 490 Minmax is outperformed by our Greedier parser. The last one reaches a number of phrases
 491 that is only 1% over the optimal for c as low as 20–30, which is 0.7–1.1 times $\log_2 n$.

492 We also show in the figures the balanced grammar method, using the values of Table 1.⁵

⁴ From www.dcc.uchile.cl/gnavarro/software/repair.tgz, directory `bal/`.

⁵ For a fair comparison of space, we consider a tight space needed to support fast extraction: For each of the z phrases we count $\log_2 n$ bits to point to the source, $\log_2(n/z)$ bits for the length (as there are z




Figure 4 Overhead factor of number of BAT-LZ versus LZ phrases as a function of the maximum length c of a chain, for our different BAT-LZ parsers and a balanced grammar.

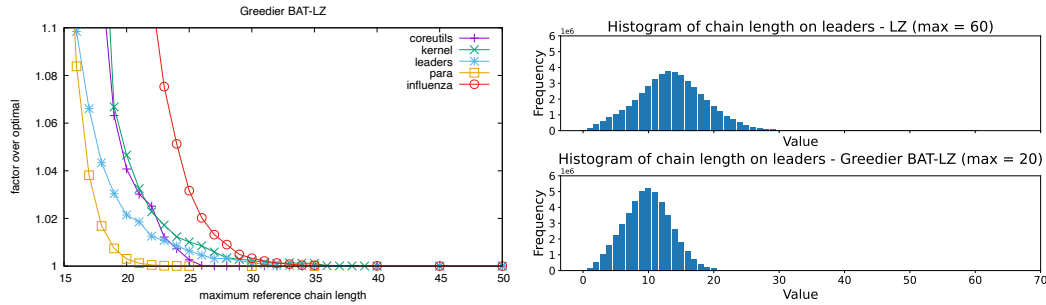
493 We can see that grammars are competitive, in some cases, with the simple baselines, but
 494 not with our new algorithms, which yield much better tradeoffs. The only exception to this
 495 analysis is `einstein`, which features a huge maximum c value of 1,736 and whose (extremely
 496 low) z value is approached only with c values near 700 using our BAT-LZ parsers. On this
 497 text, the balanced grammar offers an access time that is not achievable with our techniques.

498 Fig. 5 (left) zooms in the area where Greedier BAT-LZ reaches less than 10% extra space
 499 on top of standard LZ (excluding `einstein`).

9 Discussion and Future Work

501 A first question is whether a Greedy BAT-LZ parsing can be produced in $o(n \log^3 n)$ time
 502 within linear space, either by solving our geometric problem faster or without recasting

lengths adding up to n ; the cumulative sequence of lengths also allow finding the desired phrase using Elias-Fano codes [14, 16]), and 8 bits for the final symbol. For a grammar of size g and r symbols, we count $g \log_2 r$ bits for the right-hand sides, $g \log_2 n$ bits for the expansion lengths (cumulative on the right-hand sides to binary search them), and $r \log(g/r)$ bits to encode the rule lengths with Elias-Fano.



■ **Figure 5** Left: detail of Fig. 4, for the Greedier BAT-LZ parser, focusing on the overheads below 10% over the LZ parse. Right: a comparison of histograms with shared x and y axis representing the chain length values on **leaders**; LZ on top and Greedier BAT-LZ with $c = 20$ on the bottom.

503 the parse into a geometric problem. This question seems to be answered in a very recent
 504 work, simultaneous with ours, that gives an $O(n \log \sigma)$ -time greedy algorithm [2] based
 505 on simulating a suffix tree construction.⁶ This algorithm is likely to be faster than ours
 506 in practice, but also to use much more space, which is relevant when compressing large
 507 repetitive texts. They also propose a parse similar to our BAT-LZ2, along with others that
 508 are incomparable to ours (in particular to Greedier, our best performing BAT-LZ parse).

509 Besides our reduction to a geometric problem being of independent interest, we believe
 510 that its flexibility can be exploited to compute more sophisticated parses in $O(n \log^3 n)$ time.
 511 For example, it might compute the Greedier parse if we extend the RMQ data structure to
 512 incorporate the additional optimization criterion (minmax of sources).

