# String Attractors: Verification and Optimization 

Dominik Kempa

Department of Computer Science, University of Helsinki, Finland
dkempa@cs.helsinki.fi
(D) https://orcid.org/0000-0003-2286-7417

Alberto Policriti<br>Department of Computer Science, University of Udine, Italy alberto.policriti@uniud.it<br>(D) https://orcid.org/0000-0001-8502-5896<br>Nicola Prezza<br>Department of Computer Science, University of Pisa, Italy nicola.prezza@di.unipi.it<br>(D) https://orcid.org/0000-0003-3553-4953<br>\section*{Eva Rotenberg}<br>DTU Compute, Technical University of Denmark, Denmark erot@dtu.dk

(D) https://orcid.org/0000-0001-5853-7909


#### Abstract

String attractors [STOC 2018] are combinatorial objects recently introduced to unify all known dictionary compression techniques in a single theory. A set $\Gamma \subseteq[1 . . n]$ is a $k$-attractor for a string $S \in \Sigma^{n}$ if and only if every distinct substring of $S$ of length at most $k$ has an occurrence crossing at least one of the positions in $\Gamma$. Finding the smallest $k$-attractor is NP-hard for $k \geq 3$, but polylogarithmic approximations can be found using reductions from dictionary compressors. It is easy to reduce the $k$-attractor problem to a set-cover instance where the string's positions are interpreted as sets of substrings. The main result of this paper is a much more powerful reduction based on the truncated suffix tree. Our new characterization of the problem leads to more efficient algorithms for string attractors: we show how to check the validity and minimality of a $k$-attractor in near-optimal time and how to quickly compute exact solutions. For example, we prove that a minimum 3 -attractor can be found in $\mathcal{O}(n)$ time when $|\Sigma| \in \mathcal{O}(\sqrt[3+\varepsilon]{\log n})$ for some constant $\epsilon>0$, despite the problem being NP-hard for large $\Sigma$.


2012 ACM Subject Classification Theory of computation $\rightarrow$ Data compression
Keywords and phrases Dictionary compression, String attractors, Set cover
Digital Object Identifier 10.4230/LIPIcs.ESA.2018.52
Related Version A full version of the paper is available at [15], https://arxiv.org/abs/1803. 01695.

## 1 Introduction

The goal of dictionary compression is to reduce the size of an input string by exploiting its repetitiveness. In the last decades, several dictionary compression techniques - some more powerful than others - were developed to achieve this goal: Straight-Line programs [17] (context-free grammars generating the string), Macro schemes [23] (a set of substring equations having the string as unique solution), the run-length Burrows-Wheeler transform [4] (a string permutation whose number of equal-letter runs decreases as the string's repetitiveness

© Dominik Kempa, Alberto Policriti, Nicola Prezza, and Eva Rotenberg; licensed under Creative Commons License CC-BY
26th Annual European Symposium on Algorithms (ESA 2018).
Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 52; pp. 52:1-52:13
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany
increases), and the compact directed acyclic word graph $[3,6]$ (the minimization of the suffix tree). Each scheme from this family comes with its own set of algorithms and data structures to perform compressed-computation operations - e.g. random access - on the compressed representation. Despite being apparently unrelated, in [16] all these compression schemes were proven to fall under a common general scheme: they all induce a set $\Gamma \subseteq[1 . . n]$ whose cardinality is bounded by the compressed representation's size and with the property that each distinct substring has an occurrence crossing at least one position in $\Gamma$. A set with this property is called a string attractor. Intuitively, positions in a string attractor capture "interesting" regions of the string; a string of low complexity (that is, more compressible), will generate a smaller attractor. Surprisingly, given such a set one can build a data structure of size $\mathcal{O}(|\Gamma| \operatorname{polylog}(n))$ supporting random access queries in optimal time [16]: string attractors therefore provide a universal framework for performing compressed computation on top of any dictionary compressor (and even optimally for particular queries such as random access).

These premises suggest that an algorithm computing the smallest string attractor for a given string would be a valuable tool for designing better compressed data structures. Unfortunately, computing a minimum string attractor is NP-hard. The problem remains NP-hard even under the restriction that only substrings of length at most $k$ are captured by $\Gamma$, for any $k \geq 3$ and on large alphabets. In this case, we refer to the problem as $k$-attractor. Not all hope is lost, though: as shown in [16], dictionary compressors are actually heuristics for computing a small $n$-attractor (with polylogarithmic approximation rate w.r.t. the minimum), and, more generally, $k$-attractor admits a $\mathcal{O}(\log k)$-approximation based on a reduction to set cover. It is actually easy to find such a reduction: choose as universe the set of distinct substrings and as set collection the string's positions (i.e., set $s_{i}$ contains all substrings crossing position $i$ ). The main limitation of this approach is that the set of distinct substrings could be quadratic in size; this makes the strategy of little use in cases where the goal is to design usable (i.e., as close as possible to linear-time) algorithms on string attractors.

The core result of this paper is a much more powerful reduction from $k$-attractor to set-cover: the universe $\mathcal{U}$ of our instance is equal to the set of edges of the $k$-truncated suffix tree, while the size of the set collection $\mathcal{S} \subseteq 2^{\mathcal{U}}$ is bounded by the size of the ( $2 k-1$ )-truncated suffix tree. First of all, we obtain a universe that is always at least $k$ times smaller than the naive approach. Moreover, the size of our set-cover instance does not depend on the string's length $n$, unless $\sigma$ and $k$ do. This allows us to show that $k$-attractor is actually solvable in polynomial time for small values of $k$ and $\sigma$, and leads us to efficient algorithms for a wide range of different problems on string attractors.

The paper is organized as follows. In Section 1.1 we describe the notation used throughout the paper and we report the main notions related to $k$-attractors. In Section 1.2 we give the main theorem stating our reduction to set-cover (Theorem 5) and briefly discuss the results that we obtain in the rest of the paper by applying it. Theorem 5 itself is proven in Section 2 and is used in Section 3 to provide fast algorithms on string attractors. Finally, in Section 3.4 we introduce and study the complexity of the closely-related sharp-k-attractor problem: to capture all distinct substrings of length exactly $k$. The full version [15] of this paper covers additional material related to the approximation of minimum $k$-attractors.

### 1.1 Notation and definitions

Throughout we consider a string $S[1 . . n]$ of $n$ symbols. We assume the reader to be familiar with the notions of suffix tree [24], suffix array [18], and wavelet tree [10, 21]. By $S T^{k}(S)$ we denote a $k$-truncated suffix tree of $S$, i.e., a compact trie containing all substrings of $S$
of length at most $k$. $\mathcal{E}(\mathcal{T})$ denotes the set of edges of the compact trie $\mathcal{T} . \mathcal{L}(\mathcal{T})$ denotes the set of leaves at maximum string depth of the compact trie $\mathcal{T}$ (i.e., leaves whose string depth is equal to the maximum string depth among all leaves). Let $e=\langle u, v\rangle$ be an edge in the (truncated) suffix tree of $S$. With $s(e)$ we denote the string read from the suffix tree root to (and including) the first character in the label of $e . \lambda(e)=|s(e)|$ is the length of this string. We will also refer to $\lambda(e)$ as the string depth of $e$. Note that edges $e_{1}, \ldots, e_{t}$ of the $k$-truncated suffix tree have precisely the same labels $s\left(e_{1}\right), \ldots, s\left(e_{t}\right)$ of the suffix tree edges $e_{1}^{\prime}, \ldots, e_{t}^{\prime}$ at string depth $\lambda\left(e_{i}^{\prime}\right) \leq k$. It follows that we can use these two edge sets interchangeably when we are only interested in their labels (this will be the case in our results). Let $S A[1 . . n]$ denote the suffix array of $S$. $\left\langle l_{e}, r_{e}\right\rangle$, with $e \in \mathcal{E}\left(S T^{k}(S)\right)$ being an edge in the $k$-truncated suffix tree, will denote the suffix array range corresponding to the string $s(e)$, i.e., $S A\left[l_{e} . . r_{e}\right]$ contains all suffixes prefixed by $s(e)$.

