The Greedy Superstring Conjecture Is False

Nikolai Chukhin* Rishikesh Gajjala† Juha Harviainen‡
Alexander S. Kulikov* Ivan Mihajlin* Hiroki Shibata§ Alexander V. Smal*

*JetBrains Research †Pramaana Labs ‡University of Helsinki §Kyushu University

arXiv:2609.01365v2 [cs.DS] 6 Nov 2026

Abstract

In the Shortest Common Superstring problem (SCS), one is given a finite set of strings and is asked to find a shortest string containing every input string as a substring. Since the problem is NP-hard, various approximation algorithms have been proposed for it. The currently best known approximation ratio is 2.466 due to Englert, Matsakis, and Veselý; a better approximation ratio 2 has been announced by OpenAI just recently.

In 1988, Tarhio and Ukkonen introduced a simple greedy algorithm that iteratively takes two strings with the longest overlap and merges them together. They showed that the greedy algorithm may produce a superstring that is about two times longer than the optimal one and conjectured that this bound is tight. This has become known as the Greedy Superstring Conjecture and has remained open for almost four decades. The best known upper bound on the approximation ratio of the greedy algorithm is 3.396 due to Englert, Matsakis, and Veselý.

In this paper, we present a series of results tightening the bounds of the greedy algorithm, with the main contributions being as follows:

1. We refute the Greedy Conjecture by constructing a family of instances for which a lower bound on the approximation ratio tends to 2.4.
2. We prove that the greedy algorithm is 3-approximate.
3. For input strings of fixed length k, we prove that the approximation guarantee is exactly $9/5$ for $k = 3$ and at least 2 for all $k \geq 6$.

1 Introduction

In the Shortest Common Superstring problem (SCS), one is given a finite set of strings and asked to find a shortest string containing every input string as a substring. The problem is a classical model of assembling highly overlapping data and appears in data compression and genome assembly [GP14], where one obtains short subsequences of the DNA and wants to reconstruct the whole sequence under the assumption that the shortest superstring is the most plausible true sequence. SCS is NP-hard already for strings of length three [GMS80], and, for this reason, polynomial-time approximation algorithms are actively studied. The first constant-factor algorithm, with approximation ratio 3, was given by Blum et al. [BJL+91]. Subsequent work improved this ratio through increasingly delicate combinations of cycle covers and maximum asymmetric traveling salesperson algorithms. The best published approximation ratio is 2.466 [EMV23]; recently, OpenAI announced a 2-approximation [Ope26].

There however exists also a simple, almost four-decade-old greedy algorithm that is conjectured [TU88, Sto88, Tur89, BJL+91] to be equally good: while more than one string remains, choose an ordered pair with maximum suffix–prefix overlap, merge the pair, and iterate. The conjecture that the greedy algorithm is 2-approximate has then become known as the Greedy Superstring Conjecture. What makes this algorithm particularly appealing is that it admits a linear-time implementation over a constant-size alphabet [Ukk90] and is known to behave well in practice [RBT04, Ma09, SVB23]. Whereas it is not difficult to show that its approximation ratio is at least 2, the best known upper bound is 3.396 [EMV23]. Figure 1 summarizes the advancements on bounding the approximation guarantee of the greedy algorithm and the best known algorithms for SCS.

For many other NP-hard problems, the behavior of a natural greedy algorithm has been understood for decades. Take, for instance, the classical Vertex Cover problem: the algorithm "while there is an uncovered edge, add both of its endpoints to the cover" has approximation ratio 2 (folklore, see [Vaz01, Algorithm 1.2]). Similarly, for Set Cover, the algorithm "while there are uncovered elements, add a set covering the largest number of them" has approximation ratio $\ln n$ [Joh74]. Finally, for List Scheduling, the algorithm "while there are unscheduled jobs, assign the next job to the least loaded machine" has approximation ratio $2 - 1/\mu$ on μ machines [Gra66].

For this reason, it is surprising that the approximation ratio of such a remarkably simple greedy algorithm for SCS is still unknown. The Greedy Conjecture has been open for almost four decades, and even in the special case of strings of a fixed length the behavior of the greedy algorithm is poorly understood.

Our Contributions

Our first main contribution is to refute the Greedy Conjecture[1]. We construct a family of instances of increasing size for which a lower bound on the approximation ratio of the greedy algorithm tends arbitrarily close to 2.4. For this construction depending on the choices of two parameters, the strings use a binary alphabet and a particular execution of the greedy algorithm results in a superstring whose length is a function of the two parameters. Since the worst-case approximation ratio of the greedy algorithm is independent of the tie-breaking rule [Nik21], this lower bound holds for every tie-breaking rule. We then give an explicit construction for a string containing all the input strings as a substring. Although this is not necessarily the shortest superstring, it upper bounds the optimal solution and thus yields a lower bound for the approximation ratio. Increasing the values of parameters then brings the lower bound towards 2.4. The strings themselves contain only four

[1] Some preliminary versions of the results of our manuscript are contained in the preprints [CKMS26a], [CKMS26b], and [Shi26].

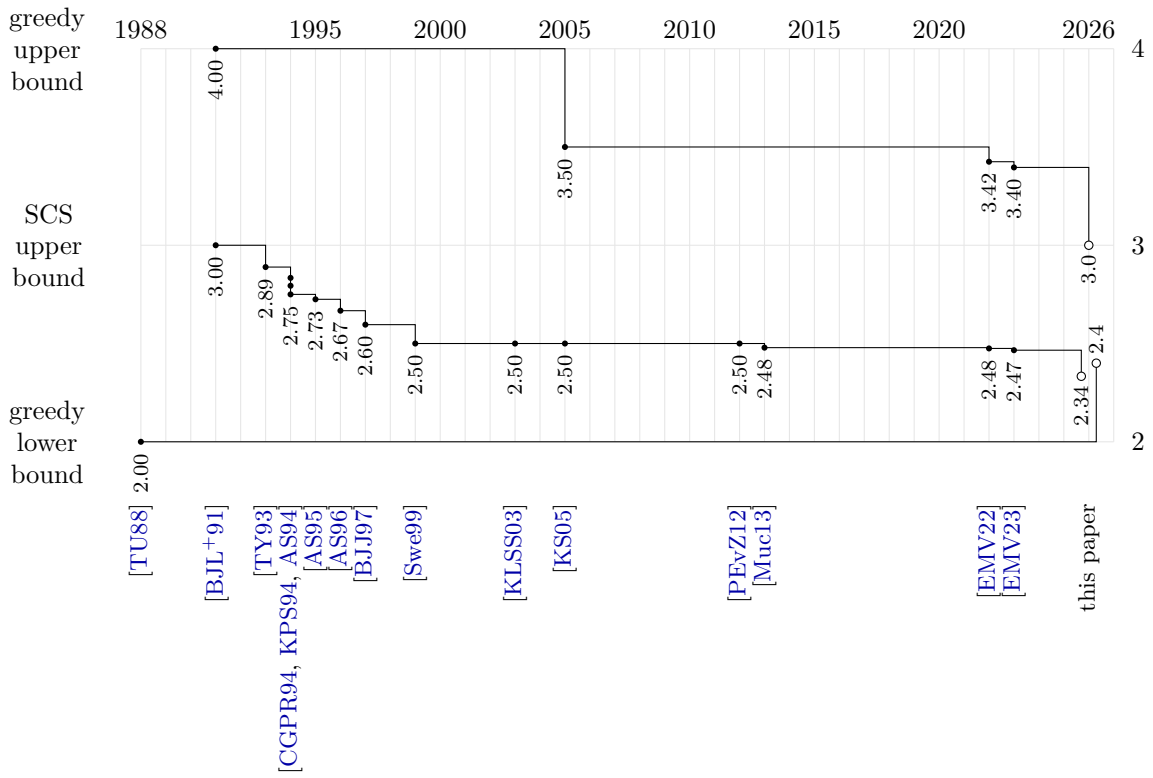


Figure 1: Timeline of approximation ratios for SCS (middle curve), upper bounds on the approximation ratio of the greedy algorithm (top curve), and lower bounds on it (bottom curve).

or five ones, which are separated from each other by gaps of varying numbers of consecutive zeros. These gaps are carefully designed to enforce which strings get merged together such that afterwards the greedy algorithm has only merges with overlap 0 remaining, with the intermediate strings starting with zeros and ending with ones.

Theorem 1.1 (Informal version of Theorem 4.1)**.** *For every $\varepsilon > 0$, there is a set of binary strings on which the greedy algorithm may output a common superstring of length at least* $(12/5 - \varepsilon)\,\mathrm{OPT}$.

To complement our lower bound, we also tighten the upper bound for the approximation guarantee of the greedy algorithm to 3. For this purpose, we start from the framework of Blum et al. [BJL+91] for proving upper bounds for SCS. The approach constructs the *overlap graph* of the input strings, which is a complete directed graph (with self-loops) where the weight of an edge from a string to another tells the number of symbols by which they overlap. Solving SCS reduces to finding the maximum-weight Hamiltonian path in the overlap graph, since to minimize the length of the superstring we want to maximize the sum of overlaps in merging the strings. The greedy algorithm for SCS then constructs a Hamiltonian path by iteratively adding edges of maximum weight such that adding the edge does not create a cycle nor increase the in- or out-degrees of a vertex to more than one. The latter constraint corresponds to each string being merged at most once on its left and right side. Letting the greedy algorithm add also edges resulting in cycles yields an optimal cycle cover of the overlap graph. This cycle cover can then be turned into a Hamiltonian path by removing one edge from each cycle with minimum overlap and merging the paths together. Our key contribution for tightening the upper bound is tightening a certain inequality from the

framework of Blum et al. [BJL+91] to its limit by proving the exact value of a particular coefficient appearing in the proof. New approaches would thus be needed to push the upper bound further.

Corollary 1.2. *The greedy algorithm for SCS is a 3-approximation algorithm.*

MGREEDY [BJL+91] concatenates the strings obtained by opening these cycles; TGREEDY instead applies the greedy algorithm to those strings. The original bound of 4 on MGREEDY also yielded the first 4-approximation guarantee for the greedy algorithm. We prove that MGREEDY has approximation ratio exactly 3 for every tie-breaking rule (Corollary 3.4). For TGREEDY, we prove an upper bound of 5/2 for every execution and, on the same binary instances used for the greedy algorithm, a lower bound of 12/5 for some choices of ties in both stages (Corollaries 3.5 and 4.4).

Combining Theorem 3.1 with the 2/3-approximation algorithm for MAX-ATSP [KLSS03, PEvZ12] gives the following corollary.

Corollary 1.3. *There is a polynomial-time 7/3-approximation algorithm for SCS.*

Finally, we consider the performance of the greedy algorithm on strings of fixed length, asking if the performance improves on shorter strings. Let ρ_k be the approximation ratio for strings of length k, that is, the supremum of $|g|/\operatorname{OPT}$ over all sets of strings of length k and all outputs g of the greedy algorithm on them. Cazaux and Rivals [CR18] proved that $2 - 1/k \le \rho_k \le (k+1)/2$, which gives $\rho_2 = 3/2$, and they conjectured that $\rho_k = 2 - 1/k$ for every k. For $k = 4$, Kulikov, Savinov, and Sluzhaev [KSS15] proved that the greedy algorithm is a 2-approximation if it breaks ties in favor of strings of the form $aaaa$, whereas for arbitrary tie-breaking the bounds of [CR18] give $7/4 \le \rho_4 \le 5/2$. We prove that $\rho_3 = 9/5$ and $\rho_k \ge 2$ for $k \ge 6$, which disproves the conjecture of Cazaux and Rivals for $k = 3$ and for every $k \ge 6$. Together with Corollary 1.2, this gives $2 \le \rho_k \le 3$ for every $k \ge 6$.

For $k = 3$, we determine ρ_3 exactly, which gives the first exact value of ρ_k for any $k \ge 3$. For the lower bound $\rho_3 \ge 9/5$, we use ℓ blocks of five strings; in the de Bruijn graph on strings of length two, the strings of each block form two cycles with two and three edges that share a vertex, and all blocks share a common vertex. Hence a single path visits all strings, whereas the greedy algorithm may merge the strings of each cycle separately, after which no two remaining strings overlap.

The proof for the upper bound is more involved: first, the greedy algorithm merges strings with overlap 2, after which the remaining overlaps are of 1 symbol, and we analyze these two phases separately. For overlaps of length 2, we again consider the de Bruijn graph to obtain its trail cover through the merges. We then consider the overlap-one merges by studying a multigraph constructed from the trail cover. Finally, we analyze the types of trails to bound the length of the greedy output.

Theorem 1.4 (Informal version of Theorem 5.1 and Proposition 5.2)**.** *The approximation ratio of the greedy algorithm on strings of length three is $\rho_3 = 9/5$.*

Similarly, for every $k \ge 6$, our instance consists of ℓ blocks of $2k-3$ strings of length k; in the de Bruijn graph on strings of length $k-1$, the strings of each block form two cycles with k and $k-3$ edges that share a vertex, and all blocks share a common vertex. Hence a single path visits all $n = (2k-3)\ell$ strings with overlaps $k-1$, which gives a shortest common superstring of length $n+k-1$. The greedy algorithm, however, may merge the strings of each cycle separately, after which the remaining strings overlap by at most one symbol, and it outputs a superstring of length $2n$.

Theorem 1.5 (Informal version of Theorem 6.1)**.** *For every $k \ge 6$, the approximation ratio of the greedy algorithm on strings of length k is at least 2, that is, $\rho_k \ge 2$.*

The initial proofs for the above results were obtained by using GPT-5.6 Sol Pro, GPT-6 Astra, and Claude Opus-5.5. We have rewritten them with the goal of simplifying them and providing more intuition. See AI methodology for further details.

2 Known Technical Results

We start by listing important known results on SCS; we introduce related concepts along the way. Gray boxes are used to visualize these concepts and facts; they are intended to help the reader and may be safely skipped.

In the rest of the paper, $\mathcal{S} = \{s_1, \ldots, s_n\}$ is an input to SCS. We assume that $\mathcal{S}$ is *substring-free*, that is, no string in $\mathcal{S}$ is a substring of another string in $\mathcal{S}$. Strings contained in other input strings can be removed during preprocessing. By OPT we denote an optimal common superstring of $\mathcal{S}$ or its length.

2.1 Prefixes and Overlaps

For nonempty strings s and t, their *overlap* $\mathrm{ov}(s,t)$ is the longest (possibly empty) string q such that $s = p \circ q$ and $t = q \circ r$, for some *nonempty* strings p and r that are in turn denoted by $\mathrm{pref}(s,t)$ and $\mathrm{suff}(s,t)$. For example, $\mathrm{ov}(\texttt{ababa}, \texttt{babc}) = \texttt{ba}$ and $\mathrm{ov}(\texttt{ababa}, \texttt{ababa}) = \texttt{aba}$. We call $\mathrm{ov}(s,s)$ the *self-overlap* of s. By $\langle s,t\rangle$ we denote the *overlap merge* of s and t:

$$\langle s,t\rangle = s \circ \mathrm{suff}(s,t) = \mathrm{pref}(s,t) \circ t = \mathrm{pref}(s,t) \circ \mathrm{ov}(s,t) \circ \mathrm{suff}(s,t).$$

Although the overlap and the prefix are defined as strings, we use these two words to denote the lengths of the corresponding strings. The notion of the overlap merge extends to more than two strings in a natural way: for $t_1, \ldots, t_k$,

$$\langle t_1, \ldots, t_k\rangle = \mathrm{pref}(t_1,t_2) \circ \mathrm{pref}(t_2,t_3) \circ \cdots \circ \mathrm{pref}(t_{k-1},t_k) \circ t_k.$$

Clearly,

$$|\langle t_1, \ldots, t_k\rangle| = \sum_{i\in[k-1]} |\mathrm{pref}(t_i,t_{i+1})| + |t_k| = \sum_{i\in[k]} |t_i| - \sum_{i\in[k-1]} |\mathrm{ov}(t_i,t_{i+1})|.$$

Hence, the smaller the sum of adjacent prefixes, the larger the sum of adjacent overlaps.

2.2 Permutations and Related Graphs

The goal in the SCS problem is to find a permutation of $\mathcal{S}$ whose overlap merge length is minimum. This is equivalent to maximizing the total overlap or minimizing the total prefix (up to the length of the last string in the permutation). The greedy algorithm tries to construct a permutation with large total overlap by taking, at every step, the largest available overlap. Formally, it maintains a set of *current strings*, which is initially $\mathcal{S}$. While more than one current string remains, it chooses an ordered pair of distinct current strings with the longest overlap, breaking ties arbitrarily, and replaces them by their overlap merge. The last remaining current string is the output.

For a set of strings $\{\texttt{abcdab}, \texttt{cdabc}, \texttt{dabcd}, \texttt{ababa}, \texttt{aaa}, \texttt{babab}\}$, a naive superstring results from concatenating them and has length 6+5+5+5+3+5 = 29: abcdabcdabcdabcdababaaaababab. An optimal superstring has length 14: aaababababcdabcd. It corresponds to a permutation $(\texttt{aaa}, \texttt{ababa}, \texttt{babab}, \texttt{abcdab}, \texttt{cdabc}, \texttt{dabcd})$. Its total overlap is $1 + 4 + 2 + 4 + 4 = 15$, exactly the number of symbols saved relative to the naive concatenation of length 29.
The largest overlap is 4; therefore, the greedy algorithm may start by merging $\langle \texttt{dabcd}, \texttt{abcdab}, \texttt{cdabc}\rangle$, after which it may merge $\langle \texttt{babab}, \texttt{ababa}\rangle$. Afterwards, there are no overlaps larger than 1, and the algorithm may finish with the following superstring

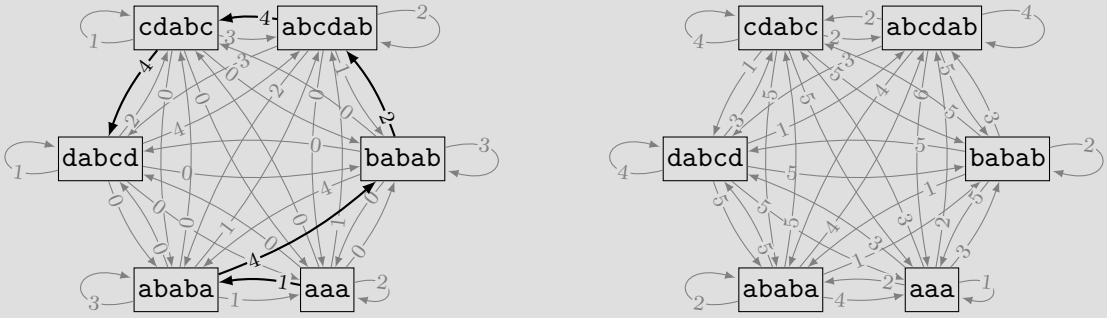
of length 16:

$$\langle \texttt{babab}, \texttt{ababa}, \texttt{aaa}, \texttt{dabcd}, \texttt{abcdab}, \texttt{cdabc} \rangle = \texttt{bababaaadabcdabc}.$$

All pairwise overlaps between input strings can be conveniently represented in an *overlap graph*: it is a complete directed graph with vertices $\mathcal{S}$ and edges $\mathcal{S} \times \mathcal{S}$ (thus, self-loops are included) where the length of an edge (s, t) is $|\mathrm{ov}(s, t)|$. A *prefix graph* is defined similarly with the only difference that the lengths are $|\mathrm{pref}(s, t)|$.

