

Efficient Terabyte-Scale Text Compression via Stable Local Consistency and Parallel Grammar Processing

Diego Díaz-Domínguez  

University of Helsinki, Finland

Abstract

We present compression algorithms designed to process terabyte-sized datasets in parallel. Our approach builds on locally consistent grammars, a lightweight form of compression, combined with simple post-processing techniques to achieve further space reductions. Locally consistent grammar algorithms are suitable for scaling as they need minimal satellite information to compact the text, but they are not inherently parallel. To enable parallelisation, we introduce a novel concept that we call *stable local consistency*. A grammar algorithm ALG is stable if for any pattern P occurring in a collection $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$, instances $\text{ALG}(T_1), \text{ALG}(T_2), \dots, \text{ALG}(T_k)$ *independently* produce cores for P with the same topology. In a locally consistent grammar, the core of P is a subset of nodes and edges in the parse tree of \mathcal{T} that remains the same in all the occurrences of P . This feature enables compression, but it only holds if ALG defines a common set of nonterminal symbols for the strings. Stability removes this restriction, allowing us to run $\text{ALG}(T_1), \text{ALG}(T_2), \dots, \text{ALG}(T_k)$ in parallel and subsequently merge their grammars into a single output equivalent to that of $\text{ALG}(\mathcal{T})$. We implemented our ideas and tested them on massive datasets. Our experiments showed that our method could process 7.9 TB of bacterial genomes in around nine hours, using 16 threads and 0.43 bits/symbol of working memory, achieving a compression ratio of 85x.

2012 ACM Subject Classification Theory of computation \rightarrow Data compression

Keywords and phrases Grammar compression, locally consistent parsing, hashing

Digital Object Identifier 10.4230/LIPIcs.SEA.2025.14

Funding This project has received funding from the European Union’s Horizon Europe research and innovation programme under grant agreement No 101060011.

1 Introduction

Classical dictionary-based compression methods such as Lempel-Ziv (LZ) [27, 39] or grammar compression [23, 4] achieve significant space reductions, but often require extensive resources, limiting their practicality for large datasets. Tools like `gzip` and `zstd` provide resource-saving simplifications of LZ that offer acceptable trade-offs for smaller inputs, but still struggle with massive repositories.

Recent heuristics have been developed for large-scale applications. For example, Deorowicz et al. [8] compress pangenomes by partitioning strings and compressing similar segments together using `zstd`. Other approaches, Hunt et al. [17], reorder genomes to improve LZ compression. Grammar algorithms like RePair [26] and SEQUITUR [32] achieve high compression ratios, but quickly exceed the available memory as the input grows. Gagie et al. [14] introduced a method using prefix-free parsing [3] to scale RePair for large inputs.

Locally consistent grammars [15, 35, 5, 10, 24] is a technique that performs rounds of locally consistent parsing [29, 38, 30, 18, 3, 19, 5, 10] to compress a text $T[1..n]$. This approach recursively segments T based on sequence patterns, producing nearly identical phrases for matching substrings. In the parse tree of a locally consistent grammar, the nodes that cover the occurrences of a pattern P share an area with identical topology and



© Diego Díaz-Domínguez;

licensed under Creative Commons License CC-BY 4.0

23rd International Symposium on Experimental Algorithms (SEA 2025).

Editors: Petra Mutzel and Nicola Prezza; Article No. 14; pp. 14:1–14:20

Leibniz International Proceedings in Informatics



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

labels. This area is the *core* of P [38], and is what makes compression possible. Locally consistent grammars are simple to construct as, unlike LZ or RePair, they only need local information to break T consistently. However, they are not only useful for compression, they also help scaling the processing of large string collections. In fact, they have been used to speed up the computation of the Burrows–Wheeler transform [9], perform pattern matching in grammar-based self-indexes [5], find maximal exact matches [11, 31], among other things (see [1, 2, 20, 21, 22] for more applications).

Although it is possible to build from T a locally consistent grammar \mathcal{G} of size $O(\delta \log \frac{n \log \sigma}{\delta \log n})$ in $O(n)$ expected time [25, 24], δ being the string complexity [37] and σ being the alphabet size of T , these techniques probably yield less impressive compression ratios in practice than LZ or RePair. However, simple transformations to reduce the size of \mathcal{G} can yield further space reductions. In this regard, Ochoa et al. [36] showed that any irreducible grammar can reach the k th-order empirical entropy of a string. Their result suggests that building \mathcal{G} and then transforming it might be an efficient alternative to greedy approaches in large datasets.

Parallelising the grammar construction in massive collections is desirable to leverage multi-core architectures. Given an input $\mathcal{T} = \{T_a, T_b\}$, an efficient solution would be to split \mathcal{T} into chunks, say $\{T_a\}$ and $\{T_b\}$, compress the chunks in different instances $\text{ALG}(\{T_a\}) = \mathcal{G}_a$ and $\text{ALG}(\{T_b\}) = \mathcal{G}_b$, and merge the resulting (small) grammars \mathcal{G}_a and \mathcal{G}_b . However, ensuring the local consistency of the merged grammar is difficult without synchronising the instances. Most locally consistent algorithms assign random fingerprints to the grammar symbols to perform the parsing. However, when there is no synchronisation, different metasymbols emitted by $\text{ALG}(\{T_a\})$ and $\text{ALG}(\{T_b\})$ that expand to equal sequences of \mathcal{T} could have different fingerprints, thus producing an inconsistent parsing of T_a and T_b . Therefore, new locally consistent schemes are necessary to make the parallelisation possible.

Our contribution. We present a parallel grammar compression method that scales to terabytes of data. Our framework consists of two operations. Let $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ be a string collection with $\sum_{T_j \in \mathcal{T}} |T_j| = n$ symbols, and let $\mathcal{H} = \{h^0, h^1, \dots, h^l\}$ be a set of hash functions. The operation $\text{BUILDGRAM}(\mathcal{T}, \mathcal{H}) = \mathcal{G}$ produces a locally consistent grammar generating strings in \mathcal{T} . Furthermore, let \mathcal{T}_a and \mathcal{T}_b be two collections with $\text{BUILDGRAM}(\mathcal{T}_a, \mathcal{H}) = \mathcal{G}_a$ and $\text{BUILDGRAM}(\mathcal{T}_b, \mathcal{H}) = \mathcal{G}_b$. The operation $\text{MERGEGRAMS}(\mathcal{G}_a, \mathcal{G}_b)$ builds a locally consistent grammar \mathcal{G}_{ab} for the collection \mathcal{T}_{ab} that combines \mathcal{T}_a and \mathcal{T}_b . BUILDGRAM uses \mathcal{H} to induce a *stable* locally consistent parsing. The stable property means that $\text{BUILDGRAM}(\mathcal{T}_a, \mathcal{H})$ and $\text{BUILDGRAM}(\mathcal{T}_b, \mathcal{H})$ *independently* produce cores with the same topology for identical patterns. The set \mathcal{H} assigns random fingerprints to the metasymbols of the grammar under construction to guide the locally consistent parsing, with the fingerprint of a metasymbol X depending on the sequence of its expansion. This feature ensures that metasymbols from different grammars that expand to matching sequences get the same fingerprints. $\text{MERGEGRAMS}(\mathcal{G}_a, \mathcal{G}_b)$ leverages the stable property to produce a grammar equivalent to that of $\text{BUILDGRAM}(\mathcal{T}_{ab}, \mathcal{H})$, thus allowing us to parallelise the compression. We show that $\text{BUILDGRAM}(\mathcal{T}, \mathcal{H})$ runs in $O(n)$ time w.h.p. and uses $O(G \log G)$ bits of working space, G being the grammar size. Similarly, $\text{MERGEGRAMS}(\mathcal{G}_a, \mathcal{G}_b)$ runs in $O(G_a + G_b)$ time and uses $O(G_a \log g_a + G_b \log g_b)$ bits, with g_a and g_b being the number of symbols in \mathcal{G}_a and \mathcal{G}_b , respectively. The parsing that we use in BUILDGRAM is inspired by the concept of induced suffix sorting [34], which has been shown to be effective for processing strings [9, 11]. In future work, we plan to use our parallel compressor not only to reduce space usage but also to process large inputs. However, we note that the concept of stability is compatible with any locally consistent grammar that uses hashing to break the text. Our experiments showed that our strategy can efficiently compress several terabytes of data.

2 Notation and basic concepts

We consider the RAM model of computation. Given an input of n symbols, we assume our procedures run in random-access memory, where the machine words use $w = \Theta(\log n)$ bits and can be manipulated in constant time. We use the big- O notation to denote time and space complexities (in bits), and the term \log to express the logarithms of base two.

2.1 Strings

A string $T[1..n]$ is a sequence of n symbols over an alphabet $\Sigma = \{1, 2, \dots, \sigma\}$. We use $T[j]$ to refer to the j th symbol in T from left to right, and $T[a..b]$ to refer to the substring starting in $T[a]$ and ending in $T[b]$. An equal-symbol run $T[a..b] = s^c$ is a substring storing c consecutive copies of $s \in \Sigma$, with $a = 1$ or $T[a - 1] \neq s$; and $b = n$ or $T[b + 1] \neq s$.

We consider a collection $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ of k strings as a multiset where each element $T_j \in \mathcal{T}$ has an arbitrary order $j \in [1..k]$. In addition, we use the operator $\|\mathcal{T}\| = \sum_{T_j \in \mathcal{T}} |T_j|$ to express the total number of symbols. We also use subscripts to differentiate collections (e.g. \mathcal{T}_a and \mathcal{T}_b). The expression $\mathcal{T}_{ab} = \mathcal{T}_a \circ \mathcal{T}_b$ denotes the combination of \mathcal{T}_a and \mathcal{T}_b into a new collection \mathcal{T}_{ab} . We assume that all collections have the same constant-size alphabet Σ .