Unless otherwise specified, we give the space of our data structures in words of $\Theta(\log n)$ bits each.

With the following definition we recall the notion of $k$-attractor of a string [16].

- Definition 1. A $k$-attractor of a string $S \in \Sigma^{n}$ is a set of positions $\Gamma \subseteq[1 . . n]$ such that every substring $S[i . . j]$ with $i \leq j<i+k$ has at least one occurrence $S\left[i^{\prime} . . j^{\prime}\right]=S[i . . j]$ with $j^{\prime \prime} \in\left[i^{\prime} . . j^{\prime}\right]$ for some $j^{\prime \prime} \in \Gamma$.

When $k=n$, we simply call $\Gamma$ an attractor of $S$.

- Definition 2. A minimal $k$-attractor of a string $S \in \Sigma^{n}$ is a $k$-attractor $\Gamma$ such that $\Gamma-\{j\}$ is not a $k$-attractor of $S$ for any $j \in \Gamma$.
- Definition 3. A minimum $k$-attractor of a string $S \in \Sigma^{n}$ is a $k$-attractor $\Gamma^{*}$ such that, for any $k$-attractor $\Gamma$ of $S,\left|\Gamma^{*}\right| \leq|\Gamma|$.
- Theorem 4. [16, Thm. 4.2] The problem of deciding whether a string $S$ admits a $k$-attractor of size at most $t$ is NP-complete for $k \geq 3$.


### 1.2 Overview of the contributions

Our main theorem is a reduction to set-cover based on the notion of truncated suffix tree:

- Theorem 5. Let $S^{\prime}=\#^{k-1} S \#^{k-1}$, with $\# \notin \Sigma$. An instance of $k$-attractor can be reduced to a set-cover instance with universe equal to the set of edges of the $k$-truncated suffix tree of $S$ and set collection with one element for each leaf of the $(2 k-1)$-truncated suffix tree of $S^{\prime \prime}$.

Figure 1 depicts the main technique (Lemma 11) standing at the core of our reduction: a set $\Gamma$ is a valid attractor if and only if it marks (or colors, in the picture), all suffix tree edges.

Using the reduction of Theorem 5, we obtain the following results. First, we present efficient algorithms to check the validity and minimality of a $k$-attractor. Note that it is trivial to perform these checks in cubic time (or quadratic, with little more care). In Theorem 18 we show that we can check whether a set $\Gamma \subseteq[1 . . n]$ is a valid $k$-attractor for $S$ in $\mathcal{O}(n)$ time and $\mathcal{O}(n)$ words of space. Using recent advances in compact data structures, we show how to further reduce this working space to $\mathcal{O}(n \log \sigma)$ bits without affecting query times when $k \leq \sigma^{\mathcal{O}(1)}$, or with a small time penalty in the general case. In particular, when $k$ is polynomial in the alphabet size, we can always check the correctness of a $k$-attractor in $\mathcal{O}(n)$ time and $\mathcal{O}(n \log \sigma)$ bits of space. With similar techniques, in Theorem 22 we show how to verify that $\Gamma$ is a minimal $k$-attractor for $S$ in near-optimal $\mathcal{O}(n \log n)$ time. To conclude, in Theorem 24 we show that a minimum $k$-attractor can be found in $\mathcal{O}(n)+\exp \left(\mathcal{O}\left(\sigma^{k} \log \sigma^{k}\right)\right)$ time. In particular, this result yields the following corollaries:

(a) The suffix tree of the string BBBABA, colored corresponding to attractor positions 2,5,6. This attractor is minimal: removing any position leaves some edge uncolored.

(b) The suffix tree of the string BBBABA, colored corresponding to attractor positions 3,4 . This attractor is minimum: it is minimal and of minimum size.

Figure 1 A position $i$ "marks" (or, here, colors) a suffix tree edge $e$ if and only if it crosses an occurrence of the string read from the root to the first letter in the label of $e$. A set of positions forms a $k$-attractor if and only if they color all edges of the $k$-truncated suffix tree (in this figure, $k=6$ and we color the whole suffix tree). Dashed lines indicate that the edge has multiple colors. The string terminator $\$$ (and edges labeled with $\$$ ) is omitted for simplicity.

- Corollary 6. $k$-attractor is in $P$ when $\sigma^{k} \log \sigma^{k} \in \mathcal{O}(\log n)$.
- Corollary 7. For constant $k$, a minimum $k$-attractor can be found in $\mathcal{O}(n)$ time if there exists $\epsilon>0$ such that $\sigma^{k+\epsilon} \in \mathcal{O}(\log n)$.

Proof. Let $\sigma^{k+\epsilon}=\sigma^{k\left(1+\epsilon^{\prime}\right)}$, where $\epsilon^{\prime}=\epsilon / k>0$ is a constant. For any constant $\epsilon^{\prime}>0$ we have that $\sigma^{k} \log \sigma^{k} \in o\left(\sigma^{k} \cdot\left(\sigma^{k}\right)^{\epsilon^{\prime}}\right)=o\left(\sigma^{k\left(1+\epsilon^{\prime}\right)}\right)$. It follows that $\sigma^{k} \log \sigma^{k} \in o(\log n)$, i.e., by Theorem 24 we can find a minimum $k$-attractor in linear time.

With the above result we can, for example, find a minimum 3-attractor in $\mathcal{O}(n)$ time when $\sigma \in \mathcal{O}(\sqrt[3+\varepsilon]{\log n})$, for some $\epsilon>0$ (keep in mind that 3 -attractor is NP-complete for large alphabets).

## 2 A better reduction to set-cover

In this section we give our main result: a reduction from $k$-attractor to set-cover using a universe smaller than the one used in [16]. We start with an alternative characterization of $k$-attractors based on the $k$-truncated suffix tree.

- Definition 8 (Marker). $j \in \Gamma$ is a marker for a suffix tree edge $e$ if and only if

$$
\exists i \in S A\left[l_{e} . . r_{e}\right]: i \leq j<i+\lambda(e)
$$

Equivalently, we say that $j$ marks $e$ (see Figure 1).

- Definition 9 (Edge marking). $\Gamma \subseteq[1 . . n]$ marks a suffix tree edge $e$ if and only if there exists a $j \in \Gamma$ that marks $e$.
- Definition 10 (Suffix tree $k$-marking). $\Gamma \subseteq[1 . . n]$ is a suffix tree $k$-marking if and only if it marks every edge $e$ such that $\lambda(e) \leq k$ (equivalently, every $e \in \mathcal{E}\left(S T^{k}(S)\right)$ ).

When $k=n$ we simply say suffix tree marking (since all edges satisfy $\lambda(e) \leq n$ ). We now show that the notions of $k$-attractor and suffix tree $k$-marking are equivalent.

- Lemma 11. $\Gamma$ is a $k$-attractor if and only if it is a suffix tree $k$-marking.