513 Other heuristics may also be of interest: there may be better ways to rank sources, other
 514 than their maximum chain length. Further, we have so far focused on reducing the worst
 515 case access time, but we might prefer to reduce the *average* access time. Our parsings do
 516 reduce it (Fig. 5 right), but this is just a side effect and has not been our main aim. So we
 517 pose as an open problem to efficiently build a leftward parse with bounded average reference
 518 chain length whose number of phrases is minimal, or in practice close to that of classical LZ.

519 Another intriguing line of work is to study the compression performance of BAT-LZ.
 520 An important result by Bannai et al. [2] shows that, letting g_{rl} be the size of the smallest
 521 run-length context-free grammar that generates a text T , there exists a BAT-LZ parse for T
 522 of size $O(g_{rl})$ if we let $c = \Theta(\log n)$ with some convenient multiplying constant. This bound
 523 is nearly optimal, because existing bounds [51] forbid the existence of BAT-LZ parses of
 524 size $O(g)$ —where $g \geq g_{rl}$ is the size of the smallest context-free grammar—with access time
 525 $c = O(\log^{1-\epsilon} n)$ for a constant $\epsilon > 0$. A relevant question is whether there is a BAT-LZ parse
 526 of size $O(z)$ —where $z \leq g_{rl}$ is the size of the Lempel-Ziv parse of T —with $c = \Theta(\log n)$.

527 Finally, from an application viewpoint, it would be interesting to incorporate BAT-LZ in
 528 the construction of the LZ-index [34] and measure how much its time performance improves
 529 at the price of an insignificant increase in space. Obtaining an efficient bounded version of
 530 the LZ-End parsing described in the same article [34] is also an interesting problem since
 531 efficient parsings for unrestricted LZ-End have appeared only recently [29, 28].

⁶ They use a slightly modified definition of Lempel-Ziv parses, which has no explicit character at the end of the phrases. The precise consequences of this difference are not totally clear to us.