2.2 Grammar compression


Grammar compression consists in representing a string $T[1..n] \in \Sigma^*$ as a small context-free grammar that generates only T [23, 4]. Formally, a grammar $\mathcal{G} = \{\Sigma, V, \mathcal{R}, S\}$ is a tuple where Σ is the alphabet of T (the *terminals*), V is the set of *nonterminals*, $\mathcal{R} \subseteq V \times (\Sigma \cup V)^*$ is the set of rewriting rules in the form $X \rightarrow Q[1..q]$. The symbol $S \in V$ is the grammar's start symbol. Given two strings $w_a = A \cdot X \cdot B, w_b = A \cdot Q[1..q] \cdot B \in (\Sigma \cup V)^*$, w_b rewrites w_a (denoted $w_a \Rightarrow w_b$) if $X \rightarrow Q[1..q]$ exists in \mathcal{R} . Furthermore, w_a derives w_b , denoted $w_a \Rightarrow^* w_b$, if there is a sequence u_1, u_2, \dots, u_x such that $u_1 = w_a, u_x = w_b$, and $u_j \Rightarrow u_{j+1}$ for $1 \leq j < x$. The string $\text{exp}(X) \in \Sigma^*$ resulting from $X \Rightarrow^* \text{exp}(X)$ is the *expansion* of X , with the decompression of T expressed as $S \Rightarrow^* \text{exp}(S) = T$. Compression algorithms ensure that every $X \in V$ occurs only once on the left-hand sides of \mathcal{R} . In this way, there is only one possible string $\text{exp}(X)$ for each X . This type of grammar is referred to as straight-line. The sum of the lengths of the right-hand sides of \mathcal{R} is the grammar size.

The *parse tree* $PT(X)$ of $X \in V$ represents $X \Rightarrow^* \text{exp}(X)$. Given the rule $X \rightarrow Q[1..q]$, the root of $PT(X)$ is a node r labelled X that has q children, which are labelled from left to right with $Q[1], Q[2], \dots, Q[q]$, respectively. The j th child of r , labelled $Q[j]$, is a leaf if $Q[j] \in \Sigma$; otherwise, it is an internal node whose subtree is recursively defined according to $Q[j]$ and its rule in \mathcal{R} .

Post-processing a grammar consists of capturing the remaining repetitions in its rules. For instance, if $XY \in V^*$ appears multiple times on the right-hand sides of \mathcal{R} , one can create a new rule $Z \rightarrow XY$ and replace the occurrences of XY with Z . *Run-length compression* encapsulates each distinct equal-symbol $X^\ell \in (\Sigma \cup V)^*$ appearing in the right-hand sides of \mathcal{R} as a constant-size rule $X' \rightarrow (X, \ell)$. *Grammar simplification* removes every rule $X \rightarrow Q[1..q]$ whose symbol X appears once on the right-hand sides, replacing its occurrence with $Q[1..q]$.

2.3 Locally-consistent parsing and grammars

A *parsing* is a partition of a string $T[1..n]$ into a sequence of phrases $T = T[1..j_1 - 1]T[j_1..j_2 - 1] \dots T[j_x..n]$, where the indices $j_1 < j_2 < \dots < j_x$ are *breaks*. Let $\text{par}(o, o') = T[j_y..j_{y+1} - 1]T[j_{y+1}..j_{y+2} - 1] \dots T[j_{y+u-1}..j_{y+u} - 1]$ denote the u phrases that cover a substring $T[o..o']$,



■ **Figure 1** Locally consistent grammar compression of $P[1..33]$. The first row (bottom-up) is P , and the next rows are the metasymbols for parsing rounds. The grey boxes are breaks. The boxes below the thick black are the core of P . The values of A^i and Z^i change if the context of P changes.

with $o \in [j_y..j_{y+1})$ and $o' \in [j_{y+u-1}, j_{y+u})$. A parsing is *locally consistent* [6] iff, for any pair of equal substrings $T[a..b] = T[a'..b']$, $\text{par}(a, b)$ and $\text{par}(a', b')$ differ in $O(1)$ phrases at the beginning and $O(1)$ at the end, with their internal phrase sequences identical.

A locally consistent parsing scheme relevant to this work is that of Nong et al. [35]. They originally described their idea to perform induced suffix sorting, but it has been shown that it is also locally consistent [10, 7]. They define a *type* for each position $T[\ell]$:

$$\begin{aligned} \text{L-type} &\iff T[\ell] > T[\ell + 1] \text{ or } T[\ell] = T[\ell + 1] \text{ and } T[\ell + 1] \text{ is L-type.} \\ \text{S-type} &\iff T[\ell] < T[\ell + 1] \text{ or } T[\ell] = T[\ell + 1] \text{ and } T[\ell + 1] \text{ is S-type.} \\ \text{LMS-type} &\iff T[\ell] \text{ is S-type and } T[\ell - 1] \text{ is L-type.} \end{aligned} \quad (1)$$

Their method scans T from right to left and sets a break at each LMS-type position.

One can create a grammar $\mathcal{G} = \{\Sigma, V, \mathcal{R}, S\}$ that only generates T by applying successive rounds of locally consistent parsing. Different works present this idea slightly differently (see [38, 30, 18, 3, 5, 10]), but the procedure is similar: in every round i , the algorithm receives a string T^i as input ($T^i = T$ with $i = 1$) and performs the following steps:

1. Break T^i into phrases using a locally consistent parsing.
2. Assign a nonterminal X to each distinct sequence $Q[1..q]$ that is a phrase in T^i .
3. Store every X in V and its associated rule $X \rightarrow Q[1..q]$ in \mathcal{R} .
4. Replace the phrases in T^i by their nonterminals to form the string T^{i+1} for $i + 1$.

The process ends when T^i does not have more breaks, in which case it creates the rule $S \rightarrow T^i$ for the start symbol and returns the resulting grammar \mathcal{G} . The phrases have a length of at least two, so the length of T^{i+1} is half the length of T^i in the worst case. Consequently, the algorithm incurs in $O(\log n)$ rounds and runs in $O(n)$ time.

The output grammar \mathcal{G} is locally consistent because it compresses the occurrences of a pattern $P[1..m]$ largely in the same way. The first parsing round transforms the occurrences into substrings in the form $A^2[1..a^2] \cdot P^2[1..m^2] \cdot Z^2[1..z^2] \in V^*$, where the superscripts indicate symbols in T^2 . The blocks $A^2[1..a^2]$ and $Z^2[1..z^2]$ have $O(1)$ variable nonterminals that change with P 's context, while $P^2[1..m^2]$ remains the same in all occurrences. In the second round, $P^2[1..m^2]$ yields strings in the form $A^3[1..a^3] \cdot P^3[1..m^3] \cdot Z^3[1..z^3]$ that recursively have the same structure. The substring $P^i[1..m^i]$ remains non-empty during the first $O(\log m)$ rounds. The substrings P^i , with $i = 1, \dots, O(\log m)$, conform the *core* of P [38] (see Figure 1).

2.4 Hashing and string fingerprints

Hashing refers to the idea of using a function $h : \mathcal{U} \rightarrow [m]$ to map elements in a universe \mathcal{U} to integers in a range $[m] = \{0, 1, \dots, m - 1\}$ uniformly at random. When the universe \mathcal{U} is large and only an unknown subset $\mathcal{K} \subset \mathcal{U}$ requires fingerprints over a range $[m]$ with $|\mathcal{K}| < m \ll |\mathcal{U}|$, the typical solution is to use a *universal* hash function. For any pair $x, y \in \mathcal{U}$,

a universal hash function ensures that the collision probability is $\Pr[h(x) = h(y)] = 1/m$. Let $\mathcal{U} = \{0, 1, 2, \dots, \sigma\}$ be the universe, $p > \sigma$ a prime number, and $a \in [p]_+ = \{1, 2, \dots, p-1\}$ and $b \in [p] = \{0, 1, \dots, p-1\}$ integers chosen uniformly at random. One can make $h : \mathcal{U} \rightarrow [m]$ universal using the formula $h(x) = ((ax + b) \bmod p) \bmod m$, with $p > m$.

These ideas can be adapted to produce fingerprints for a set $\mathcal{K} \subset \Sigma^*$ of strings over an alphabet $\Sigma = \{1, 2, \dots, \sigma\}$ [12]. Pick a prime number $p > \sigma$ and choose an integer $c \in [p]_+$ uniformly at random. Then, build the degree q polynomial $h(Q[1..q]) = (\sum_{i=1}^q Q[i] \cdot c^{i-1}) \bmod p$ over $[p]$ and regard the symbols $Q[1], Q[2], \dots, Q[q]$ in $Q[1..q] \in \mathcal{U}$ as the polynomial's coefficients. Additionally, compose h with a universal hash function $h' : [p] \rightarrow [m]$ to obtain a fingerprint in $[m]$. Let $a, b, c \in [p]_+$ be three integers chosen uniformly at random. Then the function becomes $h'(Q[1..q]) = ((a (\sum_{i=1}^q Q[i] \cdot c^{i-1}) + b) \bmod p) \bmod m$.

3 Our methods

3.1 The grammar model

We first introduce the features of the locally consistent grammar $\mathcal{G} = \{\Sigma, V, \mathcal{R}, S\}$ we build with BUILDGRAM and MERGEGRAMS. Given a rule $X \rightarrow Q[1..q] \in \mathcal{R}$, the operator $rhs(X)$ is an alias for the string $Q[1..q]$, with $rhs(s) = s$ when $s \in \Sigma$. The grammar \mathcal{G} is *fully balanced*, which means that, for any $X \in V \setminus \{S\}$, all root-to-leaf paths in $PT(X)$ have the same number e of edges. We refer to e as the *level* of X , and extend the use of level to the rule $X \rightarrow Q[1..q] \in \mathcal{R}$ associated with X . We indistinctly use the terms nonterminal and metasymbol to refer to $X \in V$.

The start symbol S is associated with a rule $S \rightarrow C[1..k]$ where $C[1..k]$ is a sequence such that each $exp(C[j]) = T_j[1..n_j] \in \mathcal{T}$ and $PT(C[j])$ has height $O(\log n_j)$. The symbols in $C[1..k]$ can have different heights and, hence, different levels (this idea will become clear in Section 3.3). Let l_{max} be the highest level among the elements in $C[1..k]$. We set the level of S equal to $l = l_{max} + 1$, which we regard as the height of \mathcal{G} .

We define the partitions $\mathcal{R} = \{\mathcal{R}^1, \dots, \mathcal{R}^{l-1}\}$ and $V = \{V^1, V^2, \dots, V^{l-1}\}$, where every pair (\mathcal{R}^i, V^i) is the set of rules and nonterminals (respectively) with level $i \in [1..l-1]$. In each subset \mathcal{R}^i , the left-hand sides are symbols over the alphabet V^i , while the right-hand sides are strings over V^{i-1} . Further, we consider $V^0 = \Sigma$ to be the set of terminals.