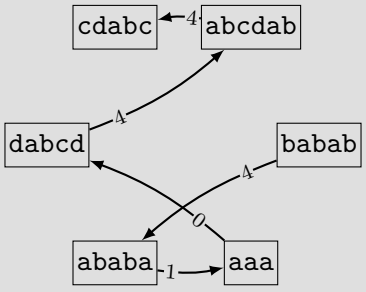
Solving SCS is the same as solving MAX-ATSP in the overlap graph, that is, finding a Hamiltonian path of maximum total length. At the same time, approximation algorithms for MAX-ATSP do not give approximation guarantees for SCS directly since the total overlap may be much larger than OPT.

The overlap and prefix graphs for our running example are shown below. In the overlap graph, a Hamiltonian path of maximum total overlap 15 is highlighted.



The greedy superstring algorithm can be conveniently represented in terms of the overlap graph: it constructs a Hamiltonian path by scanning the list of edges in the order of nonincreasing length and takes the current edge into a solution if it does not create a node of in-degree or out-degree more than one and does not close a cycle. This way, at any point of the execution of the greedy algorithm, its solution is a collection of disjoint paths.

The greedy algorithm starts by considering edges of length 4. It may take edges `dabcd` $\to$ `abcdab` and `abcdab` $\to$ `cdabc`. Then, it skips the edge `cdabc` $\to$ `dabcd` as it would create a cycle. The greedy algorithm then takes an edge `babab` $\to$ `ababa`. Finally, by skipping a bunch of edges, it takes edges `ababa` $\to$ `aaa` and `aaa` $\to$ `dabcd`.



This representation relies on the following fact, which we will use repeatedly in our analysis of the greedy algorithm.

Lemma 2.1 ([BJL+94])**.** *Throughout any execution of the greedy algorithm, the current strings remain substring-free. Moreover, if two distinct current strings are formed by merging the input strings $t_1, \ldots, t_k$ and $t'_1, \ldots, t'_\ell$, respectively, in these orders, then*

$$\mathrm{ov}(\langle t_1, \ldots, t_k \rangle, \langle t'_1, \ldots, t'_\ell \rangle) = \mathrm{ov}(t_k, t'_1).$$

Consequently, the overlap lengths selected by the greedy algorithm are nonincreasing.

The *compression* of a common superstring of $\mathcal{S}$ is the total length of the strings in $\mathcal{S}$ minus the length of the superstring. We use the following classical result.

Theorem 2.2 ([TU88, Tur89, KN24])**.** *The compression of every output of the greedy algorithm on $\mathcal{S}$ is at least half of the compression of every common superstring of $\mathcal{S}$.*

2.3 Cycle Covers

If one allows the greedy algorithm to produce cycles (without violating the in-degree and out-degree constraints), then it constructs an optimal *cycle cover* $\mathcal{C}$, that is, a collection of cycles that cover every node of the graph and has maximum total overlap. This is proved in [BJL+94, Theorem 10]. The constructed cycle cover may contain self-loops. It is not difficult to see that $\mathcal{C}$ is also a cycle cover of minimum total length in the prefix graph. Indeed, consider a cycle $C = (t_0, t_1, \ldots, t_{r-1}, t_0)$ of $\mathcal{C}$. Its total overlap is equal to $\sum_{i\in[r]} |\mathrm{ov}(t_{i-1}, t_i)|$ and its total prefix is $\sum_{i\in[r]} |\mathrm{pref}(t_{i-1}, t_i)|$ (here, indices are taken modulo r). Note, however, that $|\mathrm{pref}(t_{i-1}, t_i)| + |\mathrm{ov}(t_{i-1}, t_i)| = |t_{i-1}|$. Consequently, the total overlap of C and its total prefix sum to the total length of the strings on C. Thus, maximizing total overlap minimizes total prefix.

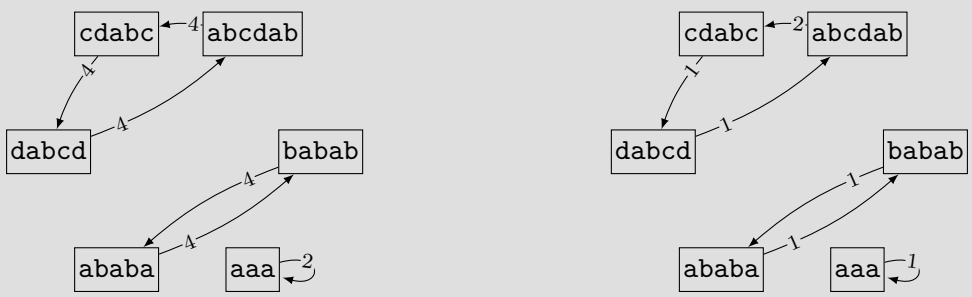
For $C \in \mathcal{C}$, let $w(C)$ be the total prefix of the cycle C and $W = \sum_{C\in\mathcal{C}} w(C)$ be the total prefix of $\mathcal{C}$. Then,

$$W \le \mathrm{OPT}\,. \tag{1}$$

Indeed, let $(t_1, \ldots, t_n)$ be a permutation of the input strings whose overlap merge has length OPT. Then, the total prefix of the cycle $(t_1, \ldots, t_n, t_1)$ is at most OPT, but at least W (since this single cycle is also a cycle cover):

$$\mathrm{OPT} = \sum_{i\in[n-1]} |\mathrm{pref}(t_i, t_{i+1})| + |t_n| \ge \sum_{i\in[n-1]} |\mathrm{pref}(t_i, t_{i+1})| + |\mathrm{pref}(t_n, t_1)| \ge W.$$

The figure below shows the same optimal cycle cover in the overlap and prefix graphs. Its total overlap and total prefix are 22 and $W = 7$, respectively; their sum is 29, the total length of the input strings.



2.4 Opening the Cycles

To turn a cycle cover into a Hamiltonian path, one may proceed as follows: for each cycle, remove its edge of minimum overlap; then, merge the resulting paths into a Hamiltonian path.

Formally, consider a cycle $C = (t_1, \ldots, t_r, t_1)$ of $\mathcal{C}$ and assume that the edge (t_r, t_1) was the last edge of this cycle added by the algorithm. We call it a *cycle-closing edge*; it has minimum overlap among the edges of the cycle. Denote its length by $o(C)$ and let $O = \sum_{C \in \mathcal{C}} o(C)$. Now, let $\sigma(C) = \langle t_1, \ldots, t_r \rangle$ be a *representative string* for C. Then,

$$\sigma(C) = \mathrm{pref}(t_1, t_2) \circ \cdots \circ \mathrm{pref}(t_{r-1}, t_r) \circ t_r = \mathrm{pref}(t_1, t_2) \circ \cdots \circ \mathrm{pref}(t_{r-1}, t_r) \circ \mathrm{pref}(t_r, t_1) \circ \mathrm{ov}(t_r, t_1).$$

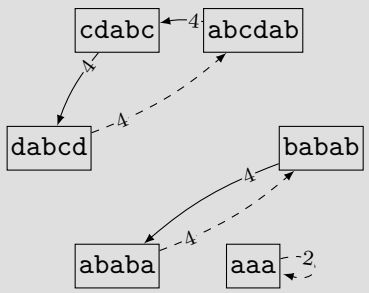
Hence,

$$|\sigma(C)| = w(C) + o(C).$$

The algorithm MGREEDY [BJL+91] constructs the cycle cover greedily and concatenates its representative strings. More precisely, it maintains a set of current strings, initially $\mathcal{S}$, and an initially empty set of representatives. While the current set is nonempty, it selects an ordered pair of current strings with maximum overlap, allowing the two selected strings to coincide and breaking ties arbitrarily. If the strings are distinct, it replaces them by their overlap merge; otherwise, it removes the selected string from the current set and adds it unchanged to the representatives. The extracted strings are precisely the strings $\sigma(C)$ obtained by opening the cycles at their cycle-closing edges. MGREEDY concatenates the representatives. TGREEDY uses this construction of the representatives as its first stage and, as its second stage, applies the greedy algorithm to them, again breaking ties arbitrarily. The superstring returned by MGREEDY has length

$$\sum_{C \in \mathcal{C}} (w(C) + o(C)) = W + O. \tag{2}$$

The cycle-closing edges of the optimal cycle cover are shown as dashed edges below. Their total overlap is $O = 4 + 4 + 2 = 10$. To the right, we show the corresponding representative strings. By concatenating the representative strings, one gets a superstring `abcdabcdbababaaaa` of length $W + O = 7 + 10 = 17$.



$\sigma(C_1) =$ `abcdabcd`
$\sigma(C_2) =$ `bababa`
$\sigma(C_3) =$ `aaa`

2.5 Upper Bounding the Overlap of Cycle-Closing Edges

Since we already know that $W \leq \mathrm{OPT}$ (see (1)), to turn an upper bound $W + O$ (see (2)) into an approximation guarantee, we need to upper bound O in terms of OPT. This is exactly what many previous papers do. Namely, they prove an upper bound of the form

$$O \leq \mathrm{OPT} + \gamma W, \tag{3}$$

for a constant γ. Combined with (1) and (2), this inequality immediately implies that MGREEDY is a $(2+\gamma)$-approximation. As shown by [BJL+94, Section 5], this implies that the approximation guarantee of the greedy algorithm is also at most $(2+\gamma)$.

The MGREEDY algorithm simply concatenates the representative strings. If, instead, one applies an α-approximation algorithm of MAX-ATSP in the overlap graph of the representative strings and merges them in the found order, the resulting algorithm has approximation guarantee $(2+(1-\alpha)\gamma)$. This was shown in [BJJ97, Muc07, Muc13]; a proof can also be found in [EMV22, Theorem 3.1]. The history of improvements of upper bounds on γ is the following: [BJL+94] proved the inequality with $\gamma = 2$, [KS05] improved it to 3/2, [EMV22] pushed it to 1.425, and [EMV23] improved it further to 1.396.

We prove the inequality (3) for $\gamma = 1$ (see Theorem 3.1). This implies that the greedy algorithm is 3-approximate. Plugging in the best known approximation guarantee $\alpha = 2/3$ [KLSS03, PEvZ12] for MAX-ATSP gives a polynomial-time 7/3-approximation for SCS. Our proof of a stronger upper bound on γ does not come at a cost of a more involved analysis: in fact, our proof is much simpler. We achieve this by considering global properties of the overlap graph instead of considering each cycle of the optimal cycle cover separately. Then, in Theorem 3.3, we prove that $\gamma = 1$ is the optimal endpoint of this approach: for every $\gamma < 1$, there exists an input to SCS where the inequality (3) fails.

2.6 Properties of Cycle Strings

A cycle with small total prefix may be an indicator of the high repetitive structure of the strings of this cycle. For example, a cycle through strings

$$\texttt{abcdabcdabcd}, \texttt{cdabcdabcdab}, \texttt{dabcdabcdabc}$$

has total prefix $2+1+1=4$, whereas its minimum edge overlap is 10. In particular, O may be much larger than W.

Two strings are *equivalent* if one is a cyclic rotation of the other. A string z is *primitive* if there is no string v and an integer $k \geq 2$ such that $z = v^k$. For a cycle $C = (t_1, \ldots, t_r, t_1)$ of $\mathcal{C}$, define its *cycle string* by

$$\tau(C) = \operatorname{pref}(t_1, t_2) \circ \cdots \circ \operatorname{pref}(t_r, t_1),$$

hence $|\tau(C)| = w(C)$.

Lemma 2.3 ([Muc13, Lemma 2.5])**.** *Every cycle string $\tau(C)$ is primitive, and the cycle strings of distinct cycles of $\mathcal{C}$ are pairwise non-equivalent.*

In the running example, the three cycles C_1, C_2, C_3 have the following cycle strings:

$$\tau(C_1) = \texttt{abcd}, \qquad \tau(C_2) = \texttt{ba}, \qquad \tau(C_3) = \texttt{a}.$$

Then, $\sigma(C_1) = \tau(C_1)^2, \sigma(C_2) = \tau(C_2)^3$, and $\sigma(C_3) = \tau(C_3)^3$ (in general, $\sigma(C)$ is not always a power of $\tau(C)$).

Fix an arbitrary total order on the alphabet. A *Lyndon word* [Shi53, Lyn54] is a string that is strictly smaller than each of its proper suffixes; equivalently, it is primitive and lexicographically smaller than all its nontrivial cyclic rotations. Every equivalence class of primitive strings has a unique Lyndon word. Moreover, a Lyndon word z has no period $p < |z|$: otherwise its suffix

of length $|z| - p$ would equal a proper prefix of z and hence would be smaller than z, contrary to the definition.

For a string x and an integer $m \geq 0$, let $\mathrm{Sub}_m(x)$ be the set of length-m substrings of x. For a string z and an integer $m \geq 0$, let $\mathrm{CSub}_m(z)$ be the set of length-m substrings of z^∞. For every cycle $C \in \mathcal{C}$, let $\tau'(C)$ be the unique Lyndon word equivalent to $\tau(C)$.

Lemma 2.4 ([BJL+94, Claim 2]). *For every cycle $C \in \mathcal{C}$,*

$$\mathrm{CSub}_{o(C)+1}(\tau'(C)) = \mathrm{CSub}_{o(C)+1}(\tau(C)) \subseteq \mathrm{Sub}_{o(C)+1}(\mathrm{OPT}).$$

Consider the optimal superstring $\mathrm{OPT} = \texttt{aaababababcdabcd}$. Using the alphabet order $\texttt{a} < \texttt{b} < \texttt{c} < \texttt{d}$, the Lyndon words corresponding to the three cycle strings are

$$\tau'(C_1) = \texttt{abcd}, \qquad \tau'(C_2) = \texttt{ab}, \qquad \tau'(C_3) = \texttt{a}.$$

The minimum overlaps of the three cycles are $o(C_1) = o(C_2) = 4$ and $o(C_3) = 2$. Then,

$$\begin{aligned}
\mathrm{CSub}_5(\tau'(C_1)) &= \{\texttt{abcda}, \texttt{bcdab}, \texttt{cdabc}, \texttt{dabcd}\},\\
\mathrm{CSub}_5(\tau'(C_2)) &= \{\texttt{ababa}, \texttt{babab}\},\\
\mathrm{CSub}_3(\tau'(C_3)) &= \{\texttt{aaa}\}.
\end{aligned}$$

Thus, every length-$(o(C)+1)$ substring of $\tau(C)^\infty$ occurs as a substring of OPT.

3 Optimal Bound on the Cycle-Cover Inequality

In this section, we give a proof of our main technical contribution: the smallest value of γ where the inequality (3) holds is equal to 1.

3.1 Upper Bound

Theorem 3.1. *For every instance of SCS, it holds that $O \leq \mathrm{OPT} + W$.*

Corollary 1.2. *The greedy algorithm for SCS is a 3-approximation algorithm.*

Lemma 2.4 shows that, for every cycle C, all length-$(o(C)+1)$ substrings of $\tau(C)^\infty$ occur in OPT. The lemma below bounds the number of cycles that can use the same superstring.