532

References

- 533 1 Alberto Apostolico. The myriad virtues of subword trees. In *Combinatorial Algorithms on*
534 *Words*, NATO ISI Series, pages 85–96. Springer-Verlag, 1985.
- 535 2 Hideo Bannai, Mitsuru Funakoshi, Diptarama Hendrian, Myuji Matsuda, and Simon J. Puglisi.
536 Height-bounded Lempel-Ziv encodings. *CoRR*, abs/2403.08209, 2024.
- 537 3 Djamel Belazzougui and Simon J. Puglisi. Range predecessor and Lempel-Ziv parsing. In
538 *Proc. 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2053–2071,
539 2016.
- 540 4 Djamel Belazzougui and Simon J. Puglisi. Range predecessor and Lempel-Ziv parsing. In
541 *Proc. 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2053–2071,
542 2016.
- 543 5 Philip Bille, Gad M. Landau, Rajeev Raman, Kunihiko Sadakane, Srinivasa Rao Satti, and
544 Oren Weimann. Random access to grammar-compressed strings and trees. *SIAM Journal on*
545 *Computing*, 44(3):513–539, 2015.
- 546 6 Timothy M. Chan, Yakov Nekrich, Saladi Rahul, and Konstantinos Tsakalidis. Orthogonal
547 point location and rectangle stabbing queries in 3-d. *Journal of Computational Geometry*,
548 13(1), 2022.
- 549 7 Timothy M. Chan and Konstantinos Tsakalidis. Dynamic orthogonal range searching on the
550 RAM, revisited. *Journal of Computational Geometry*, 9(2):45–66, 2018.
- 551 8 Moses Charikar, Eric Lehman, Ding Liu, Rina Panigrahy, Manoj Prabhakaran, Amit Sahai,
552 and Abhi Shelat. The smallest grammar problem. *IEEE Transactions on Information Theory*,
553 51(7):2554–2576, 2005.
- 554 9 Gang Chen, Simon J. Puglisi, and William F. Smyth. Lempel-Ziv factorization using less time
555 & space. *Mathematics in Computer Science*, 1:605–623, 2008.
- 556 10 Ferdinando Cicalese and Francesca Ugazio. On the complexity and approximability of bounded
557 access Lempel Ziv coding. *CoRR*, abs/2403.15871, 2024. Submitted.
- 558 11 David R. Clark. *Compact PAT Trees*. PhD thesis, University of Waterloo, Canada, 1996.
- 559 12 Francisco Claude, Gonzalo Navarro, and Alejandro Pacheco. Grammar-compressed indexes
560 with logarithmic search time. *Journal of Computer and System Sciences*, 118:53–74, 2021.
- 561 13 Francisco Claude, Gonzalo Navarro, and Alberto Ordóñez Pereira. The wavelet matrix: An
562 efficient wavelet tree for large alphabets. *Information Systems*, 47:15–32, 2015.
- 563 14 P. Elias. Efficient storage and retrieval by content and address of static files. *Journal of the*
564 *ACM*, 21:246–260, 1974.
- 565 15 Jonas Ellert, Johannes Fischer, and Max Rishøj Pedersen. New advances in rightmost Lempel-
566 Ziv. In *Proc. 30th International Symposium on String Processing and Information Retrieval*
567 *(SPIRE)*, pages 188–202, 2023.
- 568 16 R. Fano. On the number of bits required to implement an associative memory. Memo 61,
569 Computer Structures Group, Project MAC, Massachusetts, 1971.
- 570 17 Martin Farach. Optimal suffix tree construction with large alphabets. In *Proc. 38th Annual*
571 *Symposium on Foundations of Computer Science (FOCS)*, pages 137–143. IEEE Computer
572 Society, 1997.
- 573 18 J. Fischer. Combined data structure for previous- and next-smaller-values. *Theoretical*
574 *Computer Science*, 412(22):2451–2456, 2011.
- 575 19 Johannes Fischer and Volker Heun. Space-efficient preprocessing schemes for range minimum
576 queries on static arrays. *SIAM Journal on Computing*, 40(2):465–492, 2011.
- 577 20 Johannes Fischer, Tomohiro I, and Dominik Köppl. Lempel Ziv computation in small space
578 (LZ-CISS). In *Proc. 26th Annual Symposium on Combinatorial Pattern Matching (CPM)*,
579 LNCS 9133, pages 172–184, 2015.
- 580 21 Johannes Fischer, Tomohiro I, Dominik Köppl, and Kunihiko Sadakane. Lempel-Ziv
581 factorization powered by space efficient suffix trees. *Algorithmica*, 80(7):2048–2081, 2018.
- 582 22 Moses Ganardi, Artur Jez, and Markus Lohrey. Balancing straight-line programs. *Journal of*
583 *the ACM*, 68(4):article 27, 2021.