3.2 Fingerprints for the grammar symbols

In this section, we describe the set \mathcal{H} of hash functions that assign fingerprints in BUILDGRAM. The universal hash function $h^0 : \Sigma \rightarrow [m_0]$ maps terminal symbols to integers over an arbitrary range $[0, 1, \dots, m_0 - 1]$, with $m_0 > \sigma$. Furthermore, each function h^i , with $1 \leq i < l$, recursively assigns fingerprints to the right-hand sides of \mathcal{R}^i . Let $[m_{i-1}]$ be the integer range for the fingerprints emitted by h^{i-1} . We choose a random prime number $p_i > m_{i-1}$, three integer values $a_i, b_i, c_i \in [p_i]_+$, and a new integer m_i . Now, given a rule $X \rightarrow Q[1..q] \in \mathcal{R}^i$, we compute the fingerprint for $Q[1..q]$ as

$$h^i(Q[1..q]) = \left(\left(a_i \left(\sum_{j=1}^q h^{i-1}(rhs(Q[j])) \cdot c_i^{j-1} \right) + b_i \right) \bmod p_i \right) \bmod m_i. \quad (2)$$

Although $h^i : [m_{i-1}]^* \rightarrow [m_i]$ computes a fingerprint for a string, we associate this fingerprint with $X \in V^i$ because each $Q[1..q]$ has one possible X . Notice that the recursive definition of h^i implicitly traverses $PT(X)$ and ignores the nonterminals labelling $PT(X)$.

As a result, the value that h^i assigns to $Q[1..q]$ (or equivalently, to X) depends on $\text{exp}(X)$, the functions h^0, h^1, \dots, h^i , and the topology of $PT(X)$. In practice, we avoid traversing $PT(X)$ by operating bottom-up over \mathcal{G} : When we process \mathcal{R}^i , the fingerprints in h^{i-1} for V^{i-1} that we require to obtain the fingerprints in h^i are available.

BUILDGRAM does not know a priori the number of hash functions \mathcal{H} needs to compress an input. However, the locally consistent grammar algorithm that we use requires $O(\log n)$ rounds of parsing to compress $T[1..n]$ (Section 3.3). If we consider the number of rounds plus the function h^0 for the terminals, then $|\mathcal{H}| \geq \lceil \log n \rceil + 1$ is enough to process $T[1..n]$.

3.3 Our grammar algorithm

Our procedure $\text{BUILDGRAM}(\mathcal{T}, \mathcal{H})$ receives as input a collection $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ of k strings and a set \mathcal{H} of hash functions, and returns a locally consistent grammar $\mathcal{G} = \{\Sigma, V, \mathcal{R}, S\}$ that only generates strings in \mathcal{T} . We assume \mathcal{H} has at least $\lceil \log n_{max} \rceil + 1$ elements (see Section 3.2), where n_{max} is the length of the longest string in \mathcal{T} .

The algorithm of BUILDGRAM is inspired by the parsing of Nong et al. [28], which has been used not only for compression [35], but also for the processing of strings [9] (see Section B). However, as noted in the Introduction, the ideas we present here and in the next sections are compatible with any locally consistent grammar that uses hashing.

Overview of the algorithm. BUILDGRAM constructs \mathcal{G} in successive rounds of locally consistent parsing. In each round i , we run steps 1-4 of Section 2.3, breaking the strings of \mathcal{T} individually, but collapsing the rules in the same grammar \mathcal{G} . When we finish round i , we flag each string $T_j^{i+1} \in \mathcal{T}^{i+1}$ with length one as inactive (that is, fully compressed) and stop the compression in round $i + 1$ if there are no active strings. Subsequently, we create the sequence $C[1..k]$ with compressed strings (that is, symbols we marked as inactive), create the start symbol $S \in V$ with the corresponding rule $S \rightarrow C[1..k] \in \mathcal{R}$, and finish BUILDGRAM.

3.3.1 Parsing mechanism

We parse the active strings of \mathcal{T}^i (step 1 of the round) using a variant of the parsing of Nong et al. [33] (Section 2.3) that employs Equation 2 to randomise the sequences that trigger breaks. We refer to this modification as RANDLMSPAR.

Let $X, Y \in V^{i-1}$ be any pair of nonterminals. We define the partial order \prec as follows:

$$X \prec Y \iff \begin{cases} h^{i-1}(\text{rhs}(X)) < h^{i-1}(\text{rhs}(Y)) & \text{if } h^{i-1}(\text{rhs}(X)) \neq h^{i-1}(\text{rhs}(Y)), \\ \text{undefined} & \text{if } h^{i-1}(\text{rhs}(X)) = h^{i-1}(\text{rhs}(Y)). \end{cases}$$

Additionally, we define the equivalence relation:

$$X \sim Y \iff h^{i-1}(\text{rhs}(X)) = h^{i-1}(\text{rhs}(Y)),$$

to cover the cases where $X = Y$ or $X \neq Y$ and their fingerprints collide. Now, let $T_j^i[\ell], T_j^i[\ell + 1] \in V^{i-1}$ be two adjacent positions in some string $T_j^i \in \mathcal{T}^i$ during round i . We redefine the types of Equation 1 for $T_j^i[\ell]$ as follows:

$$\begin{aligned} \text{L-type} &\iff T_j^i[\ell] \succ T_j^i[\ell + 1] \text{ or } T_j^i[\ell] \sim T_j^i[\ell + 1] \text{ and } T_j^i[\ell + 1] \text{ is L-type.} \\ \text{S-type} &\iff T_j^i[\ell] \prec T_j^i[\ell + 1] \text{ or } T_j^i[\ell] \sim T_j^i[\ell + 1] \text{ and } T_j^i[\ell + 1] \text{ is S-type.} \\ \text{LMS-type} &\iff T_j^i[\ell] \text{ is S-type and } T_j^i[\ell - 1] \text{ is L-type.} \end{aligned} \tag{3}$$

Let $rhs(X) = Q[1..q]$ be the replacement for X . We modify Equation 2 as follows:

$$F[X] = h^i(Q[1..q], F) = \left(\left(a_i \left(\sum_{j=1}^q F[Q[j]] \cdot c_i^{j-1} \right) + b_i \right) \bmod p_i \right) \bmod m_i. \quad (4)$$

This operation is valid because the alphabet of $Q[1..q]$ is V^{i-1} , and F already has its fingerprints.

A note on collisions. The consecutive positions $T_j^i[\ell] \neq T_j^i[\ell + 1]$ with $h^{i-1}(rhs(T^i[\ell])) = h^i(rhs(T^i[\ell + 1]))$ (i.e. collisions) never cause a break because $T^i[\ell] \sim T^i[\ell + 1]$. Intuitively, the more contiguous colliding symbols we have, the less breaks the parsing produces, and the more inefficient the compression becomes. The chances of this situation are small if the hash function $h^{i-1} \in \mathcal{H}$ emits fingerprints in the range $[m^{i-1}]$ with $|V^{i-1}| < m^{i-1}$. However, we do not know a priori $|V^{i-1}|$, so we have to choose a very large m^{i-1} . This decision has a trade-off because a large m^{i-1} means that cells of F require more bits, and hence more working memory. In Section 4, we investigate suitable values for m^i in large collections.

Now, we present the theoretical cost of BUILDGRAM.

► **Theorem 2.** *Let \mathcal{T} be a collection of k strings and $|\mathcal{T}| = n$ symbols, where the longest string has length n_{max} . Additionally, let \mathcal{H} be a set of hash functions with $|\mathcal{H}| \geq \lceil \log n_{max} \rceil + 1$ elements. $BUILDGRAM(\mathcal{T}, \mathcal{H}) = \mathcal{G}$ runs in $O(n)$ time w.h.p. and requires $O(G \log w)$ bits on top of \mathcal{T} , where G is the grammar size of \mathcal{G} .*

Proof. Calculating the type of each $T_j^i[\ell]$ in $T_j^i[1..n_j] \in \mathcal{T}^i$ takes $O(n_j)$ time if we have the array F with precomputed fingerprints of V^{i-1} . In addition, we can use a hash table to record the parsing phrases in T_j^i , which takes $O(n_x)$ time w.h.p. If we consider all the strings in \mathcal{T}^i , the running time of the parsing round i is $\mathcal{O}(|\mathcal{T}^i|)$ in expectation. On the other hand, all the phrases in T_j^i have length > 1 , except (possibly) one phrase at each end of T_j^i . Therefore, \mathcal{T}^{i+1} has $\frac{n}{2^i} + 2k$ symbols in the worst case. Considering that BUILDGRAM requires $O(\log n_{max})$ rounds, the cumulative length of $\mathcal{T}^1, \mathcal{T}^2, \dots, \mathcal{T}^{\log n_{max}}$ is at most $n + \sum_{i=1}^{\log n_{max}} \frac{n}{2^i} + 2k \leq 2n + 2k \log n_{max}$, with $2k \log n_{max}$ being the contribution of the phrases with length one. However, BUILDGRAM stops processing a string $T_j \in \mathcal{T}$ as soon as it is fully compressed, meaning that length-one phrases contribute $\sum_{T_j \in \mathcal{T}} \log |T_j| \leq k \log n_{max}$ elements in the worst case. Therefore, as $\sum_{T_j \in \mathcal{T}} \log |T_j| < \sum_{T_j \in \mathcal{T}} |T_j| = n$, the running time of BUILDGRAM is $O(n)$ w.h.p.

Let $g = |\Sigma| + |V|$ be the number symbols in \mathcal{G} . The $O(G \log G) + g \log w + |\mathcal{H}|w = O(G \log w)$ bits of working space in BUILDGRAMS represent the $O(G \log G)$ bits of the hash tables, the array F that stores the g fingerprints, and the hash functions in \mathcal{H} . ◀

3.5 Merging grammars

We now present our algorithm for merging grammars. Let \mathcal{T}_a and \mathcal{T}_b be two collections, with n_a and n_b being the lengths of the longest strings in \mathcal{T}_a and \mathcal{T}_b , respectively, and $\mathcal{T}_{ab} = \mathcal{T}_a \circ \mathcal{T}_b$ being their union (Section 2.1). Furthermore, let $\mathcal{G}_a = BUILDGRAM(\mathcal{T}_a, \mathcal{H})$ and $\mathcal{G}_b = BUILDGRAM(\mathcal{T}_b, \mathcal{H})$ be grammars that only generate strings in \mathcal{T}_a and \mathcal{T}_b , respectively. We assume that \mathcal{H} has $|\mathcal{H}| \geq \lceil \log \max(n_a, n_b) \rceil + 1$ elements (see Section 3.2). The instance MERGEGRAMS($\mathcal{G}_a, \mathcal{G}_b$) returns a locally consistent grammar \mathcal{G}_{ab} that only generates strings in \mathcal{T}_{ab} , and that is equivalent to the output of BUILDGRAM($\mathcal{T}_{ab}, \mathcal{H}$).

