

Synchronization Strings: Highly Efficient Deterministic Constructions over Small Alphabets*

Kuan Cheng[†] Bernhard Haeupler[‡] Xin Li[§] Amirbehshad Shahrabi[¶] Ke Wu^{||}

Abstract

Synchronization strings are recently introduced by Haeupler and Shahrabi [1] in the study of codes for correcting insertion and deletion errors (insdel codes). A synchronization string is an encoding of the indices of the symbols in a string, and together with an appropriate decoding algorithm it can transform insertion and deletion errors into standard symbol erasures and corruptions. This reduces the problem of constructing insdel codes to the problem of constructing standard error correcting codes, which is much better understood. Besides this, synchronization strings are also useful in other applications such as synchronization sequences and interactive coding schemes. For all such applications, synchronization strings are desired to be over alphabets that are as small as possible, since a larger alphabet size corresponds to more redundant information added.

Haeupler and Shahrabi [1] showed that for any parameter $\varepsilon > 0$, synchronization strings of arbitrary length exist over an alphabet whose size depends only on ε . Specifically, [1] obtained an alphabet size of $O(\varepsilon^{-4})$, which left an open question on where the minimal size of such alphabets lies between $\Omega(\varepsilon^{-1})$ and $O(\varepsilon^{-4})$. In this work, we partially bridge this gap by providing an improved lower bound of $\Omega(\varepsilon^{-3/2})$, and an improved upper bound of $O(\varepsilon^{-2})$. We also provide fast explicit constructions of synchronization strings over small alphabets.

Further, along the lines of previous work on similar combinatorial objects, we study the extremal question of the smallest possible alphabet size over which syn-

chronization strings can exist for some constant $\varepsilon < 1$. We show that one can construct ε -synchronization strings over alphabets of size four while no such string exists over binary alphabets. This reduces the extremal question to whether synchronization strings exist over ternary alphabets.

1 Introduction

This paper focuses on the study of a combinatorial object called *synchronization string*. Intuitively, a synchronization string is a (finite or infinite) string that avoids similarities between pairs of intervals in the string. Such nice properties and synchronization strings themselves can actually be motivated from at least two different aspects: coding theory and pattern avoidance. We now discuss the motivations and previous work in each aspect below.

1.1 Motivation and Previous Work in Coding

Theory The general and most important goal of coding theory is to obtain reliable transmissions of information in the presence of noise or adversarial error. Starting from the pioneering works of Shannon, Hamming, and many others, coding theory has evolved into an extensively studied field, with applications found in various areas in computer science. Regarding the general goal of correcting errors, we now have a very sophisticated and almost complete understanding of how to deal with Hamming errors, i.e., symbol erasures and substitutions. On the other hand, the knowledge of codes for edit errors (i.e., symbol insertions and deletions) has lagged far behind despite also being studied intensively since the 1960s. In practice, this is one of the main reasons why communication systems require a lot of effort and resources to maintain synchronization strictly.

One major difficulty in designing codes for insertion and deletion errors is that in the received codeword, the positions of the symbols may have changed. This is in contrast to standard symbol erasures and substitutions, where the positions of the symbols always stay the same. Thus, many of the known techniques in designing codes for standard Hamming errors cannot be directly utilized to protect against insertion and deletion errors.

*Supported in part by NSF grants CCF-1527110, CCF-1618280, CCF-1617713, CCF-1814603 and NSF CAREER award CCF-1750808.

[†]kcheng17@jhu.edu. Department of Computer Science, Johns Hopkins University.

[‡]haeupler@cs.cmu.edu. Department of Computer Science, Carnegie Mellon University.

[§]lixint@cs.jhu.edu. Department of Computer Science, Johns Hopkins University.

[¶]shahrabi@cs.cmu.edu. Department of Computer Science, Carnegie Mellon University.

^{||}kwu35@jhu.edu. Department of Computer Science, Johns Hopkins University.

In this context, a recent work of Haeupler and Shahrasbi [1] introduced *synchronization strings*, which enable a black-box transformation of Hamming-type error correcting codes to codes that protect against edit errors (insdel codes for short). Informally, a synchronization string of length n is an encoding of the indices of the n positions into a string over some alphabet Σ , such that, despite some insertion and deletion errors, one can still recover the correct indices of many symbols. With these correct indices, a standard error correcting code can then be used to recover the original message. This then gives a code for insertion and deletion errors, which is the combination of a standard error correcting code and a synchronization string.

The simplest example of a synchronization string is just to record the index of each symbol, i.e., the string $1, 2, \dots, n$. It can be easily checked that even if $(1 - \varepsilon)$ fraction of these indices are deleted, one can still correctly recover the positions of the remaining εn symbols. However, this synchronization string uses an alphabet whose size grows with the length of the string. The main contribution of [1] is to show that under a slight relaxation, there exist synchronization strings of arbitrary length n over an alphabet of fixed size. Further, [1] provided efficient and streaming algorithms to correctly recover the indices of many symbols from a synchronization string after being altered by insertion and deletion errors. Formally, [1] defines *ε -synchronization strings* as follows. A string S is an ε -synchronization string if the edit distance (the number of insertions and deletions needed to transform one string into another) of any two consecutive substrings $S[i, j)$ and $S[j, k)$ is at least $(1 - \varepsilon)(k - i)$.

DEFINITION 1.1. (ε -SYNCHRONIZATION STRING) A string S is an ε -synchronization string if $\forall 1 \leq i < j < k \leq |S| + 1$, $ED(S[i, j), S[j, k)) > (1 - \varepsilon)(k - i)$. A string S is an infinite ε -synchronization string if it has infinite length and $\forall 1 \leq i < j < k$, $ED(S[i, j), S[j, k)) > (1 - \varepsilon)(k - i)$. Here $ED(,)$ stands for the edit distance.

Using the construction and decoding algorithms for ε -synchronization strings, [1] gives a code that for any $\delta \in (0, 1)$ and $\varepsilon > 0$, can correct δ fraction of insertion and deletion errors with rate $1 - \delta - \varepsilon$. Besides this, synchronization strings have found a variety of applications, such as in synchronization sequences [2], interactive coding schemes [3, 4, 5, 6, 7, 8, 9, 10], coding against synchronization errors [10, 11, 12], and edit distance tree codes [9, 13].

For all such applications, synchronization strings are desired to be over alphabets that are as small as pos-

sible, since a larger alphabet size corresponds to more redundant information added. Thus a natural question here is how small the alphabet size can be. In [1], Haeupler and Shahrasbi showed that ε -synchronization strings with arbitrary length exist over an alphabet of size $O(\varepsilon^{-4})$, they also gave a randomized polynomial time algorithm to construct such strings. In a very recent work [10], they further gave various efficient deterministic constructions for finite/infinite ε -synchronization strings, which have alphabets of size $\text{poly}(\varepsilon^{-1})$ for some unspecified large polynomial. On the other hand, the definition of synchronization strings implies that any ε^{-1} consecutive symbols in an ε -synchronization string have to be distinct—providing an $\Omega(\varepsilon^{-1})$ lower-bound for the alphabet size.

Haeupler and Shahrasbi [10] further introduced another notion called long distance synchronization strings. Besides requiring that the edit distance of any adjacent intervals should be large, long distance synchronization strings also require that any large, nearby intervals should have large edit distance.

DEFINITION 1.2. ($f(l)$ -DISTANCE ε -SYNCHRONIZATION STRING)

A string $S \in \Sigma^n$ is a $f(l)$ -distance ε -synchronization string if for every $1 \leq i < j \leq i' < j' \leq n + 1$, $ED(S[i, j), S[i', j')) > (1 - \varepsilon)(l)$ for $i' - j \leq f(l)$ where $l = j + j' - i - i'$.

As a special case, 0-distance synchronization strings are standard synchronization strings. Similar to [10], we focus on $f(l) = n \cdot \mathbb{1}_{l > c \log n}$ where $\mathbb{1}_{l > c \log n}$ is the indicator function for $l > c \log n$. This definition considers the edit distance of all pairs of large intervals and adjacent small intervals.

DEFINITION 1.3. (c -LONG-DISTANCE ε -SYNCHRONIZATION STRING)

We call $n \cdot \mathbb{1}_{l > c \log n}$ -distance ε -synchronization strings c -long-distance ε -synchronization strings.

Haeupler and Shahrasbi [10] showed that long distance synchronization strings are locally decodable. This crucial property is used in [10] to construct codes for edit errors with block transpositions and block repetitions. Long distance synchronization strings can also be applied to achieve infinite channel simulations with exponentially smaller memory and decoding time requirements [10]. We note that the constructions in [10] have alphabets of size $\text{poly}(\varepsilon^{-1})$ for some unspecified large polynomial, which motivates us to reduce the alphabet size.

1.2 Motivation and Previous Work in Pattern Avoidance

Apart from applications in coding theory and other communication problems involving insertions

and deletions, synchronization strings are also interesting combinatorial objects from a mathematical perspective. As a matter of fact, plenty of very similar combinatorial objects have been studied prior to this work.

A classical work of Thue [14] introduces and studies *square-free* strings, i.e., strings that do not contain two identical consecutive substrings. Thue shows that such strings exist over alphabets of size three and provides a fast construction of such strings using *morphisms*. The seminal work of Thue inspired further works on the same problem [15, 16, 17, 18, 19, 20] and problems with a similar pattern avoidance theme.

Krieger et. al. [21] study strings that satisfy relaxed variants of square-freeness, i.e., strings that avoid *approximate squares*. Their study provides several results on strings that avoid consecutive substrings of equal length with small additive or multiplicative Hamming distance in terms of their length. In each of these regimes, [21] gives constructions of approximate square free strings over alphabets with small constant size for different parameters.

Also, Camungol and Rampersad [22] study *approximate squares with respect to edit distance*, which is equivalent to the ε -synchronization string notion except that the edit distance property is only required to hold for pairs of consecutive substrings of equal length. [22] employs a technique based on entropy compression to prove that such strings exist over alphabets that are constant in terms of string length but exponentially large in terms of ε^{-1} . We note that the previous result of Haeupler and Shahrabi [1] already improves this dependence to $O(\varepsilon^{-4})$.

Finally, Beck [31] uses the Lovász local lemma to show the existence of words where subsequent occurrences of identical substrings of length n occur exponentially (in n) far away from each other, which is very similar to long distance synchronization string. The difference is that Beck mainly care about identical substrings being widely separated.

Again, a main question addressed in most of the above-mentioned previous work on similar mathematical objects is how small the alphabet size can be.

1.3 Our Results In this paper we study the question of how small the alphabet size of an ε -synchronization string can be. We address this question both for a specified ε and for unspecified ε . In the first case we try to bridge the gap between the upper bound of $O(\varepsilon^{-4})$ provided in [1] and the lower bound of $\Omega(\varepsilon^{-1})$. In the second case we study the question of how small the alphabet size can be to ensure the existence of an ε -synchronization string for some constant $\varepsilon < 1$. In both

cases we also give efficient constructions that improve previous results.

1.3.1 New Bounds on Minimal Alphabet Size for a given ε

Our first theorem gives improved upper bound and lower bound for the alphabet size of an ε -synchronization string for a given ε .

THEOREM 1.1. *For any $0 < \varepsilon < 1$, there exists an alphabet Σ of size $O(\varepsilon^{-2})$ such that an infinite ε -synchronization string exists over Σ . In addition, $\forall n \in \mathbb{N}$, a randomized algorithm can construct an ε -synchronization string of length n in expected time $O(n^5 \log n)$. Further, the alphabet size of any ε -synchronization string has to be at least $\Omega(\varepsilon^{-3/2})$.*

REMARK 1.1. *Here ε can be any number in $(0, 1)$, not just a constant. The lower bound is conditioned on that the synchronization string is long enough in terms of ε , as short synchronization strings can exist over small alphabets. For example, a binary string 01 is a synchronization string, but we are not very interested in such cases.*

Next, we provide efficient and even linear-time constructions of ε -synchronization strings over drastically smaller alphabets than the efficient constructions in [10]. The construction is highly-explicit in the sense that the i 'th position of a synchronization string can be computed in time $O(\log i)$. Moreover, our construction also works for the notion of long distance synchronization strings.

THEOREM 1.2. *For every $n \in \mathbb{N}$ and any constant $\varepsilon \in (0, 1)$, there is a deterministic construction of a (long-distance) ε -synchronization string of length n over an alphabet of size $O(\varepsilon^{-2})$ that runs in $\text{poly}(n)$ time. Further, there is a highly-explicit linear time construction of such strings over an alphabet of size $O(\varepsilon^{-3})$.*

In addition, in Section 5.3, we present a method to construct infinite synchronization strings using constructions of finite ones, which only increases the alphabet size by a constant factor—as opposed to the construction in [10] that increases the alphabet size quadratically.

THEOREM 1.3. *For any constant $0 < \varepsilon < 1$, there exists an explicit construction of an infinite ε -synchronization string S over an alphabet of size $O(\varepsilon^{-2})$. Further, there exists a highly-explicit construction of an infinite ε -synchronization string S over an alphabet of size $O(\varepsilon^{-3})$ such that for any $i \in \mathbb{N}$, the first i symbols can be computed in $O(i)$ time and $S[i, i + \log i]$ can be computed in $O(\log i)$ time.*

1.3.2 Minimal Alphabet Size for Unspecified ε : Three or Four? One interesting question that has been commonly addressed by previous work on similar combinatorial objects is the size of the smallest alphabet over which one can find such objects. Along the lines of [14, 20, 21, 22], we study the existence of synchronization strings over alphabets with minimal constant size.