- 584 23 Keisuke Goto and Hideo Bannai. Simpler and faster Lempel Ziv factorization. In *Proc. 23rd*
585 *Data Compression Conference (DCC)*, pages 133–142, 2013.
- 586 24 D. Gusfield. *Algorithms on Strings, Trees and Sequences: Computer Science and Computational*
587 *Biology*. Cambridge University Press, 1997.
- 588 25 Juha Kärkkäinen, Dominik Kempa, and Simon J. Puglisi. Lightweight Lempel-Ziv parsing. In
589 *Proc. 12th International Symposium on Experimental Algorithms (SEA)*, pages 139–150, 2013.
- 590 26 Juha Kärkkäinen, Dominik Kempa, and Simon J. Puglisi. Linear time Lempel-Ziv factorization:
591 Simple, fast, small. In *Proc. 24th Annual Symposium on Combinatorial Pattern Matching*
592 *(CPM)*, LNCS 7922, pages 189–200, 2013.
- 593 27 Juha Kärkkäinen, Dominik Kempa, and Simon J. Puglisi. Lazy Lempel-Ziv factorization
594 algorithms. *ACM Journal of Experimental Algorithmics*, 21(1):2.4:1–2.4:19, 2016.
- 595 28 Dominik Kempa and Dmitry Kosolobov. LZ-End parsing in compressed space. In *Proc. 27th*
596 *Data Compression Conference (DCC)*, pages 350–359, 2017.
- 597 29 Dominik Kempa and Dmitry Kosolobov. LZ-End parsing in linear time. In *Proc. 25th Annual*
598 *European Symposium on Algorithms (ESA)*, pages 53:1–53:14, 2017.
- 599 30 Dominik Kempa and Barna Saha. An upper bound and linear-space queries on the LZ-
600 End parsing. In *Proc. 33rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages
601 2847–2866, 2022.
- 602 31 John C. Kieffer and En-Hui Yang. Grammar-based codes: A new class of universal lossless
603 source codes. *IEEE Transactions on Information Theory*, 46(3):737–754, 2000.
- 604 32 Dominik Köppl and Kunihiko Sadakane. Lempel-Ziv computation in compressed space
605 (LZ-CICS). In *Proc. 26th Data Compression Conference (DCC)*, pages 3–12, 2016.
- 606 33 Sebastian Kreft and Gonzalo Navarro. Lz77-like compression with fast random access. In
607 *Proc. 20th Data Compression Conference (DCC)*, pages 239–248, 2010.
- 608 34 Sebastian Kreft and Gonzalo Navarro. On compressing and indexing repetitive sequences.
609 *Theoretical Computer Science*, 483:115–133, 2013.
- 610 35 J. Larsson and A. Moffat. Off-line dictionary-based compression. *Proceedings of the IEEE*,
611 88(11):1722–1732, 2000.
- 612 36 Abraham Lempel and Jacob Ziv. On the complexity of finite sequences. *IEEE Transactions*
613 *on Information Theory*, 22(1):75–81, 1976.
- 614 37 M. Lothaire. *Algebraic Combinatorics on Words*. Cambridge University Press, 2002.
- 615 38 Edward M. McCreight. A space-economical suffix tree construction algorithm. *J. ACM*,
616 23(2):262–272, 1976.
- 617 39 J. Ian Munro. Tables. In *Proc. 16th Conference on Foundations of Software Technology and*
618 *Theoretical Computer Science (FSTTCS)*, LNCS 1180, pages 37–42, 1996.
- 619 40 J. Ian Munro, Yakov Nekrich, and Jeffrey Scott Vitter. Fast construction of wavelet trees.
620 *Theoretical Computer Science*, 638:91–97, 2016.
- 621 41 Gonzalo Navarro. *Compact Data Structures – A practical approach*. Cambridge University
622 Press, 2016.
- 623 42 Gonzalo Navarro. Indexing highly repetitive string collections, part I: Repetitiveness measures.
624 *ACM Computing Surveys*, 54(2):article 29, 2021.
- 625 43 Gonzalo Navarro. Indexing highly repetitive string collections, part II: Compressed indexes.
626 *ACM Computing Surveys*, 54(2):article 26, 2021.
- 627 44 Yakov Nekrich. Orthogonal range searching in linear and almost-linear space. *Computational*
628 *Geometry*, 42(4):342–351, 2009.
- 629 45 Ge Nong, Sen Zhang, and Wai Hong Chan. Two efficient algorithms for linear time suffix
630 array construction. *IEEE Transactions on Computers*, 60(10):1471–1484, 2011.
- 631 46 Enno Ohlebusch and Simon Gog. Lempel-Ziv factorization revisited. In *Proc. 22nd Annual*
632 *Symposium on Combinatorial Pattern Matching (CPM)*, LNCS 6661, pages 15–26, 2011.
- 633 47 Simon J. Puglisi, William F. Smyth, and Andrew Turpin. A taxonomy of suffix array
634 construction algorithms. *ACM Computing Surveys*, 39(2):article 4, 2007.

- 635 48 Michael Rodeh, Vaughan R. Pratt, and Shimon Even. Linear algorithm for data compression
636 via string matching. *Journal of the ACM*, 28(1):16–24, 1981.
- 637 49 Wojciech Rytter. Application of Lempel-Ziv factorization to the approximation of grammar-
638 based compression. *Theoretical Computer Science*, 302(1-3):211–222, 2003.
- 639 50 Esko Ukkonen. On-line construction of suffix trees. *Algorithmica*, 14(3):249–260, 1995.
- 640 51 Elad Verbin and Wei Yu. Data structure lower bounds on random access to grammar-
641 compressed strings. In *Proc. 24th Annual Symposium on Combinatorial Pattern Matching*
642 (*CPM*), LNCS 7922, pages 247–258, 2013.
- 643 52 Peter Weiner. Linear pattern matching algorithms. In *Proc. 14th Annual Symposium on*
644 *Switching and Automata Theory (SWAT)*, pages 1–11. IEEE Computer Society, 1973.