Lemma 3.2. *Let u be a string, $m \in \{1, \ldots, |u|\}$, and let $z_1, \ldots, z_a$ be distinct Lyndon words for some $a \geq 0$. Suppose that $|z_j| \leq m$ and $\mathrm{CSub}_m(z_j) \subseteq \mathrm{Sub}_m(u)$ for every j. Then*

$$a \leq |\mathrm{Sub}_m(u)| - |\mathrm{Sub}_{m-1}(u)| + 1.$$

Consider the difference between Sub_m and Sub_{m-1}. Every distinct length-$(m-1)$ substring of u, except possibly the one occurring only at the end of u, can be extended to the right by one symbol. Therefore

$$|\mathrm{Sub}_m(u)| \geq |\mathrm{Sub}_{m-1}(u)| - 1.$$

Lemma 3.2 shows that each Lyndon word whose complete set of length-m periodic substrings occurs in u forces one additional length-m substring. To give an example, let $u = \texttt{aababbabaab}$

and $m = 4$. Then

$$\begin{aligned}\mathrm{Sub}_3(u) &= \{\mathtt{aab}, \mathtt{aba}, \mathtt{abb}, \mathtt{baa}, \mathtt{bab}, \mathtt{bba}\},\\ \mathrm{Sub}_4(u) &= \{\mathtt{aaba}, \mathtt{abaa}, \mathtt{abab}, \mathtt{abba}, \mathtt{baab}, \mathtt{baba}, \mathtt{babb}, \mathtt{bbab}\}.\end{aligned}$$

By the extension argument above, $|\,\mathrm{Sub}_4(u)| \geq |\,\mathrm{Sub}_3(u)| - 1 = 5$.
However, consider the Lyndon words $z_1 = \mathtt{aab}$, $z_2 = \mathtt{ab}$, and $z_3 = \mathtt{abb}$. Their length-4 cyclic substrings are

$$\begin{aligned}\mathrm{CSub}_4(z_1) &= \{\mathtt{aaba}, \mathtt{abaa}, \mathtt{baab}\},\\ \mathrm{CSub}_4(z_2) &= \{\mathtt{abab}, \mathtt{baba}\},\\ \mathrm{CSub}_4(z_3) &= \{\mathtt{abba}, \mathtt{babb}, \mathtt{bbab}\}.\end{aligned}$$

They are contained in $\mathrm{Sub}_4(u)$, so Lemma 3.2 implies that these three Lyndon words account for three additional length-4 substrings, which shows that

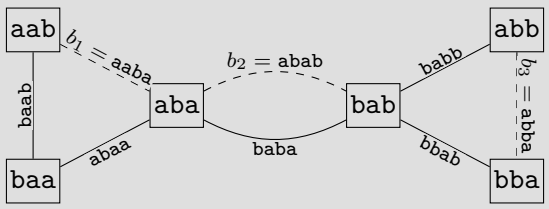
$$|\,\mathrm{Sub}_4(u)| \geq a + |\,\mathrm{Sub}_3(u)| - 1 = 8.$$

Proof. Let G_m be the following undirected multigraph: its vertices are the strings in $\mathrm{Sub}_{m-1}(u)$, and each string in $\mathrm{Sub}_m(u)$ is an edge joining its length-$(m-1)$ prefix and its length-$(m-1)$ suffix (it is an undirected version of the order-$(m-1)$ Rauzy graph [Rau82] of u, which is in turn a subgraph of the de Bruijn graph [DB46]). Thus, G_m has $|\mathrm{Sub}_{m-1}(u)|$ vertices and $|\mathrm{Sub}_m(u)|$ edges. It is connected because the consecutive length-$(m-1)$ substrings of u form a walk that visits every vertex.

For every j, let b_j be the length-m prefix of z_j^∞. Since $m \geq |z_j|$ and z_j is a Lyndon word, b_j is the unique lexicographically smallest string in $\mathrm{CSub}_m(z_j)$. The strings $b_1, \ldots, b_a$ are pairwise distinct. Indeed, suppose $b_i = b_j$. If $|z_i| = |z_j|$, since $m \geq |z_i|$, the length-$|z_i|$ prefixes of b_i and b_j are z_i and z_j, respectively, so $z_i = z_j$. If, say, $|z_i| < |z_j|$, then z_j is a prefix of z_i^∞ and hence has period $|z_i| < |z_j|$, which is impossible for a Lyndon word. After relabeling, assume $b_1 < b_2 < \cdots < b_a$ and delete the edges $b_1, b_2, \ldots, b_a$ from G_m in that order. We show that after each deletion the graph remains connected.

Consider the deletion of the edge b_j, and, for $0 \leq i < |z_j|$, let c_i and u_i be the length-m and length-$(m-1)$ substrings of z_j^∞ beginning at offset i, respectively. Then $c_0 = b_j$, and each c_i is an edge of G_m connecting u_i and u_{i+1}. If $|z_j| > 1$, deleting the edge $c_0 = b_j$ keeps the graph connected, since the alternate walk $c_1, c_2, \ldots, c_{|z_j|-1}$ is still present: each c_i with $1 \leq i < |z_j|$ is larger than b_j, whereas every previously deleted edge is smaller than b_j. If $|z_j| = 1$, then b_j is a loop, so its deletion is harmless.

For the same string $u = \mathtt{aababbabaab}$ and $m = 4$, the graph G_4 is shown below.



After all a deletions, the remaining graph is connected and has $|\mathrm{Sub}_m(u)| - a$ edges; consequently,

$|\mathrm{Sub}_m(u)| - a \geq |\mathrm{Sub}_{m-1}(u)| - 1$. □

Proof of Theorem 3.1. Enumerate the cycles $C \in \mathcal{C}$ satisfying $o(C) + 1 \geq w(C)$ as $C_1, \ldots, C_h$. For every index j, put $z_j = \tau'(C_j)$ and $\ell_j = o(C_j) + 1$. The Lyndon words $z_1, \ldots, z_h$ are distinct by Lemma 2.3, and $|z_j| = |\tau(C_j)| = w(C_j) \leq \ell_j$ for every j. Moreover, Lemma 2.4 gives $\mathrm{CSub}_{\ell_j}(z_j) \subseteq \mathrm{Sub}_{\ell_j}(\mathrm{OPT})$ for every j. This inclusion implies $\ell_j \leq \mathrm{OPT}$, since $\mathrm{CSub}_{\ell_j}(z_j)$ is nonempty, and $\mathrm{CSub}_m(z_j) \subseteq \mathrm{Sub}_m(\mathrm{OPT})$ for every $m \leq \ell_j$, since every length-m substring of z_j^∞ is a prefix of a length-ℓ_j one. For $1 \leq m \leq \mathrm{OPT}$, let a_m be the number of indices j satisfying $|z_j| \leq m \leq \ell_j$. Lemma 3.2 gives

$$a_m \leq |\mathrm{Sub}_m(\mathrm{OPT})| - |\mathrm{Sub}_{m-1}(\mathrm{OPT})| + 1.$$

Each index j is counted by a_m for the values $m = |z_j|, \ldots, \ell_j$, hence exactly $\ell_j - |z_j| + 1 = o(C_j) - w(C_j) + 2$ times. Since $|\mathrm{Sub}_0(\mathrm{OPT})| = |\mathrm{Sub}_{\mathrm{OPT}}(\mathrm{OPT})| = 1$, summing the inequalities over m gives

$$\begin{aligned}
\sum_{\substack{C \in \mathcal{C} \\ o(C)+1 \geq w(C)}} (o(C) - w(C) + 2) &= \sum_{j=1}^{h} (\ell_j - |z_j| + 1) \\
&= \sum_{m=1}^{\mathrm{OPT}} a_m \\
&\leq \sum_{m=1}^{\mathrm{OPT}} \left(|\mathrm{Sub}_m(\mathrm{OPT})| - |\mathrm{Sub}_{m-1}(\mathrm{OPT})| + 1\right) \\
&= \mathrm{OPT}\,.
\end{aligned}$$

For each cycle in the first sum, $o(C) - w(C) \leq o(C) - w(C) + 2$, whereas every omitted cycle satisfies $o(C) - w(C) < 0$. Therefore,

$$O - W \leq \mathrm{OPT}\,.$$ □

3.2 Lower Bound

In this section, we prove that the coefficient of W in Theorem 3.1 cannot be decreased further.

Theorem 3.3. *For every $\gamma < 1$, there is a substring-free set $\mathcal{S}$ of binary strings such that*

$$O > \mathrm{OPT} + \gamma W.$$

Proof. For $p \geq 2$, put $s_p = \mathtt{0}^{p-1}\mathtt{1}\mathtt{0}^{p-1}\mathtt{1}\mathtt{0}^{p-2}$.

$$\begin{aligned}
s_2 &= \mathtt{0101} \\
s_3 &= \mathtt{0010010} \\
s_4 &= \mathtt{0001000100} \\
s_5 &= \mathtt{0000100001000}
\end{aligned}$$

For $p \neq q$, the string s_p is not a substring of s_q, since each string contains exactly two $\mathtt{1}$'s, whose positions differ by p and q, respectively. Furthermore,

$$|\mathrm{ov}(s_p, s_p)| = 2p - 2 \qquad \text{and} \qquad |\mathrm{ov}(s_p, s_q)| \leq 2p - 3 \quad \text{for } p \neq q.$$

Indeed, s_p starts and ends with $\mathtt{0}^{p-1}\mathtt{1}\mathtt{0}^{p-2}$, giving a self-overlap of length at least $2p-2$. An overlap of length $2p-2$ from s_p to s_q forces $q=p$, because the first $\mathtt{1}$ in this suffix of s_p is at position p, whereas the first $\mathtt{1}$ in s_q is at position q. A longer proper overlap contains both $\mathtt{1}$'s of s_p, so matching their distance forces $q=p$, and matching their positions then forces a zero shift, which is impossible for a proper overlap.

For $m \geq 2$, let $\mathcal{S}_m = \{s_2, s_3, \ldots, s_m\}$. By the discussion above, $\mathcal{S}_m$ is substring-free and its cycle cover consisting of all self-loops has minimum total prefix. Then,

$$W = \sum_{p=2}^{m} (3p - 2 - (2p-2)) = \sum_{p=2}^{m} p = m(m+1)/2 - 1,$$
$$O = \sum_{p=2}^{m} (2p-2) = 2W - 2(m-1).$$

Since $|\mathrm{ov}(s_p, s_{p-1})| = 2p-3$,

$$\mathrm{OPT} \leq |\langle s_m, s_{m-1}, \ldots, s_2 \rangle| = |s_2| + \sum_{p=3}^{m} (|s_p| - (2p-3)) = 4 + \sum_{p=3}^{m} (p+1) = W + m.$$

Consequently,

$$O - \mathrm{OPT} \geq 2W - 2(m-1) - (W+m) = W - 3m + 2.$$

Since $W = \Theta(m^2)$, for every fixed $\gamma < 1$ and all sufficiently large m, $O - \mathrm{OPT} > \gamma W$. □

3.3 Consequences for MGREEDY and TGREEDY

The cycle-cover inequality also gives the exact approximation ratio of MGREEDY and an improved upper bound for TGREEDY.

Corollary 3.4. *For every tie-breaking rule, the approximation ratio of MGREEDY is exactly* 3.

Proof. Every output of MGREEDY has length $W + O$, so Theorem 3.1 and (1) give

$$W + O \leq \mathrm{OPT} + 2W \leq 3\,\mathrm{OPT}\,.$$

For the matching lower bound, consider the sets $\mathcal{S}_m = \{s_2, \ldots, s_m\}$ from the proof of Theorem 3.3. While s_p is the longest current string, (s_p, s_p) is the unique pair of maximum overlap, with overlap $2p-2$, since every other pair has overlap at most $2p-3$. Hence, for every tie-breaking rule, MGREEDY extracts $s_m, \ldots, s_2$ unchanged, in this order. Its final concatenation consequently has length

$$|g_m| = \sum_{p=2}^{m} (3p-2) = 3W - 2(m-1),$$

where $W = m(m+1)/2 - 1$. By the proof of Theorem 3.3, $\mathrm{OPT} \leq W + m$, so

$$\frac{|g_m|}{\mathrm{OPT}} \geq \frac{3W - 2(m-1)}{W+m} = 3 - \frac{5m-2}{W+m}.$$

The right-hand side tends to 3 as $m \to \infty$, proving the corollary. □

Corollary 3.5. *TGREEDY is a* 5/2*-approximation algorithm.*

Proof. For a set of strings $\mathcal{A}$, let $\mathrm{OPT}(\mathcal{A})$ denote the length of a shortest common superstring. Let g be any output of TGREEDY on $\mathcal{S}$. Let $\mathcal{R} = \{\sigma(C) : C \in \mathcal{C}\}$ be the set of representative strings obtained in the first stage, and let $L = \sum_{C \in \mathcal{C}} |\sigma(C)| = W + O$. The set $\mathcal{R}$ is substring-free. Indeed, if $\sigma(D)$ were a substring of $\sigma(C)$ for distinct cycles C and D, then every input string on either cycle would occur in $\tau(C)^\infty$ by [BJL+94, Claims 2 and 4]. By [BJL+94, Claim 3], the vertices of both cycles could then be covered by one cycle of total prefix at most $w(C)$. Replacing C and D by this cycle would decrease the total prefix of $\mathcal{C}$ by at least $w(D) > 0$, contradicting its optimality.

We next recall the bound

$$\mathrm{OPT}(\mathcal{R}) \le \mathrm{OPT}(\mathcal{S}) + W,$$

used in the proof of [BJL+94, Theorem 11]. For every cycle $C = (t_1, \ldots, t_r, t_1)$, let $s_C = t_1$ and $u_C = \langle t_1, \ldots, t_r, t_1 \rangle$. The string u_C contains $\sigma(C)$, starts and ends with s_C, and has length $|s_C| + w(C)$. Choose an ordering of the strings s_C whose overlap merge has minimum length. Merging the strings u_C in the same order preserves at least the overlap of each adjacent pair of strings s_C, so it gives a common superstring of $\mathcal{R}$ of length at most $\mathrm{OPT}(\{s_C : C \in \mathcal{C}\}) + W$. Since $\{s_C : C \in \mathcal{C}\} \subseteq \mathcal{S}$, this proves the bound.

Since $\mathcal{R}$ is substring-free, Theorem 2.2 applies to the second stage of TGREEDY and gives

$$L - |g| \ge \frac{1}{2}(L - \mathrm{OPT}(\mathcal{R})).$$

Combining this with the bound on $\mathrm{OPT}(\mathcal{R})$, Theorem 3.1, and (1), we obtain

$$\begin{aligned} |g| &\le \frac{1}{2}(L + \mathrm{OPT}(\mathcal{R})) \\ &\le \mathrm{OPT}(\mathcal{S}) + \frac{3}{2}W \\ &\le \frac{5}{2}\,\mathrm{OPT}(\mathcal{S}). \end{aligned}$$

□

4 Lower Bound for the Greedy Algorithm

In this section, we show that the approximation ratio of the greedy algorithm is at least $12/5 = 2.4$. Specifically, we construct a family of sets of binary strings for which the ratio of the length of the output to that of a common superstring tends to $12/5$ under an execution that selects a maximum overlap at every step. Since the worst-case approximation ratio of the greedy algorithm is independent of the tie-breaking rule [Nik21], this lower bound holds for every tie-breaking rule.

Our main result is as follows:

Theorem 4.1 (Formal version of Theorem 1.1)**.** *For all integers $r \ge 1$ and $P \ge r - 1$, there is a set of $5r$ binary strings that has a common superstring of length $5Pr + 27r^2 + 4P + 19r + 1$, whereas the greedy algorithm may output a common superstring of length $12Pr + (115r^2 + 7r)/2$.*

Therefore, the ratio of the length of the output to OPT is at least

$$\frac{12Pr + (115r^2 + 7r)/2}{5Pr + 27r^2 + 4P + 19r + 1}.$$

For $P = r^2$, this bound tends to $12/5$ as $r \to \infty$.

4.1 Gap Words

We represent binary strings containing ones as *gap words*. For nonnegative integers k, L, and R and positive integers $g_1, \ldots, g_k$, we denote

$$[L \mid g_1\ g_2\ \cdots\ g_k \mid R] = 0^L 1\, 0^{g_1-1} 1\, 0^{g_2-1} 1 \cdots 0^{g_k-1} 1\, 0^R.$$

The sequence $(g_1, \ldots, g_k)$ is called a *gap sequence*. For $k = 0$, the gap sequence is empty and the gap word is $0^L 1 0^R$. The gap word contains exactly $k+1$ ones. For every $i \in \{1, \ldots, k\}$, g_i is the distance between the i-th one and the $(i+1)$-st one. We call $g_1, \ldots, g_k$ the *gaps* and L and R the left and right *pads*, respectively.

The overlap of two gap words is determined as follows.

Lemma 4.2. *Let $S = [L \mid u_1\ \cdots\ u_k \mid R]$ and $T = [L' \mid v_1\ \cdots\ v_\ell \mid R']$ be gap words such that neither is a substring of the other. Let J be the set of integers j with $0 \le j \le \min\{k, \ell\}$ that satisfy the following conditions:*

1. *$u_{k-j+i} = v_i$ for every $i \in \{1, \ldots, j\}$.*
2. *$L' < u_{k-j}$ (or $L' \le L$ if $j = k$).*
3. *$R < v_{j+1}$ (or $R \le R'$ if $j = \ell$).*

Then

$$\operatorname{ov}(S, T) = \begin{cases} [L' \mid v_1\ \cdots\ v_{j_{\max}} \mid R] & \text{if } J \neq \emptyset, \\ 0^{\min\{R, L'\}} & \text{otherwise,} \end{cases}$$

where $j_{\max} = \max J$ in the first case.

Proof. A gap word that is a suffix of S and a prefix of T must have the last j gaps of S and the first j gaps of T for some $0 \le j \le \min\{k, \ell\}$. Its left pad must be L' and its right pad must be R, so its gap word representation is $[L' \mid v_1\ \cdots\ v_j \mid R]$. This gap word is a suffix of S and a prefix of T precisely when the shared gaps agree and both pads fit into the neighboring gaps or pads, as specified by the three conditions. Its length is $L' + 1 + \sum_{i=1}^{j} v_i + R$, which increases strictly with j because each v_i is positive. A suffix of S that is a prefix of T and has no gap word representation consists entirely of zeros, and the longest such string is $0^{\min\{R, L'\}}$. Every gap word satisfying the three conditions is longer than this string, so $j_{\max}$ gives the longest common suffix and prefix when J is nonempty. Since neither S nor T is a substring of the other, every common suffix and prefix is proper in both strings, so the longest of these strings is $\operatorname{ov}(S, T)$. □

For the gap words S and T in Lemma 4.2, let $Q(S, T)$ be the largest integer q with $0 \le q \le \min\{k, \ell\}$ such that $u_{k-q+i} = v_i$ for every $i \in \{1, \ldots, q\}$. Since every $j \in J$ satisfies $j \le Q(S, T)$ and all gap values are positive, Lemma 4.2 gives

$$|\operatorname{ov}(S, T)| \le L' + R + 1 + \sum_{i=1}^{Q(S,T)} v_i. \tag{4}$$

Figure 2 shows an example of gap words and their overlaps.