Overview of the Algorithm. The merging consists of making \mathcal{G}_a absorb the content that is unique to \mathcal{G}_b . Specifically, we discard rules from \mathcal{G}_b whose expansions occur in \mathcal{T}_a , and for those expanding to sequences not in \mathcal{T}_a , we add them as new rules in \mathcal{G}_a .

3.6 The merge algorithm

For the grammar \mathcal{G}_a , let \mathcal{R}_a be its set of rules, let V_a be its set of nonterminals, and let l_a be its height. The symbols \mathcal{R}_b , V_b , and l_b denote equivalent information for \mathcal{G}_b . We consider the partitions $\mathcal{R}_a = \{\mathcal{R}_a^1, \mathcal{R}_a^2, \dots, \mathcal{R}_a^{l_a-1}\}$ and $V_a = \{V_a^1, V_a^2, \dots, V_a^{l_a-1}\}$, where every (\mathcal{R}_a^i, V_a^i) is the set of rules and nonterminals (respectively) that BUILDGRAM($\mathcal{T}_a, \mathcal{H}$) produced during the parsing round $i \in [1..l_a - 1]$. The elements $\mathcal{R}_b = \{\mathcal{R}_b^1, \mathcal{R}_b^2, \dots, \mathcal{R}_b^{l_b-1}\}$ and $V_b = \{V_b^1, V_b^2, \dots, V_b^{l_b-1}\}$ are equivalent partitions for \mathcal{G}_b .

MERGEGRAMS processes the grammar levels $1, 2, \dots, \max(l_a, l_b) - 1$ in increasing order. In each round i , we keep the invariant that the right-hand sides of \mathcal{R}_a^i and \mathcal{R}_b^i are comparable. That is, given two rules $X_a \rightarrow Q_a[1..q_a] \in \mathcal{R}_a^i$ and $X_b \rightarrow Q_b[1..q_b] \in \mathcal{R}_b^i$, $Q_a[1..q_a] = Q_b[1..q_b]$ implies $\text{exp}(X_a) = \text{exp}(X_b) \in \Sigma^*$. When $i = 1$, the invariant holds as the right-hand sides of \mathcal{R}_a^1 and \mathcal{R}_b^1 are over Σ , which is the same for \mathcal{T}_a and \mathcal{T}_b (Section 2.1).

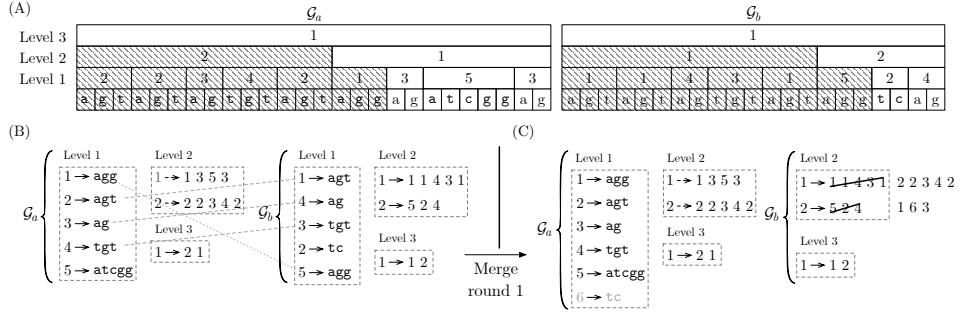
We begin the algorithm by creating an array $L_a[0..l_a]$ that stores in $L_a[i] = \sum_{j=0}^{i-1} |V_a^j|$ the number of nonterminals with level $< i$. Observe that $L_a[0] = 0$ and $L_a[1] = |\Sigma|$ because $V_a^0 = \Sigma$. We create an equivalent array L_b for \mathcal{G}_b . We also initialize the sets $E^1, E^2, \dots, E^{l_b-1}$ where every $E^i \subseteq [1..k_b]$ keeps the indexes of $C_b[1..k_b]$ with level- i symbols.

The first step of the merge round i is to scan the right-hand sides of \mathcal{R}_a^i , and for each $X_a \rightarrow Q_a[1..q_a]$, we modify $X_a = X_a - L_a[i]$ and $Q_a[j] = Q_a[j] - L_a[i - 1]$, with $j \in [1..q_a]$. The change $X_a = X_a - L_a[i]$ allows us to append new elements to V_a^i and \mathcal{R}_a^i while maintaining the validity of the symbols on the right-hand sides of \mathcal{R}_a^{i+1} , whose alphabet is V_a^i . Then, we create a hash table H_a that stores every modified rule $X_a \rightarrow Q[1..q_a] \in \mathcal{R}_a^i$ as a key-value pair $(Q_a[1..q_a], X_a)$, and an empty array $M_b[1..|V_b^i|]$ to update the right-hand sides of \mathcal{R}_b^{i+1} .

We check which right-hand sides in \mathcal{R}_b^i occur as keys in H_a . Recall that, when $i = 1$, the strings in \mathcal{R}_a^i are already comparable to the keys in H_a because they are over Σ and the subtraction of $L_a[i - 1 = 0]$ does not change their values. For $i > 1$, we make these strings comparable during the previous round $i - 1$. If the string $Q_b[1..q_b]$ of a rule $X_b \rightarrow Q_b[1..q_b]$ occurs in H_a as a key, we extract the associated value X_a from the hash table. On the other hand, if $Q_b[1..q_b]$ does not exist in H_a , we create a new symbol $X_a = |V_a^i| + 1$ and set $V_a^i = V_a^i \cup \{X_a\}$. Subsequently, we record the new rule $X_a \rightarrow Q_b[1..q_b]$ in \mathcal{R}_a^i and store $M[X_b - L_b[i]] = X_a$. Once we process \mathcal{R}_b^i , we scan the right-hand sides of \mathcal{R}_b^{i+1} and use M to update their symbols. We also use E^i to update each position $j \in E^i$ as $C_b[j] = M[C_b[j] - L_b[j]]$. Now, we discard \mathcal{R}_b^i and continue to the next merge round $i + 1$.

When we finish round i , the right-hand sides of \mathcal{R}_b^{i+1} are over $[1..|V_a^i|]$, and the right-hand sides of \mathcal{R}_a^{i+1} will be over $[1..|V_a^i|]$ after we update their values with L_a . These modifications will make both strings sets comparable, and our invariant will hold for $i + 1$.

After $\min(l_a, l_b) - 1$ rounds of merge, one of the input grammars will run out of levels. The remaining rounds will skip the creation and query of H_a , and will append new rules directly to \mathcal{R}_a^i (if any). After we finish the rounds, we concatenate the compressed strings to form $C_{ab}[1..k_a + k_b] = C_a[1..k_a] \cdot C_b[1..k_b]$ and update the starting rule $S_a \rightarrow C_{ab}[1..k_a + k_b]$. The last step of MERGEGRAMS is to modify the symbols in \mathcal{G}_a . For that purpose, we recompute L_a , and for every level- i rule $X_a \rightarrow Q_a[1..q_a]$, we set $X_a = X_a + L_a[i]$ and $Q_a[j] = Q_a[j] + L_a[i - 1]$, with $j \in [1..q_a]$. Figure 3 shows an example of MERGEGRAMS.



■ **Figure 3** Example of MERGEGRAMS($\mathcal{G}_a, \mathcal{G}_b$). As the parsing is stable, $T_a[1..19] = T_b[1..19]$ have cores (dashed boxes) with the same topology in \mathcal{G}_a and \mathcal{G}_b . In (B-C), the nonterminals represent the relative position of their rules in their corresponding levels. For example, the left-hand side of $2 \rightarrow 524$ in \mathcal{G}_b (side B) is 2 because that rule is the second in level 2. On the other hand, the symbol 2 in 524 refers to the second rule of level 1. In the first merge round, MERGEGRAMS checks which right-hand sides in level one of \mathcal{G}_b are also right-hand sides in level one of \mathcal{G}_a (dashed lines in side B). Only tc is not in \mathcal{G}_a , so the algorithm appends it at the end of level one in \mathcal{G}_a and assigns it the new metasymbol 6 (side C). Subsequently, it discards level one in \mathcal{G}_b and updates the right-hand sides of level two in \mathcal{G}_b according to their corresponding metasymbols in \mathcal{G}_a . In (C), the rule $1 \rightarrow 11431$ becomes $1 \rightarrow 22342$ and $2 \rightarrow 524$ becomes $2 \rightarrow 163$. For example, 5 becomes 1 on the right-hand side of $2 \rightarrow 524$ because the level one rule $5 \rightarrow \text{agg}$ in \mathcal{G}_b matches $1 \rightarrow \text{agg}$ in \mathcal{G}_a (see dashed lines in side B). After the update, MERGEGRAMS goes to the next round and operates recursively.