It is easy to observe that no such string can exist over a binary alphabet since any binary string of length four either contains two consecutive identical symbols or two consecutive identical substrings of length two. On the other hand, one can extract constructions over constant-sized alphabets from the existence proofs in [1, 10], but the unspecified constants there would be quite large. In Section 7.2, for some $\varepsilon < 1$, we provide a construction of arbitrarily long ε -synchronization strings over an alphabet of size four. This narrows down the question to whether such strings exist over alphabets of size three.

To construct such strings, we introduce the notion of *weak synchronization string*, which requires substrings to satisfy a similar property as that of an ε -synchronization string, except that the lower bound on edit distance is rounded down to the nearest integer. We show that weak synchronization strings exist over binary alphabets and use one such string to modify a ternary square-free string ([14]) into a synchronization string over an alphabet of size four.

1.3.3 Constructing Synchronization Strings Using Uniform Morphisms Morphisms have been widely used in previous work as a tool to construct similar combinatorial objects. A uniform morphism of rank r over an alphabet Σ is a function $\phi : \Sigma \rightarrow \Sigma^r$ that maps any symbol of an alphabet Σ to a string of length r over the same alphabet. Using this technique, some similar combinatorial objects in previous work have been constructed by taking a symbol from the alphabet and then repeatedly using an appropriate morphism to replace each symbol with a string [20, 21]. Here we investigate whether such tools can also be utilized to construct synchronization strings. In Section 7.1, we show that no such morphism can construct arbitrarily long ε -synchronization strings for any $\varepsilon < 1$.

1.4 Applications Our new constructions of synchronization strings can be used in all known applications of these objects, including converting standard error correcting codes into error correcting codes for edit errors that approach the singleton bound, giving locally decodable codes for edit errors, constructing codes for block transpositions and replications, obtaining near-

linear time interactive coding schemes for edit errors, achieving exponentially more efficient infinite channel simulations and so on.

For all these applications, our constructions can reduce the alphabet size of the code by at least a ε^{-2} factor where ε is the error parameter of the synchronization string. This also further improves the alphabet size of an approximate square free string with respect to edit distance in [22] to $O(\varepsilon^{-2})$.

In addition, our construction relies on a novel implementation of a deterministic Lovász Local Lemma. This may have further applications in other related problems.

1.5 Overview of Our Techniques Our constructions differ from previous constructions in several key aspects. Here we first highlight the new ideas we use to achieve the improvements.

1. Non-uniform Lovász Local Lemma. To show the existence of ε -synchronization strings over an alphabet of size $O(\varepsilon^{-2})$, as in previous works we use Lovász Local Lemma. However, instead of using a uniform sample space, we use a *non-uniform* sample space. We carefully set up this sample space so that the dependency graph is still sparse, and thus the Local Lemma can be used to establish the existence of synchronization strings.

2. Adapted algorithmic Lovász Local Lemma. To get an explicit construction of our synchronization strings, we need to use an algorithmic version of Lovász Local Lemma. However, since our sample space is non-uniform, we also need to adapt the algorithm accordingly. Specifically, we need to adjust the sampling process to generate the non-uniform sample space very carefully.

3. Adapted deterministic Lovász Local Lemma. To get a deterministic construction of our synchronization strings, we further need a deterministic algorithmic Lovász Local Lemma. A direct application of previous known results such as that of of Chandrasekaran et al. [23] is problematic, since it will make the dependency graph too dense. Thus, we need to carefully *redefine* the bad events in the Local Lemma to make this work.

4. Synchronization circles. To protect small intervals in a synchronization string, we introduce a new combinatorial object called *synchronization circle*. A synchronization circle is a generalization of synchronization string, which can be viewed as a circle of symbols. It has the property that no matter

where one cuts the circle, it becomes a synchronization string. We show that a synchronization circle can be constructed by concatenating two synchronization strings over two disjoint alphabets, which only increases the alphabet size by a factor of two. This idea is also used in our construction of infinite synchronization strings to reduce the alphabet size.

5. **Reverse connection between insdel codes and synchronization strings.** We give a new way to construct synchronization strings, by concatenating different codewords from an appropriately designed insdel code. This reduces the task of constructing synchronization strings to a much smaller size, which is important to achieve our highly explicit constructions. This also complements the connection established in [1], which shows that insdel codes can be constructed from synchronization strings.

We now explain our techniques in more details.

1.5.1 Existence of ε -synchronization strings over small alphabets Haeupler and Shahrabi [1] showed the existence of ε -synchronization strings over an alphabet of size $O(\varepsilon^{-4})$. By using the Local Lemma, they first show that with positive probability, a uniformly random string S over an alphabet of size $O(\varepsilon^{-2})$ satisfies the requirement of a synchronization string when the intervals have size at least t for some $t = \Omega(\varepsilon^{-2})$. To take care of the small intervals, they do a symbol-wise concatenation of S and another string R whose i 'th position is $R[i] = (i \bmod t)$. This results in an alphabet size of $O(\varepsilon^{-4})$.

To get rid of the concatenation, we directly use a non-uniform sample space in the Local Lemma. In our construction, we randomly sample each symbol (say the i 'th symbol) in the string S from an alphabet Σ of size $O(\varepsilon^{-2})$, conditioned on that this symbol is different from the previous $t - 1$ or $i - 1$ symbols, whichever is smaller. We then define the bad events in the following way: an interval $S[i, k]$ is bad if there is a $j, i < j < k$, s.t. $ED(S[i, j], S[j, k]) \leq (1 - \varepsilon)(k - i)$. At a first glance, as each symbol now depends on the previous $t - 1$ symbols, we may not directly apply Lovász Local Lemma. However, our key observation is that, although the symbols are not independent of each other, the badness of an interval I is mutually independent of the badness of all intervals that do not intersect with I . We show this by a direct computation of the probabilities. Now we can apply Lovász Local Lemma to show that S is an ε -synchronization string with positive probability.

1.5.2 Constructions of ε -synchronization strings over small alphabets

Randomized construction. In order to apply the algorithmic framework of Lovász Local Lemma in [24] and [25], we need the probability space in the Local Lemma to be sampled from n independent random variables $\mathcal{P} = \{P_1, \dots, P_n\}$, and that each event in the collection of bad events $\mathcal{A} = \{A_1, \dots, A_m\}$ is determined by some subsets of \mathcal{P} . However our non-uniform sample space discussed above does not satisfy these conditions.

To rectify this, we implement our previous sampling process as follows (which gives an alternative view of the previous sampling procedure): we initially fix an arbitrary order for the symbols in the alphabet Σ . When we sample the i 'th symbol of the string, we reorder the alphabet such that the first $t - 1$ symbols in the alphabet are exactly the previous $t - 1$ symbols in the string, and the rest of the symbols are placed in the current order as the last $|\Sigma| - t + 1$ symbols in the alphabet. Then we uniformly randomly choose an integer P_i in $\{1, 2, \dots, |\Sigma| - t + 1\}$ and let the i 'th symbol be the $(t - 1 + P_i)$ 'th symbol in the new order of Σ . In this way, the random variables $\{P_1, \dots, P_n\}$ are indeed independent, and the i 'th symbol is different from the previous $t - 1$ symbols. Furthermore, the event of any interval $S[i, k]$ being bad depends only on $\{P_i, \dots, P_k\}$. This enables us to follow the framework in [24] and [25] and give a randomized construction of synchronization strings over small alphabets.

Deterministic polynomial time construction. We combine the sampling algorithm described above with the deterministic Lovász local lemma of Chandrasekaran et al. [23] to give a polynomial-time construction of c -long distance ε -synchronization strings. In addition to the requirement on edit distance of adjacent intervals, long distance synchronization strings also require that any pair of large disjoint intervals $S[i_1, i_1 + l_1], S[i_2, i_2 + l_2]$, where $i_1 + l_1 \leq i_2$ and $l_1 + l_2 \geq c \log n$, has large enough edit distance. However, it's problematic to define the bad event corresponding to intervals $S[i_1, i_1 + l_1]$ and $S[i_2, i_2 + l_2]$ in the usual way: $ED(S[i_1, i_1 + l_1], S[i_2, i_2 + l_2]) \leq (1 - \varepsilon)(l_1 + l_2)$. This is because under this definition, the event not only depends on the random variables $\{P_j\}$ within these intervals, but also depends on the random variables between these two intervals. This will make the dependency graph too dense to apply the Local Lemma.

To fix this, we redefine the badness of the pair of intervals $S[i_1, i_1 + l_1]$ and $S[i_2, i_2 + l_2]$ as follows: the pair is bad if there exists a choice of the random variables in the middle, such that according to the sampling algorithm, $ED(S[i_1, i_1 + l_1], S[i_2, i_2 + l_2]) \leq (1 - \varepsilon)(l_1 + l_2)$.

This definition removes the dependence on the random variables between the two intervals, and we show that under this definition, the badness of a pair of intervals is still mutually independent of the badness of all disjoint pairs of intervals. Fortunately, the probability of the bad events under this new definition can also be bounded appropriately. Thus we can apply the deterministic Lovász Local Lemma of Chandrasekaran et al. [23].

Deterministic highly explicit construction.

A long distance synchronization string of length n considers the edit distance of all pairs of large intervals and adjacent small intervals. Our construction of long distance synchronization strings uses concatenations of different codewords of a carefully designed insdel code \mathcal{C} where each codeword has length $m = O(\log n)$, and we show that this construction can handle both large intervals and small intervals.

For large intervals the intuition is the following. If the codeword has length $m = c \log n$, then after concatenations of different codewords, any pair of intervals whose total length is larger than m must have large edit distance, since the intervals must contain parts from at least two different codewords.

We are now left with small adjacent intervals. This is where we need to carefully design the code \mathcal{C} . [1] showed that an insdel code can be constructed by doing a symbol-wise concatenation of a standard error correcting code with a synchronization string. If we construct \mathcal{C} in this way, each codeword in \mathcal{C} is a synchronization string itself, which can protect adjacent small intervals if they lie within one codeword. However, we have no guarantee for small intervals crossing two codewords. To address this problem, we introduce a new combinatorial object called synchronization circle, which is a generalization of synchronization string. A synchronization circle S of length n requires that S is a synchronization string beginning from any position i . That is, for any $1 \leq i \leq n$, $S[i], \dots, S[n], S[1], \dots, S[i-1]$ is a synchronization string. We give explicit constructions for synchronization circles by concatenating two synchronization strings on disjoint alphabets. Now by doing a symbol-wise concatenation of a standard error correcting code with a synchronization circle, we have that every codeword in \mathcal{C} is itself a synchronization circle. This guarantees that any two adjacent small intervals in our synchronization string have large edit distance.

Note that our construction is highly-explicit because we can construct each block (each codeword of \mathcal{C} in the construction) independently.

Explicit construction of infinite ε -synchronization strings To construct an infinite ε -synchronization string, we first construct a sequence of synchronization strings whose lengths increase

exponentially. That is, for some $k = O(1/\varepsilon)$, we construct synchronization strings S_1, S_2, \dots such that S_i has length k^i , and concatenate them sequentially. For the analysis, consider any pair of intervals whose total length is between k^t and k^{t+1} for some t , then we can ignore the part of these two intervals before S_t . This is because the number of symbols before S_t is quite small compared to the rest of the string. The remaining part either lies in a synchronization string S_i , or crosses two synchronization strings S_{i-1} and S_i . The first case is already taken care of by the property of a synchronization string. To take care of the second case, we use two disjoint alphabets Σ_1, Σ_2 to construct these synchronization strings. That is, we construct S_i for odd i 's over Σ_1 and S_i for even i 's over Σ_2 . If we use our previous highly-explicit constructions of synchronization strings to construct these S_i 's, then this construction is also highly-explicit.

1.5.3 Existence of ε -synchronization string over alphabet of size 4 To construct a synchronization string over an alphabet of size four, we introduce the notion of *weak synchronization string*, which requires substrings to satisfy a similar property as that of an ε -synchronization string, except that the lower bound on edit distance is rounded down to the nearest integer. We show that weak synchronization strings exist over binary alphabets and use one such string to modify a ternary square-free string ([14]) into a synchronization string over an alphabet of size four.

2 Discussions and Open Problems

In this paper we give highly efficient constructions of various ε -synchronization strings over small alphabets. In particular, we reduce the alphabet size from $O(\varepsilon^{-4})$ or $\text{poly}(1/\varepsilon)$ in previous works to $O(\varepsilon^{-2})$. We also obtain an improved lower bound of $\Omega(\varepsilon^{-3/2})$. This leaves the open problem of what is the tight bound. Similarly, for synchronization strings with unspecified ε , we show that such strings exist over an alphabet of size 4 but cannot exist over an alphabet of size 2, which leaves the open problem of whether synchronization strings can exist over an alphabet of size 3. Experimental results seem to suggest that the answer is yes.

It would be interesting to find other applications of synchronization strings, and we believe our new way of using the deterministic Lovász Local Lemma may also find other applications.

3 Notations and Definitions

Usually we use Σ (probably with some subscripts) to denote the alphabet and Σ^* to denote all strings over alphabet Σ .