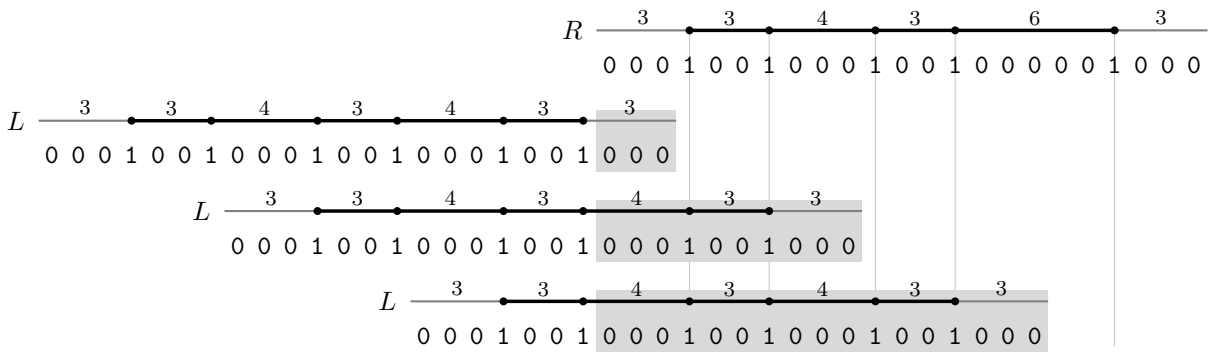


Figure 2: An illustration of the gap words $L = [3 \mid 3\ 4\ 3\ 4\ 3 \mid 3]$ and $R = [3 \mid 3\ 4\ 3\ 6 \mid 3]$ and three possible overlaps of L with R. Each dot on a horizontal line represents an occurrence of one in the corresponding string. The numbers above the black and gray segments indicate the gap values and pad lengths, respectively. The corresponding binary string is shown directly below each horizontal line. The overlapping portions are shaded. In the three overlaps shown, the strings share only zeros, one gap, and three gaps, respectively. The longest overlap is $\mathrm{ov}(L, R) = [3 \mid 3\ 4\ 3 \mid 3]$.

4.2 Construction and Overview

We now introduce the construction of our lower-bound instances. Fix integers $r \geq 1$ and $P \geq r - 1$. Let $M = P + 8r$ and $p_i = P + r + i$ for $1 \leq i \leq r$. We call M the *common gap* and p_i the i-th *private gap* for $1 \leq i \leq r$. Let $\mathcal{S}$ consist of the following five binary strings for each $i \in \{1, \ldots, r\}$:

$$X_i = [P + 2r \mid p_i\ M\ p_i\ M \mid P], \qquad Y_i = [P \mid M\ p_i\ M\ p_i \mid P + 4r],$$
$$U_i = [P \mid M\ M\ p_i\ M \mid P], \qquad V_i = [P \mid M\ p_i\ M\ M \mid P + r],$$
$$Z_i = [P + r \mid M\ M\ p_i \mid 0].$$

By construction, the set $\mathcal{S}$ is substring-free.

Overview We now explain how the construction yields a short common superstring and a much longer string produced by the greedy algorithm. In both constructions, we first merge the five input strings for each i into a single string. For the short common superstring and the long common superstring produced by the greedy algorithm, we write the strings obtained by merging five strings as

$$S_i = \langle Z_i, U_i, Y_i, X_i, V_i \rangle,$$
$$T_i = \langle X_i, Y_i, U_i, V_i, Z_i \rangle,$$

respectively. We choose the lengths of the left and right pads of the input strings so that the greedy algorithm has an execution that produces all the strings T_i. As shown in Figure 3, the sum of the lengths of the prefixes of the four adjacent pairs is $3P + O(r)$ for S_i and $8P + O(r)$ for T_i. Since the lengths of the last strings are $|V_i| = 6P + O(r)$ and $|Z_i| = 4P + O(r)$, respectively, we obtain $|S_i| = 9P + O(r)$ and $|T_i| = 12P + O(r)$.

We next merge the strings obtained for different indices i. For $1 \leq i < r$, the strings S_i and S_{i+1} share a long overlap of length $4P + O(r)$, so the length of the corresponding prefix is $5P + O(r)$. In contrast, T_i and T_j have empty overlap whenever $i \neq j$ because their private gaps differ and we set the lengths of their left and right pads to prevent overlaps. Thus, merging $S_1, \ldots, S_r$ in this

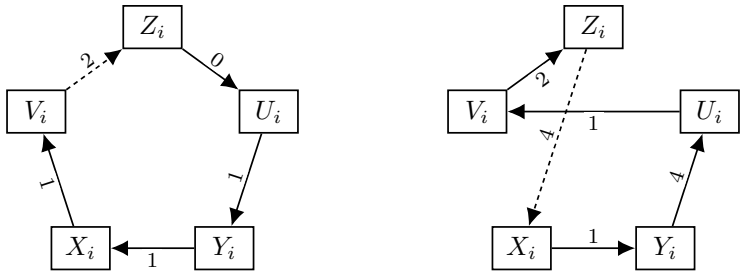


Figure 3: Illustrations of prefix graphs for the short common superstring (left) and the long string produced by the greedy algorithm (right). An integer label x on an edge indicates that the length of the corresponding prefix is $Px + O(r)$. For $i < r$, the dashed arrows represent $V_i \to Z_{i+1}$ on the left and $Z_i \to X_{i+1}$ on the right, and their labels indicate the lengths of the corresponding prefixes.

order gives a short common superstring, whereas merging $T_1, \ldots, T_r$ gives their concatenation, with lengths

$$|\langle S_1, \ldots, S_r \rangle| = \sum_{i=1}^{r-1} |\operatorname{pref}(S_i, S_{i+1})| + |S_r| = 5Pr + O(r^2 + P),$$
$$|\langle T_1, \ldots, T_r \rangle| = \sum_{i=1}^{r} |T_i| = 12Pr + O(r^2).$$

Consequently, the ratio of the length of the string produced by the greedy algorithm to the length of the short common superstring approaches $12/5$ when we choose $P = r^2$ and take the limit as $r \to \infty$.

4.3 Constructing a Short Common Superstring

We first give an explicit form of the short common superstring described in the overview and verify that it contains every input string. Let

$$H = [P + r \mid M\ M\ p_1\ M\ p_1\ \cdots\ M\ M\ p_r\ M\ p_r\ M\ M \mid P + r].$$

The gap sequence of H contains $3r + 2$ occurrences of M and two occurrences of each p_i. The sum of these gap values is the distance between the first and last ones of H. Adding one for the first one and the lengths of the two pads gives

$$\begin{aligned} |H| &= (3r+2)M + 1 + 2(P+r) + 2\sum_{i=1}^{r} p_i \\ &= (3r+2)(P+8r) + 1 + 2(P+r) + 2\sum_{i=1}^{r}(P+r+i) \\ &= 5Pr + 27r^2 + 4P + 19r + 1. \end{aligned}$$

The gap sequence of each string in $\mathcal{S}$ appears as a consecutive segment of that of H. For each such occurrence in H, the runs of zeros immediately before the first one and after the last one are at least as long as the corresponding left and right pads, respectively (see Figure 4). Thus, H is a common superstring of $\mathcal{S}$, which gives

$$\mathrm{OPT} \le |H| = 5Pr + 27r^2 + 4P + 19r + 1.$$

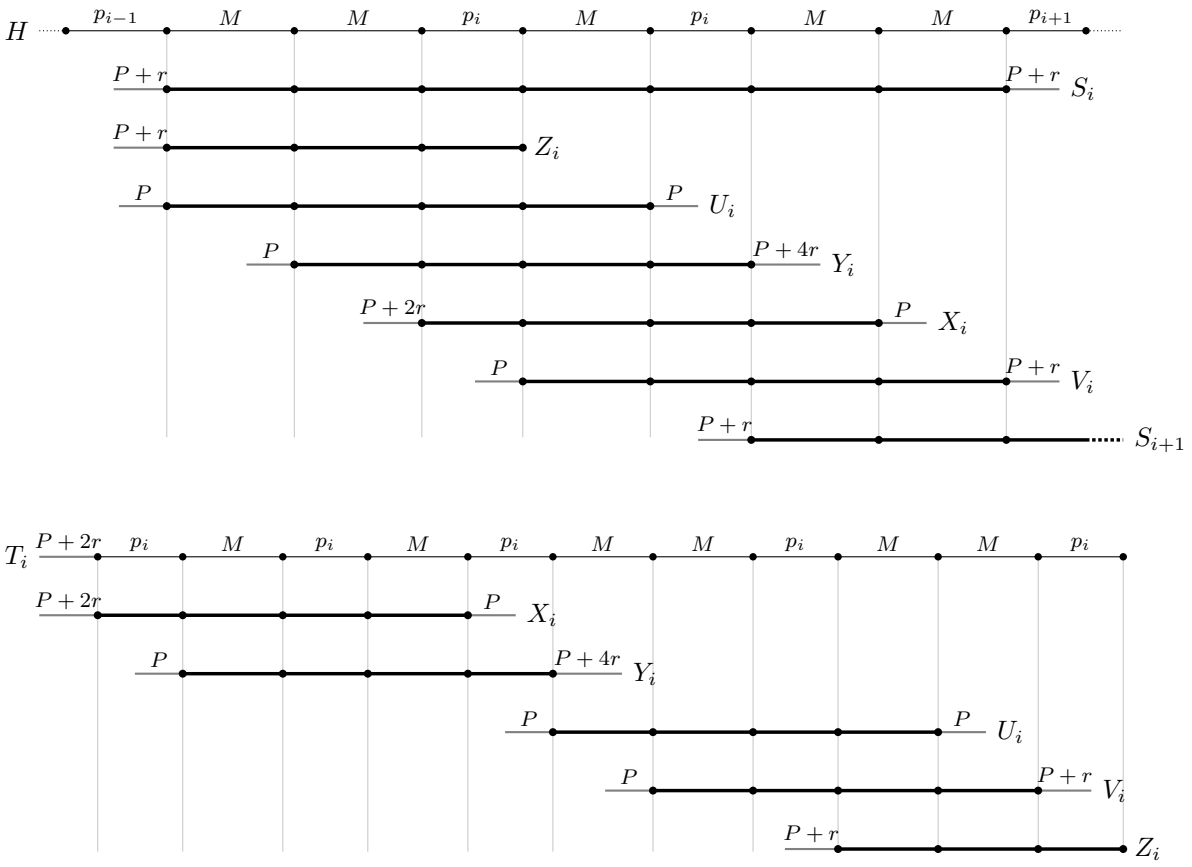


Figure 4: An illustration of the occurrences of strings in H and T_i. The top diagram shows S_i, S_{i+1}, and the five input strings with index i in H, and the bottom diagram shows the five input strings in T_i. Each dot on a horizontal line represents an occurrence of one. The numbers above the black and gray segments indicate the gap values and pad lengths, respectively.

4.4 Constructing a Long Common Superstring

We now introduce how the greedy algorithm can produce a long common superstring of $\mathcal{S}$.

For each $i \in \{1, \ldots, r\}$, we write the intermediate strings obtained as we merge the input strings as follows.

$$\begin{aligned}
A_i &= \langle X_i, Y_i \rangle = [P + 2r \mid p_i \ M \ p_i \ M \ p_i \mid P + 4r], \\
B_i &= \langle U_i, V_i \rangle = [P \mid M \ M \ p_i \ M \ M \mid P + r], \\
C_i &= \langle B_i, Z_i \rangle = \langle U_i, V_i, Z_i \rangle = [P \mid M \ M \ p_i \ M \ M \ p_i \mid 0], \\
T_i &= \langle A_i, C_i \rangle = \langle X_i, Y_i, U_i, V_i, Z_i \rangle \\
&= [P + 2r \mid p_i \ M \ p_i \ M \ p_i \ M \ M \ p_i \ M \ M \ p_i \mid 0].
\end{aligned}$$

We will show that the following procedure is a valid execution of the greedy algorithm.

Step (i) For $i = r, r-1, \ldots, 1$, we first merge X_i with Y_i to obtain A_i and then merge U_i with V_i to obtain B_i.

Step (ii) For each $i \in \{1, \ldots, r\}$, we merge B_i with Z_i to obtain C_i.

Step (iii) For each $i \in \{1, \ldots, r\}$, we merge A_i with C_i to obtain T_i.

Step (iv) We concatenate $T_1, \ldots, T_r$ without overlap.

The overlaps used in these merge procedures are

$$\begin{aligned}
\operatorname{ov}(X_i, Y_i) &= \operatorname{ov}(U_i, V_i) = [P \mid M\ p_i\ M \mid P], \\
\operatorname{ov}(B_i, Z_i) &= [P + r \mid M\ M \mid P + r], \\
\operatorname{ov}(A_i, C_i) &= [P \mid\ \mid P + 4r],
\end{aligned}$$

and their lengths are

$$\begin{aligned}
\alpha_i &= |\operatorname{ov}(X_i, Y_i)| = |\operatorname{ov}(U_i, V_i)| = 2P + 2M + p_i + 1 \\
&= 5P + 17r + i + 1, \\
\beta &= |\operatorname{ov}(B_i, Z_i)| = 2P + 2r + 2M + 1 \\
&= 4P + 18r + 1, \\
\gamma &= |\operatorname{ov}(A_i, C_i)| \\
&= 2P + 4r + 1.
\end{aligned}$$

Since $P \geq r - 1$, we have $\alpha_r \geq \cdots \geq \alpha_1 \geq \beta > \gamma$. The overlap between T_i and T_j is empty for $i \neq j$. The correctness of these overlaps follows from Lemma 4.2 and the definitions of the strings. The bottom diagram of Figure 4 illustrates the occurrences of the five input strings in T_i.

Since the gap sequence of T_i contains six occurrences of M and five occurrences of p_i, we have

$$|T_i| = 1 + P + 2r + 6M + 5p_i = 12P + 55r + 5i + 1.$$

The strings $T_1, \ldots, T_r$ are concatenated without overlap, so the length of the resulting string is

$$|T_1 \cdots T_r| = \sum_{i=1}^{r} (12P + 55r + 5i + 1) = 12Pr + \frac{115r^2 + 7r}{2}.$$

Thus, it remains to show that this procedure is a valid execution of the greedy algorithm. Together with the bound $\mathrm{OPT} \leq |H|$, this will establish the bound in Theorem 4.1:

$$\frac{|T_1 \cdots T_r|}{\mathrm{OPT}} \geq \frac{|T_1 \cdots T_r|}{|H|} = \frac{12Pr + (115r^2 + 7r)/2}{5Pr + 27r^2 + 4P + 19r + 1}.$$

4.5 Validity of the Greedy Procedure

We show that the procedure in Steps (i)–(iv) is a valid execution of the greedy algorithm. Throughout this subsection, we use without further mention the overlap formula in Lemma 2.1 and the fact that the maximum overlap length never increases.

We first state a lemma listing the possible pairs of input strings with an overlap longer than β.

Lemma 4.3. *If a pair of distinct strings $S, T \in \mathcal{S}$ satisfies $|\operatorname{ov}(S, T)| > \beta$, then*

$$(S, T) \in \{(X_i, Y_i), (X_i, V_i), (U_i, Y_i), (U_i, V_i), (Y_i, X_i)\}$$

holds for some $i \in \{1, \ldots, r\}$.

Sequence	Prefix	Suffix
(p_i, M)	X_i	X_i, U_i
(M, p_i)	Y_i, V_i	Y_i, Z_i
(M, M)	U_i, Z_i	V_i
(M, p_i, M)	Y_i, V_i	X_i, U_i
(p_i, M, p_i)	X_i	Y_i
(M, M, p_i)	U_i, Z_i	Z_i
(p_i, M, M)	—	V_i

Table 1: The list of strings in $\mathcal{S}$ whose gap sequences begin or end with the sequence in the first column. In the row for (M, M), the index i ranges from 1 to r in both columns. In every other row, each string has the same index as p_i.

(S, T)	$\operatorname{ov}(S, T)$	$\lvert\operatorname{ov}(S, T)\rvert$
(X_i, Y_i)	$[P \mid M\ p_i\ M \mid P]$	$5P + 17r + i + 1$
(X_i, V_i)	$[P \mid M\ p_i\ M \mid P]$	$5P + 17r + i + 1$
(U_i, Y_i)	$[P \mid M\ p_i\ M \mid P]$	$5P + 17r + i + 1$
(U_i, V_i)	$[P \mid M\ p_i\ M \mid P]$	$5P + 17r + i + 1$
(Y_i, X_i)	$[P + 2r \mid p_i\ M\ p_i \mid P + 4r]$	$5P + 16r + 2i + 1$

Table 2: The overlaps and their lengths for the pairs in Lemma 4.3, where $1 \le i \le r$.

Proof. For every string in $\mathcal{S}$, the left and right pads have lengths at most $P + 2r$ and $P + 4r$, respectively, and every gap has value at most M. If $Q(S, T) \le 1$, then the inequality in (4) gives

$$\lvert\operatorname{ov}(S, T)\rvert \le (P + 2r) + (P + 4r) + M + 1 = 2P + 6r + M + 1 = 3P + 14r + 1 < \beta.$$

The gap sequences cannot share four or more entries, so $Q(S, T) \le 3$. Table 1 lists the first and last two or three entries of the gap sequences of the input strings.

Suppose first that $Q(S, T) = 2$. By Table 1, the two shared gaps are (p_i, M), (M, p_i), or (M, M). In the first two cases, $p_i \le P + 2r$, so the inequality in (4) gives

$$\lvert\operatorname{ov}(S, T)\rvert \le (P + 2r) + (P + 4r) + M + p_i + 1 \le 4P + 16r + 1 < \beta.$$

For (M, M), the possible pairs are $(S, T) = (V_i, U_j)$ and $(S, T) = (V_i, Z_j)$ for $i, j \in \{1, \ldots, r\}$. The right pad of S has length $P + r$ and the left pad of T has length at most $P + r$, so the inequality in (4) gives

$$\lvert\operatorname{ov}(S, T)\rvert \le 2P + 2r + 2M + 1 = \beta.$$

Finally, suppose that $Q(S, T) = 3$. Since $p_1, \ldots, p_r$ are different, Table 1 shows that the five types of pairs in the statement and (Z_i, U_i) are the only possibilities for distinct S and T. For the latter pair, the inequality in (4) gives

$$\lvert\operatorname{ov}(Z_i, U_i)\rvert \le P + 2M + p_i + 1 = 4P + 17r + i + 1 \le \beta.$$

This proves the claim. □

Table 2 shows the overlaps and their lengths for the pairs in Lemma 4.3. The first four pairs have overlaps of length α_i, and the last pair has an overlap of length $\alpha_i - r + i \le \alpha_i$.