► **Theorem 3.** Let \mathcal{G}_a (respectively, \mathcal{G}_b) be a locally consistent grammar that generates strings in a collection \mathcal{T}_a of k_a strings (respectively, a collection \mathcal{T}_b with k_b strings). The size of \mathcal{G}_a is G_a and the size of \mathcal{G}_b is G_b . Similarly, let $g_a = |\Sigma| + |V_a|$ be the number of grammar symbols (respectively, $g_b = |\Sigma| + |V_b|$). MERGEGRAMS($\mathcal{G}_a, \mathcal{G}_b$) = \mathcal{G}_{ab} builds a locally consistent grammar generating strings in $\mathcal{T}_{ab} = \mathcal{T}_a \circ \mathcal{T}_b$ in $O(G_a + G_b)$ time w.h.p. and $O(G_a \log g_a + G_b \log g_b)$ bits of space

Proof. We obtain L_a and L_b in one scan of the nonterminal sets, which takes $O(g_a + g_b)$ time. It is also possible to obtain the sets E^1, E^2, \dots, E^{l_b} in $O(G_b)$ time. Let G_a^i (respectively, G_b^i) be the number symbols on the right-hand sides of \mathcal{R}_a^i (respectively, \mathcal{R}_b^i). Filling in the hash table H_a requires a linear scan of \mathcal{R}_a^i , which runs in $O(G_a^i)$ time w.h.p. On the other hand, scanning \mathcal{R}_b^i and querying its right-hand sides in H_a takes $O(G_b^i)$ time w.h.p. In addition, modifying the right-hand sides of \mathcal{R}_b^{i+1} with M takes $O(G_b^{i+1})$ time. If we transfer the cost of updating \mathcal{R}_b^{i+1} to the next round $i + 1$, performing the merge round i takes $O(G_a^i + G_b^i)$ time w.h.p. Now, let $g_a^i = |V_a^i|$ and $g_b^i = |V_b^i|$ be the number of level- i symbols. We require $G_a^i \log g_a^{i-1} + G_b^i \log g_b^{i-1}$ bits to encode \mathcal{R}_a^i and \mathcal{R}_b^i , $O(G_a^i \log g_a^{i-1})$ bits for H_a , and $g_b^i \log g_b^i$ bits for M . Consequently, the cost of the round i is $O(G_a^i \log g_a^{i-1} + G_b^i \log g_b^{i-1})$ bits. As $G_a = k_a + |\Sigma| + \sum_{j=1}^{l_a-1} G_a^j$ and $G_b = k_b + |\Sigma| + \sum_{j=1}^{l_b-1} G_b^j$, then MERGEGRAMS runs in $O(G_a + G_b)$ time w.h.p and uses $O(G_a \log g_a + G_b \log g_b)$ bits of space. ◀

4 Experiments

Implementation details. We implemented our framework in C++ in a tool called LCG (<https://github.com/ddiazdom/lcg>). We support parallel compression by interleaving executions of BUILDGRAM and MERGEGRAMS as follows: given a collection \mathcal{T} and an integer p , LCG uses p threads that execute BUILDGRAM in parallel to compress different subsets of \mathcal{T} into different buffer grammars. When the combined space usage of the buffer grammars

■ **Table 1** Performance of the competitor tools. The best result of each column is in bold. Cells with a dash are experiments that crashed or are incompatible with the input.

Tool	Compression ratio plain/compressed				Compression speed MB/s				Working memory bps			
	ATB	HUM	COVID	KERNEL	ATB	HUM	COVID	KERNEL	ATB	HUM	COVID	KERNEL
LCG	85.26	135.54	328.10	99.99	232.26	244.73	506.44	237.91	0.43	0.29	0.36	2.05
agc	–	144.90	237.93	–	–	120.42	53.71	–	–	0.15	0.18	–
zstd	58.19	4.72	344.99	38.23	95.51	27.21	442.79	226.45	0.004	0.01	0.10	0.53
RePair	–	–	778.62	127.03	–	–	1.89	2.21	–	–	30.63	150.11
BigRePair	–	–	586.12	109.87	–	–	47.89	21.04	–	–	2.28	3.69

exceeds a given threshold, LCG merges them into a sink grammar using MERGEGRAMS and resets the buffer grammars. Section A explains this idea in more detail. We refer to this strategy as PBUILDGRAM to differentiate it from our description in Section 3.3. After running PBUILDGRAM, LCG run-length compresses the output (RL step), and then removes unique nonterminals from the output of PBUILDGRAM + RL (SIMP step).

Experimental setup and inputs. We compared LCG against other state-of-the-art compressors, measuring the compression speed in MB per second (MB/s), the peak of the working memory in bits per input symbol (bps), and the compression ratio. Furthermore, we assessed the amount of compression LCG achieves, its resource usage, and how it scales with the number of threads. We conducted the experiments on a machine with AlmaLinux 8.4, 3 TiB of RAM, and processor Intel(R) Xeon(R) CPU E7-8890 v4 @ 2.20GHz, with 192 cores. We tested four collections. HUM: all the human genome assemblies available in NCBI up to August 27, 2024 (3.46 TB, $\sigma=16$). ATB: release 2.0 of the AllTheBacteria dataset [17], which contains genomes of bacteria and archaea (7.9 TB, $\sigma = 5$). COVID: all the SARS-CoV-2 genomes in NCBI up to November 5 (267.4 GB, $\sigma = 16$). KERNEL: last 40 versions of the Linux kernel (.h, .c, .txt, and .rst files) up to December 13, 2024 (54.4 GB, $\sigma = 190$).

Competitor tools. zstd (<https://github.com/facebook/zstd>) is an efficient tool that uses a simplified version of LZ and encodes the output using Huffman [16] and ANS [13]. agc (<https://github.com/refresh-bio/agc>) is a compressor for highly similar genomic sequences by Deorowicz et al. [8] that breaks strings into segments and groups segments into blocks according to sequence similarity. RePair (<https://users.dcc.uchile.cl/~gnavarro/software/repair.tgz>) is a popular grammar compression algorithm by Larsson and Moffat [26] that recursively replaces the most common pair of symbols in the text. BigRePair (<https://gitlab.com/manzai/bigrepair>) is a RePair variant by Gagie et al. [14] that scales the compression by using prefix-free parsing [3]. LCG, agc, zstd, and BigRePair support multithreading, so we used 16 threads in each. RePair does not support multithreading. For zstd, we used compression level 15 and a window size of 2 GiB—the tool does not allow longer windows with compression level ≥ 15 .

4.1 Results and discussion

4.1.1 Comparison with other text compressors

LCG was the fastest tool, with a speed ranging 232.26 – 506.44 MB/s (Table 1). The speed of the other tools varied with the input, but RePair remained the slowest (1.89-2.21 MB/s). The fact that RePair used one thread and the other tools 16 alone does not explain these results. For instance, LCG was 268 times faster than RePair in COVID (506.44 MB/s versus 1.89

■ **Table 2** Compression statistics. The value $g = |V|$ is number of nonterminals, $G = \sum_{X \in V} |rhs(X)|$ is the grammar size, $max(g^i) = \max_{i \in [1..l-1]} |V^i|$ is the maximum number of nonterminals in a level, $max(G^i) = \max_{i \in [1..l-1]} \sum_{X \in V^i} |rhs(X)|$ is the maximum level size. “Size reduction” is the percentage of decrease for G relative to previous grammar. The column “% nonter. increase” is the percentage of increase for g relative to the grammar of PBUILDGRAM.

Dataset	RePair		LCG							
	g	G	PBUILDGRAM				RL		SIMP	
			g	G	$max(g^i)$	$max(G^i)$	% nonter. increase	Size reduction	% deleted rules	Size reduction
ATB	–	–	9.5×10^9	2.9×10^{10}	3.3×10^9	1.0×10^{10}	6.68×10^{-3}	0.69	83.35	27.02
HUM	–	–	2.4×10^9	1.2×10^{10}	4.5×10^8	1.5×10^9	2.16×10^{-2}	36.73	72.30	22.31
COVID	1.7×10^7	1.1×10^8	1.1×10^8	5.1×10^8	1.5×10^7	1.8×10^8	5.22×10^{-3}	35.16	88.71	29.11
KERNEL	6.6×10^7	1.3×10^8	6.4×10^7	2.2×10^8	2.4×10^7	7.6×10^7	2.70×10^{-2}	1.76	82.25	24.62


MB/s, respectively). We also observed that LCG achieved higher speeds in more compressible inputs (e.g. COVID), probably because the hash tables recording phrases from the text (see encoding in Section A.1) have to perform more lookups than insertions –lookups are cheaper.

The most space-efficient tool was `zstd`, with a working memory usage of 0.004 – 0.53 bps. This result is due to the cap of 2 GiB that `zstd` uses for the LZ window, regardless of the input size. This threshold keeps memory usage low, but limits compression in large datasets where repetitiveness is spread (HUM or KERNEL). On the other hand, `zstd` with ATB yielded important space reductions probably because Hunt et al. [17] preprocessed ATB to place similar strings close to each other to improve LZ compression. LCG used far less working memory than the other grammar compressors (0.29 – 2.05 bps versus 30.63 – 150.11 of RePair and 2.28 – 3.69 of BigRePair), though it is still high compared to `zstd`.

RePair obtained the best compression ratios and is likely to outperform the other tools in ATB and HUM –where we could not run RePair. Compared to LCG, RePair achieved 2.37 times more compression in COVID and 1.27 times more in KERNEL. The difference was smaller with BigRePair, with RePair achieving 1.32 times more compression in COVID and 1.15 times in KERNEL. Despite LCG did not obtain the best ratio, it still achieves important reductions, and its trade-off between compression and resource usage seems to be the best. Besides, it is still possible to further compress in LCG by applying RePair or Huffman encoding. We think these additional steps would be fast, as they operate over a small grammar.

4.1.2 Breakdown of our method

Compression. PBUILDGRAM produced substantially larger grammars than RePair, being 69.32% larger in KERNEL and 363.63% larger in COVID (see Table 2). However, we expected this result as PBUILDGRAM is not greedy. Interestingly, PBUILDGRAM produced fewer nonterminals than RePair in KERNEL (6.4×10^7 versus 6.6×10^7 , respectively). The maximum number of nonterminals $max(g^i)$ produced in a round was $3.3 \times 10^9 < 2^{32} - 1$ (ATB), indicating that 32 bit fingerprints (i.e. $m^i < 2^{32}$ in \mathcal{H}) are likely to be enough to keep the number of collisions low in repetitive collections with dozens of TB, although nonrepetitive collections of this size might require larger fingerprints. Recall from Section 3.4 that the more bits we use for fingerprints, the less collisions we have, and hence less impact on the compression, but it also means more working memory. The effect of RL varied with the input: the grammar size reductions ranged between 0.69% and 36.73%, with a negligible increase in the number of nonterminals. RL performed well in HUM and COVID because they have long equal-symbol runs of Ns. On the other hand, SIMP removed 81.65% of the nonterminals and reduced the size of the grammar by 25.76%, on average.



■ **Figure 4** Performance of LCG. (A) Running time breakdown. (B) Memory peak breakdown. “Fps” are the fingerprints in PBUILDGRAM, and “Sat. data” are arrays and grammars of the buffers (Section A.1). (C) Performance of LCG in HUM relative to the number of threads. The left y axis is the compression speed and the right y axis is the memory peak.