DEFINITION 3.1. (SUBSEQUENCE) *The subsequence of a string S is any sequence of symbols obtained from S by deleting some symbols. It doesn't have to be continuous.*

DEFINITION 3.2. (EDIT DISTANCE) *For every $n \in \mathbb{N}$, the edit distance $ED(S, S')$ between two strings $S, S' \in \Sigma^n$ is the minimum number of insertions and deletions required to transform S into S' .*

DEFINITION 3.3. (LONGEST COMMON SUBSEQUENCE) *For any strings S, S' over Σ , the longest common subsequence of S and S' is the longest pair of subsequence that are equal as strings. We denote by $LCS(S, S')$ the length of the longest common subsequence of S and S' .*

Note that $ED(S, S') = |S| + |S'| - 2LCS(S, S')$ where $|S|$ denotes the length of S .

DEFINITION 3.4. (SQUARE-FREE STRING) *A string S is a square free string if $\forall 1 \leq i < i + 2l \leq |S| + 1$, $(S[i, i + l])$ and $S[i + l, i + 2l])$ are different as words.*

We also introduce the following generalization of a synchronization string, which will be useful in our deterministic constructions of synchronization strings.

DEFINITION 3.5. (ε -SYNCHRONIZATION CIRCLE) *A string S is an ε -synchronization circle if $\forall 1 \leq i \leq |S|$, $S_i, S_{i+1}, \dots, S_{|S|}, S_1, S_2, \dots, S_{i-1}$ is an ε -synchronization string.*

Note that in a communication that suffers from symbol corruptions and erasures, any erasure leads to half a unit of Hamming distance from original string and any corruption results into one. To easily quantify errors that consist of both types, we define *half-errors* as follows. Here, the name half-error comes from the fact that a symbol erasure is half as bad as a symbol corruption.

DEFINITION 3.6. (HALF-ERRORS) *Any set of a symbol substitutions and b symbol erasures is referred to as $2a + b$ half-errors, where one symbol substitution is counted as two half errors, and one symbol erasure is counted as one half error.*

4 ε -synchronization Strings and Circles with Alphabet Size $O(\varepsilon^{-2})$

In this section we show that by using a non-uniform sample space together with the Lovász local lemma, we can have a randomized polynomial time construction of an ε -synchronization string with alphabet size $O(\varepsilon^{-2})$. We then use this to give a simple construction of an ε -synchronization circle with alphabet size $O(\varepsilon^{-2})$ as well. Although the constructions here are randomized,

the parameter ε can be anything in $(0, 1)$ (even sub-constant), while our deterministic constructions in later sections usually require ε to be a constant in $(0, 1)$.

We first recall the general Lovász local lemma, [30].

LEMMA 4.1. (General Lovász Local Lemma) *Let A_1, \dots, A_n be a set of bad events. $G(V, E)$ is a dependency graph for this set of events if $V = \{1, \dots, n\}$ and each event A_i is mutually independent of all the events $\{A_j : (i, j) \notin E\}$.*

If there exists $x_1, \dots, x_n \in [0, 1)$ such that for all i we have

$$\Pr[A_i] \leq x_i \prod_{(i,j) \in E} (1 - x_j)$$

Then the probability that none of these events happens is bounded by

$$\Pr\left[\bigwedge_{i=1}^n \bar{A}_i\right] \geq \prod_{i=1}^n (1 - x_i) > 0$$

Using this lemma, we have the following theorem showing the existence of ε -synchronization strings over an alphabet of size $O(\varepsilon^{-2})$.

THEOREM 4.1. $\forall \varepsilon \in (0, 1)$ and $\forall n \in \mathbb{N}$, there exists an ε -synchronization string S of length n over alphabet Σ of size $\Theta(\varepsilon^{-2})$.

Proof. Suppose $|\Sigma| = c_1 \varepsilon^{-2}$ where c_1 is a constant. Let $t = c_2 \varepsilon^{-2}$ and $0 < c_2 < c_1$. The sampling algorithm is as follows:

1. Randomly pick t different symbols from Σ and let them be the first t symbols of S . If $t \geq n$, we just pick n different symbols.
2. For $t + 1 \leq i \leq n$, we pick the i th symbol $S[i]$ uniformly randomly from $\Sigma \setminus \{S[i-1], \dots, S[i-t+1]\}$

Now we prove that there's a positive probability that S contains no *bad* interval $S[i, k]$ which violates the requirement that $ED(S[i, j], S[j+1, k]) > (1 - \varepsilon)(k - i)$ for any $i < j < k$. This requirement is equivalent to $LCS(S[i, j], S[j+1, k]) < \frac{\varepsilon}{2}(k - i)$.

Notice that for $k - i \leq t$, the symbols in $S[i, k]$ are completely distinct. Hence we only need to consider the case where $k - i > t$. First, let's upper bound the probability that an interval is bad:

$$\begin{aligned} & \Pr[\text{interval I of length } l \text{ is bad}] \\ & \leq \binom{l}{\varepsilon l} (|\Sigma| - t)^{-\frac{\varepsilon l}{2}} \end{aligned}$$

The first inequality holds because if the interval is bad, then it has to contain a repeating sequence $a_1a_2\dots a_p a_1a_2\dots a_p$ where p is at least $\frac{\varepsilon l}{2}$. Such sequence can be specified via choosing εl positions in the interval and the probability that a given sequence is valid for the string in this construction is at most $(|\Sigma| - t)^{-\frac{\varepsilon l}{2}}$. The second inequality comes from Stirling's inequality.

The inequality above indicates that the probability that an interval of length l is bad can be upper bounded by $C^{-\varepsilon l}$, where C is a constant and can be arbitrarily large by modifying c_1 and c_2 .

Now we use general Lovász local lemma to show that S contains no bad interval with positive probability. First we'll show the following lemma.

CLAIM 1. *The badness of interval $I = S[i, j]$ is mutually independent of the badness of all intervals that do not intersect with I .*

Proof. Suppose the intervals before I that do not intersect with I are I_1, \dots, I_m , and those after I are $I'_1, \dots, I'_{m'}$. We denote the indicator variables of each interval being bad as b , b_k and $b'_{k'}$.

First we prove that there exists $p \in (0, 1)$ such that $\forall x_1, x_2, \dots, x_m \in \{0, 1\}$,

$$\Pr[b = 1 | b_k = x_k, k = 1, \dots, m] = p$$

According to our construction, we can see that for any fixed prefix $S[1, i-1]$, the probability that I is bad is a fixed real number p' . That is,

$$\forall \text{ valid } \tilde{S} \in \Sigma^{i-1}, \Pr[b = 1 | S[1, i-1] = \tilde{S}] = p'$$

This comes from the fact that, the sampling of the symbols in $S[i, k]$ only depends on the previous $h = \min\{i-1, t-1\}$ different symbols, and up to a relabeling these h symbols are the same h symbols (e.g., we can relabel them as $\{1, \dots, h\}$ and the rest of the symbols as $\{h+1, \dots, |\Sigma|\}$). On the other hand the probability that $b = 1$ remains unchanged under any relabeling of the symbols, since if two sampled symbols are the same, they will stay the same; while if they are different, they will still be different. Thus we have:

$$\begin{aligned} & \Pr[b = 1 | b_k = x_k, i = 1, \dots, m] \\ &= \frac{\sum_{\tilde{S}} \Pr[b = 1, S[1, i-1] = \tilde{S}]}{\sum_{\tilde{S}} \Pr[S[1, i-1] = \tilde{S}]} \\ &= p' \sum_{\tilde{S}} \frac{\Pr[S[1, i-1] = \tilde{S}]}{\sum_{\tilde{S}'} \Pr[S[1, i-1] = \tilde{S}']} \\ &= p'. \end{aligned}$$

In the equations, \tilde{S} indicates all valid string that prefix $S[1, i-1]$ can be such that $b_k = x_k, k = 1, \dots, m$. Hence, b is independent of $\{b_k, k = 1, \dots, m\}$. Similarly, we can prove that the joint distribution of $\{b'_{k'}, k' = 1, \dots, m'\}$ is independent of that of $\{b, b_k, k = 1, \dots, m\}$. Hence b is independent of $\{b_k, b'_{k'}, k = 1, \dots, m, k' = 1, \dots, m'\}$, which means, the badness of interval I is mutually independent of the badness of all intervals that do not intersect with I .

Obviously, an interval of length l intersects at most $l + l'$ intervals of length l' . To use Lovász local lemma, we need to find a sequence of real numbers $x_{i,k} \in [0, 1)$ for intervals $S[i, k]$ for which

$$\Pr[S[i, k] \text{ is bad}] \leq x_{i,k} \prod_{S[i,k] \cap S[i',k'] \neq \emptyset} (1 - x_{i',k'})$$

The rest of the proof is the same as that of Theorem 5.7 in [1].

We propose $x_{i,k} = D^{-\varepsilon(k-i)}$ for some constant $D \geq 1$. Hence we only need to find a constant D such that for all $S[i, k]$,

$$C^{-\varepsilon(k-i)} \leq D^{-\varepsilon(k-i)} \prod_{l=t}^n [1 - D^{-\varepsilon l}]^{l+(k-i)}$$

That is, for all $l' \in \{1, \dots, n\}$,

$$\begin{aligned} C^{-l'} &\leq D^{-l'} \prod_{l=t}^n [1 - D^{-\varepsilon l}]^{\frac{l+l'}{\varepsilon}} \\ \Rightarrow C &\geq \frac{D}{\prod_{l=t}^n [1 - D^{-\varepsilon l}]^{\frac{l+l'+1}{\varepsilon}}} \end{aligned}$$

Notice that the righthand side is maximized when $n = \infty, l' = 1$. Hence it's sufficient to show that

$$C \geq \frac{D}{\prod_{l=t}^{\infty} [1 - D^{-\varepsilon l}]^{\frac{l+1}{\varepsilon}}}$$

Let $L = \max_{D>1} \frac{D}{\prod_{l=t}^{\infty} [1 - D^{-\varepsilon l}]^{\frac{l+1}{\varepsilon}}}$. We only need to guarantee that $C > L$.

We claim that $L = \Theta(1)$. Since $t = c_2 \varepsilon^{-2} = \omega\left(\frac{\log \frac{1}{\varepsilon}}{\varepsilon}\right)$,

$$\frac{D}{\prod_{l=t}^{\infty} [1 - D^{-\varepsilon l}]^{\frac{l+1}{\varepsilon}}} \leq \frac{D}{1 - \frac{2}{\varepsilon^3} \frac{D^{-\frac{1}{\varepsilon}}}{(1 - D^{-\varepsilon})^2}}$$

We can see that for $D = 7$, $\max_{\varepsilon} \left\{ \frac{2}{\varepsilon^3} \frac{D^{-\frac{1}{\varepsilon}}}{(1 - D^{-\varepsilon})^2} \right\} < 0.9$. Therefore (5) is bounded by a constant, which means $L = \Theta(1)$ and the proof is complete.

Using a modification of an argument in [1], we can also obtain a randomized construction.

LEMMA 4.2. *There exists a randomized algorithm which for any $\varepsilon \in (0, 1)$ and any $n \in \mathbb{N}$, constructs an ε -synchronization string of length n over alphabet of size $O(\varepsilon^{-2})$ in expected time $O(n^5 \log n)$.*

Proof. The algorithm is similar to that of Lemma 5.8 in [1], using algorithmic Lovász Local lemma [24] and the extension in [25].

THEOREM 4.2. (THEOREM 3.1 IN [25]) *Suppose there is an $\xi \in [0, 1)$ and an assignment of reals $x : \mathcal{A} \mapsto (0, 1)$ such that:*

$$\forall A \in \mathcal{A} : \Pr[A] \leq (1 - \xi)x(A) \prod_{B \in \Gamma(A)} (1 - x(B))$$

Here, $T := \sum_{A \in \mathcal{A}} x(A)$, where $\Gamma(A) = \{B \in \mathcal{A} : B \neq A \text{ and } vbl(B) \cap vbl(A) \neq \emptyset\}$, in which $vbl(A)$ is the smallest subset of random variables determining event A .

Furthermore:

1. if $\xi = 0$, then the expected number of re-samplings done by the MT algorithm is at most $v_1 = T \max_{A \in \mathcal{A}} \frac{1}{1 - x(A)}$, and for any parameter $\lambda \geq 1$, the MT algorithm terminates within λv_1 re-samplings with probability at least $1 - 1/\lambda$.
2. if $\xi > 0$, then the expected number of re-samplings done by the MT algorithm is at most $v_2 = O(\frac{n}{\xi} \log \frac{T}{\xi})$, and for any parameter $\lambda \geq 1$, the MT algorithm terminates within λv_2 re-samplings with probability $1 - \exp(-\lambda)$.

Here, MT algorithm is a sampling algorithm introduced in [24], which can quickly terminates with high probability.

It starts with a string sampled according to the sampling algorithm in the proof of Theorem 4.1, over alphabet Σ of size $C\varepsilon^{-2}$ for some large enough constant C . Then the algorithm checks all $O(n^2)$ intervals for a violation of the requirements for ε -synchronization string. If a bad interval is found, this interval is resampled by randomly choosing every symbol s.t. each one of them is different from the previous $t - 1$ symbols, where $t = c'\varepsilon^{-2}$ with c' being a constant smaller than C .