Proof. $(\Rightarrow)$ Let $\Gamma$ be a $k$-attractor. Pick any suffix tree edge $e$ such that $\lambda(e) \leq k$. Then, $\lambda(e)=|s(e)| \leq k$ and, by definition of $k$-attractor, there exists a $j \in \Gamma$ and an $i$ such that $s(e)=S[i . . i+|s(e)|-1]$ and $i \leq j \leq i+|s(e)|-1$. We also have that $i \in S A\left[l_{e} . . r_{e}\right]$ (being $\left\langle l_{e}, r_{e}\right\rangle$ precisely the suffix array range of suffixes prefixed by $\left.s(e)\right)$. Putting these results together, we found an $i \in S A\left[l_{e} . . r_{e}\right]$ such that $i \leq j \leq i+\lambda(e)-1$ for some $j \in \Gamma$, which by Definition 9 means that $\Gamma$ marks $e$. Since the argument works for any edge $e$ at string depth at most $k$, we obtain that $\Gamma$ is a suffix tree $k$-marking.
$(\Leftarrow)$ Let $\Gamma$ be a suffix tree $k$-marking. Let, moreover, $s$ be a substring of $S$ of length at most $k$. Consider the lowest suffix tree edge $e$ (i.e., the $e$ with maximum $\lambda(e)$ ) such that $s(e)$ prefixes $s$. In particular, $\lambda(e) \leq k$. Note that, by definition of suffix tree, every occurrence $S[i . . i+|s(e)|-1]=s(e)$ of $s(e)$ in $S$ prefixes an occurrence of $s: S[i . . i+|s|-1]=s$. By definition of $k$-marking, there exists a $j \in \Gamma$ such that $j$ is a marker for $e$, which means (by Definition 8) that $\exists i \in S A\left[l_{e} . . r_{e}\right]: i \leq j<i+\lambda(e)$. Since $i \in S A\left[l_{e} . . r_{e}\right], S A[i]$ is an occurrence of $s(e)$, and therefore of $s$. But then, we have that $i \leq j<i+\lambda(e)=i+|s(e)| \leq$ $i+|s|$, i.e., $S[S A[i] . . S A[i]+|s|-1]$ is an occurrence of $s$ crossing $j \in \Gamma$. Since the argument works for every substring $s$ of $S$ of length at most $k$, we obtain that $\Gamma$ is a $k$-attractor.

An equivalent formulation of Lemma 11 is that $\Gamma$ is a $k$-attractor if and only if it marks all edges of the $k$-truncated suffix tree. In particular (case $k=n$ ), $\Gamma$ is an attractor if and only if it is a suffix tree marking.

Lemma 11 will be used to obtain a smaller universe $\mathcal{U}$ in our set-cover reduction. With the following lemmas we show that also the size of the set collection $\mathcal{S}$ can be considerably reduced when $k$ and $\sigma$ are small.

- Definition 12 ( $k$-equivalence). Two positions $i, j \in[1 . . n]$ are $k$-equivalent, indicated as $i \equiv_{k} j$, if and only if

$$
S^{\prime}[i-k+1 . . i+k-1]=S^{\prime}[j-k+1 . . j+k-1]
$$

where $S^{\prime}[i]=\#$ if $i<1$ or $i>n$ (note that we allow negative positions) and $S^{\prime}[i]=S[i]$ otherwise, and $\# \notin \Sigma$ is a new character.

It is easy to see that $k$-equivalence is an equivalence relation. First, we bound the size of the distinct equivalence classes of $\equiv_{k}$ (i.e., the size of the quotient set $[1 . . n] / \equiv_{k}$ ).

- Lemma 13. $\left|[1 . . n] / \equiv_{k}\right|=\left|\mathcal{L}\left(S T^{2 k-1}\left(S^{\prime}\right)\right)\right| \leq \min \left\{n, \sigma^{2 k-1}+2 k-2\right\}$

Proof. By definition of $\equiv_{k}$, the set [1..n]/ $\equiv_{k}$ has one element per distinct substring of length $(2 k-1)$ in $S^{\prime}$, that is, per distinct path from the suffix tree root to each of the nodes in $\mathcal{L}\left(S T^{2 k-1}\left(S^{\prime}\right)\right)$. Clearly, $\left|\mathcal{L}\left(S T^{2 k-1}\left(S^{\prime}\right)\right)\right| \leq n$. On the other hand, there are at most $\sigma^{2 k-1}$ distinct substrings of length $2 k-1$ on $\Sigma$. Moreover, there are $2 k-2$ additional substrings to consider on the borders of $S^{\prime}$ (to include the runs of symbol \#). It follows that the cardinality of $\mathcal{L}\left(S T^{2 k-1}\left(S^{\prime}\right)\right)$ is upper-bounded also by $\sigma^{2 k-1}+2 k-2$.

We now show that any minimal $k$-attractor can have at most one element from each equivalence class of $\equiv_{k}$.

- Lemma 14. If $\Gamma$ is a minimal $k$-attractor, then for any $1 \leq i \leq n$ it holds $\left|\Gamma \cap[i]_{\equiv_{k}}\right| \leq 1$.

Proof. Suppose, by contradiction, that $\left|\Gamma \cap[i]_{\equiv_{k}}\right|>1$ for some $i$. Then, let $j, j^{\prime} \in \Gamma \cap[i]_{\equiv_{k}}$, with $j \neq j^{\prime}$. By definition of $\equiv{ }_{k}, S^{\prime}[j-k+1 . . j+k-1]=S^{\prime}\left[j^{\prime}-k+1 . . j^{\prime}+k-1\right]$. This means that if a substring of $S$ of length at most $k$ has an occurrence crossing position $j$
in $\Gamma$ then it has also one occurrence crossing position $j^{\prime} \in(\Gamma-\{j\})$. On the other hand, any other substring occurrence crossing any position $j^{\prime \prime} \neq j, j^{\prime}$ is also captured by $\Gamma-\{j\}$ since $j^{\prime \prime}$ belongs to this set. This implies that $\Gamma-\{j\}$ is a $k$-attractor, which contradicts the minimality of $\Gamma$.

Moreover, if we swap any element of a $k$-attractor with an equivalent element then the resulting set is still a $k$-attractor:

- Lemma 15. Let $\Gamma$ be a $k$-attractor. Then, $(\Gamma-\{j\}) \cup\left\{j^{\prime}\right\}$ is a $k$-attractor for any $j \in \Gamma$ and any $j^{\prime} \equiv_{k} j$.

Proof. Pick any occurrence of a substring $s,|s| \leq k$, crossing position $j$. By definition of $\equiv_{k}$, since $j^{\prime} \equiv_{k} j$ there is also an occurrence of $s$ crossing $j^{\prime}$. This implies that $\Gamma^{\prime}=(\Gamma-\{j\}) \cup\left\{j^{\prime}\right\}$ is a $k$-attractor.

Lemmas 14 and 15 imply that we can reduce the set of candidate positions from [1..n] to $\mathcal{C}=\left\{\min (I) \mid I \in[1 . . n] / \equiv_{k}\right\}$ (that is, an arbitrary representative - in this case, the minimum - from any class of $\equiv_{k}$ ), and still be able to find a minimal/minimum $k$-attractor. Note that, by Lemma $13,\left|[1 . . n] / \equiv_{k}\right| \leq \min \left\{n, \sigma^{2 k-1}+2 k-2\right\}$.