We now verify that every merge in Steps (i)–(iv) uses a maximum overlap among the current strings. For each step, we compare the overlap of the prescribed pair with all current overlaps.

Step (i) We first prove the validity of Step (i) by induction over $i = r, r-1, \ldots, 1$. Suppose that we have completed the merges for indices $r, r-1, \ldots, i+1$ using maximum overlaps. The set of current strings is

$$\{A_j, B_j, Z_j \mid i < j \le r\} \cup \{X_j, Y_j, U_j, V_j, Z_j \mid 1 \le j \le i\}.$$

We show that the maximum length of an overlap among the current strings is α_i. By Lemma 4.3, overlaps between strings with different indices have length at most β. For each $j \le i$, Lemma 4.3 and Table 2 show that the maximum length of an overlap between distinct strings with index j is α_j. For $j > i$, the current strings with index j are $A_j = \langle X_j, Y_j \rangle$, $B_j = \langle U_j, V_j \rangle$, and Z_j. Since no pair of distinct strings in $\{A_j, B_j, Z_j\}$ has S as a suffix of one string and T as a prefix of the other for any pair (S, T) in Table 2, every overlap involving A_j, B_j, or Z_j with $j > i$ has length at most $\beta \le \alpha_i$, by Lemma 4.3. Thus, the inequalities $\beta \le \alpha_i$ and $\alpha_j < \alpha_i$ for $j < i$ show that the maximum length is α_i.

We can therefore merge X_i with Y_i to obtain A_i using an overlap of length α_i. The pair (U_i, V_i) still has overlap of length α_i after this merge, so it can be selected next. Hence, every merge in Step (i) uses a maximum overlap by induction.

Step (ii) After Step (i), no pair in Table 2 occurs as a suffix and a prefix of two distinct current strings, so all overlaps have length at most β. Each pair (B_i, Z_i) in Step (ii) has overlap of length β. Since the maximum length of an overlap cannot increase, each of these pairs has maximum overlap when it is selected.

Step (iii) At the start of Step (iii), the current strings are $A_i = \langle X_i, Y_i \rangle$ and $C_i = \langle U_i, V_i, Z_i \rangle$ for $1 \le i \le r$. By Lemma 4.2, we have

$$\begin{aligned}
&|\operatorname{ov}(A_i, A_j)| = |\operatorname{ov}(Y_i, X_j)| = P + 2r < \gamma && \text{for } i \ne j,\\
&|\operatorname{ov}(C_i, C_j)| = |\operatorname{ov}(Z_i, U_j)| = P + 1 < \gamma && \text{for } i \ne j,\\
&|\operatorname{ov}(A_i, C_j)| = |\operatorname{ov}(Y_i, U_j)| = 2P + 4r + 1 = \gamma,\\
&|\operatorname{ov}(C_i, A_j)| = |\operatorname{ov}(Z_i, X_j)| \le 2P + 3r + i + 1 \le \gamma.
\end{aligned}$$

Thus, every overlap has length at most γ. Each pair (A_i, C_i) still has overlap of length γ when it is selected, so Step (iii) is valid.

Step (iv) At the start of Step (iv), the current strings are $T_i = \langle X_i, Y_i, U_i, V_i, Z_i \rangle$ for $1 \le i \le r$. For $i \ne j$, we have $\operatorname{ov}(T_i, T_j) = \operatorname{ov}(Z_i, X_j)$. We show that this overlap is empty. Recall that

$$Z_i = [P + r \mid M\ M\ p_i \mid 0], \qquad X_j = [P + 2r \mid p_j\ M\ p_j\ M \mid P].$$

Z_i ends with the suffix $\mathtt{1}\mathtt{0}^{p_i - 1}\mathtt{1}$ and it does not occur in X_j because p_i differs from both p_j and M. Hence, any nonempty overlap must have the form $\mathtt{0}^a\mathtt{1}$ with $0 \le a \le p_i - 1$. However, X_j begins with $\mathtt{0}^{P+2r}\mathtt{1}$, so such an overlap would require $a = P + 2r$. Since $a \le p_i - 1 < P + 2r$, no such overlap exists. Thus, $|\operatorname{ov}(T_i, T_j)| = |\operatorname{ov}(Z_i, X_j)| = 0$ for $i \ne j$. Therefore, concatenating $T_1, \ldots, T_r$ without overlap in Step (iv) is a valid execution of the greedy algorithm. This completes the proof that the procedure is a valid execution of the greedy algorithm.

4.6 A Lower Bound for TGREEDY

The same construction also gives a lower bound for TGREEDY. Its first stage can extract the strings A_i and C_i as representatives, and its second stage can reproduce Steps (iii) and (iv).

Corollary 4.4. *For every $\varepsilon > 0$, there is a set of binary strings on which TGREEDY may output a common superstring of length at least $(12/5 - \varepsilon)\,\mathrm{OPT}$.*

Here, "may" means that TGREEDY produces this output for some way of breaking ties in both stages.

Proof. Fix $r \geq 1$ and set $P = r^2$ in the construction above. We first verify that the first stage can extract exactly the representatives $\{A_i, C_i \mid 1 \leq i \leq r\}$. The gap representations give the following self-overlaps:

$$\begin{aligned}
|\mathrm{ov}(X_i, X_i)| &= 4P + 11r + i + 1,\\
|\mathrm{ov}(Y_i, Y_i)| &= 4P + 13r + i + 1,\\
|\mathrm{ov}(U_i, U_i)| &= 3P + 8r + 1,\\
|\mathrm{ov}(V_i, V_i)| &= 3P + 9r + 1,\\
|\mathrm{ov}(Z_i, Z_i)| &= P + r + 1,\\
|\mathrm{ov}(A_i, A_i)| &= \delta_i = 5P + 16r + 2i + 1,\\
|\mathrm{ov}(B_i, B_i)| &= 4P + 17r + 1,\\
|\mathrm{ov}(C_i, C_i)| &= 4P + 17r + i + 1.
\end{aligned}$$

In particular, every input string and every B_i has self-overlap smaller than β. For the strings A_i, we have

$$\delta_i = \alpha_i - r + i \leq \alpha_i, \qquad \delta_i \geq \delta_1 = \beta + (r-1)^2 + 1 > \beta.$$

Before each merge of Step (i), extract every current A_j whose self-overlap exceeds that merge's overlap, in decreasing order of self-overlap. Deleting such strings cannot increase any remaining overlap or remove either endpoint of a later prescribed merge. The proof of Step (i) therefore still bounds every overlap between distinct current strings by the next prescribed value α_i. After these extractions, every remaining A_j has self-overlap at most α_i, while every other current string has self-overlap below $\beta \leq \alpha_i$. Thus, each prescribed merge is valid. After these merges, all overlaps between distinct current strings are at most β, as in Step (ii). Extract any remaining A_i in decreasing order of δ_i; these extractions are valid because $\delta_i > \beta$. The current strings are now precisely B_i and Z_i for $1 \leq i \leq r$.

Next perform the merges $B_i \to Z_i$ of Step (ii), each with overlap β. The new string C_i has self-overlap at most β, so these merges remain valid when ties between an extraction and a merge are broken in favor of the merge. The remaining strings are the C_i, whose mutual overlaps have length $P + 1$ for distinct indices. Since every C_i has self-overlap larger than $P + 1$, extract them in decreasing order of self-overlap. This completes the first stage with exactly the representatives $\{A_i, C_i \mid 1 \leq i \leq r\}$.

These representatives are substring-free. Different indices have distinct private gaps, and the gap sequence (p_i, M, p_i, M, p_i) of A_i is not a consecutive segment of (M, M, p_i, M, M, p_i), the gap sequence of C_i. Since C_i has more ones than A_i, the reverse containment is also impossible. Thus, the greedy algorithm in the second stage can follow Steps (iii) and (iv), merging each A_i with C_i

and concatenating the resulting strings T_i. Its output g has the same length as before, and the same common superstring H gives

$$\frac{|g|}{\mathrm{OPT}} \geq \frac{|g|}{|H|} = \frac{12Pr + (115r^2 + 7r)/2}{5Pr + 27r^2 + 4P + 19r + 1}.$$

With $P = r^2$, this lower bound tends to $12/5$ as $r \to \infty$, proving the corollary. □

5 Tight Bound for Strings of Length Three

In this section, we consider inputs in which all strings have the same length k. Let ρ_k be the approximation ratio of the greedy algorithm on such inputs, that is, the supremum of $|g|/\mathrm{OPT}$ over all finite alphabets, all nonempty sets $\mathcal{S}$ of strings of length k, and all outputs g of the greedy algorithm on $\mathcal{S}$. Cazaux and Rivals [CR18] proved that $5/3 \leq \rho_3 \leq 2$. We show that $\rho_3 = 9/5$, starting with the upper bound.

Theorem 5.1 (Formal version of Theorem 1.4, upper bound)**.** *For every nonempty set $\mathcal{S}$ of strings of length three, every output g of the greedy algorithm on $\mathcal{S}$ satisfies*

$$|g| < \frac{9}{5}\,\mathrm{OPT}\,.$$

Consequently, $\rho_3 \leq 9/5$.

Then, we show that this bound is tight.

Proposition 5.2 (Formal version of Theorem 1.4, lower bound)**.** *For every integer $\ell \geq 1$, there is a set of 5ℓ strings of length three that has a common superstring of length $5\ell + 2$, whereas the greedy algorithm may output a common superstring of length 9ℓ. Consequently, $\rho_3 \geq 9/5$.*

Here, "may" means that the greedy algorithm produces this output for some way of breaking ties. Unlike the bound of Theorem 4.1, this bound does not extend to every tie-breaking rule by [Nik21] as the transformation used there is not length-preserving.

Throughout this section, $\mathcal{S}$ is a set of n strings of length three, and g is the output of an arbitrary run of the greedy algorithm on $\mathcal{S}$. By Lemma 2.1, the greedy algorithm first merges strings with overlap two, which we call the overlap-two stage, then strings with overlap one, which we call the overlap-one stage, and then concatenates the remaining strings. Let κ_2 and κ_1 be the numbers of strings after the overlap-two and the overlap-one stage, respectively, so that $\kappa_1 \leq \kappa_2$. There are $n - \kappa_2$ merges with overlap two and $\kappa_2 - \kappa_1$ merges with overlap one, so

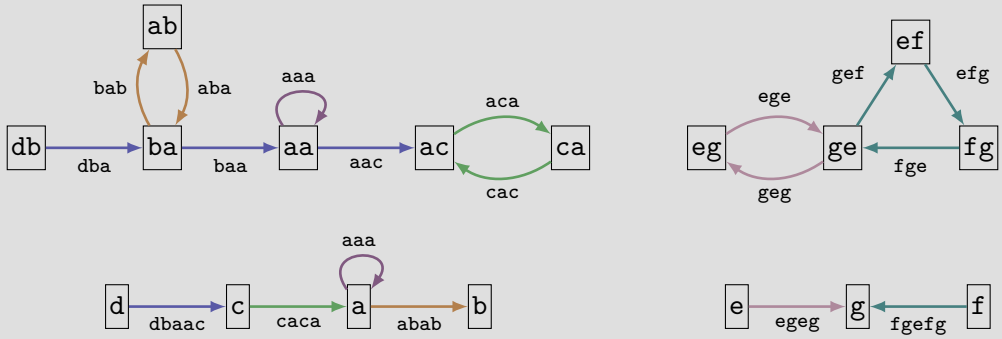
$$|g| = 3n - 2(n - \kappa_2) - (\kappa_2 - \kappa_1) = n + \kappa_1 + \kappa_2.$$

We describe the overlap-two stage in terms of the directed graph G_2 whose vertices are the strings of length two that are prefixes or suffixes of strings in $\mathcal{S}$, and in which every string $abc \in \mathcal{S}$ is an edge from ab to bc; it is a subgraph of the de Bruijn graph [DB46]. Thus, G_2 has n edges, and two input strings overlap by two symbols exactly when the first one ends at the vertex where the second one starts, so the overlap-two stage merges them along trails of G_2. For the overlap-one stage, we create a multigraph G_1 on symbols, where every string left by the overlap-two stage corresponds to an edge from its first to its last symbol. Thus, G_1 is built like G_2, but for the strings left by the overlap-two stage and with their prefixes and suffixes of length one.

As a running example, consider the $n = 13$ strings

$$\mathcal{S} = \{\texttt{aaa}, \texttt{aac}, \texttt{aba}, \texttt{aca}, \texttt{baa}, \texttt{bab}, \texttt{cac}, \texttt{dba}, \texttt{efg}, \texttt{fge}, \texttt{geg}, \texttt{ege}, \texttt{gef}\}.$$

The graph G_2 is shown below at the top, with every edge labeled by its string. The greedy algorithm may first merge `dba`, `baa`, and `aac` into `dbaac`, `aba` and `bab` into `abab`, and `cac` and `aca` into `caca`, as well as `fge`, `gef`, and `efg` into `fgefg` and `ege` and `geg` into `egeg`. Then no two of the current strings `dbaac`, `abab`, `caca`, `aaa`, `fgefg`, and `egeg` overlap by two symbols, so $\kappa_2 = 6$. The edges of G_1, shown at the bottom, correspond to the same six strings and are drawn in the same colors. In the overlap-one stage, the greedy algorithm may merge `caca` and `abab` into `cacabab`, and then `dbaac` and `cacabab` into `dbaacacabab`. After that, no two current strings overlap, so $\kappa_1 = 4$, and the greedy algorithm outputs $g =$ `dbaacacababaaafgefgegeg` of length $23 = n + \kappa_1 + \kappa_2$.



Before giving the details, we explain the idea of the proof of Theorem 5.1. The overlap-two stage leaves κ_2 strings, which correspond to κ_2 trails of G_2 that cover all its edges. We say that a trail is *closed* if it ends at the vertex where it starts, and *open* otherwise. Let $\pi(G_2)$ be the minimum number of trails that cover all edges of G_2, then Lemma 5.5 shows that $\mathrm{OPT} \geq n+\pi(G_2)+1$. No trail formed by the overlap-two stage ends at the vertex where another one starts, since otherwise their strings would overlap by two, so at most $\pi(G_2)$ of these trails are open by Lemma 5.4. It remains to bound the number of closed trails. If all closed trails had at least three edges, then the n edges of G_2 would lie on κ_2 trails, at most $\pi(G_2)$ of which are open and have at least one edge each, while the others have at least three edges each, so $n \geq 3\kappa_2 - 2\pi(G_2)$, and using the trivial bound $\kappa_1 \leq \kappa_2$ would give $|g| \leq n + 2\kappa_2 \leq \frac{5n+4\pi(G_2)}{3} < \frac{5}{3}\,\mathrm{OPT}$, which is the value $\rho_3 = 2 - 1/3$ conjectured by Cazaux and Rivals [CR18]. Hence, the main difficulty lies in closed trails with one or two edges, formed by a string aaa or by two strings aba and bab for distinct symbols a and b. Such trails make κ_2 large, and our main lemma, Lemma 5.7, shows that when κ_2 is large, the edges of G_1 can be covered by few trails, so the overlap-one stage merges many strings and κ_1 is small.

The following example shows that short closed trails make Theorem 5.1 tight: the overlap-two stage may form only closed trails with two and three edges, after which no two current strings overlap.

Proof of Proposition 5.2. Let $\mathtt{a}, \mathtt{b}, \mathtt{x}_1, \ldots, \mathtt{x}_\ell$ be distinct symbols, and let $\mathcal{S}$ be the set of the length-three substrings of

$$h = \mathtt{abx}_1\mathtt{ax}_1\mathtt{ab} \cdots \mathtt{x}_\ell\mathtt{ax}_\ell\mathtt{ab}.$$

It consists of the 5ℓ distinct strings $\mathtt{abx}_i$, $\mathtt{bx}_i\mathtt{a}$, $\mathtt{x}_i\mathtt{ax}_i$, $\mathtt{ax}_i\mathtt{a}$, and $\mathtt{x}_i\mathtt{ab}$ for $i \in [\ell]$, and $|h| = 5\ell + 2$. For every i, the edges $\mathtt{ax}_i\mathtt{a}$ and $\mathtt{x}_i\mathtt{ax}_i$ of G_2 form a closed trail, and so do the edges $\mathtt{bx}_i\mathtt{a}$, $\mathtt{x}_i\mathtt{ab}$,

and $\mathtt{abx}_i$. Consecutive edges of a trail overlap by two symbols, so the overlap-two stage may merge the strings along each of these 2ℓ trails, which gives the current strings $\mathtt{ax}_i\mathtt{ax}_i$ and $\mathtt{bx}_i\mathtt{abx}_i$. No two of these strings overlap. Indeed, each of them starts and ends with the same string, $\mathtt{ax}_i$ or $\mathtt{bx}_i$, which is different for different strings, and contains $\mathtt{x}_i$ only at its second and last positions. Hence, if a nonempty proper suffix of one of them is a proper prefix of another one, then it ends with some $\mathtt{x}_i$, so it consists of the first two symbols of the latter string, and the former string ends with these two symbols, which is impossible. Therefore, $\kappa_1 = \kappa_2 = 2\ell$, and $|g| = n + \kappa_1 + \kappa_2 = 9\ell$. Since $\mathrm{OPT} \le |h|$ and $9\ell/(5\ell+2)$ tends to $9/5$ as $\ell \to \infty$, we obtain $\rho_3 \ge 9/5$, and Theorem 5.1 gives $\rho_3 \le 9/5$. □

5.1 A Trail Cover of G_2

A *trail* P in a directed multigraph is a walk without repeated edges. A trail P with $m \ge 1$ edges visits vertices $P_0, P_1, \ldots, P_m$. We call every index i with $0 \le i \le m$ a position of P; it is *internal* if $0 < i < m$. The trail P is *closed* if $P_0 = P_m$ and *open* otherwise. A *trail cover* of a directed multigraph H is a set of edge-disjoint trails that contain all edges of H, and $\pi(H)$ denotes the minimum size of a trail cover of H. The *weak components* of H are the connected components of the undirected multigraph obtained by ignoring the directions of the edges, and H is weakly connected if it has one weak component. For a vertex v of H, let $\deg^+(v)$ and $\deg^-(v)$ be its out-degree and in-degree. The *excess* of H is $\mathrm{exc}(H) = \sum_v \max\{\deg^+(v) - \deg^-(v), 0\}$, the total number of outgoing edges in excess of incoming ones.