One subtle point of our algorithm is the following. Note that in order to apply the algorithmic framework of [24] and [25], one needs the probability space to be sampled from n independent random variables $\mathcal{P} = \{P_1, \dots, P_n\}$ so that each event in the collection $\mathcal{A} =$

$\{A_1, \dots, A_m\}$ is determined by some subset of \mathcal{P} . Then, when some bad event A_i happens, one only resamples the random variables that decide A_i . Upon first look, it may appear that in our application of the Lovász Local lemma, the sampling of the i 'th symbol depends on the the previous $h = \min\{i - 1, t - 1\}$ symbols, which again depend on previous symbols, and so on. Thus the sampling of the i 'th symbol depends on the sampling of all previous symbols. However, we can implement our sampling process as follows: for the i 'th symbol we first independently generate a random variable P_i which is uniform over $\{1, 2, \dots, |\Sigma| - h\}$, then we use the random variables $\{P_1, \dots, P_n\}$ to decide the symbols, in the following way. Initially we fix some arbitrary order of the symbols in Σ , then for $i = 1, \dots, n$, to get the i 'th symbol, we first reorder the symbols Σ so that the previous h chosen symbols are labeled as the first h symbols in Σ , and the rest of the symbols are ordered in the current order as the last $|\Sigma| - h$ symbols. We then choose the i 'th symbol as the $(h + P_i)$ 'th symbol in this new order. In this way, the random variables $\{P_1, \dots, P_n\}$ are indeed independent, and the i 'th symbol is indeed chosen uniformly from the $|\Sigma| - h$ symbols excluding the previous h symbols. Furthermore, the event of any interval $S[i, k]$ being bad only depends on the random variables (P_i, \dots, P_k) since no matter what the previous h symbols are, they are relabeled as $\{1, \dots, h\}$ and the rest of the symbols are labeled as $\{h + 1, \dots, |\Sigma|\}$. From here, the same sequence of (P_i, \dots, P_k) will result in the same behavior of $S[i, k]$ in terms of which symbols are the same. We can thus apply the same algorithm as in [1].

Note that the time to get the i 'th symbol from the random variables $\{P_1, \dots, P_n\}$ is $O(n \log \frac{1}{\varepsilon})$ since we need $O(n)$ operations each on a symbol of size $C\varepsilon^{-2}$. Thus resampling each interval takes $O(n^2 \log \frac{1}{\varepsilon})$ time since we need to resample at most n symbols. For every interval, the edit distance can be computed using the Wagner-Fischer dynamic programming within $O(n^2 \log \frac{1}{\varepsilon})$ time. Theorem 4.2 shows that the expected number of re-sampling is $O(n)$. The algorithm will repeat until no bad interval can be found. Hence the overall expected running time is $O(n^5 \log \frac{1}{\varepsilon})$.

Note that without loss of generality we can assume that $\varepsilon > 1/\sqrt{n}$ because for smaller errors we can always use the indices directly, which have alphabet size n . So the overall expected running time is $O(n^5 \log n)$.

We can now construct an ε -synchronization circle using Theorem 4.1.

THEOREM 4.3. *For every $\varepsilon \in (0, 1)$ and every $n \in \mathbb{N}$, there exists an ε -synchronization circle S of length n over alphabet Σ of size $O(\varepsilon^{-2})$.*

Proof. First, by Theorem 4.1, we can have two ε -synchronization strings: S_1 with length $\lceil \frac{n}{2} \rceil$ over Σ_1 and S_2 with length $\lfloor \frac{n}{2} \rfloor$ over Σ_2 . Let $\Sigma_1 \cap \Sigma_2 = \emptyset$ and $|\Sigma_1| = |\Sigma_2| = O(\varepsilon^{-2})$. Let S be the concatenation of S_1 and S_2 . Then S is over alphabet $\Sigma = \Sigma_1 \cup \Sigma_2$ whose size is $O(\varepsilon^{-2})$. Now we prove that S is an ε -synchronization circle.

$\forall 1 \leq m \leq n$, consider string $S' = s_m, s_{m+1}, \dots, s_n, s_1, s_2, \dots, s_{m-1}$. Notice that for two strings T and T' over alphabet Σ , $LCS(T, T') \leq \frac{\varepsilon}{2}(|T| + |T'|)$ is equivalent to $ED(T, T') \geq (1 - \varepsilon)(|T| + |T'|)$. For any $i < j < k$, we call an interval $S'[i, k]$ good if $LCS(S'[i, j], S'[j+1, k]) \leq \frac{\varepsilon}{2}(k-i)$. It suffices to show that $\forall 1 \leq i, k \leq n$, the interval $S'[i, k]$ is good.

Without loss of generality let's assume $m \in [\lceil \frac{n}{2} \rceil, n]$.

Intervals which are substrings of S_1 or S_2 are good intervals, since S_1 and S_2 are ε -synchronization strings.

We are left with intervals crossing the ends of S_1 or S_2 .

If $S'[i, k]$ contains s_n, s_1 but doesn't contain $s_{\lceil \frac{n}{2} \rceil}$: If $j < n - m + 1$, then there's no common subsequence between $S'[i, j]$ and $S'[n - m + 2, k]$. Thus

$$LCS(S'[i, j], S'[j+1, k]) \leq \frac{\varepsilon}{2}(n - m + 1 - i) < \frac{\varepsilon}{2}(k - i)$$

If $j \geq n - m + 1$, then there's no common subsequence between $S'[j+1, k]$ and $S'[i, n - m + 1]$. Thus

$$LCS(S'[i, j], S'[j+1, k]) \leq \frac{\varepsilon}{2}(k - (n - m + 2)) < \frac{\varepsilon}{2}(k - i)$$

Thus intervals of this kind are good.

If $S'[i, k]$ contains $s_{\lfloor \frac{n}{2} \rfloor}, s_{\lceil \frac{n}{2} \rceil}$ but doesn't contain s_n : If $j \leq n - m + \lfloor \frac{n}{2} \rfloor + 1$, then there's no common subsequence between $S'[i, j]$ and $S'[n - m + \lceil \frac{n}{2} \rceil + 1, k]$, thus

$$LCS(S'[i, j], S'[j+1, k]) < \frac{\varepsilon}{2}(k - i)$$

If $j \geq n - m + \lfloor \frac{n}{2} \rfloor + 1$, then there's no common subsequence between $S'[j+1, k]$ and $S'[i, n - m + \lfloor \frac{n}{2} \rfloor + 1]$. Thus

$$LCS(S'[i, j], S'[j+1, k]) < \frac{\varepsilon}{2}(k - i)$$

Thus intervals of this kind are good.

If $S'[i, k]$ contains $s_{\lceil \frac{n}{2} \rceil}$ and s_n : If $n - m + 2 \leq j \leq n - m + \lfloor \frac{n}{2} \rfloor + 1$, then the common subsequence is either that of $S'[i, n - m + 1]$ and $S'[n - m + \lceil \frac{n}{2} \rceil + 1, k]$ or that of $S'[n - m + 2, j]$ and $S'[j+1, n - m + \lfloor \frac{n}{2} \rfloor + 1]$. This is because $\Sigma_1 \cap \Sigma_2 = \emptyset$. Thus

$$\begin{aligned} & LCS(S'[i, j], S'[j+1, k]) \\ & \leq \max\{LCS(S'[i, n - m + 1], S'[n - m + \lceil \frac{n}{2} \rceil + 1, k]), \\ & \quad LCS(S'[n - m + 2, j], S'[j+1, n - m + \lfloor \frac{n}{2} \rfloor + 1])\} \\ & < \frac{\varepsilon}{2}(k - i) \end{aligned}$$

If $j \leq n - m + 1$, then there's no common subsequence between $S'[i, j]$ and $S'[n - m + 2, n - m + \lfloor \frac{n}{2} \rfloor + 1]$. Thus

$$LCS(S'[i, j], S'[j+1, k]) < \frac{\varepsilon}{2}(k - i)$$

If $j \geq S'[n - m + \lceil \frac{n}{2} \rceil + 1]$, the proof is similar to the case where $j \leq n - m + 1$.

This shows that S' is an ε -synchronization string. Thus by the definition of synchronization circle, the construction gives an ε -synchronization circle.

5 Deterministic Constructions of Long-Distance Synchronization Strings

In this section, we give deterministic constructions of synchronization strings. In fact, we consider a generalized version of synchronization strings, i.e., $f(l)$ -distance ε -synchronization strings first defined by Haeupler and Shahrasbi [10]. Throughout this section, ε is considered to be a constant in $(0, 1)$.

5.1 Polynomial Time Constructions of Long-Distance Synchronization Strings

Here, by combining the deterministic Lovász local lemma of Chandrasekaran et al. [23] and the non-uniform sample space used in Theorem 4.1, we give a deterministic polynomial-time construction of c -long ε -synchronization strings over an alphabet of size $O(\varepsilon^{-2})$.

THEOREM 5.1. (THEOREM 2 IN [23]) *Let the time needed to compute the conditional probability $\Pr[A | \forall i \in I : P_i = v_i]$ for any $A \in \mathcal{A}$ and any partial evaluation $(v_i \in D_i)_{i \in I}$ where $I \subset [n]$ be at most t_C . If there is an $\epsilon \in (0, 1)$ and an assignment of reals $x : \mathcal{A} \mapsto (0, 1)$ such that*

$$\forall A \in \mathcal{A} : \Pr[A] \leq x'(A)^{1+\epsilon} = \left(x(A) \prod_{B \in \Gamma(A)} (1 - x(B)) \right)^{1+\epsilon}$$

Then, a deterministic algorithm can find a good evaluation in time

$$O(t_C) \cdot \frac{\hat{D} M^{2(1+1/\epsilon)} \log M}{\epsilon w_{\min}} = \tilde{O} \left(t_C \cdot \hat{D} \cdot M^{3+2/\epsilon} \right)$$

Here,

- $x'(A) = x(A) \prod_{B \in \Gamma(A)} (1 - x(B))$
- $\hat{D} := \max_{P \in \mathcal{P}} \{ \text{domain}(P) \}$ where \mathcal{P} is the set of random variables.
- $M := \max \left\{ n, 4m, 4 \sum_{a \in \bar{\mathcal{A}}} \frac{2|vbl(A)|}{x'(A)} \cdot \frac{x(A)}{1 - x(A)} \right\}$ where $\bar{\mathcal{A}} = \{ A \in \mathcal{A} | x'(A) \geq \frac{1}{4m} \}$, n is the number of random variables and m is the number of events in \mathcal{A} .

- $w_{min} := \min_{A \in \mathcal{A}} \{-\log x'(A)\}$
- $\Gamma(A) = \Gamma_{\mathcal{A}}(A)$ for the set of all events $B \neq A$ in \mathcal{A} with $vbl(B) \cap vbl(A) \neq \emptyset$, in which $vbl(A)$ is the smallest subset of random variables determining event A .

First we recall the following property of c -long synchronization strings.

LEMMA 5.1. (COROLLARY 4.4 OF [10]) *If S is a string which satisfies the edit distance requirement of c -long-distance ε -synchronization property stated in Definition 1.2 for any two non-adjacent intervals of total length $2c \log n$ or less, then it satisfies such requirement for all pairs of non-adjacent intervals.*

We now have the following theorem.

THEOREM 5.2. *For any $n \in \mathbb{N}$ and any constant $\varepsilon \in (0, 1)$, there is a deterministic construction of an $O(1/\varepsilon)$ -long-distance ε -synchronization string of length n , over an alphabet of size $O(\varepsilon^{-2})$, in time $\text{poly}(n)$.*

Proof. To prove this, we will use the Lovász local lemma and its deterministic algorithm in [23]. Suppose the alphabet is Σ with $|\Sigma| = q = c_1 \varepsilon^{-2}$ where c_1 is a constant. Let $t = c_2 \varepsilon^{-2}$ and $0 < c_2 < c_1$. We denote $|\Sigma| - t$ as q . The sampling algorithm of string S (1-index based) is as follows:

- Initialize an arbitrary order for Σ .
- For i th symbol:
 - Denote $h = \min\{t - 1, i - 1\}$. Generate a random variable P_i uniformly over $\{1, 2, \dots, |\Sigma| - h\}$.
 - Reorder Σ such that the previous h chosen symbols are labeled as the first h symbols in Σ , and the rest are ordered in the current order as the last $|\Sigma| - h$ symbols.
 - Choose the $(P_i + h)$ 'th symbol in this new order as $S[i]$.

Define the bad event A_{i_1, l_1, i_2, l_2} as intervals $S[i_1, i_1 + l_1]$ and $S[i_2, i_2 + l_2]$ violating the $c = O(1/\varepsilon)$ -long-distance synchronization string property for $i_1 + l_1 \leq i_2$. In other words, A_{i_1, l_1, i_2, l_2} occurs if and only if $ED(S[i_1, i_1 + l_1], S[i_2, i_2 + l_2]) \leq (1 - \varepsilon)(l_1 + l_2)$, which is equivalent to $LCS(S[i_1, i_1 + l_1], S[i_2, i_2 + l_2]) \geq \frac{\varepsilon}{2}(l_1 + l_2)$.

Note that according to the definition of c -long distance ε -synchronization string and Lemma 5.1, we only need to consider A_{i_1, l_1, i_2, l_2} where $l_1 + l_2 < c \log n$

and $c \log n \leq l_1 + l_2 \leq 2c \log n$. Thus we can upper bound the probability of A_{i_1, l_1, i_2, l_2} ,

$$\Pr[A_{i_1, l_1, i_2, l_2}] \leq \left(\frac{el}{\varepsilon l \sqrt{|\Sigma| - t}} \right)^{\varepsilon l} = \left(\frac{e}{\varepsilon \sqrt{|\Sigma| - t}} \right)^{\varepsilon l} = \hat{C}^{\varepsilon l},$$

where $l = l_1 + l_2$ and \hat{C} is a constant which depends on c_1 and c_2 .