We can now prove our main theorem.
Proof of Theorem 5. We build our set-cover instance $\langle\mathcal{U}, \mathcal{S}\rangle$ as follows. We choose $\mathcal{U}=$ $\mathcal{E}\left(S T^{k}(S)\right)$, i.e., the set of edges of the $k$-truncated suffix tree. The set collection $\mathcal{S}$ is defined as follows: let $s_{i}=\left\{e \in \mathcal{E}\left(S T^{k}(S)\right) \mid i\right.$ marks $\left.e\right\}, \mathcal{C}=\left\{\min (I) \mid I \in[1 . . n] / \equiv_{k}\right\}$ and put:

$$
\mathcal{S}=\left\{s_{i} \mid i \in \mathcal{C}\right\}
$$

By definition of $\equiv_{k}$, each $I \in[1 . . n] / \equiv_{k}$ is unambiguously identified by a substring of length $2 k-1$ of the string $S^{\prime}=\#^{k-1} S \#^{k-1}$. We therefore obtain $|\mathcal{S}|=\left|\mathcal{L}\left(S T^{2 k-1}\left(S^{\prime}\right)\right)\right|$. We now prove the correctness and completeness of the reduction.

Correctness. By the definition of our reduction, a solution $\left\{s_{i_{1}}, \ldots, s_{i_{\gamma}}\right\}$ to $\langle\mathcal{U}, \mathcal{S}\rangle$ yields a set $\Gamma=\left\{i_{1}, \ldots, i_{\gamma}\right\}$ of positions marking every edge in $\mathcal{E}\left(S T^{k}(S)\right)$. Then, Lemma 11 implies that $\Gamma$ is a $k$-attractor.

Completeness. Let $\Gamma=\left\{i_{1}, \ldots, i_{\gamma}\right\}$ be a minimal $k$-attractor. Then, Lemmas 14 and 15 imply that the following set is also a minimal $k$-attractor of the same size: $\Gamma^{\prime}=\left\{j_{1}=\right.$ $\left.\min \left(\left[i_{1}\right]_{\equiv_{k}}\right), \ldots, j_{\gamma}=\min \left(\left[i_{\gamma}\right]_{\equiv_{k}}\right)\right\}$. Note that $\Gamma^{\prime} \subseteq\left\{\min (I) \mid I \in[1 . . n] / \equiv_{k}\right\}$. By Lemma 11, $\Gamma^{\prime}$ marks every edge in $\mathcal{E}\left(S T^{k}(S)\right)$. Then, by definition of our reduction the collection $\left\{s_{j_{1}}, \ldots, s_{j_{\gamma}}\right\}$ covers $\mathcal{U}=\mathcal{E}\left(S T^{k}(S)\right)$.

In the rest of the paper, we use the notation $\mathcal{U}=\mathcal{E}\left(S T^{k}(S)\right)$ and $\mathcal{C}=\{\min (I) \mid I \in$ $\left.[1 . . n] / \equiv_{k}\right\}$ to denote the universe to be covered (edges of the $k$-truncated suffix tree) and the candidate attractor positions, respectively. Recall, moreover, that $|\mathcal{U}| \leq \min \left\{n, \sigma^{k}\right\}$ and $|\mathcal{C}| \leq \min \left\{n, \sigma^{2 k-1}+2 k-2\right\}$.

## 3 Faster algorithms

In this section we use properties of our reduction to provide faster algorithms for a range of problems: (i) checking that a given set $\Gamma \subseteq[1 . . n]$ is a $k$-attractor, (ii) checking that a given set is a minimal $k$-attractor, and (iii) finding a minimum $k$-attractor. We note that problems (i)-(ii) admit naive cubic solutions, while problem (iii) is NP-hard for $k \geq 3$ [16].

We assume that the input string $S$ is over the integers alphabet [1.. $\sigma$ ] where $\sigma \in n^{\mathcal{O}(1)}$. Note that suffix-sorting can be performed in linear time and space on such alphabets [14].

### 3.1 Checking the attractor property

Given a string $S$, a set $\Gamma \subseteq[1 . . n]$, and an integer $k \geq 1$, is $\Gamma$ a $k$-attractor for $S$ ? We show that this question can be answered in $\mathcal{O}(n)$ time.

The main idea is to use Lemma 11 and check, for every suffix tree edge $e$ at string depth at most $k$, if $\Gamma$ marks $e$. Consider the suffix array $S A[1 . . n]$ of $S$ and the array $D[1 . . n]$ defined as follows: $D[i]=\operatorname{succ}(\Gamma, S A[i])-S A[i]$, where $\operatorname{succ}(X, x)$ returns the smallest element larger than or equal to $x$ in the set $X$ (i.e., $D[i]$ is the distance between $S A[i]$ and the element of $\Gamma$ following - and possibly equal to $-S A[i])$. $D$ can be built in linear time and space by creating a bit-vector $B[1, n]$ such that $B[i]=1 \mathrm{iff} i \in \Gamma$ and pre-processing $B$ for constant-time successor queries [13,5]. We build a range-minimum data structure (RMQ) on $D(\mathcal{O}(n)$ bits of space, constant query time [9]). Then for every suffix tree edge $e$ such that $\lambda(e) \leq k$, we check (in constant time) that $\lambda(e)>\min \left(D\left[l_{e} . . r_{e}\right]\right)$. The following lemma ensures that this is equivalent to checking whether $\Gamma$ marks $e$.

- Lemma 16. $\lambda(e)>\min \left(D\left[l_{e} . . r_{e}\right]\right)$ if and only if $\Gamma$ marks $e$.

Proof. $(\Rightarrow)$ Assume that $\lambda(e)>\min \left(D\left[l_{e} . . r_{e}\right]\right)$. By definition of $D$, this means that there exist an index $i^{\prime} \in\left[l_{e} . . r_{e}\right]$ and a $j \in \Gamma$, with $j \geq i=S A\left[i^{\prime}\right]$, such that $j-i=D\left[i^{\prime}\right]<\lambda(e)$. Equivalently, $i \leq j<i+\lambda(e)$, i.e., $\Gamma$ marks $e$.
$(\Leftarrow)$ Assume that $\Gamma$ marks $e$. Then, by definition, there exist an index $i^{\prime} \in\left[l_{e} . . r_{e}\right]$ and a $j \in \Gamma$ such that $S A\left[i^{\prime}\right]=i \leq j<i+\lambda(e)$. Then, $j-i<\lambda(e)$. Since $D\left[i^{\prime}\right]$ is computed taking the $j \in \Gamma, j \geq S A\left[i^{\prime}\right]$, minimizing $j-S A\left[i^{\prime}\right]$, it must be the case that $D\left[i^{\prime}\right] \leq j-i<\lambda(e)$. Since $i^{\prime} \in\left[l_{e} . . r_{e}\right]$, this implies that $\min \left(D\left[l_{e} . . r_{e}\right]\right)<\lambda(e)$.

Together, Lemmas 11 and 16 imply that, if $\lambda(e)>\min \left(D\left[l_{e} . . r_{e}\right]\right)$ for every edge at string depth at most $k$, then $\Gamma$ is a $k$-attractor for $S$. Since the suffix tree, as well as the other structures used by our algorithm, can be built in linear time and space on alphabet [1..n] [7] and checking each edge takes constant time, we obtain that the problem of checking whether a set $\Gamma \subseteq[1 . . n]$ is a valid $k$-attractor can be solved in $\mathcal{O}(n)$ time and $\mathcal{O}(n)$ words of space. We now show how to improve upon this working space by using recent results in the field of compact data structures. In the following result, we assume that the input string is packed in $\mathcal{O}(n \log \sigma)$ bits (that is, $\mathcal{O}\left(n / \log _{\sigma} n\right)$ words).

We first need the following Lemma from [2]:

- Lemma 17. [2, Thm. 3] In $\mathcal{O}(n)$ time and $\mathcal{O}(n \log \sigma)$ bits of space we can enumerate the following information for each suffix tree edge e:
- The suffix array range $\left\langle l_{e}, r_{e}\right\rangle$ of the string $s(e)$, and
- the length $\lambda(e)$ of $s(e)$.