Resource usage. The grammar encoding we chose (Section A.1) had an important effect on the usage of resources. LCG spent 93.2% of its running time, on average, executing PBUILDGRAM (Figure 4A). The bottleneck was the lookup/insertion of phrases in the hash tables H^1, H^2, \dots, H^{l-1} of the grammars when BUILDGRAM (subroutine of PBUILDGRAM) parsed its input text. These hash table operations are costly because of the high number of cache misses and string comparisons. In particular, the first three parsing rounds of BUILDGRAM are the most expensive because the text is not small enough and produces a large set of phrases (see Figure 6). Consequently, BUILDGRAM has to hash more frequently and in larger hash tables in those rounds. The impact of the other steps is negligible, with RL and SIMP accounting (on average) for 1.82% and 4.98% of the running time, respectively. The working memory usage (Figure 4B) varied with the input: in ATB and HUM, the sink grammar and the fingerprints of PBUILDGRAM accounted for 75% – 82% of the usage. The rest of the memory was satellite data: the arrays and grammars of the thread buffers. The results are different in COVID and KERNEL, where the memory usage is dominated by satellite data. The hash tables H^1, H^2, \dots, H^{l-1} imposed a considerable space overhead as they store their keys (i.e. parsing phrases) in arrays of 32-bit cells to perform fast lookups (via `memcmp`). We can reduce the memory cost by using hash tables that store keys using VBytes instead. In this way, the keys still use an integral number of bytes, and we can still use `memcmp` for lookups. Our preliminary tests (not shown here) suggest that using VBytes in the hash tables (and 32-bit fingerprints) reduces the space of the sink grammar by 39% in HUM, 47% in COVID, and 45% in KERNEL. If we consider that the buffer grammars use the same encoding, the change to VBytes could drastically reduce the peak of working memory in LCG. We can decrease working memory even further by keeping parts of the sink grammar on disk.

Effect of parallelism. The compression speed of LCG in HUM increased steadily with the threads (Figure 4C), while the working memory peak remained stable in 0.29 bps until 16 threads. After that, the peak increased to 0.31 bps with 20 threads, and 0.33 with 24.

5 Conclusions and further work

We presented a parallel grammar compressor that processes texts at high speeds while achieving high compression ratios. Our working memory usage is still high compared to popular general-purpose compressors like `zstd`, but we can greatly reduce the gap by using VByte encoding or keeping some parts of the grammar on disk. On the other hand, we use substantially less memory than popular grammar compressors. In fact, to our knowledge,

LCG is the only grammar-based tool that scales to terabytes of data. Furthermore, our simple strategy captures repetitions from distant parts of the text, making it more robust than other widely spread compression heuristics. Additional reductions in LCG are possible by using greedy methods, such as RePair, or statistical compression on top of the output grammar, but it will slow down the analysis of strings. As mentioned, our goal is not only to compress but also to scale string processing algorithms in massive collections. In the literature, it has been shown that the use of locally consistent grammars can speed up those algorithms [9, 11] but the efficient computation of the grammar remains a bottleneck. We solved that problem in this work. Integrating our scheme with those algorithms could enable the processing of an unprecedented volume of strings.

References

- 1 Tuğkan Batu and S Cenk Sahinalp. Locally consistent parsing and applications to approximate string comparisons. In *Proc. 9th International Conference on Developments in Language Theory (DLT)*, pages 22–35, 2005.
- 2 Or Birenzweige, Shay Golan, and Ely Porat. Locally consistent parsing for text indexing in small space. In *Proc. 31th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 607–626, 2020. doi:10.1137/1.9781611975994.37.
- 3 Christina Boucher, Travis Gagie, Alan Kuhnle, Ben Langmead, Giovanni Manzini, and Taher Mun. Prefix-free parsing for building big BWTs. *Algorithms for Molecular Biology*, 14:1–15, 2019. doi:10.1186/S13015-019-0148-5.
- 4 Moses Charikar, Eric Lehman, Ding Liu, Rina Panigrahy, Manoj Prabhakaran, Amit Sahai, and Abhi Shelat. The smallest grammar problem. *IEEE Transactions on Information Theory*, 51(7):2554–2576, 2005. doi:10.1109/TIT.2005.850116.
- 5 Anders Roy Christiansen, Mikko Berggren Ettiienne, Tomasz Kociumaka, Gonzalo Navarro, and Nicola Prezza. Optimal-time dictionary-compressed indexes. *ACM Transactions on Algorithms (TALG)*, 17(1):1–39, 2020. doi:10.1145/3426473.
- 6 Richard Cole and Uzi Vishkin. Deterministic coin tossing and accelerating cascades: micro and macro techniques for designing parallel algorithms. In *Proc. 18th Annual ACM Symposium on Theory of Computing (STOC)*, pages 206–219, 1986. doi:10.1145/12130.12151.
- 7 Jin-Jie Deng, Wing-Kai Hon, Dominik Köppl, and Kunihiko Sadakane. FM-indexing grammars induced by suffix sorting for long patterns. In *Proc. 22nd Data Compression Conference (DCC)*, pages 63–72. IEEE, 2022. doi:10.1109/DCC52660.2022.00014.
- 8 Sebastian Deorowicz, Agnieszka Danek, and Heng Li. AGC: compact representation of assembled genomes with fast queries and updates. *Bioinformatics*, 39(3):btad097, 2023. doi:10.1093/BIOINFORMATICS/BTAD097.
- 9 Diego Díaz-Domínguez and Gonzalo Navarro. Efficient construction of the BWT for repetitive text using string compression. *Information and Computation*, 294:105088, 2023. doi:10.1016/J.IC.2023.105088.
- 10 Diego Díaz-Domínguez, Gonzalo Navarro, and Alejandro Pacheco. An LMS-based grammar self-index with local consistency properties. In *Proc. 28th International Symposium on String Processing and Information Retrieval (SPIRE)*, pages 100–113, 2021. doi:10.1007/978-3-030-86692-1_9.
- 11 Diego Díaz-Domínguez and Leena Salmela. Computing all-vs-all MEMs in grammar-compressed text. In *Proc. 30th International Symposium on String Processing and Information Retrieval (SPIRE)*, pages 157–170. Springer, 2023. doi:10.1007/978-3-031-43980-3_13.
- 12 Martin Dietzfelbinger, Joseph Gil, Yossi Matias, and Nicholas Pippenger. Polynomial hash functions are reliable. In *Proc. 19th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 235–246, 1992.
- 13 Jarek Duda. Asymmetric numeral systems: entropy coding combining speed of huffman coding with compression rate of arithmetic coding. *arXiv preprint arXiv:1311.2540*, 2013.

- 14 Travis Gagie, Tomohiro I, Giovanni Manzini, Gonzalo Navarro, Hiroshi Sakamoto, and Yoshimasa Takabatake. Rpair: Rescaling RePair with rsync. In *Proc. 26th International Symposium on String Processing and Information Retrieval (SPIRE)*, pages 35–44, 2019. doi:10.1007/978-3-030-32686-9_3.
- 15 Paweł Gawrychowski, Adam Karczmarsz, Tomasz Kociumaka, Jakub Łącki, and Piotr Sankowski. Optimal dynamic strings. In *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1509–1528, 2018.
- 16 David A. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, 1952.
- 17 Martin Hunt, Leandro Lima, Wei Shen, John Lees, and Zamin Iqbal. AllTheBacteria—all bacterial genomes assembled, available and searchable, 2024. bioRxiv preprint. doi:10.1101/2024.03.08.584059.
- 18 Artur Jeż. A really simple approximation of smallest grammar. *Theoretical Computer Science*, 616:141–150, 2016. doi:10.1016/J.TCS.2015.12.032.
- 19 Dominik Kempa and Tomasz Kociumaka. String synchronizing sets: sublinear-time BWT construction and optimal LCE data structure. In *Proc. 51st Annual ACM Symposium on Theory of Computing (STOC)*, pages 756–767, 2019. doi:10.1145/3313276.3316368.
- 20 Dominik Kempa and Tomasz Kociumaka. Resolution of the burrows-wheeler transform conjecture. In *Proc. 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1002–1013, 2020. doi:10.1109/FOCS46700.2020.00097.
- 21 Dominik Kempa and Tomasz Kociumaka. Dynamic suffix array with polylogarithmic queries and updates. In *Proc. 54th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1657–1670, 2022. doi:10.1145/3519935.3520061.
- 22 Dominik Kempa and Tomasz Kociumaka. Collapsing the hierarchy of compressed data structures: Suffix arrays in optimal compressed space. In *Proc. 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1877–1886, 2023. doi:10.1109/FOCS57990.2023.00114.
- 23 John C. Kieffer and En Hui Yang. Grammar-based codes: a new class of universal lossless source codes. *IEEE Transactions on Information Theory*, 46(3):737–754, 2000. doi:10.1109/18.841160.
- 24 Tomasz Kociumaka, Gonzalo Navarro, and Francisco Olivares. Near-optimal search time in δ -optimal space, and vice versa. *Algorithmica*, 86(4):1031–1056, 2024. doi:10.1007/S00453-023-01186-0.
- 25 Tomasz Kociumaka, Gonzalo Navarro, and Nicola Prezza. Toward a definitive compressibility measure for repetitive sequences. *IEEE Transactions on Information Theory*, 69(4):2074–2092, 2022. doi:10.1109/TIT.2022.3224382.
- 26 N. Jesper Larsson and Alistair Moffat. Off-line dictionary-based compression. *Proceedings of the IEEE*, 88(11):1722–1732, 2000. doi:10.1109/5.892708.
- 27 Abraham Lempel and Jacob Ziv. On the complexity of finite sequences. *IEEE Transactions on Information Theory*, 22(1):75–81, 1976. doi:10.1109/TIT.1976.1055501.
- 28 Heng Li and Richard Durbin. Fast and accurate long-read alignment with Burrows-Wheeler Transform. *Bioinformatics*, 26(5):589–595, 2010. doi:10.1093/BIOINFORMATICS/BTP698.
- 29 Kurt Mehlhorn, Rajamani Sundar, and Christian Uhrig. Maintaining dynamic sequences under equality tests in polylogarithmic time. *Algorithmica*, 17:183–198, 1997. doi:10.1007/BF02522825.
- 30 S. Muthukrishnan and Süleyman Cenk Sahinalp. Approximate nearest neighbors and sequence comparison with block operations. In *Proc of 32nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 416–424, 2000. doi:10.1145/335305.335353.
- 31 Gonzalo Navarro. Computing MEMs and relatives on repetitive text collections. *ACM Transactions on Algorithms*, 21(1):1–33, 2024. doi:10.1145/3701561.
- 32 Craig G. Nevill-Manning and Ian H. Witten. Compression and explanation using hierarchical grammars. *The Computer Journal*, 40(2_and_3):103–116, 1997.