However, to apply the deterministic Lovász local lemma (LLL), we need to have two additional requirements. The first requirement is that each bad event depends on up to logarithmically many variables, and the second is that the inequalities in the Lovász local lemma hold with a constant exponential slack [10].

The first requirement may not be true under the current definition of badness. Consider for example the random variables $P_{i_1}, \dots, P_{i_1 + l_1 - 1}, P_{i_2}, \dots, P_{i_2 + l_2 - 1}$ for a pair of split intervals $S[i_1, i_1 + l_1], S[i_2, i_2 + l_2]$ where the total length $l_1 + l_2$ is at least $2c \log n$. The event A_{i_1, l_1, i_2, l_2} may depend on too many random variables (i.e., $P_{i_1}, \dots, P_{i_2 + l_2 - 1}$).

To overcome this, we redefine the badness of the split interval $S[i_1, i_1 + l_1]$ and $S[i_2, i_2 + l_2]$ as follows: let B_{i_1, l_1, i_2, l_2} be the event that there exists $P_{i_1 + l_1}, \dots, P_{i_2 - 1}$ (i.e., the random variables chosen between the two intervals) such that the two intervals generated by $P_{i_1}, \dots, P_{i_1 + l_1 - 1}$ and $P_{i_2}, \dots, P_{i_2 + l_2 - 1}$ (together with $P_{i_1 + l_1}, \dots, P_{i_2 - 1}$) makes $LCS(S[i_1, i_1 + l_1], S[i_2, i_2 + l_2]) \geq \frac{\varepsilon}{2}(l_1 + l_2)$ according to the sampling algorithm. Note that if B_{i_1, l_1, i_2, l_2} does not happen, then certainly A_{i_1, l_1, i_2, l_2} does not happen.

Notice that with this new definition of badness, B_{i_1, l_1, i_2, l_2} is independent of $\{P_{i_1 + l_1}, \dots, P_{i_2 - 1}\}$ and only depends on $\{P_{i_1}, \dots, P_{i_1 + l_1 - 1}, P_{i_2}, \dots, P_{i_2 + l_2 - 1}\}$. In particular, this implies that B_{i_1, l_1, i_2, l_2} is independent of the badness of all other intervals which have no intersection with $(S[i_1, i_1 + l_1], S[i_2, i_2 + l_2])$.

We now bound $\Pr[B_{i_1, l_1, i_2, l_2}]$. When considering the two intervals $S[i_1, i_1 + l_1]$, $S[i_2, i_2 + l_2]$ and their edit distance under our sampling algorithm, without loss of generality we can assume that the order of the alphabet at the point of sampling $S[i_1]$ is $(1, 2, \dots, q)$ just by renaming the symbols. Now, if we fix the order of the alphabet at the point of sampling $S[i_2]$ in our sampling algorithm, then $S[i_2, i_2 + l_2]$ only depends on $\{P_{i_2}, \dots, P_{i_2 + l_2}\}$ and thus $LCS(S[i_1, i_1 + l_1], S[i_2, i_2 + l_2])$ only depends on $\{P_{i_1}, \dots, P_{i_1 + l_1 - 1}, P_{i_2}, \dots, P_{i_2 + l_2}\}$.

Conditioned on any fixed order of the alphabet at the point of sampling $S[i_2]$, we have that $LCS(S[i_1, i_1 + l_1], S[i_2, i_2 + l_2]) \geq \frac{\varepsilon}{2}(l_1 + l_2)$ happens with probability at most $\hat{C}^{\varepsilon l}$ by the same computation as we upper bound $\Pr[A_{i_1, l_1, i_2, l_2}]$. Note that there are at most $q!$ different

orders of the alphabet. Thus by a union bound we have

$$\Pr[B_{i_1, l_1, i_2, l_2}] \leq \hat{C}^{\varepsilon l} \times q! = C^{\varepsilon l},$$

for some constant C .

In order to meet the second requirement of the deterministic algorithm of Lovász local lemma, we also need to find real numbers $x_{i_1, i_1 + l_1, i_2, i_2 + l_2} \in [0, 1]$ such that for any B_{i_1, l_1, i_2, l_2} ,

$$\Pr[B_{i_1, l_1, i_2, l_2}] \leq \left[x_{i_1, l_1, i_2, l_2} \prod_X (1 - x_{i'_1, l'_1, i'_2, l'_2}) \right]^{1.01}.$$

where X denotes the condition that $[S[i_1, i_1 + l_1] \cup S[i_2, i_2 + l_2]] \cap [S[i'_1, i'_1 + l'_1] \cup S[i'_2, i'_2 + l'_2]] \neq \emptyset$. We propose $x_{i_1, l_1, i_2, l_2} = D^{-\varepsilon(l_1 + l_2)}$ for some $D > 1$ to be determined later. D has to be chosen such that for any i_1, l_1, i_2, l_2 and $l = l_1 + l_2$:

$$\begin{aligned} & \left(\frac{e}{\varepsilon \sqrt{|\Sigma| - t}} \right)^{\varepsilon l} \\ & \leq \left[D^{-\varepsilon l} \prod_X (1 - D^{-\varepsilon(l'_1 + l'_2)}) \right]^{1.01} \end{aligned}$$

Notice that

$$\begin{aligned} & D^{-\varepsilon l} \prod_X (1 - D^{-\varepsilon(l'_1 + l'_2)}) \\ & \geq D^{-\varepsilon l} \\ & \times \prod_{l' = c \log n}^{2c \log n} \prod_{l'_1 = 1}^{l'} (1 - D^{-\varepsilon l'})^{\tilde{\ell}_n} \times \prod_{l'' = t}^{c \log n} (1 - D^{-\varepsilon l''})^{l+l''} \\ & = D^{-\varepsilon l} \prod_{l' = c \log n}^{2c \log n} \prod_{l'_1 = 1}^{l'} (1 - D^{-\varepsilon l'})^{4(l+l')n} \\ & \times \prod_{l'' = t}^{c \log n} (1 - D^{-\varepsilon l''})^{l+l''} \\ & \geq D^{-\varepsilon l} (1 - 32c^3 n \log^3 n D^{-\varepsilon c \log n}) \left[1 - \frac{D^{-c_2 \varepsilon^{-1}}}{1 - D^{-\varepsilon}} \right]^l \\ & \times \left(1 - \frac{D^{-c_2/\varepsilon} (D^{-\varepsilon} + c_2/\varepsilon^2 - c_2 D^{-\varepsilon}/\varepsilon^2)}{(1 - D^{-\varepsilon})^2} \right) \end{aligned}$$

where in the inequalities,

$$\tilde{\ell} = [(l_1 + l'_1) + (l_1 + l'_2) + (l_2 + l'_1) + (l_2 + l'_2)]$$

For $D = 2$ and $c = 2/\varepsilon$, $\lim_{\varepsilon \rightarrow 0} \frac{2^{-c_2/\varepsilon}}{1 - 2^{-\varepsilon}} = 0$. Thus, for sufficiently small ε , $\frac{2^{-c_2/\varepsilon}}{1 - 2^{-\varepsilon}} < \frac{1}{2}$. Moreover,

$32c^2 n \log^2 n D^{-\varepsilon l'} = \frac{2^8 \log^3 n}{\varepsilon^3 n} = o(1)$. Finally, for sufficiently small ε , $1 - \frac{D^{-c_2/\varepsilon} (D^{-\varepsilon} + c_2/\varepsilon^2 - c_2 D^{-\varepsilon}/\varepsilon^2)}{(1 - D^{-\varepsilon})^2} > 2^{-\varepsilon}$. Therefore, for sufficiently small ε and sufficiently large n , the requirement for LLL is satisfied under the condition:

$$\begin{aligned} & D^{-\varepsilon l} \prod_X (1 - D^{-\varepsilon(l'_1 + l'_2)}) \\ & \geq 2^{-\varepsilon l} \left(1 - \frac{1}{2} \right) (2^{-\varepsilon})^l \left(1 - \frac{1}{2} \right) \geq \frac{4^{-\varepsilon l}}{4} \end{aligned}$$

So for LLL to work, the following should be guaranteed:

$$\left(\frac{e}{\varepsilon \sqrt{|\Sigma| - t}} \right)^{\frac{\varepsilon l}{1.01}} \leq \frac{4^{-\varepsilon l}}{4} \Leftrightarrow \frac{4^{2.02(1+\varepsilon)e^2}}{\varepsilon^2} \leq |\Sigma| - t$$

Hence the second requirement holds for $|\Sigma| - t = \frac{4^{4.04} \varepsilon^2}{\varepsilon^2} = O(\varepsilon^{-2})$.

COROLLARY 5.1. *For any $n \in \mathbb{N}$ and any constant $\varepsilon \in (0, 1)$, there is a deterministic construction of an ε -synchronization string of length n , over an alphabet of size $O(\varepsilon^{-2})$, in time $\text{poly}(n)$.*

By a similar concatenation construction used in the proof of Theorem 4.3, we also have a deterministic construction for synchronization circles.

COROLLARY 5.2. *For any $n \in \mathbb{N}$ and any constant $\varepsilon \in (0, 1)$, there is a deterministic construction of an ε -synchronization circle of length n , over an alphabet of size $O(\varepsilon^{-2})$, in time $\text{poly}(n)$.*

5.2 Deterministic linear time constructions of c -long distance ε -synchronization string Here we give a much more efficient construction of a c -long distance ε -synchronization string, using synchronization circles and standard error correcting codes. We show that the following algorithm gives a construction of c -long distance synchronization strings.

ALGORITHM 5.1. [Explicit Linear Time Construction of c -long distance ε -synchronization string]

Input:

- An error correcting code $\hat{\mathcal{C}} \subset \Sigma_{\hat{\mathcal{C}}}^m$, with distance δm and block length $m = c \log n$.
- An ε_0 -synchronization circle $SC = (sc_1, \dots, sc_m)$ of length m over alphabet Σ_{SC} .

Operations:

- Construct a code $\mathcal{C} \subset \Sigma^m$ such that

$$\mathcal{C} = \{((\hat{c}_1, sc_1), \dots, (\hat{c}_m, sc_m)) | (\hat{c}_1, \dots, \hat{c}_m) \in \hat{\mathcal{C}}\}$$

where $\Sigma = \Sigma_{\hat{\mathcal{C}}} \times \Sigma_{SC}$.

- Let S be concatenation of all codewords $\mathcal{C}_1, \dots, \mathcal{C}_N$ from \mathcal{C} .

Output: S .

To prove the correctness, we first recall the following theorem from [1].

THEOREM 5.3. (THEOREM 4.2 OF [1]) *Given an ε_0 -synchronization string S with length n , and an efficient error correcting code \mathcal{C} with block length n , that corrects up to $n\delta^{\frac{1+\varepsilon_0}{1-\varepsilon_0}}$ half errors, one can obtain an insertion/deletion code \mathcal{C}' that can be decoded from up to $n\delta$ deletions, where $\mathcal{C}' = \{(c'_1, \dots, c'_n) \mid \forall i \in [n], c'_i = (c_i, S[i])\}, (c_1, \dots, c_n) \in \mathcal{C}\}.$*

We have the following property of longest common subsequence.

LEMMA 5.2. Suppose T_1 is the concatenation of ℓ_1 strings, $T_1 = S_1 \circ \dots \circ S_{\ell_1}$ and T_2 is the concatenation of ℓ_2 strings, $T_2 = S'_1 \circ \dots \circ S'_{\ell_2}$. If there exists an integer t such that for all i, j , we have $LCS(S_i, S'_j) \leq t$, then we have $LCS(T_1, T_2) \leq (\ell_1 + \ell_2)t$.

Proof. We rename the strings in T_2 by $S_{\ell_1+1}, \dots, S_{\ell_1+\ell_2}$. Suppose the longest common subsequence between T_1 and T_2 is \tilde{T} , which can be viewed as a matching between T_1 and T_2 .

we can divide \tilde{T} sequentially into disjoint intervals, where each interval corresponds to a common subsequence between a different pair of strings (S_i, S_j) , where S_i is from T_1 and S_j is from T_2 . In addition, if we look at the intervals from left to right, then for any two consecutive intervals and their corresponding pairs (S_i, S_j) and $(S_{i'}, S_{j'})$, we must have $i' \geq i$ and $j' \geq j$ since the matchings which correspond to two intervals cannot cross each other. Furthermore either $i' > i$ or $j' > j$ as the pair (S_i, S_j) is different from $(S_{i'}, S_{j'})$.

Thus, starting from the first interval, we can label each interval with either i or j such that every interval receives a different label, as follows. We label the first interval using either i or j . Then, assuming we have already labeled some intervals and now look at the next interval. Without loss of generality assume that the previous interval is labeled using i , now if the current $i' > i$ then we can label the current interval using i' ; otherwise we must have $j' > j$ so we can label the current interval using j' . Thus the total number of the labels is at most $l_1 + l_2$, which means the total number of the intervals is also at most $l_1 + l_2$. Note that each interval has length at most t , therefore we can upper bound $LCS(T_1, T_2)$ by $(l_1 + l_2)t$.