Proof. In [2, Thm. 3] (see also [1]) the authors show how to enumerate the following information for each right-maximal substring $W$ of $S$ in $\mathcal{O}(n)$ time and $\mathcal{O}(n \log \sigma)$ bits of space: $|W|$ and the suffix array range $\operatorname{range}(W b)$ of the string $W b$, for all $b \in \Sigma$ such that $W b$ is a substring of $S$. Since $W$ is right-maximal, those $W b$ are equal to our strings $s(e)$ (for every edge $e$ ). It follows that our problem is solved by outputting all range $(W b)$ and $|W b|$ returned by the algorithm in [2, Thm. 3].

We can now prove our theorem. Note that the input set $\Gamma \subseteq[1 . . n]$ can be encoded in $n$ bits, so also the input fits in $\mathcal{O}(n \log \sigma)$ bits.

- Theorem 18. Given a string $S \in[1 . . \sigma]^{n}$, a set $\Gamma \subseteq[1 . . n]$, and an integer $k \geq 1$, we can check whether $\Gamma$ is a $k$-attractor for $S$ in:
- Optimal $\mathcal{O}(n \log \sigma)$ bits of space and $\mathcal{O}\left(n \log ^{\epsilon} n\right)$ time, for any constant $\epsilon>0$, or
- $\mathcal{O}(n(\log \sigma+\log k))$ bits of space and $\mathcal{O}(n)$ time.

Proof. To achieve the first trade-off we will replace the $D$ array (occupying $\mathcal{O}(n \log n)$ bits) with a smaller data structure supporting random access to $D$. We start by replacing the standard suffix array with a compressed suffix array (CSA) [8, 11]. Given a text stored in $\mathcal{O}(n \log \sigma)$ bits, the CSA can be built in deterministic $\mathcal{O}(n)$ time and optimal $\mathcal{O}(n \log \sigma)$ bits of space [20], and supports access queries to the suffix array $S A$ in $\mathcal{O}\left(\log ^{\epsilon} n\right)$ time [11], for any constant $\epsilon>0$ chosen at construction time. Given that $D[i]=\operatorname{succ}(\Gamma, S A[i])-S A[i]$ and we can compute the successor function in constant time using a $\mathcal{O}(n)$-bit data structure (array $B), D[i]$ can be computed in $\mathcal{O}\left(\log ^{\epsilon} n\right)$ time. Using access to $D$, the RMQ data structure (occupying $\mathcal{O}(n)$ bits) can be built in $\mathcal{O}\left(n \log ^{\epsilon} n\right)$ time and $\mathcal{O}(n)$ bits of space [9, Thm. 5.8]. At this point, observe that the order in which we visit suffix tree edges does not affect the correctness of our algorithm. By using Lemma 17 we can enumerate $\lambda(e)$ and $\left\langle l_{e}, r_{e}\right\rangle$ for every suffix tree edge $e$ in linear time and compact space, and check $\lambda(e)>\min \left(D\left[l_{e} . . r_{e}\right]\right)$, whenever $\lambda(e) \leq k$ (Lemma 16).

To achieve the second trade-off we observe that in our algorithm we only explore the suffix tree up to depth $k$ (i.e., we only perform the check of Lemma 16 when $\lambda(e) \leq k$ ), hence any $D[i]>k$ can be replaced with $D[i]=k+1$ without affecting the correctness of the verification procedure. In this way, array $D$ can be stored in just $\mathcal{O}(n \log k)$ bits. To compute the $D$ array in $\mathcal{O}(n)$ time and compact space we observe that it suffices to access all pairs $\langle i, S A[i]\rangle$ in any order (not necessarily $\langle 1, S A[1]\rangle,\langle 2, S A[2]\rangle, \ldots$ ). From [2, Thm. 10], in $\mathcal{O}(n)$ time and $\mathcal{O}(n \log \sigma)$ bits of space we can build a compressed suffix array supporting constant-time LF function computation. By repeatedly applying LF from the first suffix array position, we enumerate entries of the inverse suffix array $I S A$ in right-to-left order in $\mathcal{O}(n)$ time [2, Lem. 1]. This yields the sequence of pairs $\langle I S A[i], i\rangle=\langle j, S A[j]\rangle$, for $i=n, \ldots, 1$ and $j=I S A[i]$, which can be used to compute $D$ in linear time and compact space. As in the first trade-off, we use Lemma 17 to enumerate $\lambda(e)$ and $\left\langle l_{e}, r_{e}\right\rangle$ for every suffix tree edge $e$, and check $\lambda(e)>\min \left(D\left[l_{e} . . r_{e}\right]\right)$, whenever $\lambda(e) \leq k$ (Lemma 16).

Note that with the second trade-off of Theorem 18 we achieve $\mathcal{O}(n)$ time and optimal $\mathcal{O}(n \log \sigma)$ bits of space when $k \leq \sigma^{\mathcal{O}(1)}$ (in particular, this is always the case when $k$ is constant). Note also that, since we now assume that the input string is packed in $\mathcal{O}\left(n / \log _{\sigma} n\right)$ words, the running time is not optimal (as $\Omega\left(n / \log _{\sigma} n\right)$ is a lower-bound in this model).

### 3.2 Checking minimality

Given a string $S$, a set $\Gamma \subseteq[1 . . n]$, and an integer $k \geq 1$, is $\Gamma$ a minimal $k$-attractor for $S$ ? The main result of this section is that this question can be answered in near-optimal $\mathcal{O}(n \log n)$ time.

We first show that minimal $k$-attractors admit a convenient characterization based on the concept of suffix tree $k$-marking.

- Definition 19 ( $k$-necessary position). $j \in \Gamma$ is $k$-necessary with respect to $\Gamma$ if and only if there is at least one suffix tree edge $e$ such that:

1. $\lambda(e) \leq k$,
2. $j$ marks $e$, and
3. If $j^{\prime} \in \Gamma$ marks $e$, then $j^{\prime}=j$

- Definition 20 ( $k$-necessary set). $\Gamma$ is $k$-necessary if and only if all its elements are $k$ necessary with respect to $\Gamma$.

When $\Gamma$ is clear from the context, we just say $k$-necessary (referring to some $j \in \Gamma$ ) instead of $k$-necessary with respect to $\Gamma$.

- Lemma 21. $\Gamma$ is a minimal $k$-attractor if and only if:

1. It is a $k$-attractor, and
2. it is $k$-necessary.

Proof. $(\Rightarrow)$ Let $\Gamma$ be a minimal $k$-attractor. Let $j \in \Gamma$. Since $\Gamma$ is minimal, $\Gamma-\{j\}$ is not a $k$-attractor. From Theorem 11, this implies that $\Gamma-\{j\}$ is not a $k$-marking, i.e., there exists a suffix tree edge $e$, with $\lambda(e) \leq k$, that is not marked by $\Gamma-\{j\}$. On the other hand, the fact that $\Gamma$ is a $k$-attractor implies (Theorem 11) that $\Gamma$ is a $k$-marking, i.e., it also marks edge $e$. This, in particular, implies that $j$ marks $e$. Now, let $j^{\prime} \in \Gamma$ be a position that marks $e$. Assume, by contradiction, that $j^{\prime} \neq j$. Then, $j^{\prime} \in \Gamma-\{j\}$, which implies that $\Gamma-\{j\}$ marks $e$. This is a contradiction, therefore it must be the case that $j^{\prime}=j$, i.e., $j$ is $k$-necessary. Since the argument works for any $j \in \Gamma$, we obtain that all $j \in \Gamma$ are $k$-necessary.
$(\Leftarrow)$ Assume that $\Gamma$ is a $k$-attractor and all $j \in \Gamma$ are $k$-necessary. Then, choose an arbitrary $j \in \Gamma$. By Definition 19, there exists an edge $e$ that is only marked by $j$, i.e., for every $j^{\prime} \in \Gamma-\{j\}, j^{\prime}$ does not mark $e$. This implies (Theorem 11) that $\Gamma-\{j\}$ is not a $k$-attractor. Since the argument works for any $j \in \Gamma$, we obtain that $\Gamma$ is a minimal $k$-attractor.