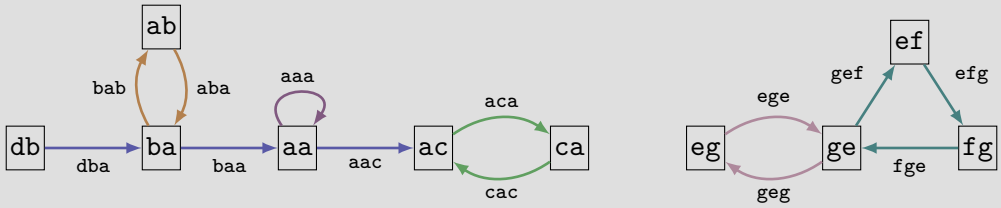
Lemma 5.3 ([GMS80, Lemma 2])**.** *If H is weakly connected and has at least one edge, then $\pi(H) = \max\{\mathrm{exc}(H), 1\}$.*

We can view the overlap-two stage of the greedy algorithm as building a trail cover of G_2, in which every current string is the overlap merge of the edges of a trail in their order on the trail. Let $\mathcal{T}$ be the resulting trail cover, which consists of κ_2 trails. Consecutive vertices of a trail P of G_2 with m edges overlap in one symbol, so there are symbols $a_0, \ldots, a_{m+1}$, which we call the *symbols of* P, such that $P_i = a_i a_{i+1}$ for $0 \le i \le m$. Note that no trail of $\mathcal{T}$ ends at the first vertex of another trail of $\mathcal{T}$, since otherwise their current strings would overlap by two after the overlap-two stage. The next lemma shows that $\mathcal{T}$ has the minimum possible number of open trails.

Lemma 5.4. *Let ν be the number of weak components of G_2 that contain no open trail of $\mathcal{T}$. Then $\mathcal{T}$ has exactly $\pi(G_2) - \nu$ open trails.*

Proof. Since no trail of $\mathcal{T}$ ends at the first vertex of another one, for every vertex v of G_2, either no trail of $\mathcal{T}$ ends at v, or no trail of $\mathcal{T}$ starts at v, or the only trail of $\mathcal{T}$ that starts or ends at v is a closed trail that starts at v. The number of trails of $\mathcal{T}$ that start at v exceeds the number of those that end at v by $\deg^+(v) - \deg^-(v)$, so in each case exactly $\max\{\deg^+(v) - \deg^-(v), 0\}$ open trails of $\mathcal{T}$ start at v. Thus, every weak component H of G_2 contains exactly $\mathrm{exc}(H)$ open trails of $\mathcal{T}$, and summing $\pi(H) = \max\{\mathrm{exc}(H), 1\}$ over all weak components H of G_2 proves the lemma. □

Recall the graph G_2 of the running example, in which the trails of $\mathcal{T}$ are drawn in different colors.



The graph G_2 has two weak components. In the first one, `db` is the only vertex with more outgoing than incoming edges, so its excess is 1; in the second one, every vertex has as many outgoing as incoming edges, so its excess is 0, and $\pi(G_2) = 1 + 1 = 2$ by Lemma 5.3. The first component contains the open trail `db` $\to$ `ba` $\to$ `aa` $\to$ `ac`, whereas the second one consists only of closed trails, so $\nu = 1$, and $\mathcal{T}$ indeed has exactly $\pi(G_2) - \nu = 1$ open trails.

We now bound OPT from below.

Lemma 5.5.

$$\mathrm{OPT} \geq n + \pi(G_2) + 1.$$

Proof. Let $\mathrm{OPT} = |\langle t_1, \ldots, t_n \rangle|$ for an ordering $t_1, \ldots, t_n$ of $\mathcal{S}$, and split this ordering into maximal blocks in which consecutive strings overlap by two symbols. The strings of every block are the consecutive edges of a trail of G_2, and every input string lies in exactly one block, so these trails form a trail cover of G_2. Hence, the number r of blocks is at least $\pi(G_2)$. Since $n - r$ pairs of consecutive strings overlap by two symbols and the other $r - 1$ pairs by at most one symbol,

$$\mathrm{OPT} \geq 3n - 2(n - r) - (r - 1) = n + r + 1 \geq n + \pi(G_2) + 1. \qquad \square$$

Let G_1 be the directed multigraph whose vertices are symbols and which has an edge e_P from a_0 to a_{m+1} for every trail $P \in \mathcal{T}$ with symbols $a_0, \ldots, a_{m+1}$. In other words, e_P goes from the first to the last symbol of the current string of P. Now, we analyze the overlap-one stage.

Lemma 5.6.

$$2\kappa_1 \leq \kappa_2 + \pi(G_1).$$

Proof. Let $\mathcal{R}$ be the set of the κ_2 current strings after the overlap-two stage. The set $\mathcal{R}$ is substring-free, and the rest of the run of the greedy algorithm is a run of the greedy algorithm on $\mathcal{R}$ with output g. This run performs $\kappa_2 - \kappa_1$ merges with overlap one, so the compression of g, that is, the total length of the strings in $\mathcal{R}$ minus $|g|$, is $\kappa_2 - \kappa_1$. Every string of $\mathcal{R}$ is the current string of a unique trail $P \in \mathcal{T}$ and thus corresponds to the edge e_P of G_1. Take a minimum trail cover of G_1 and order $\mathcal{R}$ by listing, trail after trail, the strings that correspond to the edges of each trail in their order on the trail. If two edges are consecutive on a trail of G_1, then the first ends at the symbol at which the second starts, so the corresponding strings overlap by at least one symbol. Hence, all pairs of consecutive strings in this order except at most $\pi(G_1) - 1$ overlap by at least one symbol, so the overlap merge of $\mathcal{R}$ in this order has compression at least $\kappa_2 - \pi(G_1)$. By Theorem 2.2 applied to the input $\mathcal{R}$, $\kappa_2 - \kappa_1 \geq (\kappa_2 - \pi(G_1))/2$. $\square$

Recall that $|g| = n + \kappa_1 + \kappa_2$ and that $\mathrm{OPT} \geq n + \pi(G_2) + 1$ by Lemma 5.5. Since $\frac{9}{5}n + \frac{8}{5}\pi(G_2) < \frac{9}{5}(n + \pi(G_2) + 1)$, it suffices to show that $|g| \leq \frac{9}{5}n + \frac{8}{5}\pi(G_2)$, that is,

$$\kappa_1 + \kappa_2 \leq \frac{4}{5}n + \frac{8}{5}\pi(G_2). \tag{5}$$

We have two upper bounds on κ_1: the trivial bound $\kappa_1 \le \kappa_2$ and the bound $\kappa_1 \le \frac{1}{2}(\kappa_2 + \pi(G_1))$ of Lemma 5.6. Adding one fifth of the first bound to four fifths of the second one, we get

$$\kappa_1 + \kappa_2 \le \frac{1}{5}\kappa_2 + \frac{2}{5}(\kappa_2 + \pi(G_1)) + \kappa_2 = \frac{8}{5}\kappa_2 + \frac{2}{5}\pi(G_1),$$

so (5) holds whenever

$$\pi(G_1) \le 2n + 4\pi(G_2) - 4\kappa_2. \tag{6}$$

To understand the right-hand side of (6), consider a trail $P \in \mathcal{T}$ with m_P edges. Its edges are m_P input strings, so they add $2m_P$ to the right-hand side through n. The trail itself is one of the κ_2 current strings after the overlap-two stage, so it subtracts four through κ_2. If P is open, then it also adds four through $\pi(G_2)$, since $\pi(G_2)$ is the number of open trails of $\mathcal{T}$ plus ν by Lemma 5.4. Accordingly, we define the *contribution* of P as $\mathrm{contr}(P) = 2m_P - 4$ if P is closed and $\mathrm{contr}(P) = 2m_P$ if P is open, so that

$$2n + 4\pi(G_2) - 4\kappa_2 = 4\nu + \sum_{P \in \mathcal{T}} \mathrm{contr}(P),$$

and (6) becomes

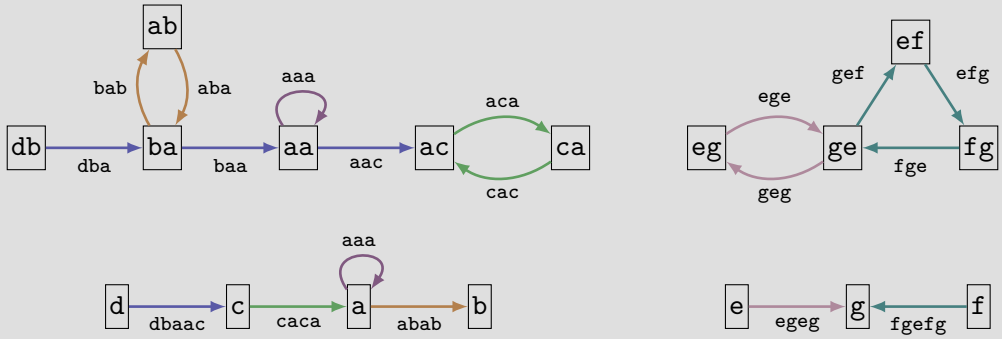
$$\pi(G_1) \le 4\nu + \sum_{P \in \mathcal{T}} \mathrm{contr}(P).$$

Thus, Theorem 5.1 follows from the following main lemma, which we prove in Section 5.2.

Lemma 5.7.

$$\pi(G_1) \le 4\nu + \sum_{P \in \mathcal{T}} \mathrm{contr}(P).$$

Recall the graphs G_2 and G_1 of the running example.



Here, $\mathcal{T}$ consists of six trails: in the first component, the open trail `db` → `ba` → `aa` → `ac` with three edges, the loop `aaa`, and the closed trails `ab` → `ba` → `ab` and `ca` → `ac` → `ca` with two edges each, and in the second component, the closed trails `eg` → `ge` → `eg` and `fg` → `ge` → `ef` → `fg` with two and three edges. Their contributions are 6, −2, 0, 0, 0, and 2, respectively. Since $\nu = 1$, the right-hand side of Lemma 5.7 is $4 + 6 = 10$. In fact, one can cover G_1 with three trails, so $\pi(G_1) = 3$.

5.2 A Trail Cover of G_1

In this subsection, we prove Lemma 5.7, that is, we construct a trail cover of G_1 with at most $4\nu + \sum_{P \in \mathcal{T}} \mathrm{contr}(P)$ trails. Recall that every trail $P \in \mathcal{T}$ corresponds to a single edge e_P of G_1.

The contribution $\mathrm{contr}(P)$ is positive for open trails and for closed trails with at least three edges, but it is 0 for closed trails with two edges and -2 for loops. We split the trails of $\mathcal{T}$, and thus the edges of G_1, into groups such that, except for at most ν groups, every group has positive total contribution and its edges in G_1 can be covered by at most this many trails. We form the groups as follows. First, a loop aaa that shares its weak component of G_2 with other edges lies on another trail of $\mathcal{T}$ that passes through aa, and we put the loop into the group of this trail, which absorbs the loop; the remaining contribution $\mathrm{rem}(P) = \mathrm{contr}(P) - 2\,\mathrm{loops}(P)$ of a trail P, where $\mathrm{loops}(P)$ is the number of loops that P absorbs, is still nonnegative (Lemma 5.8). Second, a closed trail with $\mathrm{rem}(P) = 0$ uses only two symbols a and b, and all such trails on a and b together with their loops form a *bundle*, whose edges in G_1 form a single trail of G_1 (Lemmas 5.10 and 5.11); the total contribution of a bundle is 0, so we put it into the group of a trail with positive remaining contribution whenever possible. Third, the other trails, that is, the open trails and the closed trails with $\mathrm{rem}(P) > 0$, are called *connectors*; we attach every bundle that shares its weak component of G_2 with a connector to one such connector, and the group of a connector Q, formed by Q, its loops, and the attached bundles, has total contribution $\mathrm{rem}(Q)$ and can be covered by $\mathrm{rem}(Q)$ trails of G_1 (Lemma 5.12). Finally, a weak component of G_2 without connectors is a single loop or a single bundle; it has no open trails, so there are at most ν such components, and we cover each of them by one trail.

A loop is a trail of $\mathcal{T}$ that consists of a single edge from aa to aa, that is, of the string aaa for a symbol a; its edge in G_1 is a loop at a. Consider a loop aaa whose weak component in G_2 has another edge. Then another trail $P \in \mathcal{T}$ visits aa, and since no trail of $\mathcal{T}$ ends at the first vertex of another one, it neither starts nor ends at aa, so $P_i = aa$ for an internal position i of P. Several trails, or several positions of one trail, may visit aa; we fix one such trail P and position i and say that P *absorbs* the loop aaa at position i, so every loop is absorbed by at most one trail. Note that $a_{i-1} \neq a \neq a_{i+2}$, since otherwise P would contain the edge aaa. For a trail $P \in \mathcal{T}$ that is not a loop, let $\mathrm{loops}(P)$ be the number of loops that P absorbs, and let $\mathrm{rem}(P) = \mathrm{contr}(P) - 2\,\mathrm{loops}(P)$ be the *remaining contribution* of P.

By the following lemma, to prove Lemma 5.7, it suffices to construct a trail cover of G_1 of size at most $\sum_P \mathrm{rem}(P) + \nu$, where the sum is over the trails $P \in \mathcal{T}$ that are not loops.

Lemma 5.8.

$$\sum_P \mathrm{rem}(P) + \nu \leq 4\nu + \sum_{P \in \mathcal{T}} \mathrm{contr}(P),$$

where the sum is over the trails $P \in \mathcal{T}$ that are not loops.

Proof. By the definition of $\mathrm{rem}(P)$, the sum $\sum_P \mathrm{rem}(P)$ equals the sum of the contributions of the trails of $\mathcal{T}$ that are not loops and of the loops that are absorbed. Every loop that is not absorbed forms a weak component of G_2 without open trails, so these loops contribute at least -2ν, and therefore $\sum_P \mathrm{rem}(P) \leq 2\nu + \sum_{P \in \mathcal{T}} \mathrm{contr}(P)$, which implies the lemma. □

A closed trail P with $\mathrm{rem}(P) = 0$ that is not a loop adds nothing to the bound $\sum_P \mathrm{rem}(P) + \nu$, so its edge e_P must be covered by a trail of G_1 that is accounted for by other trails of $\mathcal{T}$. To describe such trails, let $\mathrm{free}(P)$ be the number of internal positions at which P absorbs no loop; these positions include all internal positions i with $a_i \neq a_{i+1}$. The trail P has $m_P - 1$ internal positions, so $\mathrm{free}(P) = m_P - 1 - \mathrm{loops}(P)$, that is, $\mathrm{rem}(P) = 2\,\mathrm{free}(P) - 2$ if P is closed and $\mathrm{rem}(P) = 2\,\mathrm{free}(P) + 2$ if P is open. We start with a simple property of closed trails.

Lemma 5.9. *Let $P \in \mathcal{T}$ be a closed trail with $m \geq 2$ edges and symbols $a_0, \ldots, a_{m+1}$. Then at least two indices $0 \leq i < m$ satisfy $a_i \neq a_{i+1}$.*

Proof. Since P is closed, $a_m = a_0$. Since P has $m \geq 2$ distinct edges, the sequence $a_0, \ldots, a_m$ is not constant but starts and ends with the same symbol, so $a_i \neq a_{i+1}$ for at least two indices $0 \leq i < m$. □

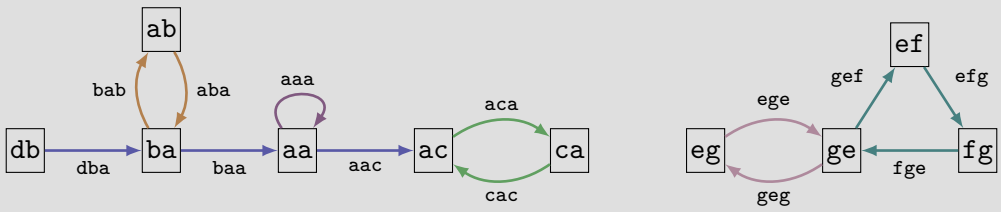
The next lemma describes the closed trails with $\mathrm{rem}(P) = 0$.

Lemma 5.10. *Every closed trail $P \in \mathcal{T}$ that is not a loop satisfies* $\mathrm{free}(P) \geq 1$*, that is,* $\mathrm{rem}(P) \geq 0$*. If* $\mathrm{rem}(P) = 0$*, then the current string of P is $ab^j a^\ell b$ for some distinct symbols a and b and some $j, \ell \in \{1, 2\}$, and P absorbs a loop at every visit to aa or bb (if such exists).*

Proof. Let P have $m \geq 2$ edges and symbols $a_0, \ldots, a_{m+1}$. By Lemma 5.9, at least two indices $0 \leq i < m$ satisfy $a_i \neq a_{i+1}$, and all of them except possibly $i = 0$ are internal positions at which P absorbs no loop. Hence, $\mathrm{free}(P) \geq 1$, and if $\mathrm{free}(P) = 1$, then there are exactly two such indices, 0 and some j, and P absorbs loops at all internal positions except j. With $a = a_0$ and $b = a_1$, this means that $a_1 = \cdots = a_j = b$ and $a_{j+1} = \cdots = a_m = a$, and $a_{m+1} = a_1 = b$. At every internal position i with $a_i = a_{i+1}$, P absorbs a loop, so $a_{i-1} \neq a_i \neq a_{i+2}$; hence, $j \leq 2$ and $m - j \leq 2$, and the current string $a_0 a_1 \cdots a_{m+1}$ of P is $ab^j a^{m-j} b$, which proves the second statement. □

By Lemma 5.10, every closed trail $P \in \mathcal{T}$ with $\mathrm{rem}(P) = 0$ that is not a loop has current string $ab^j a^\ell b$ for two distinct symbols a and b, so we can group these trails by their pair of symbols. Namely, for distinct symbols a and b, the *bundle* $Z_{a,b}$ consists of the closed trails $P \in \mathcal{T}$ with $\mathrm{rem}(P) = 0$ whose symbols are a and b, together with the loops that these trails absorb; thus, $Z_{a,b} = Z_{b,a}$. The trails of $\mathcal{T}$ that are neither loops nor in bundles, that is, the open trails and the closed trails P with $\mathrm{rem}(P) > 0$, are called *connectors*.