LEMMA 5.3. *The output S in Algorithm 5.1 is an ε_1 -synchronization circle, where $\varepsilon_1 \leq 10(1 - \frac{1-\varepsilon_0}{1+\varepsilon_0}\delta)$.*

Proof. Suppose $\hat{\mathcal{C}}$ can correct up to δm half errors. Then according to Theorem 5.3, \mathcal{C} can correct up to $\frac{1-\varepsilon_0}{1+\varepsilon_0}\delta m$ deletions.

Let $\alpha = 1 - \frac{1-\varepsilon_0}{1+\varepsilon_0} \delta$. Notice that \mathcal{C} has the following properties:

1. $LCS(\mathcal{C}) = \max_{c_1, c_2 \in \mathcal{C}} LCS(c_1, c_2) \leq \alpha m$
2. Each codeword in \mathcal{C} is an ε -synchronization circle over Σ .

Consider any shift of the start point of S , we only need to prove that $\forall 1 \leq i < j < k \leq n, \text{LCS}(S[i, j], S[j + 1, k]) < \frac{\varepsilon_1}{2}(k - i)$.

Suppose $S_1 = S[i, j]$ and $S_2 = S[j + 1, k]$. Let $\varepsilon_1 = 10\alpha$.

Case 1: $k - i > m$. Let $|S_1| = s_1$ and $|S_2| = s_2$, thus $s_1 + s_2 > m$. If we look at each S_h for $h = 1, 2$, then S_h can be divided into some consecutive codewords, plus at most two incomplete codewords at both ends. In this sense each S_h is the concatenation of ℓ_h strings with $\ell_h < \frac{s_h}{m} + 2$. An example of the worst case appears in Figure 1.

$$\begin{array}{ccccccc}
 & c_1 & & c_0 & & c_3 & \\
 \text{S} & 1 & 2 & 3 & \overset{5}{\cancel{6}} & 1 & 5 & 6 & 2 & 4 & 3 & \overset{1}{\cancel{4}} & 5 & 6 \\
 & i & & & & j & & & & & k & & &
 \end{array}$$

Figure 1: Example of the worst case, where j splits a codeword, and there are two incomplete codewords at both ends.

Now consider the longest common subsequence between any pair of these strings where one is from S_1 and the other is from S_2 , we claim that the length of any such longest common subsequence is at most αm . Indeed, if the pair of strings are from two different codewords, then by the property of the code \mathcal{C} we know the length is at most αm . On the other hand, if the pair of strings are from a single codeword (this happens when j splits a codeword, or when $S[i]$ and $S[k]$ are in the same codeword), then they must be two disjoint intervals within a codeword. In this case, by the property that any codeword is also a synchronization circle, the length of the longest common subsequence of this pair is at most $\frac{\varepsilon_0}{2} m$.

Note that $\alpha = 1 - \frac{1-\varepsilon_0}{1+\varepsilon_0} \delta \geq 1 - \frac{1-\varepsilon_0}{1+\varepsilon_0} = \frac{2\varepsilon_0}{1+\varepsilon_0} \geq \varepsilon_0$ (since $\delta, \varepsilon_0 \in (0, 1)$). Thus $\frac{\varepsilon_0}{2}m < \alpha m$. Therefore, by Lemma 5.2, we have

$$\begin{aligned}
& LCS(S_1, S_2) \\
& < (\frac{s_1}{m} + 2 + \frac{s_2}{m} + 2)\alpha m \\
(5.1) \quad & = \alpha(s_1 + s_2 + 4m) \\
& < 5\alpha(s_1 + s_2) \\
& = 5\alpha(k - i) = \frac{\varepsilon_1}{2}(k - i)
\end{aligned}$$

Case 2: If $k - i \leq m$, then according to the property of synchronization circle SC , we know that the longest common subsequence of S_1 and S_2 is less than $\frac{\varepsilon_0}{2}(k - i) \leq \alpha(k - i) \leq \frac{\varepsilon_1}{2}(k - i)$.

As a result, the longest common subsequence of $S[i, j]$ and $S[j + 1, k]$ is less than $\frac{\varepsilon_1}{2}(k - i)$, which means that S is an ε_1 -synchronization circle.

Similarly, we also have the following lemma.

LEMMA 5.4. *The output S of algorithm 5.1 is a c -long distance ε -synchronization string of length $n = Nm$ where N is the number of codewords in \mathcal{C} , $\varepsilon = 12(1 - \frac{1-\varepsilon_0}{1+\varepsilon_0}\delta)$.*

Proof. By Lemma 5.3, S is an ε_1 -synchronization string, thus the length of longest common subsequence for adjacent intervals S_1, S_2 with total length $l < c \log n$ is less than $\frac{\varepsilon_1}{2}l$. We only need to consider pair of intervals S_1, S_2 whose total length $l \in [c \log n, 2c \log n]$.

Notice that the total length of S_1 and S_2 is at most $2c \log n$, which means that S_1 and S_2 each intersects with at most 3 codewords from \mathcal{C} . Using Lemma 5.2, we have that $LCS(S_1, S_2) \leq 6\alpha l$.

Thus picking $\varepsilon = \max\{12\alpha, \varepsilon_1\} = 12\alpha = 12(1 - \frac{1-\varepsilon_0}{1+\varepsilon_0}\delta)$, S from algorithm 5.1 is a c -long distance ε -synchronization circle.

We need the following code constructed by Guruswami and Indyk [26].

LEMMA 5.5. (THEOREM 3 OF [26]) *For every $0 < r < 1$, and all sufficiently small $\varepsilon > 0$, there exists a family of codes of rate r and relative distance $(1 - r - \varepsilon)$ over an alphabet of size $2^{O(\varepsilon^{-4}r^{-1}\log(1/\varepsilon))}$ such that codes from the family can be encoded in linear time and can also be uniquely decoded in linear time from $2(1 - r - \varepsilon)$ fraction of half errors.*

LEMMA 5.6. [Error Correcting Code by Brute Force Search]

For any $n \in \mathbb{N}$, any $\varepsilon \in [0, 1]$, one can construct an error correcting code in time $O(2^{\varepsilon n}(\frac{2e}{\varepsilon})^n n \log(1/\varepsilon))$ and space $O(2^{\varepsilon n} n \log(1/\varepsilon))$, with block length n , number of codewords $2^{\varepsilon n}$, distance $d = (1 - \varepsilon)n$, alphabet size $2e/\varepsilon$.

Proof. We conduct a brute-force search here to find all the codewords one by one.

We denote the code as \mathcal{C} and the alphabet as Σ . Let $|\Sigma| = q$. At first, let $\mathcal{C} = \emptyset$. Then we add an arbitrary element in Σ^n to \mathcal{C} . Every time after a new element C is added to \mathcal{C} , we exclude every such element in Σ^n that has distance less than d from C . Then we pick an arbitrary one from the remaining elements, adding it to \mathcal{C} . Keep doing this until $|\mathcal{C}| = 2^{\varepsilon n}$.

Note that given $C \in \Sigma^n$, the total number of elements that have distance less than d to C , is at most $\binom{n}{d}q^d = \binom{n}{(n-d)}q^d \leq (\frac{e}{\varepsilon})^{\varepsilon n}q^{(1-\varepsilon)n}$. We have to require that $|\mathcal{C}|(\frac{e}{\varepsilon})^{\varepsilon n}q^{(1-\varepsilon)n} \leq q^n$. Let $q = 2e/\varepsilon$. So \mathcal{C} can be $2^{\varepsilon n}$.

The exclusion operation takes time $O((\frac{2e}{\varepsilon})^n n \log(1/\varepsilon))$ as we have to exhaustively search the space and for each word we have to compute its hamming distance to the new added code word. Since there are $2^{\varepsilon n}$ code words, the time complexity is as stated.

We have to record those code words, so the space complexity is also as stated.

We can now use Algorithm 5.1 to give a linear time construction of c -long distance ε -synchronization strings.

THEOREM 5.4. *For every $n \in \mathbb{N}$ and any constant $0 < \varepsilon < 1$, there is a deterministic construction of a $c = O(\varepsilon^{-2})$ -long distance ε -synchronization string $S \in \Sigma^n$ where $|\Sigma| = O(\varepsilon^{-3})$, in time $O(n)$. Moreover, $S[i, i + \log n]$ can be computed in $O(\frac{\log n}{\varepsilon^2})$ time.*

Proof. Suppose we have an error correcting code $\hat{\mathcal{C}}$ with distance rate $1 - \varepsilon' = \frac{1 + \frac{\varepsilon}{36}}{1 - \frac{\varepsilon}{36}}(1 - \frac{\varepsilon}{12})$, message rate $r_c = O(\varepsilon'^2)$, over an alphabet of size $|\Sigma_c| = O(\varepsilon'^{-1})$, with block length $m = O(\varepsilon'^{-2} \log n)$. Let $c = O(\varepsilon'^{-2}) = O(\varepsilon^{-2})$. We apply Algorithm 5.1, using $\hat{\mathcal{C}}$ and an $\frac{\varepsilon}{36}$ -synchronization circle SC of length m over an alphabet of size $O(\varepsilon^{-2})$. Here SC is constructed by Corollary 5.2 in time $\text{poly}(m) = \text{poly}(\log n)$. By Lemma 5.4, we have a c -long-distance $12(1 - \frac{1 - \varepsilon/36}{1 + \varepsilon/36}(1 - \varepsilon')) = \varepsilon$ -synchronization string of length $m \cdot |\Sigma_c|^{r_c m} \geq n$.

It remains to show that we can have such a $\hat{\mathcal{C}}$ with linear time encoding. We use the code in Lemma 5.5 as the outer code and the one in Lemma 5.6 as inner code. Let \mathcal{C}_{out} be an instantiation of the code in Lemma 5.5 with rate $r_o = \varepsilon_o = \frac{1}{3}\varepsilon'$, relative distance $d_o = (1 - 2\varepsilon_o)$ and alphabet size $2^{O(\varepsilon_o^{-5} \log(1/\varepsilon_o))}$, and block length $n_o = \frac{\varepsilon_o^4 \log n}{\log(1/\varepsilon_o)}$, which is encodable and decodable in linear time.

Further, according to Lemma 5.6 one can find a code \mathcal{C}_{in} with rate $r_i = O(\varepsilon_i)$ where $\varepsilon_i = \frac{1}{3}\varepsilon'$, relative

distance $1 - \varepsilon_i$, over an alphabet of size $\frac{2\varepsilon}{\varepsilon_i}$, and block length $n_i = O(\varepsilon_i^{-6} \log(1/\varepsilon_i))$. Note that since the block length and alphabet size are both constant because ε is a constant. So the encoding can be done in constant time.

Concatenating \mathcal{C}_{out} and \mathcal{C}_{in} gives the desire code $\hat{\mathcal{C}}$ with rate $O(\varepsilon'^2)$, distance $1 - O(\varepsilon')$ and alphabet of size $O(\varepsilon'^{-1})$ and block length $O(\varepsilon'^{-2} \log n)$. Moreover, the encoding of $\hat{\mathcal{C}}$ can be done in linear time, because the encoding of \mathcal{C}_{out} is in linear time and the encoding of \mathcal{C}_{in} is in constant time.

Note that since every codeword of $\hat{\mathcal{C}}$ can be computed in time $O(\frac{\log n}{\varepsilon^2})$, $S[i, i + \log n]$ can be computed in $O(\frac{\log n}{\varepsilon^2})$ time.

COROLLARY 5.3. *For every $n \in \mathbb{N}$ and any constant $0 < \varepsilon < 1$, there is a deterministic construction of an ε -synchronization string $S \in \Sigma^n$ where $|\Sigma| = O(\varepsilon^{-3})$, in time $O(n)$. Moreover, $S[i, i + \log n]$ can be computed in $O(\frac{\log n}{\varepsilon^2})$ time.*

5.3 Explicit Constructions of Infinite Synchronization Strings

In this section we focus on the construction of infinite synchronization strings.

To measure the efficiency of the construction of an infinite string, we consider the time complexity for computing the first n elements of that string. An infinite synchronization string is said to have an explicit construction if there is an algorithm that computes any position $S[i]$ in time $\text{poly}(i)$. Moreover, it is said to have a highly-explicit construction if there is an algorithm that computes any position $S[i]$ in time $O(\log i)$.

We have the following algorithm.

ALGORITHM 5.2. [Construction of infinite ε -synchronization string]

Input:

- A constant $\varepsilon \in (0, 1)$.

Operations:

- Let $q \in \mathbb{N}$ be the size of an alphabet large enough to construct an $\frac{\varepsilon}{2}$ -synchronization string. Let Σ_1 and Σ_2 be two alphabets of size q such that $\Sigma_1 \cap \Sigma_2 = \emptyset$.
- Let $k = \frac{4}{\varepsilon}$. For $i = 1, 2, \dots$, construct an $\frac{\varepsilon}{2}$ -synchronization string S_{k^i} of length k^i , where S_{k^i} is over Σ_1 if i is odd and over Σ_2 otherwise.
- Let S be the sequential concatenation of $S_k, S_{k^2}, S_{k^3}, \dots, S_{k^t}, \dots$

Output: S .

LEMMA 5.7. *If there is a construction of $\frac{\varepsilon}{2}$ -synchronization strings with alphabet size q , then*

Algorithm 5.2 constructs an infinite ε -synchronization string with alphabet size $2q$.

Proof. Algorithm 5.2 can be shown as in the figure below.