A naive solution for the minimality-checking problem is to test the $k$-attractor property on $\Gamma-\{i\}$ for every $i \in \Gamma$ using Theorem 18. This solution, however, runs in quadratic time. Our efficient strategy relies on colored range reporting and consists in checking, for every suffix tree edge $e$, if there is only one $j \in \Gamma$ marking it. In this case, we flag $j$ as necessary. If, in the end, all attractor positions are flagged as necessary, then the attractor is minimal by Lemma 21. Notice that it can turn out that a single suffix tree edge can be marked by more than one necessary $j \in \Gamma$.

- Theorem 22. Given a string $S \in[1 . . \sigma]^{n}$, a set $\Gamma \subseteq[1 . . n]$, and an integer $k \geq 1$, we can check whether $\Gamma$ is a minimal $k$-attractor for $S$ in $\mathcal{O}(n \log n)$ time and $\mathcal{O}(n \log |\Gamma|)$ space.

Proof. We associate to each element in $\Gamma$ a distinct color from the set $\Sigma_{\Gamma}=\left\{c_{i} \mid i \in \Gamma\right\}$, and we build a two-dimensional $\Sigma_{\Gamma}$-colored grid $L \subseteq[1 . . n]^{2} \times \Sigma_{\Gamma}$ (i.e., each point $\langle i, j\rangle$ in $L$ is associated with a color from $\left.\Sigma_{\Gamma}\right)$ defined as $L=\left\{\left\langle i, D[i], c_{S A[i]+D[i]}\right\rangle, i=1, \ldots, n\right\}$, that is, at coordinates $\langle i, D[i]\rangle$ we insert a point "colored" with the color associated to the attractor position immediately following - and possibly equal to $-S A[i]$. Then, for every suffix tree edge $e$ we check that $L \cap\left[l_{e} . . r_{e}\right] \times[0 . . \lambda(e)-1]$ contains at least two distinct colors. If there are at least two distinct colors $c_{k} \neq c_{k^{\prime}}$ in the range $\left[l_{e} . . r_{e}\right] \times[0 . . \lambda(e)-1]$, then we do not mark $k$ and $k^{\prime}$ as necessary (note that they could be marked later by some other edge, though). However, even if there is only one color $c_{k}$ in the range this may not be enough to mark $k$ as necessary. The reason for this is that in array $D$ we are tracking only the attractor position $i^{\prime}$ immediately following each text position $i$; it could well be that the attractor position $i^{\prime \prime}>i^{\prime}$ immediately following $i^{\prime}$ marks $e$, but we miss it because we track only $i^{\prime}$. This problem can be easily solved inserting in $L$ also a point corresponding to the second nearest attractor position following every text position (so the number of points only doubles). It is easy to see that this is sufficient to solve our problem, since we only aim at enumerating at most two distinct colors in a range.

At this point, we have reduced the problem to the so-called three-sided colored orthogonal range reporting problem in two dimensions: report the distinct colors inside a three-sided orthogonal range in a grid. For this problem, the fastest known data structure takes $\mathcal{O}(n \log n)$ space and answers queries in $\mathcal{O}\left(\log ^{2} n+i\right)$ time, where $n$ is the number of points in the grid and $i$ is the number of returned points [12]. This would result in an overall running time of $\mathcal{O}\left(n \log ^{2} n\right)$ for our algorithm. We note that our problem is, however, simpler than the general one. In our case, it is enough to list two distinct colors (if any); we are not interested in counting the total number of such colors or reporting an arbitrary number of them.

Our solution relies on wavelet trees [10]. First, we pre-process the set of $v \leq 2 n$ points so that they fit in a grid $[1 . . v] \times[1 . . v]$ such that every row and every column contain exactly one point. Mapping the original query on this grid can be easily done in constant time using well-established rank reduction techniques that we do not discuss here (see, e.g. [22]). We can view this grid as an integer vector $V \in[1 . . v]^{v}$, where each $V[i]$ is associated with a color $V[i] . c \in \Sigma_{\Gamma}$. We build a wavelet tree $W T(V)$ on $V$. Let $u_{s}$ denote the internal node of $W T(V)$ reached following the path $s \in\{0,1\}^{*}$. With $V_{s}$ we denote the subsequence of $V$ associated with $u_{s}$, i.e., the subsequence of $V$ such that the binary representation of each $V_{s}[i]$ is prefixed by $s$. For each internal node $u_{s}$ we store (i) the sequence of colors $C_{s}=V_{s}[1] . c, \ldots, V_{s}\left[\left|V_{s}\right|\right] . c$, and (ii) a bitvector $B_{s}\left[1 . .\left|V_{s}\right|\right]$ such that $B_{s}[1]=0$ and $B_{s}[i] \neq B_{s}[i-1]$ iff $C_{s}[i] \neq C_{s}[i-1]$. We pre-process each $B_{s}$ for constant-time rank and select queries [13, 5]. Overall, our data structure takes $\mathcal{O}(n \log \Gamma)$ words of space (that is, $\mathcal{O}(n \log \Gamma \log n)$ bits: at each of the $\log n$ levels of the wavelet tree we store $v \leq 2 n$ colors).

To report two distinct colors in the range $[l, r] \times\left[l^{\prime}, r^{\prime}\right]$, we find in $\mathcal{O}(\log v)$ time the $\mathcal{O}(\log v)$ nodes of $W T(V)$ covering $\left[l^{\prime}, r^{\prime}\right]$ as usually done when solving orthogonal range queries on wavelet trees (see [21] for full details). For each such node $u_{s}$, let the range $V_{s}\left[l_{s}, r_{s}\right]$ contain the elements in common between $V_{s}$ and $V[l . . r]$ (i.e., the range obtained mapping $[l . . r]$ on $V_{s}$ ). We check whether in $B_{s}\left[l_{s} . . r_{s}\right]$ there are two distinct bits at adjacent positions. If this is the case, we locate their positions $B_{s}\left[i_{0}\right]=0$ and $B_{s}\left[i_{1}\right]=1$, with $l_{s} \leq i_{0}=i_{1}-1 \leq r_{s}$ (the case $i_{1}=i_{0}-1$ is symmetric), and return the colors $C_{s}\left[i_{0}\right]$ and $C_{s}\left[i_{1}\right]$. By construction of $B_{s}, C_{s}\left[i_{0}\right] \neq C_{s}\left[i_{1}\right]$ and therefore we are done. Note that $i_{0}$ and $i_{1}$ can be easily found in constant time using rank/select queries on $B_{s}$.

If, on the other hand, all sequences $B_{s}\left[l_{s} . . r_{s}\right]$ are unary, then we just need to retrieve the $\mathcal{O}(\log v)$ colors $C_{s}\left[l_{s}\right]$, for all the $u_{s}$ covering the interval $\left[l^{\prime}, r^{\prime}\right]$, and check if any two of them are distinct (e.g. radix-sort them in linear time). Also this step runs in $\mathcal{O}(\log v)$ time.