Recall the graph G_2 of the running example.



Let Q be the open trail $\texttt{db} \to \texttt{ba} \to \texttt{aa} \to \texttt{ac}$, so its internal positions are 1 and 2. It absorbs the loop `aaa` at position 2, hence, $\mathrm{loops}(Q) = 1$, $\mathrm{free}(Q) = 1$, and $\mathrm{rem}(Q) = \mathrm{contr}(Q) - 2\,\mathrm{loops}(Q) = 4$. The closed trail $\texttt{ab} \to \texttt{ba} \to \texttt{ab}$ absorbs no loop at its only internal position 1, so $\mathrm{free}(P) = 1$ and $\mathrm{rem}(P) = 2\,\mathrm{free}(P) - 2 = 0$; similarly, the closed trail $\texttt{ca} \to \texttt{ac} \to \texttt{ca}$ has $\mathrm{rem}(P) = 0$. Their current strings `abab` and `caca` have the form $ab^j a^\ell b$ of Lemma 5.10 with $j = \ell = 1$, so they form the bundles $Z_{\texttt{a,b}}$ and $Z_{\texttt{a,c}}$, whereas Q is a connector. In the second component, the closed trail $Q' = \texttt{fg} \to \texttt{ge} \to \texttt{ef} \to \texttt{fg}$ absorbs no loops at its internal positions 1 and 2, so $\mathrm{free}(Q') = 2$ and $\mathrm{rem}(Q') = 2\,\mathrm{free}(Q') - 2 = 2$, and Q' is a connector; the closed trail $\texttt{eg} \to \texttt{ge} \to \texttt{eg}$ has $\mathrm{rem}(P) = 0$ and forms the bundle $Z_{\texttt{e,g}}$. Thus, $\sum_P \mathrm{rem}(P) = \mathrm{rem}(Q) + \mathrm{rem}(Q') = 6$, and the edges of the bundles in G_1 must be covered together with e_Q and $e_{Q'}$, by trails accounted for by $\mathrm{rem}(Q)$ and $\mathrm{rem}(Q')$.

Lemma 5.11. *Besides loops, every bundle $Z_{a,b}$ contains at most one trail that starts at ab and at most one trail that starts at ba. The edges of the trails of $Z_{a,b}$ in G_1 form a trail of G_1. Moreover, different bundles visit disjoint sets of vertices.*

Proof. Since no trail of $\mathcal{T}$ ends at the first vertex of another one, different closed trails of $\mathcal{T}$ start at different vertices, and by Lemma 5.10, every trail of $Z_{a,b}$ that is not a loop starts at ab or ba. The edges of these trails in G_1 are therefore an edge from a to b, an edge from b to a, or both, and the loops of $Z_{a,b}$ give loops at a and b, so these edges form a trail of G_1. By Lemma 5.10, the trails of $Z_{a,b}$ visit only the vertices ab, ba, aa, and bb, and they absorb the loop aaa at every visit to aa. Hence, two bundles can share only a vertex aa, but then both would absorb the loop aaa. □

A weak component of G_2 without connectors consists of a single loop or of a single bundle, since every loop in a weak component with other edges is absorbed by a trail of a bundle, and different bundles share no vertices. Such a component contains no open trail of $\mathcal{T}$, so there are at most ν such components, and the term ν in the bound $\sum_P \mathrm{rem}(P) + \nu$ accounts for the one trail that covers each of them.

Now, consider a bundle $Z = Z_{a,b}$ in a weak component of G_2 with a connector. Then some edge outside Z is incident to a vertex visited by Z. The trail of $\mathcal{T}$ that contains this edge is not a loop, since Z absorbs the loops, and it is not in another bundle by Lemma 5.11, so some connector visits a vertex visited by Z. We attach Z to one of the connectors as follows.

If some connector Q starts or ends at a vertex v visited by Z, we attach Z to Q at this endpoint. No trail of Z starts at v. Indeed, every trail of Z, including its loops, is closed, and no trail of $\mathcal{T}$ ends at the first vertex of another one. In particular, v is neither aa nor bb, since by Lemma 5.10, if Z visits aa, then it contains the loop aaa, which starts at aa, and similarly for bb. Hence, v is ab or ba, say, $v = ab$. By Lemma 5.10, every trail of Z that is not a loop starts at ab or ba, so it starts at ba, and by Lemma 5.11, Z consists of a single such trail and possibly of loops at aa and bb. Therefore, e_Q together with the edges of Z forms a trail in G_1. By Lemma 5.11, at most one bundle is attached at each endpoint of Q, and e_Q together with the bundles attached at both endpoints still forms a trail of G_1.

In the running example, the connector $Q = \mathtt{db} \to \mathtt{ba} \to \mathtt{aa} \to \mathtt{ac}$ ends at $v = \mathtt{ac}$, which the bundle $Z_{\mathtt{a},\mathtt{c}}$ visits, so $Z_{\mathtt{a},\mathtt{c}}$ is attached to Q at this endpoint. This bundle consists of the single trail $\mathtt{ca} \to \mathtt{ac} \to \mathtt{ca}$, which starts at $\mathtt{ca}$ (not at v), and its edge $\mathtt{caca}$ in G_1 goes from $\mathtt{c}$ to $\mathtt{a}$. The edge $e_Q = \mathtt{dbaac}$ goes from $\mathtt{d}$ to $\mathtt{c}$, so together they form the trail $\mathtt{d} \to \mathtt{c} \to \mathtt{a}$ of G_1.

Otherwise, no connector starts or ends at a vertex visited by Z, and we take a connector Q that visits a vertex visited by Z and attach Z to Q at an internal position i such that Z visits Q_i. The trail Q absorbs no loop at i, since either Q_i is ab or ba, or the loop at Q_i is absorbed by Z. By Lemma 5.11, different bundles are attached at different positions.

Now, every edge of G_1 outside the weak components of G_2 without connectors lies in the group of some connector Q, which consists of Q, the bundles attached to Q, and the loops absorbed by Q or by these bundles. We cover the edges of each group in G_1 separately, except for some loops: a loop aaa can be inserted into any trail of G_1 that visits a, so it needs a trail of its own only if a has no other edges in G_1. A loop aaa is *isolated in* G_1 if no other edge of G_1 is incident to a. By Lemma 5.10, the loops absorbed by bundles are not isolated in G_1.

Note that a loop becomes isolated when its symbol occurs only *inside* the current strings. For example, for the strings $\mathtt{xyy}$, $\mathtt{yyy}$, and $\mathtt{yyz}$, the greedy algorithm could merge $\mathtt{xyy}$ and $\mathtt{yyz}$ into $\mathtt{xyyz}$, so that the open trail $\mathtt{xy} \to \mathtt{yy} \to \mathtt{yz}$ absorbs the loop $\mathtt{yyy}$. In G_1, the string $\mathtt{xyyz}$ gives an edge from $\mathtt{x}$ to $\mathtt{z}$, and no other current string starts or ends with $\mathtt{y}$, so the loop $\mathtt{yyy}$ is isolated in G_1 and needs a trail of its own.

For a connector Q, let D_Q be the subgraph of G_1 formed by e_Q, by the edges in G_1 of the bundles attached to Q, and by the edges of the loops isolated in G_1 that Q absorbs. The next lemma shows that the remaining contribution of Q bounds the size of a trail cover of D_Q.

Lemma 5.12. *Every connector Q satisfies $\pi(D_Q) \le \mathrm{rem}(Q)$.*

Proof. Let B_Q be the set of internal positions at which bundles are attached to Q; since Q absorbs no loop at these positions, $|B_Q| \le \mathrm{free}(Q)$. Let λ_Q be the number of loops isolated in G_1 that Q absorbs. Suppose that Q absorbs such a loop aaa at position i, so that $Q_{i-1} = ba$ and $Q_{i+1} = ac$ for some symbols $b \neq a$ and $c \neq a$. No bundle visits ba, since a bundle that visits ba is $Z_{a,b}$, whose edges in G_1 are incident to a; similarly, no bundle visits ac. If Q is closed, then $i-1$ and $i+1$ are internal positions, since otherwise Q starts at ba or ac, and e_Q is incident to a. Since $c \neq a$, Q absorbs no loop at $i+1$. Hence, the positions $i+1$ of these λ_Q loops are distinct positions without loops outside B_Q, and they are internal, except that maybe $i+1 = m_Q$ is possible if Q is open. Therefore,

$$|B_Q| + \lambda_Q \le \begin{cases} \mathrm{free}(Q) + 1 & \text{if } Q \text{ is open,} \\ \mathrm{free}(Q) & \text{if } Q \text{ is closed.} \end{cases} \tag{7}$$

The edge e_Q and the bundles attached at the endpoints of Q form one trail of G_1, and so does every bundle attached at a position in B_Q, by Lemma 5.11. Together with the λ_Q isolated loops, this gives $\pi(D_Q) \le 1 + |B_Q| + \lambda_Q$. Recall that $\mathrm{rem}(Q) = 2\,\mathrm{free}(Q) + 2$ if Q is open and $\mathrm{rem}(Q) = 2\,\mathrm{free}(Q) - 2$ if Q is closed. Hence, by (7), $\pi(D_Q) \le \mathrm{free}(Q) + 2 \le \mathrm{rem}(Q)$ if Q is open, and $\pi(D_Q) \le \mathrm{free}(Q) + 1 \le 2\,\mathrm{free}(Q) - 2 = \mathrm{rem}(Q)$ if Q is closed and $\mathrm{free}(Q) \ge 3$.

It remains to consider a closed connector Q with $\mathrm{free}(Q) = 2$, that is, with $\mathrm{rem}(Q) = 2$. By (7), $|B_Q| + \lambda_Q \le 2$, and if $|B_Q| + \lambda_Q \le 1$, then $\pi(D_Q) \le 1 + |B_Q| + \lambda_Q \le 2 = \mathrm{rem}(Q)$. Hence, we may assume that $|B_Q| + \lambda_Q = 2$. We first show that $\lambda_Q = 0$. Suppose that Q absorbs a loop isolated in G_1 at position i. As shown above, $Q_{i-1} = ba$ and $Q_{i+1} = ac$ for symbols $b \neq a$ and $c \neq a$, and since Q is closed, $i-1$ and $i+1$ are internal positions. The vertices ba and ac consist of two distinct symbols, so Q absorbs no loop at $i-1$ and $i+1$, and since $\mathrm{free}(Q) = 2$, these are the only internal positions at which Q absorbs no loop. Every position in B_Q is such a position, and a bundle attached at a position j visits Q_j, but, as shown above, no bundle visits ba or ac; hence, $B_Q = \varnothing$. Then $\lambda_Q = 2 - |B_Q| = 2$, so Q absorbs another loop isolated in G_1 at a position $i' \neq i$. By the same argument, $i'-1$ and $i'+1$ are also internal positions at which Q absorbs no loop, so $\{i'-1, i'+1\} = \{i-1, i+1\}$, which implies $i' = i$, a contradiction.

Hence, $\lambda_Q = 0$, and $|B_Q| = 2$, that is, bundles are attached at both internal positions of Q without loops. Let Q have m edges and symbols $a_0, \ldots, a_{m+1}$. By Lemma 5.9, at least two indices $0 \le i < m$ satisfy $a_i \neq a_{i+1}$, and Q absorbs no loop at the internal ones among them; since $\mathrm{free}(Q) = 2$, there are two or three such indices.

Suppose that there are two such indices. Since $a_m = a_0$, the sequence $a_0, \ldots, a_m$ changes from some symbol a to another symbol b at one of these indices and back from b to a at the other, so Q visits ab and ba at them. A bundle attached to Q at one of these indices visits ab or ba, so this bundle is $Z_{a,b}$, which visits both ab and ba. If one of the two indices is 0, then $Z_{a,b}$ visits Q_0, so it is attached at an endpoint and not at the other index, and thus $|B_Q| \le \mathrm{free}(Q) - 1$. Otherwise, both indices are internal positions without loops, so they form B_Q, but $Z_{a,b}$ is attached at only one position. Both cases contradict $|B_Q| = 2$.

Hence, there are three such indices, and one of them is 0, since $\mathrm{free}(Q) = 2$. Let these indices be $0 < j < k < m$, and let $a = a_0$, $b = a_1$, and $c = a_{j+1}$. Then $a_1 = \cdots = a_j = b$, $a_{j+1} = \cdots = a_k = c$, and $a_{k+1} = \cdots = a_m = a$, where the last equality holds since $a_m = a_0$. Hence,

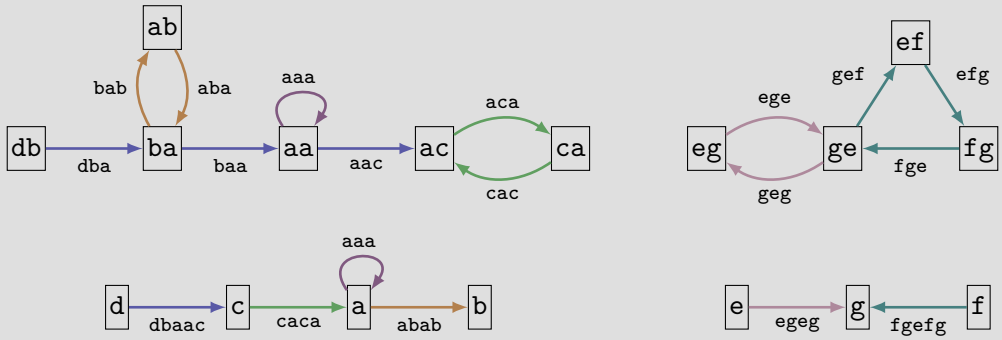
Q visits ab, bc, and ca at the positions 0, j, and k. The bundles attached at the positions j and k visit bc and ca, so they are $Z_{b,c}$ and $Z_{c,a}$.

We now list the edges of D_Q that are not loops.

- The edge e_Q, which goes from $a_0 = a$ to $a_{m+1} = a_1 = b$.
- An edge from b to a if a bundle is attached at an endpoint of Q: since Q starts and ends at ab, this bundle is $Z_{a,b}$, and as shown above, its only trail that is not a loop starts at ba.
- At most one edge from b to c and at most one edge from c to b, which come from $Z_{b,c}$, and similarly, at most one edge from c to a and at most one edge from a to c, which come from $Z_{c,a}$, by Lemma 5.11.

Since $\lambda_Q = 0$, there are no other edges, so, ignoring loops, every symbol has in-degree and out-degree at most two in D_Q, that is, $-2 \le \deg^+(v) - \deg^-(v) \le 2$ for every vertex v. These differences sum to zero, so if $\operatorname{exc}(D_Q) \ge 3$, then at least two vertices have positive differences and at least two other vertices have negative differences, which is impossible, since D_Q has only the three vertices a, b, and c. Hence, $\operatorname{exc}(D_Q) \le 2$. Finally, D_Q is weakly connected, since e_Q joins a and b, and $Z_{b,c}$ joins b and c, so Lemma 5.3 gives $\pi(D_Q) \le 2 = \operatorname{rem}(Q)$. □

Recall the graphs G_2 and G_1 of the running example.



In this example $Z_{\mathtt{a,b}}$ is attached to Q at the internal position 1, where Q visits `ba`. The loop `aaa` is not isolated in G_1, so D_Q consists of the edges `dbaac`, `caca`, and `abab` of G_1. The proof of Lemma 5.12 covers them by the trail `d` → `c` → `a` and by the edge `a` → `b` of the bundle attached at position 1, so $\pi(D_Q) \le 2 \le \operatorname{rem}(Q)$. In the second component, no connector starts or ends at `eg` or `ge`, so $Z_{\mathtt{e,g}}$ is attached to Q' at the internal position 1, where Q' visits `ge`. Then $D_{Q'}$ consists of the edges `fgefg` and `egeg` of G_1, which both end at `g`, so $\pi(D_{Q'}) = 2 = \operatorname{rem}(Q')$, and Lemma 5.12 holds with equality. Finally, as in the proof below, the loop `aaa` is inserted at `a`, which gives the trail `d` → `c` → `a` → `a` → `b`; together with the two trails that cover $D_{Q'}$, it covers G_1 with three trails, whereas $\sum_Q \operatorname{rem}(Q) + \nu = 7$.

Proof of Lemma 5.7. We cover the edges of G_1 as follows. For every weak component of G_2 without connectors, we take the trail of G_1 formed by its edges in G_1, and for every connector Q, we take a minimum trail cover of D_Q. These trails are edge-disjoint, since every bundle is attached to at most one connector and every loop is absorbed by at most one trail, and they cover all edges of G_1 except the loops absorbed by connectors that are not isolated in G_1. Such a loop aaa is the only uncovered edge at a, and another edge of G_1 is incident to a, so the loop can be inserted into a trail

that visits a. The weak components of G_2 without connectors contain no open trails, so there are at most ν of them. Therefore Lemmas 5.12 and 5.8 give

$$\pi(G_1) \le \sum_Q \mathrm{rem}(Q) + \nu = \sum_P \mathrm{rem}(P) + \nu \le 4\nu + \sum_{P \in \mathcal{T}} \mathrm{contr}(P),$$

where the first sum is over connectors Q and the second over the trails $P \in \mathcal{T}$ that are not loops; these two sums are equal, since $\mathrm{rem}(P) = 0$ for every trail P in a bundle. □

6 Lower Bound for Strings of Length at Least Six

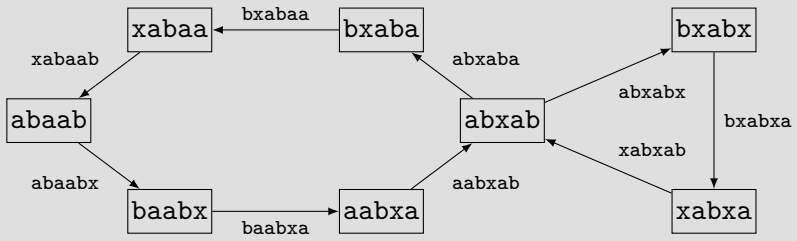
In this section, we show that $\rho_k \ge 2$ for every $k \ge 6$, where ρ_k, defined in Section 5, is the approximation ratio of the greedy algorithm on strings of length k. Thus, whereas the greedy algorithm is a 9/5-approximation on strings of length three by Theorem 5.1, its approximation ratio on strings of any fixed length $k \ge 6$ is at least 2.