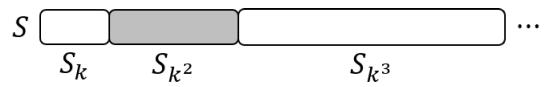


Figure 2: S_k and S_{k^3} are over alphabet Σ_1 and S_{k^2} is over Σ_2 .

Now we show that S is an infinite ε -synchronization string.

CLAIM 2. *Let $x < y < z$ be positive integers and let t be such that $k^t \leq |S[x, z]| < k^{t+1}$. Then $ED(S[x, y], S[y, z]) \geq (1 - \frac{\varepsilon}{2})(z - x)(1 - \frac{2}{k})$.*

Proof. Let l_i be the index of S where $S_{k^{i+1}}$ starts. Then $l_i = \sum_{j=1}^i k^j = \frac{k^{i+1} - k}{k - 1}$. Notice that $l_{t-1} < 2k^{t-1}$ and $|S[x, z]| \geq k^t$, one can throw away all elements of $S[x, z]$ whose indices are less than l_{t-1} without losing more than $\frac{2k^{t-1}}{k^t} = \frac{2}{k}$ fraction of the elements of $S[x, z]$. We use $S[x', z]$ to denote the substring after throwing away the symbols before l_{t-1} . Thus $x' \geq l_{t-1}$.

Since $x' \geq l_{t-1}$, $S[x', z]$ either entirely falls into a synchronization string S_{k^t} or crosses two synchronization strings S_{k^t} and $S_{k^{t+1}}$ over two entirely different alphabets Σ_1 and Σ_2 . Thus the edit distance of $S[x', y]$ and $S[y, z]$ is at least $(1 - \frac{\varepsilon}{2})(z - x)$.

Since $k = \frac{4}{\varepsilon}$, we have that

$$\begin{aligned} & ED(S[x, y], S[y, z]) \\ & \geq \left(1 - \frac{\varepsilon}{2}\right)(z - x) \left(1 - \frac{2}{k}\right) \\ & \geq (1 - \varepsilon)(z - x) \end{aligned}$$

This shows that S is an ε -synchronization string.

If we instantiate Algorithm 5.2 using Corollary 5.1, then we have the following theorem.

THEOREM 5.5. *For any constant $0 < \varepsilon < 1$, there exists an explicit construction of an infinite ε -synchronization string S over an alphabet of size $O(\varepsilon^{-2})$.*

Proof. We combine Algorithm 5.2 and Corollary 5.1. In the algorithm, we can construct every substring S_{k^i} in polynomial time with alphabet size $q = O(\varepsilon^{-2})$, by Corollary 5.1. So the first n symbols of S can be computed in polynomial time.

By Lemma 5.7, S is an infinite ε -synchronization string over an alphabet of size $2q = O(\varepsilon^{-2})$.

If we instantiate Algorithm 5.2 using Corollary 5.3, then we have the following theorem.

THEOREM 5.6. *For any constant $0 < \varepsilon < 1$, there exists a highly-explicit construction of an infinite ε -synchronization string S over an alphabet of size $O(\varepsilon^{-3})$. Moreover, for any $i \in \mathbb{N}$, the first i symbols can be computed in $O(i)$ time and $S[i, i + \log i]$ can be computed in $O(\log i)$ time.*

Proof. Combine Algorithm 5.2 and Corollary 5.3. In the algorithm, we can construct every substring S_{k^i} in linear time with alphabet size $q = O(\varepsilon^{-3})$, by Corollary 5.3. So the first i symbols can be computed in $O(i)$ time. Also any substring $S[i, i + \log i]$ can be computed in time $O(\log i)$.

By Lemma 5.7, S is an infinite ε -synchronization string over an alphabet of size $2q = O(\varepsilon^{-3})$.

6 $\Omega(\varepsilon^{-3/2})$ Lower-Bound on Alphabet Size

The *twin word* problem was introduced by Axenovich, Person, and Puzynina [27] and further studied by Bukh and Zhou [28]. Any set of two identical disjoint subsequences in a given string is called a twin word. [27, 28] provided a variety of results on the relations between the length of a string, the size of the alphabet over which it is defined, and the size of the longest twin word it contains. We will make use of the following result from [28] that is built upon Lemma 5.9 from [29] to provide a new lower-bound on the alphabet size of synchronization strings.

THEOREM 6.1. (THEOREM 3 FROM [28]) *There exists a constant c so that every word of length n over a q -letter alphabet contains two disjoint equal subsequences of length $cnq^{-2/3}$.*

Further, Theorem 6.4 of [1] states that any ε -synchronization string of length n has to satisfy ε -self-matching property which essentially means that it cannot contain two equal (not necessarily disjoint) subsequences of length εn or more. These two requirements lead to the following inequality for an ε -synchronization string of length n over an alphabet of size q .

$$cnq^{-2/3} \leq \varepsilon n \Rightarrow c' \varepsilon^{-3/2} \leq q$$

7 Synchronization Strings over Small Alphabets

In this section, we focus on synchronization strings over small constant-sized alphabets. We study the question of what is the smallest possible alphabet size over which arbitrarily long ε -synchronization strings can exist for some $\varepsilon < 1$, and how such synchronization strings can be constructed.

Throughout this section, we will make use of square-free strings introduced by Thue [14], which is a weaker notion than synchronization strings that requires all consecutive equal-length substrings to be non-identical. Note that no synchronization strings or square-free strings of length four or more exist over a binary alphabet since a binary string of length four either contains two consecutive similar symbols or two identical consecutive substrings of length two. However, for ternary alphabets, arbitrarily long square-free strings exist and can be constructed efficiently using *uniform morphism* [20]. In Section 7.1, we will briefly review this construction and show that no uniform morphism can be used to construct arbitrary long synchronization strings. In Section 7.2, we make use of ternary square-free strings to show that arbitrarily long ε -synchronization strings exist over alphabets of size four for some $\varepsilon < 1$.

7.1 Morphisms cannot Generate Synchronization Strings Previous works show that one can construct infinitely long square-free or approximate-square-free strings using *uniform morphisms*. A uniform morphism of rank r over an alphabet Σ is a function $\phi : \Sigma \rightarrow \Sigma^r$ that maps any symbol out of an alphabet Σ to a string of length r over the same alphabet. Applying the function ϕ over some string $S \in \Sigma^*$ is defined as replacing each symbol of S with $\phi(S)$.

[21, 15, 16, 17, 20] show that there are uniform morphisms that generate the combinatorial objects they study respectively. More specifically, one can start from any letter of the alphabet and repeatedly apply the morphism on it to construct those objects. For instance, using the uniform morphisms of rank 11 suggested in [20], all such strings will be square-free. In this section, we investigate the possibility of finding similar constructions for synchronization strings. We will show that no such morphism can possibly generate an infinite ε -synchronization strings for any fixed $0 < \varepsilon < 1$.

The key to this claim is that a matching between two substrings is preserved under an application of the uniform morphism ϕ . Hence, we can always increase the size of a matching between two substrings by applying the morphism sufficiently many times, and then adding new matches to the matching from previous steps.

THEOREM 7.1. *Let ϕ be a uniform morphism of rank r over alphabet Σ . Then ϕ does not generate an infinite ε -synchronization string, for any $0 < \varepsilon < 1$.*

Proof. To prove this, we show that for any $0 < \varepsilon < 1$, applying morphism ϕ sufficiently many times over any symbol of alphabet Σ produces a string that has two neighboring intervals which contradict ε -

synchronization property. First, we claim that, without loss of generality, it suffices to prove this for morphisms ϕ for which $\phi(\sigma)$ contains all elements of Σ for any $\sigma \in \Sigma$. To see this, consider the graph G with $|\Sigma|$ vertices where each vertex corresponds to a letter of the alphabet and there is a (σ_1, σ_2) edge if $\phi(\sigma_1)$ contains σ_2 . It is straightforward to verify that after applying morphism ϕ over a letter sufficiently many times, the resulting string can be split into a number of substrings so that the symbols in any of them belong to a subset of Σ that corresponds to some strongly connected component in G . As ε -synchronization string property is a hereditary property over substrings, this gives that one can, without loss of generality, prove the above-mentioned claim for morphisms ϕ for which the corresponding graph G is strongly connected. Further, let d be the greatest common divisor of the size of all cycles in G . One can verify that, for some sufficiently large k , ϕ^{kd} will be a morphism that, depending on the letter σ to perform recursive applications of the morphism on, will always generate strings over some alphabet Σ_σ and $\phi^{kd}(\sigma')$ contains all symbols of Σ_σ for all $\sigma' \in \Sigma_\sigma$. As proving the claim for ϕ^{kd} implies it for ϕ as well, the assumption mentioned above does not harm the generality.

We now proceed to prove that for any morphism ϕ of rank r as described above, any positive integer $n \in \mathbb{N}$, and any positive constant $0 < \delta < 1$, there exists $m \in \mathbb{N}$ so that

$$LCS(\phi^m(a), \phi^m(b)) \geq \left[1 - \left(1 - \frac{1}{|\Sigma|^2 r} \right)^n - \delta \right] \cdot r^m$$

for any $a, b \in \Sigma$ where ϕ^m represents m consecutive applications of morphism ϕ and $LCS(., .)$ denotes the longest common substring.

Having such claim proved, one can take $\delta = (1 - \varepsilon)/2$ and n large enough so that m applications of ϕ over any pair of symbols entail strings with a longest common substring that is of a fraction larger than $1 - (1 - \varepsilon) = \varepsilon$ in terms of the length of those strings. Then, for any string $S \in \Sigma^*$, one can take two arbitrary consecutive symbols of $\phi(S)$ like $S[i]$ and $S[i + 1]$. Applying morphism ϕ , m more times on $\phi(S)$ makes the corresponding intervals of $\phi^{m+1}(S)$ have an edit distance that is smaller than $1 - \varepsilon$ fraction of their combined lengths. This shows that $\phi^{m+1}(S)$ is not an ε -synchronization string and finishes the proof.

Finally, we prove the claim by induction on n . For the base case of $n = 1$, given the assumption of all members of Σ appearing in $\phi(\sigma)$ for all $\sigma \in \Sigma$, $\phi(a)$ and $\phi(b)$ have a non-empty common subsequence. This

gives that

$$\begin{aligned} & LCS(\phi(a), \phi(b)) \geq 1 \\ & = \left[1 - \left(1 - \frac{1}{r} \right) \right] \cdot r > \left[1 - \left(1 - \frac{1}{|\Sigma|^2 r} \right) - \delta \right] \cdot r \end{aligned}$$

Therefore, choosing $m = 1$ finishes the induction base.

We now prove the induction step. Note that by induction hypothesis, for some given n , one can find m_1 such that

$$LCS(\phi^{m_1}(a), \phi^{m_1}(b)) \geq \left[1 - \left(1 - \frac{1}{|\Sigma|^2 r} \right)^n - \frac{\delta}{2} \right] \cdot r^{m_1}.$$

Now, let $m_2 = \lceil \log_r \frac{2}{\delta} \rceil$. Consider $\phi^{m_2}(a)$ and $\phi^{m_2}(b)$. Note that among all possible pairs of symbols from Σ^2 , one appears at least $\frac{r^{m_2}}{|\Sigma|^2}$ times in respective positions of $\phi^{m_2}(a)$ and $\phi^{m_2}(b)$. Let (a', b') be such pair. As $\phi(a')$ and $\phi(b')$ contain all symbols of Σ , one can take one specific occurrence of a fixed arbitrary symbol $\sigma \in \Sigma$ in all appearances of the pair $\phi(a')$ and $\phi(b')$ to find a common subsequence of size $\frac{r^{m_2}}{|\Sigma|^2} = \frac{r^{m_2+1}}{|\Sigma|^2 r}$ or more between $\phi^{m_2+1}(a)$ and $\phi^{m_2+1}(b)$ (See Figure 3).

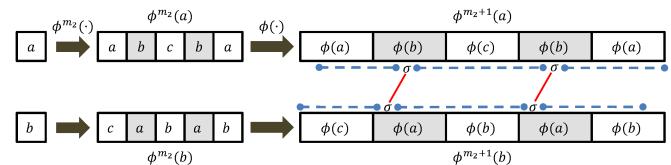


Figure 3: Induction step in Theorem 7.1; Most common pair $(a', b') = (b, a)$.

Note that one can apply the morphism ϕ further times over $\phi^{m_2+1}(a)$ and $\phi^{m_2+1}(b)$ and such common subsequence will still be preserved; However, one might be able to increase the size of it by adding new elements to the common subsequence from equal length pairs of intervals between current common subsequence elements (denoted by blue dashed line in Figure 3). The total length of such intervals is

$$1 - \frac{1}{|\Sigma|^2 r} - \frac{r}{r^{m_2+1}} = 1 - \frac{1}{|\Sigma|^2 r} - \frac{\delta}{2}$$

or more. In fact, using the induction hypothesis, by applying the morphism m_1 more times, one can get the

following for $m = m_1 + m_2 + 1$.

$$\begin{aligned} & LCS(\phi^m(a), \phi^m(b)) \\ & \geq \left[\frac{1}{|\Sigma|^2 r} + \left(1 - \left(1 - \frac{1}{|\Sigma|^2 r} \right)^n - \frac{\delta}{2} \right) \right. \\ & \quad \cdot \left. \left(1 - \frac{1}{|\Sigma|^2 r} - \frac{\delta}{2} \right) \right] r^m \\ & \geq \left[1 - \left(1 - \frac{1}{|\Sigma|^2 r} \right)^{n+1} - \delta \right] r^m \end{aligned}$$

This completes the induction step and finishes the proof.