Overall, our solution uses $\mathcal{O}(n \log |\Gamma|)$ space and runs in $\mathcal{O}(n \log v)=\mathcal{O}(n \log n)$ time.

### 3.3 Computing a minimum $\boldsymbol{k}$-attractor

Computing a minimum $k$-attractor is NP-hard for $k \geq 3$ and large $\sigma$. In this section we show that the problem is actually polynomial-time-solvable for small $k$ and $\sigma$. Our algorithm takes advantage of both our reduction to set-cover and the verification algorithm of Theorem 18.

First, we give an upper-bound to the cardinality of the set of all minimal $k$-attractors. This will speed up our procedure for finding a minimum $k$-attractor (which is, in particular, minimal). By Lemma 13, there are no more than $\exp \left(\mathcal{O}\left(\sigma^{2 k}\right)\right) k$-attractors for $S$. With the following lemma, we give a better upper-bound to the number of minimal $k$-attractors.

- Lemma 23. The number of minimal $k$-attractors is $\exp \left(\mathcal{O}\left(\sigma^{k} \log \sigma^{k}\right)\right)$.

Proof. Let minimal $(\sigma, k)$ denote the maximum number of minimal $k$-attractors on the alphabet $[1 . . \sigma]$ (independently of the string length $n$ ). Let $\Gamma$ be a minimal $k$-attractor. By Lemma 21, for every $j \in \Gamma$ there is at least one edge $e \in \mathcal{U}$ marked by $j$ only. Let
edge : $\Gamma \rightarrow \mathcal{U}$ be the function defined as follows: edge $(j)=e$ such that (i) $e$ is marked by $j$ only, and (ii) among all edges marked by $j$ only, $e$ is the one with the lexicographically smallest $s(e)$ (where, if $s\left(e^{\prime}\right)$ prefixes $s\left(e^{\prime \prime}\right)$, then we consider $s\left(e^{\prime}\right)$ smaller than $s\left(e^{\prime \prime}\right)$ in lexicographic order). Let $\mathcal{U}^{\prime}=$ edge $(\Gamma)$ be the image of $\Gamma$ through edge function. By its definition, edge is a bijection between $\Gamma$ and $\mathcal{U}^{\prime}$. This implies that $|\Gamma|=\left|\mathcal{U}^{\prime}\right| \leq|\mathcal{U}| \leq \sigma^{k}$ : a minimal $k$-attractor is a set of cardinality at most $\sigma^{k}$ chosen from a universe $\mathcal{C}$ of size at most $|\mathcal{C}| \leq \sigma^{2 k-1}+2 k-2 \leq \sigma^{2 k}$, therefore: minimal $(\sigma, k) \leq \sum_{i=1}^{\sigma^{k}}\binom{\sigma^{2 k}}{i}$. We now give an upper-bound to the function $f(N, t)=\sum_{i=1}^{t}\binom{N}{i}$, where we assume $t \geq 2$ for simplicity (the hypothesis holds in our case since $t=\sigma^{k}$ ). Then, we will plug our bound in the above inequality. Since $\binom{N}{i}<\frac{N^{i}}{i!}$, we have that

$$
\begin{aligned}
f(N, t) & <\sum_{i=1}^{t} \frac{N^{i}}{i!} \\
& =\sum_{i=1}^{t}\left(\frac{N}{t}\right)^{i} \cdot \frac{t^{i}}{i!} \\
& \leq \sum_{i=1}^{t}\left(\frac{N}{t}\right)^{i} \cdot \sum_{i=1}^{t} \frac{t^{i}}{i!} \\
& \in \mathcal{O}\left(\left(\frac{N \cdot e}{t}\right)^{t}\right)
\end{aligned}
$$

We obtain our claim: $\operatorname{minimal}(\sigma, k) \leq f\left(\sigma^{2 k}, \sigma^{k}\right) \in \mathcal{O}\left(e^{\sigma^{k}} \cdot \sigma^{k \sigma^{k}}\right) \leq \exp \left(\mathcal{O}\left(\sigma^{k} \log \sigma^{k}\right)\right)$.
Using the above lemma, we now provide a strategy to find a minimum $k$-attractor.

- Theorem 24. A minimum $k$-attractor can be found in $\mathcal{O}(n)+\exp \left(\mathcal{O}\left(\sigma^{k} \log \sigma^{k}\right)\right)$ time.

Proof. Let $c(i)=S^{\prime}[i-k+1 . . i+k-1]$, where $S^{\prime}[i]=\# \notin \Sigma$ if $i<1$ or $i>n$ and $S^{\prime}[i]=S[i]$ otherwise, be the context string associated to position $i$. Consider the string

$$
C=c\left(i_{1}\right) \$ c\left(i_{2}\right) \$ \ldots \$ c\left(i_{t}\right)
$$

where $\left\{i_{1}, i_{2}, \ldots, i_{t}\right\}=\mathcal{C}$ and $\# \neq \$ \notin \Sigma$. By our choice of $\mathcal{C}$, the length of this string is $|C|=(|\mathcal{C}| \cdot 2 k)-1 \leq\left(\sigma^{2 k-1}+2 k\right) \cdot 2 k \in \mathcal{O}\left(\sigma^{2 k} \cdot k\right) \leq \exp \left(\mathcal{O}\left(\log \sigma^{k}\right)\right)$. We can build $C$ in $\mathcal{O}(n+|C|)$ time using the suffix tree of $S$ (i.e., extracting all paths from the root to nodes at string depth at most $2 k-1$ ).

Let now $\Gamma^{\prime \prime} \subseteq\{k \cdot(2 j+1) \mid j=0, \ldots, t-1\}$. It is easy to see that $\Gamma^{\prime}=\{i \mid C[i]=$ $\$$ or $C[i]=\#\} \cup \Gamma^{\prime \prime}$ is a $k$-attractor for $C$ if and only if the set $\Gamma=\left\{i_{(x-k) /(2 k)} \mid x \in \Gamma^{\prime \prime}\right\}$ is a $k$-attractor for $S$. Suppose that $\Gamma^{\prime}$ is a $k$-attractor for $C$, and consider a substring $s$ of $S$ of length at most $k$. By construction of $C, s$ is also a substring of $C$; in particular, there is an occurrence $C[i . . i+|s|-1]=s$ crossing a position $k \cdot(2 j+1) \in \Gamma^{\prime \prime}$, for some $j$. Then, $i_{j} \in \Gamma$ and, by the way we defined $C$, there is an occurrence of $s$ crossing position $i_{j}$ in S . Conversely, suppose that $\Gamma$ is a $k$-attractor for $S$, and let $s$ be a substring of $C$ of length at most $k$. If $s$ contains either $\$$ or $\#$, then it must cross one of the positions in $\{i \mid C[i]=\$$ or $C[i]=\#\} \subseteq \Gamma^{\prime}$. Otherwise, it appears inside one of the substrings $c\left(i_{k}\right)$, for some $k \in[1 . . t]$. But then, this means that $s$ appears in $S$ and, in particular, that it has some occurrence $s^{\prime}=s$ crossing a position $i_{j} \in \Gamma$. From the way we constructed $C, s$ has an occurrence in $C$ crossing position $k \cdot(2 j+1) \in \Gamma^{\prime \prime} \subseteq \Gamma^{\prime}$.