Theorem 6.1 (Formal version of Theorem 1.5)**.** *For all integers $k \ge 6$ and $\ell \ge 1$, there is a set of $(2k-3)\ell$ strings of length k that has a common superstring of length $(2k-3)\ell + k - 1$, whereas the greedy algorithm may output a common superstring of length $(4k-6)\ell$. Consequently, $\rho_k \ge 2$ for every $k \ge 6$.*

As in Proposition 5.2, "may" means that the greedy algorithm produces this output for some way of breaking ties. The second statement follows from the first one, since $(4k-6)\ell/((2k-3)\ell + k - 1)$ tends to 2 as $\ell \to \infty$.

Before giving the details, we explain the idea of the construction. As in Section 5, consider the graph whose vertices are strings of length $k-1$ and in which every input string is an edge from its length-$(k-1)$ prefix to its length-$(k-1)$ suffix; it is a subgraph of the de Bruijn graph [DB46]. Two input strings overlap by $k-1$ symbols exactly when the first one ends at the vertex where the second one starts. The input consists of ℓ blocks, and the strings of each block form two cycles with k and $k-3$ edges that share a vertex. Moreover, the cycles with k edges of all blocks pass through one common vertex. Hence there is a common superstring that walks along both cycles of every block and passes from one block to the next through this common vertex, so that consecutive strings in it overlap by $k-1$ symbols, and every block adds only $2k-3$ symbols to it, one per string. The greedy algorithm, however, may merge the strings of each of the two cycles of a block separately, and the two resulting strings overlap by a single symbol, so every block adds $4k-6$ symbols to its output.

For $k = 6$ and $\ell = 1$, the input consists of the six rotations of `abaabx` and the three strings `abxabx`, `bxabxa`, and `xabxab`, which form the following two cycles.



The common superstring `abaabxabxabaab` of length 14 walks from `abaab` along the left cycle to `abxab`, around the right cycle, and back to `abaab`. The greedy algorithm may instead

merge the rotations of `abaabx` into `xabaabxabaa` and the other three strings into `bxabxabx`, which overlap by one symbol, and output `bxabxabxabaabxabaa` of length 18.

Proof of Theorem 6.1. Let $\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{x}_1, \ldots, \mathtt{x}_\ell$ be distinct symbols, and let $y = \mathtt{c}^{k-6}$, which is empty for $k = 6$. For every $i \in [\ell]$, let

$$u_i = \mathtt{abaabx}_i y \qquad \text{and} \qquad v_i = \mathtt{abx}_i y,$$

so that $|u_i| = k$ and $|v_i| = k - 3$. The i-th block consists of the strings of $\mathrm{CSub}_k(u_i)$, that is, the rotations of u_i, and of the strings of $\mathrm{CSub}_k(v_i)$, and the input $\mathcal{S}$ is the union of the ℓ blocks.

We first show that $|\mathcal{S}| = (2k-3)\ell$. Since $\mathtt{x}_i$ occurs in u_i and v_i only once, the position of the first $\mathtt{x}_i$ distinguishes the k rotations of u_i, as well as the length-k substrings of v_i^∞ that start at $k-3$ consecutive positions, which form $\mathrm{CSub}_k(v_i)$. Moreover, every rotation of u_i contains three $\mathtt{a}$'s, whereas every string of $\mathrm{CSub}_k(v_i)$ contains at most two, since any three occurrences of $\mathtt{a}$ in v_i^∞ span at least $2k - 5 > k$ symbols. Finally, all strings of the i-th block contain $\mathtt{x}_i$, so different blocks are disjoint.

The string

$$h = y\mathtt{abaabx}_1 y\mathtt{abx}_1 y\mathtt{abaab} \cdots \mathtt{x}_\ell y\mathtt{abx}_\ell y\mathtt{abaab}$$

has length $(2k-3)\ell + k - 1$, and it is a common superstring of $\mathcal{S}$. Indeed, for every i, the string h contains $y\mathtt{abaabx}_i y\mathtt{abx}_i y\mathtt{abaab}$. The suffix $\mathtt{abx}_i y\mathtt{abaab}$ of this string contains the rotations $\mathtt{abx}_i y\mathtt{aba}$, $\mathtt{bx}_i y\mathtt{abaa}$, and $\mathtt{x}_i y\mathtt{abaab}$ of u_i, and its prefix $y\mathtt{abaabx}_i y\mathtt{ab}$ is a substring of u_i^∞ of length $2k-4$ whose $k-3$ length-k substrings are the other rotations of u_i. The substring $\mathtt{abx}_i y\mathtt{abx}_i y\mathtt{ab}$ of h is a substring of v_i^∞ of length $2k - 4 = |v_i| + k - 1$, so its length-k substrings form $\mathrm{CSub}_k(v_i)$. Moreover, h is a shortest common superstring of $\mathcal{S}$: the strings of $\mathcal{S}$ are distinct and have length k, and hence its common superstring is of length at least $|\mathcal{S}| + k - 1 = |h|$.

We now describe a run of the greedy algorithm. For every i, let

$$u_i' = \mathtt{x}_i y\mathtt{abaabx}_i y\mathtt{abaa} \qquad \text{and} \qquad v_i' = \mathtt{bx}_i y\mathtt{abx}_i y\mathtt{abx}_i.$$

These are substrings of u_i^∞ and v_i^∞ of lengths $2k - 1 = |u_i| + k - 1$ and $2k - 4 = |v_i| + k - 1$, so their length-k substrings form $\mathrm{CSub}_k(u_i)$ and $\mathrm{CSub}_k(v_i)$, respectively. Consecutive length-k substrings of a string overlap by $k-1$ symbols, and by Lemma 2.1, no two current strings of the greedy algorithm overlap by more than $k-1$ symbols. Hence, the greedy algorithm may start by merging the length-k substrings of every u_i' and v_i' from left to right, which gives the 2ℓ current strings $u_1', v_1', \ldots, u_\ell', v_\ell'$.

We claim that $\mathrm{ov}(v_i', u_i') = \mathtt{x}_i$ for every i, and that every other overlap of two distinct strings among $u_1', v_1', \ldots, u_\ell', v_\ell'$ is empty. Note that u_i' starts with $\mathtt{x}_i$, ends with $\mathtt{aa}$, and contains $\mathtt{x}_i$ only at positions 1 and $k+1$, whereas v_i' starts with $\mathtt{b}$, ends with $\mathtt{x}_i$, and does not contain $\mathtt{aa}$. In particular, $\mathrm{ov}(v_i', u_i')$ is nonempty. Let $s \in \{u_i', v_i'\}$ and $t \in \{u_j', v_j'\}$ be distinct strings such that $\mathrm{ov}(s,t)$ is nonempty. If $\mathrm{ov}(s,t)$ ends with $\mathtt{a}$, then $s = u_i'$, and $\mathrm{ov}(s,t) \neq \mathtt{a}$, since t starts with $\mathtt{x}_j$ or $\mathtt{b}$. Then $\mathrm{ov}(s,t)$ ends with $\mathtt{aa}$, so $t = u_j'$, and $\mathrm{ov}(s,t)$ is a suffix of u_i' that starts with $\mathtt{x}_j$, so $j = i$ and $t = s$, a contradiction. Otherwise, $s = v_i'$ and $\mathrm{ov}(s,t)$ ends with $\mathtt{x}_i$, so t contains $\mathtt{x}_i$, that is, $t = u_i'$. Then $|\mathrm{ov}(s,t)|$ is 1 or $k+1$, but in the latter case, $\mathrm{ov}(s,t)$ would contain $\mathtt{aa}$, unlike v_i'. This proves the claim.

Hence, the greedy algorithm continues by merging v_i' with u_i' for every i. By Lemma 2.1, these merges do not change the other overlaps, so the resulting ℓ strings do not overlap, and the greedy algorithm concatenates them. Thus, it outputs a common superstring of length

$$\sum_{i \in [\ell]} (|v_i'| + |u_i'| - 1) = (4k-6)\ell. \qquad \square$$

AI methodology

We follow the recommendations of the working group report [AIT26] and use its template for this section.

Activity	**AI used**
Asking the research question	No
Coming up with the approach	Yes
Proof development	Yes
Writing and exposition	Yes
Checking for bugs	Yes
Other supporting tasks	Yes

Details (elaborating on the points above). We used GPT-5.6 Sol Pro to explore different approaches to prove and refute the greedy conjecture. While we were trying to formulate a self-reduction for SCS, the agent suggested a proof of the inequality in Theorem 3.1; it was not directly related to the discussed question, but the agent had found a simpler route to the statement we tried to prove. The initial proof was complicated, but working together we simplified it significantly. Then we used the agent to do mechanical work to construct the example in Theorem 3.3.

The first counterexample to the Greedy Conjecture and a series of improved lower bounds in Theorem 4.1 were obtained using GPT-5.6 Sol and GPT-6 Astra. After the initial counterexample with the lower bound 9/4, we noticed that many constructions—including it—use cactus graphs that are Eulerian. From this observation, we obtained improvements to 7/3 and then 19/8 with the help of GPT-6 Astra. Afterwards, the progress was stuck for a while, until we created a repository containing all the knowledge from the individual sessions, using which Claude Opus-5.5 finally discovered the construction yielding the lower bound 12/5.

We also used GPT-5.6 Sol Ultra to find all typos in the paper. We have written all the arguments by ourselves, checked for correctness and originality, and assume full responsibility for the manuscript.

References

[AIT26] AI and TCS: The Next Six Months. Working Group Report. Simons Institute for the Theory of Computing, https://simons.berkeley.edu/AITCSreport, 2026.

[AS94] Chris Armen and Clifford Stein. A $2\frac{3}{4}$-Approximation Algorithm for the Shortest Superstring Problem. Technical report, Dartmouth College, Hanover, NH, USA, 1994.

[AS95] Chris Armen and Clifford Stein. Improved length bounds for the shortest superstring problem (extended abstract). In *WADS*, volume 955 of *Lecture Notes in Computer Science*, pages 494–505. Springer, 1995.

[AS96] Chris Armen and Clifford Stein. A $2\frac{2}{3}$-Approximation Algorithm for the Shortest Superstring Problem. In *CPM*, volume 1075 of *Lecture Notes in Computer Science*, pages 87–101. Springer, 1996.

[BJJ97] Dany Breslauer, Tao Jiang, and Zhigen Jiang. Rotations of Periodic Strings and Short Superstrings. *Journal of Algorithms*, 24(2):340–353, 1997.

[BJL+91] Avrim Blum, Tao Jiang, Ming Li, John Tromp, and Mihalis Yannakakis. Linear approximation of shortest superstrings. In *STOC*, pages 328–336. ACM, 1991.

[BJL+94] Avrim Blum, Tao Jiang, Ming Li, John Tromp, and Mihalis Yannakakis. Linear Approximation of Shortest Superstrings. *J. ACM*, 41(4):630–647, 1994.

[CGPR94] Artur Czumaj, Leszek Gasieniec, Marek Piotrów, and Wojciech Rytter. Parallel and sequential approximations of shortest superstrings. In *SWAT*, volume 824 of *Lecture Notes in Computer Science*, pages 95–106. Springer, 1994.

[CKMS26a] Nikolai Chukhin, Alexander S. Kulikov, Ivan Mihajlin, and Alexander Smal. The greedy superstring algorithm achieves ratio 2 for strings of length 6 already. *CoRR*, abs/2608.20018, 2026.

[CKMS26b] Nikolai Chukhin, Alexander S. Kulikov, Ivan Mihajlin, and Alexander Smal. A tight cycle-cover inequality for shortest common superstring. *CoRR*, abs/2609.27921, 2026.

[CR18] Bastien Cazaux and Eric Rivals. Relationship between Superstring and Compression Measures: New Insights on the Greedy Conjecture. *Discrete Applied Mathematics*, 245:59–64, 2018.

[DB46] Nicolaas Govert De Bruijn. A combinatorial problem. In *Nederl. Akad. Wetensch. Proc.*, volume 49, pages 758–764, 1946.

[EMV22] Matthias Englert, Nicolaos Matsakis, and Pavel Veselý. Improved approximation guarantees for shortest superstrings using cycle classification by overlap to length ratios. In *STOC*, pages 317–330. ACM, 2022.

[EMV23] Matthias Englert, Nicolaos Matsakis, and Pavel Veselý. Approximation Guarantees for Shortest Superstrings: Simpler and Better. In *ISAAC*, volume 283 of *LIPIcs*, pages 29:1–29:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.

[GMS80] John Gallant, David Maier, and James A. Storer. On Finding Minimal Length Superstrings. *Journal of Computer and System Sciences*, 20(1):50–58, 1980.

[GP14] Theodoros P. Gevezes and Leonidas S. Pitsoulis. The Shortest Superstring Problem. In *Optimization in Science and Engineering*, pages 189–227. Springer, 2014.

[Gra66] Ronald L. Graham. Bounds for Certain Multiprocessing Anomalies. *Bell System Technical Journal*, 45(9):1563–1581, 1966.

[Joh74] David S. Johnson. Approximation Algorithms for Combinatorial Problems. *Journal of Computer and System Sciences*, 9(3):256–278, 1974.

[KLSS03] Haim Kaplan, Moshe Lewenstein, Nira Shafrir, and Maxim Sviridenko. Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs. In *FOCS*, pages 56–65. IEEE Computer Society, 2003.

[KN24] Pavel E. Kalugin and Maksim S. Nikolaev. The Greedy Algorithm for the Shortest Common Superstring Problem Is a ½-Approximation in Terms of Compression: A Simple Proof. In *SOSA*, pages 97–99. SIAM, 2024.

[KPS94] S Rao Kosaraju, James K Park, and Clifford Stein. Long tours and short superstrings. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 166–177. IEEE, 1994.

[KS05] Haim Kaplan and Nira Shafrir. The Greedy Algorithm for Shortest Superstrings. *Information Processing Letters*, 93(1):13–17, 2005.

[KSS15] Alexander S. Kulikov, Sergey Savinov, and Evgeniy Sluzhaev. Greedy Conjecture for Strings of Length 4. In *CPM*, volume 9133 of *Lecture Notes in Computer Science*, pages 307–315. Springer, 2015.

[Lyn54] Roger C Lyndon. On Burnside's problem. *Transactions of the American Mathematical Society*, 77(2):202–215, 1954.

[Ma09] Bin Ma. Why Greed Works for Shortest Common Superstring Problem. *Theoretical Computer Science*, 410(51):5374–5381, 2009.

[Muc07] Marcin Mucha. A tutorial on shortest superstring approximation. https://www.mimuw.edu.pl/~mucha/teaching/aa2008/ss.pdf, 2007.

[Muc13] Marcin Mucha. Lyndon Words and Short Superstrings. In *SODA*, pages 958–972. SIAM, 2013.

[Nik21] Maksim S. Nikolaev. All Instantiations of the Greedy Algorithm for the Shortest Common Superstring Problem Are Equivalent. In *SPIRE*, pages 61–67, 2021.

[Ope26] OpenAI. A Polynomial-Time 2-Approximation for Shortest Common Superstring. OpenAI Math Release preprint, 2026.

[PEvZ12] Katarzyna E. Paluch, Khaled M. Elbassioni, and Anke van Zuylen. Simpler Approximation of the Maximum Asymmetric Traveling Salesman Problem. In *STACS*, pages 501–506, 2012.

[Rau82] Gérard Rauzy. Suites à termes dans un alphabet fini. *Séminaire de théorie des nombres de Bordeaux*, pages 1–16, 1982.

[RBT04] Heidi J. Romero, Carlos A. Brizuela, and Andrei Tchernykh. An Experimental Comparison of Approximation Algorithms for the Shortest Common Superstring Problem. In *Mexican International Conference on Computer Science*, pages 27–34. IEEE Computer Society, 2004.

[Shi53] Anatolii Illarionovich Shirshov. Subalgebras of free Lie algebras. *Matematicheskii Sbornik*, 75(2):441–452, 1953.

[Shi26] Hiroki Shibata. Disproving the Greedy Superstring Conjecture. *CoRR*, abs/2609.01365, 2026.

[Sto88] James A. Storer. *Data Compression: Methods and Theory*. Computer Science Press, 1988.

[SVB23] Ondřej Sladký, Pavel Veselý, and Karel Břinda. Masked Superstrings as a Unified Framework for Textual k-mer Set Representations. *bioRxiv*, 2023.

[Swe99] Z. Sweedyk. A $2\frac{1}{2}$-Approximation Algorithm for Shortest Superstring. *SIAM Journal on Computing*, 29(3):954–986, 1999.

[TU88] Jorma Tarhio and Esko Ukkonen. A Greedy Approximation Algorithm for Constructing Shortest Common Superstrings. *Theoretical Computer Science*, 57:131–145, 1988.

[Tur89] Jonathan S. Turner. Approximation Algorithms for the Shortest Common Superstring Problem. *Information and Computation*, 83(1):1–20, 1989.

[TY93] Shang-Hua Teng and F. Frances Yao. Approximating shortest superstrings. In *FOCS*, pages 158–165. IEEE Computer Society, 1993.

[Ukk90] Esko Ukkonen. A Linear-Time Algorithm for Finding Approximate Shortest Common Superstrings. *Algorithmica*, 5(3):313–323, 1990.

[Vaz01] Vijay V. Vazirani. *Approximation Algorithms*. Springer, 2001.