7.2 Synchronization Strings over Alphabets of Size Four In this section, we show that synchronization strings of arbitrary length exist over alphabets of size four. In order to do so, we first introduce the notion of *weak ε -synchronization strings*. This weaker notion is very similar to the synchronization string property except the edit distance requirement is rounded down.

DEFINITION 7.1. (WEAK ε -SYNCHRONIZATION STRINGS)

String S of length n is a weak ε -synchronization string if for every $1 \leq i < j < k \leq n$,

$$ED(S[i, j], S[j, k]) \geq \lfloor (1 - \varepsilon)(k - i) \rfloor.$$

We start by showing that binary weak ε -synchronization strings exist for some $\varepsilon < 1$.

7.2.1 Binary Weak ε -Synchronization Strings

Here we prove that an infinite binary weak ε -synchronization string exists. The main idea is to take a synchronization string over some large alphabet and convert it to a binary weak synchronization string by mapping each symbol of that large alphabet to a binary string and separating each binary encoded block with a block of the form $0^k 1^k$.

THEOREM 7.2. *There exists a constant $\varepsilon < 1$ and an infinite binary weak ε -synchronization string.*

Proof. Take some arbitrary $\varepsilon' \in (0, 1)$. According to [1], there exists an infinite ε' -synchronization string S over a sufficiently large alphabet Σ . Let $k = \lceil \log |\Sigma| \rceil$. Translate each symbol of S into k binary bits, and separate the translated k -blocks with $0^k 1^k$. We claim that this new string T is a weak ε -synchronization binary string for some $\varepsilon < 1$.

First, call a translated k -length symbol followed by $0^k 1^k$ a *full block*. Call any other (possibly empty) substring a *half block*. Then any substring of T is a half-block followed by multiple full blocks and ends with a half block.

Let A and B be two consecutive substrings in T . Without loss of generality, assume $|A| \leq |B|$ (because edit distance is symmetric). Let M be a longest common subsequence between A and B . Partition blocks of B into the following 4 types of blocks:

1. Full blocks that match completely to another full block in A .
2. Full blocks that match completely but not to just 1 full block in A .
3. Full blocks where not all bits within are matched.
4. Half blocks.

The key claim is that the $3k$ elements in B which are matched to a type-2 block in A are not contiguous and, therefore, there is at least one unmatched symbol in B surrounded by them. To see this, assume by contradiction that all letters of some type-2 block in A are matched contiguously. The following simple analysis over 3 cases contradicts this assumption:

- Match starts at middle of some translated k -length symbol, say position $p \in [2, k]$. Then the first 1 of 1^k in A will be matched to the $(k - p + 2)$ -th 0 of 0^k in B , contradiction.
- Match starts at 0-portion of $0^k 1^k$ block, say at the p -th 0. Then the p -th 1 of 1^k in A will be matched to the first 0 of 0^k in B , contradiction.
- Match starts at 1-portion of $0^k 1^k$ block, say at the p -th 1. Then the p -th 0 of 0^k in A will be matched to the first 1 of 1^k in B , contradiction.

Let the number of type-i blocks in B be t_i . For every type-2 block, there is an unmatched letter in A between its first and last matches. Hence, $|A| \geq |M| + t_2$. For every type-3 block, there is an unmatched letter in B within. Hence, $|B| \geq |M| + t_3$. Therefore, $|A| + |B| \geq 2|M| + t_2 + t_3$.

Since $|A| \leq |B|$, the total number of full blocks in both A and B is at most $2(t_1 + t_2 + t_3) + 1$. (the additional +1 comes from the possibility that the two half-blocks in B allows for one extra full block in A) Note t_1 is a matching between the full blocks in A and the full blocks in B . So due to the ε' -synchronization property of S , we obtain the following.

$$\begin{aligned} t_1 & \leq \frac{\varepsilon'}{2} (2(t_1 + t_2 + t_3) + 1) \\ \implies t_1 & \leq \frac{\varepsilon'}{1 - \varepsilon'} (t_2 + t_3) + \frac{\varepsilon'}{2(1 - \varepsilon')} \end{aligned}$$

Furthermore, $t_1 + t_2 + t_3 + 2 > \frac{|B|}{3k} \geq \frac{|A|+|B|}{6k}$. This, along with the above inequality, implies the following.

$$\frac{1}{1-\varepsilon'}(t_2 + t_3) + \frac{4-3\varepsilon'}{2(1-\varepsilon')} > \frac{|A|+|B|}{6k}.$$

The edit distance between A and B is

$$\begin{aligned} & ED(A, B) \\ &= |A| + |B| - 2|M| \geq t_2 + t_3 \\ &> \frac{1-\varepsilon'}{6k}(|A| + |B|) - \frac{4-3\varepsilon'}{2} \\ &> \frac{1-\varepsilon'}{6k}(|A| + |B|) - 2. \end{aligned}$$

Set $\varepsilon = 1 - \frac{1-\varepsilon'}{18k}$. If $|A| + |B| \geq \frac{1}{1-\varepsilon}$, then

$$\begin{aligned} & \frac{1-\varepsilon'}{6k}(|A| + |B|) - 2 \\ & \geq \left(\frac{1-\varepsilon'}{6k} - 2(1-\varepsilon) \right) (|A| + |B|) \\ &= (1-\varepsilon)(|A| + |B|) \geq \lfloor (1-\varepsilon)(|A| + |B|) \rfloor. \end{aligned}$$

As weak ε -synchronization property trivially holds for $|A| + |B| < \frac{1}{1-\varepsilon}$, this will prove that T is a weak ε -synchronization string.

7.2.2 ε -Synchronization Strings over Alphabets of Size Four A corollary of Theorem 7.2 is the existence of infinite synchronization strings over alphabets of size four. Here we make use of the fact that infinite ternary square-free strings exist, which was proven in previous work [14]. We then modify such a string to fulfill the synchronization string property, using the existence of an infinite binary weak synchronization string.

THEOREM 7.3. *There exists some $\varepsilon \in (0, 1)$ and an infinite ε -synchronization string over an alphabet of size four.*

Proof. Take an infinite ternary square-free string T over alphabet $\{1, 2, 3\}$ [14] and some $\varepsilon \in (\frac{11}{12}, 1)$. Let S be an infinite weak binary $\varepsilon' = (12\varepsilon - 11)$ -synchronization string. Consider the string W that is similar to T except that the i -th occurrence of symbol 1 in T is replaced with symbol 4 if $S[i] = 1$. Note W is still square-free. We claim W is an ε -synchronization string as well.

Let $A = W[i, j], B = W[j, k]$ be two consecutive substrings of W . If $k - i < 1/(1-\varepsilon)$, then $ED(A, B) \geq 1 > (1-\varepsilon)(k-i)$ by square-freeness.

Otherwise, $k - i \geq 1/(1-\varepsilon) \geq 12$. Consider all occurrences of 1 and 4 in A and B , which form consecutive subsequences A_s and B_s of S respectively.

Note that $|A_s| + |B_s| \geq (k-i-3)/4$, because, by square-freeness, there cannot be a length-4 substring consisting only of 2's and 3's in W .

By weak synchronization property,

$$\begin{aligned} ED(A_s, B_s) &\geq \lfloor (1-\varepsilon')(|A_s| + |B_s|) \rfloor \\ &\geq \lfloor 3(1-\varepsilon)(k-i-3) \rfloor \\ &> 3(1-\varepsilon)(k-i) - 9(1-\varepsilon) - 1 \\ &\geq (1-\varepsilon)(k-i), \end{aligned}$$

and hence, $ED(A, B) \geq ED(A_s, B_s) \geq (1-\varepsilon)(k-i)$. Therefore, W is an ε -synchronization string.

References

- [1] Bernhard Haeupler and Amirbehshad Shahrasbi. Synchronization strings: codes for insertions and deletions approaching the singleton bound, In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, pages 33-46. ACM, 2017.
- [2] Hugues Mercier, Vijay K Bhargava, and Vahid Tarokh. A survey of error-correcting codes for channels with symbol synchronization errors, IEEE Communications Surveys & Tutorials, 12(1), 2010.
- [3] Ran Gelles. Coding for interactive communication: A survey, 2015.
- [4] Ran Gelles and Bernhard Haeupler. Capacity of interactive communication over erasure channels and channels with feedback. In Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1296-1311. Society for Industrial and Applied Mathematics, 2015.
- [5] Mohsen Ghaffari, Bernhard Haeupler, and Madhu Sudan. Optimal error rates for interactive coding i: Adaptivity and other settings. In Proceedings of the forty-sixth annual ACM symposium on Theory of computing, pages 794-803. ACM, 2014.
- [6] Mohsen Ghaffari and Bernhard Haeupler. Optimal error rates for interactive coding II: Efficiency and list decoding. In Foundations of Computer Science (FOCS), IEEE 55th Annual Symposium on, pages 394-403, 2014.
- [7] Bernhard Haeupler. Interactive channel capacity revisited. In Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on, pages 226-235. IEEE, 2014.
- [8] Gillat Kol and Ran Raz. Interactive channel capacity. In Proceedings of the forty-fifth annual ACM symposium on Theory of computing, pages 715-724. ACM, 2013.
- [9] Bernhard Haeupler, Amirbehshad Shahrasbi, and Ellen Vitercik. Synchronization strings: Channel simulations and interactive coding for insertions and deletions. In 45th International Colloquium on Automata, Languages, and Programming (ICALP), pages 75:1-75:14, 2018.

[10] Bernhard Haeupler and Amirbehshad Shahrasbi. Synchronization strings: Explicit constructions, local decoding, and applications. *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, 2018.

[11] Bernhard Haeupler, Amirbehshad Shahrasbi, and Madhu Sudan. Synchronization strings: List decoding for insertions and deletions. In *45th International Colloquium on Automata, Languages, and Programming (ICALP)*, 76:1–76:14, 2018.

[12] Bernhard Haeupler, Aviad Rubinstein, and Amirbehshad Shahrasbi. Near-linear time insertion-deletion codes and $(1+\varepsilon)$ -approximating edit distance via indexing. *arXiv preprint arXiv:1810.11863*, 2018.

[13] Mark Braverman, Ran Gelles, Jieming Mao, and Raefail Ostrovsky. Coding for interactive communication correcting insertions and deletions. *IEEE Transactions on Information Theory*, 63(10):6256–6270, 2017.

[14] Axel Thue. Über unendliche zeichenreihen (1906). *Selected Mathematical Papers of Axel Thue*. Universitetsforlaget, 1977.

[15] Axel Thue. Über die gegenseitige lage gleicher teile gewisser zeichenreihen. *Christiana Videnskabs-Selskabs Skrifter, I. Math. Naturv. Klasse*, 1, 1912.

[16] John Leech. 2726. a problem on strings of beads. *The Mathematical Gazette*, 41(338):277–278, 1957.

[17] Max Crochemore. Sharp characterizations of squarefree morphisms. *Theoretical Computer Science*, 18(2):221–226, 1982.

[18] Robert Shelton. Aperiodic words on three symbols. *Journal für die Reine und Angewandte Mathematik*, 321:195–209, 1981.

[19] Robert O Shelton and Raj P Soni. Aperiodic words on three symbols. iii. *Journal für die Reine und Angewandte Mathematik*, 330:44–52, 1982.

[20] Boris Zolotov. Another solution to the Thue problem of non-repeating words. *arXiv preprint arXiv:1505.00019*, 2015.

[21] Dalia Krieger, Pascal Ochem, Narad Rampersad, and Jeffrey Shallit. Avoiding approximate squares. In *International Conference on Developments in Language Theory*, pages 278–289. Springer, 2007.

[22] Serina Camungol and Narad Rampersad. Avoiding approximate repetitions with respect to the longest common subsequence distance. *Involve: A Journal of Mathematics*, 9(4):657–666, 2016.

[23] Karthekeyan Chandrasekaran, Navin Goyal, and Bernhard Haeupler. Deterministic algorithms for the Lovász local lemma. *SIAM Journal on Computing*, 42(6):2132–2155, 2013.

[24] Robin A Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *Journal of the ACM (JACM)*, 57(2):11, 2010.

[25] Bernhard Haeupler, Barna Saha, and Aravind Srinivasan. New constructive aspects of the Lovász local lemma. *Journal of the ACM (JACM)*, 58(6):28, 2011.

[26] V. Guruswami and P. Indyk. Linear-time encodable/decodable codes with near-optimal rate. *IEEE Transactions on Information Theory*, 51(10):3393–3400, 2005.

[27] Maria Axenovich, Yury Person, and Svetlana Puzynina. A regularity lemma and twins in words. *Journal of Combinatorial Theory, Series A*, 120(4):733–743, 2013.

[28] Boris Bukh and Lidong Zhou. Twins in words and long common subsequences in permutations. *Israel Journal of Mathematics*, 213(1):183–209, 2016.

[29] Paul Beame and Dang-Trinh Huynh-Ngoc. On the value of multiple read/write streams for approximating frequency moments. In *49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 499–508, 2008.

[30] Alon, N. and Spencer, J.H., *The Probabilistic Method*. Wiley: 2004. Wiley Series in Discrete Mathematics and Optimization.

[31] J. Beck, An application of Lovász local lemma: There exists an infinite 01-sequence containing no near identical intervals. *Colloq. Mathematica Societatis Janos Bolyai*, 37. Finite and Infinite Sets, 1981, pp. 103–107.