- 33 Ge Nong. Practical linear-time $O(1)$ -workspace suffix sorting for constant alphabets. *ACM Transactions on Information Systems*, 31(3):1–15, 2013.
- 34 Ge Nong, Sen Zhang, and Wai Hong Chan. Linear suffix array construction by almost pure induced-sorting. In *Proc. 19th Data Compression Conference (DCC)*, pages 193–202, 2009. doi:10.1109/DCC.2009.42.
- 35 Daniel Saad Nogueira Nunes, Felipe A. Louza, Simon Gog, Mauricio Ayala-Rincón, and Gonzalo Navarro. A grammar compression algorithm based on induced suffix sorting. In *Proc. 28th Data Compression Conference (DCC)*, pages 42–51, 2018. doi:10.1109/DCC.2018.00012.
- 36 Carlos Ochoa and Gonzalo Navarro. RePair and all irreducible grammars are upper bounded by high-order empirical entropy. *IEEE Transactions on Information Theory*, 65(5):3160–3164, 2018. doi:10.1109/TIT.2018.2871452.
- 37 Sofya Raskhodnikova, Dana Ron, Ronitt Rubinfeld, and Adam Smith. Sublinear algorithms for approximating string compressibility. *Algorithmica*, 65:685–709, 2013. doi:10.1007/S00453-012-9618-6.
- 38 Sühleyman Cenk Sahinalp and Uzi Vishkin. Symmetry breaking for suffix tree construction. In *Proc. 26th Annual ACM Symposium on Theory of Computing (STOC)*, pages 300–309, 1994. doi:10.1145/195058.195164.
- 39 Jacob Ziv and Abraham Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, 23(3):337–343, 1977. doi:10.1109/TIT.1977.1055714.

A Implementation details

This section presents the details of the implementation of PBUILDGRAM, the compression algorithm we implemented in LCG. We use BUILDGRAM (Section 3.3) and MERGEGRAMS (Section 3.5) as building blocks to parallelize the process without losing compression power. Section A.1 describes the encoding we use for locally consistent grammars, Section A.2 introduces some changes in BUILDGRAM to make parallel execution easier, and Section A.3 describes the steps of PBUILDGRAM.

A.1 Grammar encoding

Let $\mathcal{G} = \{\Sigma, V, \mathcal{R}, S\}$ be a locally consistent grammar of height l generating strings in $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$. Additionally, let (\mathcal{R}^i, V^i) be the i th pair, with $i \in [1..l-1]$, of the level-based partition of \mathcal{G} (see Section 3.5). A hash table H^i keeps for every $X \rightarrow Q[1..q] \in \mathcal{R}^i$ the key-value pair $(Q[1..q], X)$, while an array $F^i[1..|V^i|]$ stores the fingerprints in h^i of V^i . The array $F^0[1..\sigma]$ stores the fingerprints in h^0 for Σ . Additionally, the array $C[1..k]$ maintains the symbols representing the compressed strings of \mathcal{T} . Similarly to Section 3.6, the sets E^1, E^2, \dots, E^{l-1} store in $E^i \subseteq [1..k]$ the indexes of $C[1..k]$ with level- i symbols. Overall, the encoding of \mathcal{G} is $(\Sigma, F^0), (H^1, F^1, E^1), \dots, (H^{l-1}, F^{l-1}, E^{l-1})$, and C . We remark that the nonterminals in (\mathcal{R}^i, V^i) represent ranks. Specifically, the value $X = r$ for a rule $X \rightarrow Q[1..q]$ means that this rule is the r th in \mathcal{R}^i , while $Q[u] = r$, with $u [1..q]$, means that the rule where $Q[u] \in V^{i-1}$ is the left-hand side of the r th in \mathcal{R}^{i-1} . The use of ranks simplifies the merge with other grammars (see Section 3.6). From now on, we will use the operator $space(\mathcal{G})$ to denote the amount of bits for the encoding. We extend it so that $space(\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_p)$ represents $\sum_{j=1}^p space(\mathcal{G}_j)$. We implemented each hash table H^i using Robin Hood hashing, storing the nonterminal phrases in 32-bit cells. Finally, we implemented \mathcal{H} using the C++ library `xxHash`¹.

¹ <https://github.com/Cyan4973/xxHash>

A.2 Modification to the grammar algorithm

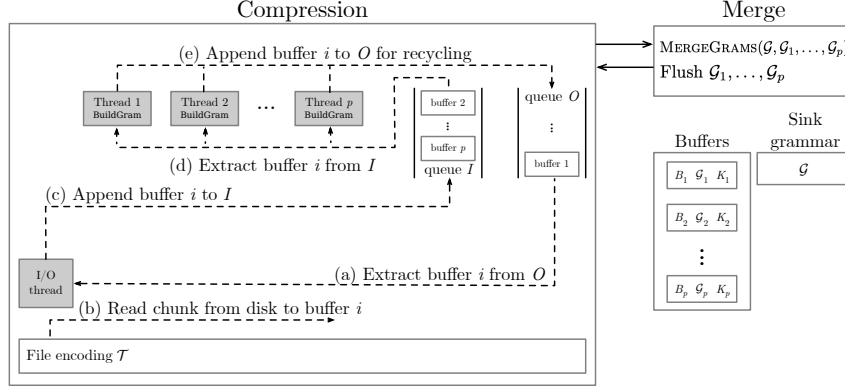
We modify BUILDGRAM to operate in parallel efficiently. The new signature of the algorithm is BUILDGRAM($\mathcal{T}, \mathcal{H}, \mathcal{G}_a, \mathcal{G}_b$) = \mathcal{G}_a , where \mathcal{G}_a and \mathcal{G}_b are two (possible nonempty) grammars, and the output is an updated version of \mathcal{G}_a . In this variant, we only record phrases in \mathcal{T} that are not in \mathcal{G}_a or \mathcal{G}_b . We store these phrases in \mathcal{G}_a and keep \mathcal{G}_b unchanged, using it as a “read-only” component. However, the general strategy to compress \mathcal{T} remains the same. The convenience of the change will become evident in the next section. We assume that \mathcal{G}_a and \mathcal{G}_b use the encoding we described in Section A.1. We add the subscript a or b to the encoding’s components to differentiate their origin, \mathcal{G}_a or \mathcal{G}_b (respectively).

BUILDGRAM now works as follows: let us assume that we are in the i th parsing round and that we receive the partially compressed collection \mathcal{T}^i as input. The alphabet of each $T_j^i \in \mathcal{T}^i$ is $V_a^{i-1} \cup V_b^{i-1}$, and we use the fingerprints in F_a^{i-1} and F_b^{i-1} to compute the types of each $T_j^i[\ell]$ (Equation 1) and thus the parsing phrases. Notice that, when $i = 1$, it holds $V_a^0 = V_b^0 = \Sigma$ and $F_a^{i-1} = F_b^{i-1}$. Let $Q[1..q] = T_j^i[\ell..\ell + q - 1]$ be the active phrase in the parsing. We first check if $Q[1..q]$ exists in H_b^i as a key. If that is the case, we get its corresponding value X in the hash table and assign it to the active phrase. On the other hand, if $Q[1..q]$ is not a key in H_b^i , we perform a lookup operation in H_a^i . Like before, if $Q[1..q]$ is a key there, we get the corresponding value X and assign it to the active phrase. Finally, if $Q[1..q]$ is not a key in H_a^i either, we insert the active phrase $T_j^i[\ell..\ell + q - 1] = Q[1..q]$ into H_a^i associated with a new metasymbol X . Let s_a and s_b be the sizes of H_a^i and H_b^i (respectively) before starting the parsing round i . The value we assign to the new metasymbol is $X = s_a + s_b + 1$. Subsequently, we replace $Q[1..q]$ by X in T_j^i and move to the next phrase. Before ending the parsing round, we store in F_a^i the fingerprints in h^i for the new phrases we inserted in H_a^i . Let $Q[1..q]$ be one of the new phrases, and let X be its metasymbol. We compute the fingerprint $h^i(Q[1..q])$ using Equation 4 and store the result in $F_a^i[X - s_b]$. Identifying the precedence of a symbol $T^{i+1}[\ell] \in V^i$ in the next round $i + 1$ is simple: $T^{i+1}[\ell] = Y > s_b$ means that the phrase associated with Y is a key in H_a^i and its fingerprint in h^i is in $F_a^i[Y - s_b]$ (s_b never changes). On the other hand, $Y \leq s_b$ means that its phrase is in H_b^i and its fingerprint in h^i is in $F_b^i[Y]$. Recall that we need F_a^i and F_b^i to compute the position types for T_x^{i+1} . Similarly, we can obtain the fingerprints in h^{i+1} for the new phrases of H_a^i by modifying Equation 4 to receive two fingerprint arrays F_a^i and F_b^i instead of one.

A.3 Parallel grammar construction

Now that we have explained our grammar encoding and the changes that we need to compress in parallel, we are ready to describe PBUILDGRAM. Figure 5 shows the steps in detail. We receive as input a string collection \mathcal{T} (a file), the number p of compression threads, and a threshold t indicating the approximate amount of working memory we can use. We first initialize p distinct buffers $(B_1, \mathcal{G}_1, K_1), (B_2, \mathcal{G}_2, K_2), \dots, (B_p, \mathcal{G}_p, K_p)$ and an empty grammar \mathcal{G} . Each buffer j is a triplet $(B_j, \mathcal{G}_j, K_j)$, where B_j is an array to store chunks of \mathcal{T} (one at a time), \mathcal{G}_j is the grammar where we compress the chunks we load into B_j , and K_j is an array of pairs $(x_1, y_1), (x_2, y_2) \dots$ indicating the chunks we have loaded into B_j . Specifically, $K_j[u] = (x, y)$ means that the u th chunk we loaded into B_j contained the subset $T_x, T_{x+1}, \dots, T_{x+y-1} \subseteq \mathcal{T}$. Notice that these strings lie contiguously in \mathcal{T} ’s file. There is also an array K with information equivalent to that of K_j but for \mathcal{G} .

PBUILDGRAM consists in a loop that interleaves two steps, compression and merge. During the compression step, p parallel threads compress the buffers B_1, B_2, \dots, B_p into the corresponding buffer grammars $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_p$, and continue doing it while $space(\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_p) < t$.