At this point, we check whether $\Gamma^{\prime}$ is a $k$-attractor for $C$ for all possible $\Gamma^{\prime \prime}$, and return the smallest such set. Instead of trying all subsets of $\mathcal{C}$, we use Lemma 23 and generate only subsets of $\mathcal{C}$ of size at most $\sigma^{k}$; these subsets will include all minimal $k$-attractors and, in particular, all minimum $k$-attractors. By Lemma 23 , there are at most $\exp \left(\mathcal{O}\left(\sigma^{k} \log \sigma^{k}\right)\right)$ such sets, and each verification takes linear $\mathcal{O}(|C|)=\exp \left(\mathcal{O}\left(\log \sigma^{k}\right)\right)$ time using Theorem 18. Overall, our algorithm for the minimum $k$-attractor runs in $\mathcal{O}(n)+\exp \left(\mathcal{O}\left(\log \sigma^{k}\right)\right) \cdot \exp \left(\mathcal{O}\left(\sigma^{k} \log \sigma^{k}\right)\right)$ time, which simplifies to $\mathcal{O}(n)+\exp \left(\mathcal{O}\left(\sigma^{k} \log \sigma^{k}\right)\right)$.

### 3.4 Sharp attractors

The complexity of $k$-attractor has been fully characterized in [16], except for the particular case $k=2$ : for $k \geq 3$, the problem has been proven to be NP-complete, while for $k=1$ it is clearly solvable in linear time. While we have not been able to settle the case $k=2$, we show a polynomial-time solution under the additional constraint that only substrings of length exactly equal to 2 are captured by the attractor set. We denote with the name $k$-sharp attractor this variant. Formally, we define a $k$-sharp attractor of a string $S \in \Sigma^{n}$ to be a set of positions $\Gamma \subseteq[1 . . n]$ such that every substring $S[i . . j]$ with $j-i+1=k$ has an occurrence $S\left[i^{\prime} . . j^{\prime}\right]=S[i . . j]$ with $j^{\prime \prime} \in\left[i^{\prime} . . j^{\prime}\right]$ for some $j^{\prime \prime} \in \Gamma$. In other words, a $k$-sharp-attractor is a subset that covers all substrings of length exactly $k$. By Minimum- $k$-Sharp-Attractor we denote the optimization problem of finding the smallest $k$-sharp attractor of a given input string, and by $k$-Sharp-Attractor $=\{\langle T, p\rangle \mid$ String $T$ has a $k$-sharp-attractor of size $\leq p\}$ we denote the corresponding decision problem. In the full version [15] of this paper we show that, for any constant $k \geq 3, k$-Sharp-Attractor is NP-complete. The proof is an adaptation of the original one proposed for $k$-attractors in [16], and is based on a reduction from set cover. Here we show that for $k=2$ the problem can be solved in polynomial time (again, the case $k=1$ is trivially solvable in linear time).

- Theorem 25. Minimum 2-sharp-attractor is in $P$.

Proof. It is easy to show that 2 -sharp-attractor is in $P$ by a reduction to edge cover. Given a string $S$, let $V \subseteq \Sigma^{2}$ be the set of strings of length 2 that occur at least once in $S$. For every substring of length 3 of the form $x y z$, add the edge $(x y, y z)$ to the edge-set $E$, and add self-loops for the first and last pair.

A position $\gamma \in \Gamma$ thus corresponds to an edge, $e_{\gamma}$, and it is easy to see that $\Gamma$ is a 2-sharp-attractor if and only if $\left\{e_{\gamma} \mid \gamma \in \Gamma\right\}$ is an edge cover.

The number of vertices and edges in this graph are both $\leq n$, so a minimum edge cover can be found in $\mathcal{O}(n \sqrt{n})$ time [19].

## References

1 Djamal Belazzougui. Linear time construction of compressed text indices in compact space. In Annual Symposium on Theory of Computing (STOC), pages 148-193. ACM, 2014.
2 Djamal Belazzougui, Fabio Cunial, Juha Kärkkäinen, and Veli Mäkinen. Linear-time string indexing and analysis in small space. arXiv preprint 1609.06378, 2016.
3 Anselm Blumer, Janet Blumer, David Haussler, Ross McConnell, and Andrzej Ehrenfeucht. Complete inverted files for efficient text retrieval and analysis. Journal of the ACM, 34(3):578-595, 1987.
4 Michael Burrows and David J. Wheeler. A block sorting lossless data compression algorithm. Technical Report 124, Digital Equipment Corporation, 1994.
5 David Clark. Compact Pat trees. PhD thesis, University of Waterloo, 1998.
6 Maxime Crochemore and Renaud Vérin. Direct construction of compact directed acyclic word graphs. In Combinatorial Pattern Matching (CPM), pages 116-129. Springer, 1997.
7 Martin Farach. Optimal suffix tree construction with large alphabets. In Annual Symposium on Foundations of Computer Science (FOCS), pages 137-143. IEEE, 1997.
8 Paolo Ferragina and Giovanni Manzini. Indexing compressed text. Journal of the ACM, 52(4):552-581, 2005.
9 Johannes Fischer and Volker Heun. Space-efficient preprocessing schemes for range minimum queries on static arrays. SIAM Journal on Computing, 40(2):465-492, 2011.

10 Roberto Grossi, Ankur Gupta, and Jeffrey Scott Vitter. High-order entropy-compressed text indexes. In Annual Symposium on Discrete Algorithms (SODA), pages 841-850. SIAM, 2003.

11 Roberto Grossi and Jeffrey Scott Vitter. Compressed suffix arrays and suffix trees with applications to text indexing and string matching. SIAM Journal on Computing, 35(2):378407, 2005.
12 Prosenjit Gupta, Ravi Janardan, and Michiel Smid. Further results on generalized intersection searching problems: counting, reporting, and dynamization. Journal of Algorithms, 19(2):282-317, 1995.
13 Guy Joseph Jacobson. Succinct static data structures. PhD thesis, Carnegie Mellon University, 1988.
14 Juha Kärkkäinen, Peter Sanders, and Stefan Burkhardt. Linear work suffix array construction. Journal of the ACM, 53(6):918-936, 2006.
15 Dominik Kempa, Alberto Policriti, Nicola Prezza, and Eva Rotenberg. String attractors: Verification and optimization. arXiv, 2018. arXiv:1803.01695.
16 Dominik Kempa and Nicola Prezza. At the roots of dictionary compression: String attractors. In Annual Symposium on Theory of Computing (STOC), pages 827-840. ACM, 2018.

17 John C. Kieffer and En-Hui Yang. Grammar-based codes: A new class of universal lossless source codes. IEEE Transactions on Information Theory, 46(3):737-754, 2000.
18 Udi Manber and Gene Myers. Suffix arrays: A new method for on-line string searches. SIAM Journal on Computing, 22(5):935-948, 1993.
19 Silvio Micali and Vijay V. Vazirani. An $O(\sqrt{|V|} \cdot|E|)$ algorithm for finding maximum matching in general graphs. In Annual Symposium on Foundations of Computer Science (SFCS), pages 17-27. IEEE, 1980.
20 J. Ian Munro, Gonzalo Navarro, and Yakov Nekrich. Space-efficient construction of compressed indexes in deterministic linear time. In Annual Symposium on Discrete Algorithms (SODA), pages 408-424. SIAM, 2017.
21 Gonzalo Navarro. Wavelet trees for all. Journal of Discrete Algorithms, 25:2-20, 2014.
22 Gonzalo Navarro. Compact data structures: A practical approach. Cambridge University Press, 2016.
23 James A. Storer and Thomas G. Szymanski. Data compression via textual substitution. Journal of the ACM, 29(4):928-951, 1982.
24 Peter Weiner. Linear pattern matching algorithms. In Annual Symposium on Switching and Automata Theory (SWAT/FOCS), pages 1-11. IEEE, 1973.