■ **Figure 5** Schematic representation of PBUILDGRAM. The steps (a-e) indicate the cycle of a buffer during the compression step.

When the space exceeds the threshold, a merge step collapses $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_p$ into the sink grammar \mathcal{G} and flushes the p buffers. The algorithm then enters a new iteration that restarts the compression from the point in \mathcal{T} where it left it the last time.

The compression step initializes an I/O thread that reads \mathcal{T} from the disk, loading the chunks sequentially from left to right in the arrays B_1, B_2, \dots, B_p . On the other hand, the p compression threads process the arrays B_1, B_2, \dots, B_p in parallel using the variant of BUILDGRAM we described in Section A.2. Every thread receives a buffer $j \in [1..p]$ and runs $\text{BUILDGRAM}(B_j, \mathcal{H}, \mathcal{G}_j, \mathcal{G}) = \mathcal{G}_j$. We synchronize the I/O thread with the compression threads using two concurrent queues I and O . The queue I keeps the chunks that are ready to be processed by the compressor threads, whereas O contains the buffers that were already processed and can be recycled by the I/O thread to insert new chunks. When the algorithm starts, O contains all buffers.

The synchronization process works as follows: let u be the next chunk of \mathcal{T} that PBUILDGRAM has to process. The I/O thread extracts the head $(B_j, \mathcal{G}_j, K_j)$ from O , reads the u th chunk $T_x, T_{x+1}, \dots, T_{x+y} \subseteq \mathcal{T}$ from disk, and loads it into B_j . Subsequently, it appends (x, y) to K_j , and finally it appends $(B_j, \mathcal{G}_j, K_j)$ to I . The I/O thread continues to process the next chunks $u+1, u+2, \dots$ in the same way as long as the compression process remains active. On the other hand, each compression thread tries to acquire the next buffer available $(B_j, \mathcal{G}_j, K_j)$ from the head of I . After the thread acquires the buffer and runs BUILDGRAM, it flushes B_j and pushes $(B_j, \mathcal{G}_j, K_j)$ into O , thus marking this buffer for recycling. Notice that a compression thread can process multiple noncontiguous chunks of \mathcal{T} and collapse their information into the same grammar \mathcal{G}_j . However, later in the execution of the algorithm, we modify \mathcal{G}_j using the information in K_j to fix this problem.

During the merge step, we execute $\text{MERGEGRAM}(\mathcal{G}, \mathcal{G}_j) = \mathcal{G}$ with each $j \in [1..p]$. It is possible to collapse the grammars in parallel in a merge sort fashion: let us assume w.l.o.g that $p+1$ is a power of two. Thus, $(p+1)/2$ threads execute in parallel processes $\text{MERGEGRAM}(\mathcal{G}_1, \mathcal{G}_2) = \mathcal{G}_1, \text{MERGEGRAM}(\mathcal{G}_3, \mathcal{G}_4) = \mathcal{G}_3, \dots, \text{MERGEGRAM}(\mathcal{G}, \mathcal{G}_p) = \mathcal{G}$ to produce the new grammars $\mathcal{G}_1, \mathcal{G}_3, \dots, \mathcal{G}$. Subsequently, the $(p+1)/4$ threads collapse the new grammars in the same way, and the process continues until only one \mathcal{G} remains. Every time we execute $\text{MERGEGRAM}(\mathcal{G}_j, \mathcal{G}_{j+1}) = \mathcal{G}_j$, we also concatenate the corresponding arrays K_j and K_{j+1} to keep track of the chunks of \mathcal{T} that the resulting \mathcal{G}_j encodes. Notice that after the merge, K has all the information of K_1, K_2, \dots, K_p . Finally, we reset the buffers and begin a new iteration of compression.

After we process all the chunks, we perform one last merge step to collapse the buffers in \mathcal{G} and then use the information in K to reorder the elements of $C[1..k]$. Once we finish, we return \mathcal{G} and complete the execution of PBUILDGRAM.

A.4 Storing the final grammar

As mentioned, we post-process the output of PBUILDGRAM using RL and then SIMP. The resulting file of this process (i.e., the output of LCG) uses $G \log g$ bits to encode \mathcal{R} , $g \log G$ bits to store pointers on the right-hand sides of \mathcal{R} , and $k \log w$ bits to store pointers to the compressed sequences of the strings in \mathcal{T} .

A.5 Advantage of our parallel scheme

Our parallel scheme can use a high number of threads with little contention (and thus achieve high compression speeds), while keeping the amount of working memory manageable. A thread executing BUILDGRAM($B_j, \mathcal{H}, \mathcal{G}_j, \mathcal{G}$) is the only one modifying B_j and \mathcal{G}_j , and although the sink \mathcal{G} can be accessed by other threads concurrently, they only read information (i.e., little to no contention). On the other hand, there is some contention when the threads modify the queues I and O concurrently to remove or insert buffers (respectively). However, the compression threads spend most of their time executing BUILDGRAM (modifying a queue is cheap), and it is unlikely that many of them attempt to access the same queue at the same time. On the other hand, the I/O thread might have more contention as we increase the number of threads because it has to compete to modify O .

As mentioned above, PBUILDGRAM keeps the consumption of working memory manageable as we add more threads. Specifically, $space(\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_p)$ is not bigger than $space(\mathcal{G})$ by a factor of p . In the first compression iteration, the sink grammar \mathcal{G} is empty and because the compression threads do not synchronize when they execute BUILDGRAM, the grammars $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_p$ will be redundant. Therefore, memory usage will grow rapidly at the beginning, exceeding the memory threshold t and trigger a merge. In this phase, we will collapse redundant content into \mathcal{G} and delete buffer grammars, thus reducing memory consumption. In the next compression iteration, \mathcal{G} will be non-empty, so every instance BUILDGRAM($B_j, \mathcal{H}, \mathcal{G}_j, \mathcal{G}$) = \mathcal{G}_j will only add to \mathcal{G}_j what is not in \mathcal{G} . In addition, \mathcal{G}_j will grow slower with every new merge iteration because there will be less “new” sequence content in \mathcal{T} . We note that memory usage still depends on $max(t, space(\mathcal{G}))$, with $space(\mathcal{G})$ depending, in turn, on the amount of repetitiveness in \mathcal{T} .

B Speeding up string processing algorithms (sketch)

In this section, we briefly explain how locally consistent grammars help speed up string processing algorithms (the idea might vary with the application). The most expensive operation in string algorithms is to find matching sequences. We use the fact that a grammar collapses redundant sequence information, so the search space in which a string algorithm has to operate is significantly smaller in the grammar than in the text.

Let $\mathcal{G} = \{\Sigma, V, \mathcal{R}, S\}$ be a locally consistent grammar generating elements in \mathcal{T} , and let ALG(\mathcal{G}) be a string processing algorithm. As before, we divide $\mathcal{R} = \{\mathcal{R}^1, \mathcal{R}^2, \dots, \mathcal{R}^{l-1}\}$ according to levels. We regard the right-hand sides of \mathcal{R}^i as a string collection and run (a section of) ALG using \mathcal{R}^i as input. The process will give some complete answers and some partial answers. We pass the complete answers to the next level $i + 1$ and use them as

satellite data. When ALG is recursive and returns from the recursion to level i , we use the new information to complete what is missing in the partial answers.

An example of this idea is the computation of maximal exact matches (MEM). Let $lcp(X, Y)$ be the longest common prefix between $exp(X)$ and $exp(Y)$, with $X, Y \in V^{i-1}$. Similarly, let $lcs(X, Y)$ be the longest common suffix of $exp(X)$ and $exp(Y)$. Assume that we have run a standard algorithm to compute the MEMs on the right-hand sides of \mathcal{R}^i and that we found a match $A[j_1..j_2] = B[k_1..k_2]$ between two right-hand sides A, B of \mathcal{R}^i . Completing the MEM requires us to obtain $lcs(A[j_1 - 1], B[k_1])$ and $lcp(A[j_2 + 1], B[k_2 + 1])$. Once we compute all the matches in \mathcal{R}^i , we use the output so that we can compute lcs and lcp at the next level $i + 1$. We still have to project the MEMs to text positions, but this operation is cheaper than computing MEMs in their plain locations, especially if the text is redundant.

The grammar we produce with PBUILDGRAM still requires some modifications to run string algorithms as described above, but the cost of this transformation is proportional to the grammar size, not the text size.

C Figures

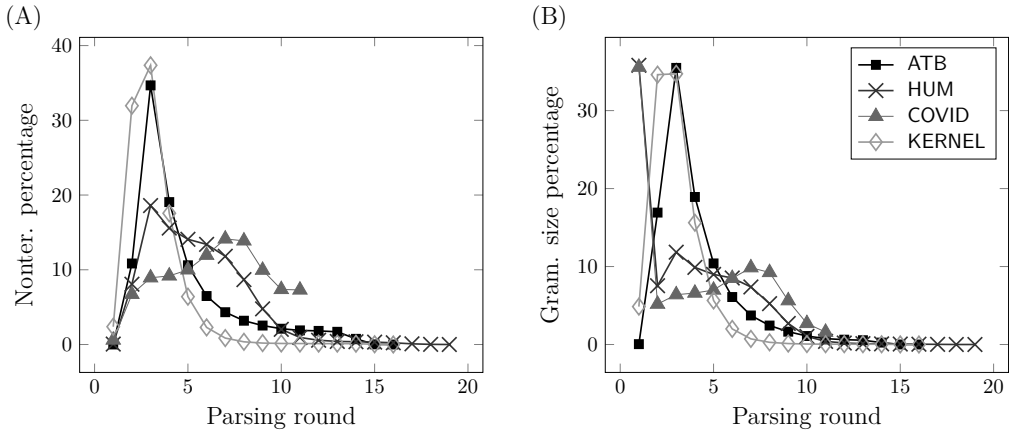


Figure 6 Number of nonterminals and phrases generated in each parsing round for the output grammar $\mathcal{G} = \{\Sigma, V, \mathcal{R}, S\}$ of PBUILDGRAM. The x-axes are the parsing rounds $i = 1, 2, \dots, l$. (A) Percentage $g^i/g \times 100$ that the number $g^i = |V^i|$ of level- i nonterminals contributes to the total number $g = |V|$ of nonterminals. (B) Percentage $G^i/G \times 100$ that the number of symbols $G^i = \sum_{X \in V^i} |rhs(X)|$ contributes to the total grammar size G . High percentages denote expensive rounds in each collection.