

Efficient Linear and Affine Codes for Correcting Insertions/Deletions*

Kuan Cheng[†]

Venkatesan Guruswami[‡]

Bernhard Haeupler[§]

Xin Li[¶]

Abstract

This paper studies *linear* and *affine* error-correcting codes for correcting synchronization errors such as insertions and deletions. We call such codes linear/affine insdel codes.

Linear codes that can correct even a single deletion are limited to have information rate at most $1/2$ (achieved by the trivial 2-fold repetition code). Previously, it was (erroneously) reported that more generally no non-trivial linear codes correcting k deletions exist, i.e., that the $(k+1)$ -fold repetition codes and its rate of $1/(k+1)$ are basically optimal for any k . We disprove this and show the existence of binary linear codes of length n and rate just below $1/2$ capable of correcting $\Omega(n)$ insertions and deletions. This identifies rate $1/2$ as a sharp threshold for recovery from deletions for linear codes, and reopens the quest for a better understanding of the capabilities of linear codes for correcting insertions/deletions.

We prove novel outer bounds and existential inner bounds for the rate vs. (edit) distance trade-off of linear insdel codes. We complement our existential results with an efficient synchronization-string-based transformation that converts any asymptotically-good linear code for Hamming errors into an asymptotically-good linear code for insdel errors. Lastly, we show that the $\frac{1}{2}$ -rate limitation does not hold for affine codes by giving an explicit affine code of rate $1 - \epsilon$ which can efficiently correct a constant fraction of insdel errors.

*A full version of this paper appears in [7]

[†]Center on Frontiers of Computing Studies, Peking University. ckkcdh@pku.edu.cn. Supported in part by a start-up funding of Peking University CFCS, a Simons Investigator Award (#409864, David Zuckerman) and NSF Award CCF-1617713.

[‡]Computer Science Department, Carnegie Mellon University. venkatg@cs.cmu.edu. Research supported in part by NSF grant CCF-1814603.

[§]Computer Science Department, Carnegie Mellon University. haeupler@cs.cmu.edu. Supported in part by NSF awards CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808 and a Sloan Research Fellowship.

[¶]Department of Computer Science, Johns Hopkins University. lixints@cs.jhu.edu. Supported by NSF Award CCF-1617713 and NSF CAREER Award CCF-1845349.

1 Introduction

Error-correcting codes resilient to synchronization errors such as insertions and deletions (insdel errors) have witnessed a lot of exciting recent progress. Efficiently constructible and decodable codes with strong, or even near-optimal, parameters have been discovered in several regimes.

For codes over the binary or other fixed constant-sized alphabets, these include (i) codes to correct a small (constant) number of insdel errors [22, 8, 4], (ii) High-rate codes to correct a small fraction of insdel errors [12, 8, 18, 11], (iii) codes of positive rate over an alphabet size q to correct a fraction of insdel errors approaching $1 - 1/(q + \sqrt{q})$ [5], (iv) near-optimal codes for block edit errors including transpositions [9], (v) list-decodable codes over large alphabets able to correct more than 100% of insertions and achieving an optimal rate independent of the fraction of insertions [17], and (vi) list-decodable codes for any alphabet of positive rate for the largest possible combination of insertions and deletions that is information-theoretically correctible [10, 19, 24, 11]. Many of these recent results build on (ideas from) synchronization strings [6, 15] introduced in [14] which elegantly reduced questions of tackling synchronization errors to the much better understood problem of correcting Hamming errors and erasures. This also led to codes over large alphabets with the optimal trade-off between rate and fraction of insdel errors corrected [14] and a near-linear encoding and decoding complexity [13].

Interestingly, not a single one of the above-mentioned constructions provides a linear code. Indeed, to the best of our knowledge, no linear insdel code constructions, beyond the trivial repetition codes, have appeared in the literature. This stands in stark contrast to error-correcting codes for Hamming errors where most important codes both in theory and practice are linear codes, including Hamming, Reed-Solomon and Reed-Muller codes, algebraic-geometry codes, polar codes, Turbo codes, expander codes, and LDPC codes.

Linear codes have many advantages: 1.) They have nice and compact representations via either their generator matrix G or their parity check matrix H . 2.) These compact representations also directly imply that

every linear codes has a simple and efficient $O(n^2)$ -time encoding procedure $\text{Enc}(x) = xG$, as well as an efficient test $Hx = 0$ for whether a given string x is a codeword. Faster encoding algorithms are often possible, and when H is sparse, even linear time decoders have been discovered. 3.) Linear codes allow for simpler analysis and allow powerful methods from linear algebra to apply. For example, many properties of a linear code can be related to its (Hamming) weight distribution, including the minimum distance of a linear code which is equal to the minimum weight of any non-zero codeword.

One of the reasons why the recent wave of results on error correcting codes for synchronization errors did not produce any results on linear insdel codes are two negative results suggesting that no such non-trivial codes exist. In particular, in 2007 [1] proved that any linear code that even corrects a single deletion must have a rate of at most $1/2$, which is the rate achieved by the trivial repetition code that outputs every symbol twice. This result is independent of the size of the alphabet and stands in stark contrast to codes for Hamming errors or non-linear insdel codes which can correct a constant fraction of errors while having a rate arbitrarily close to 1. In [4] it was claimed that the impossibility result of [1] extends to more than one deletion. In particular, the rate achievable by any linear code correcting a fixed number k of deletions is upper bound by (about) $1/(k+1)$, which is the rate of the trivial $(k+1)$ -fold repetition code.

The starting point for this paper was the realization that, while the impossibility result of [1] stands, the proof of its extension to larger k in [4] is flawed. It has since been retracted in [3], as a result of this work. This reopens the quest to better understand the potential and limitations of linear codes in tackling insertion and deletion errors.

In this paper we provide new existential results as well as explicit constructions of efficient linear codes. Our results are the first to provide any non-trivial linear insdel codes that drastically outperform simple repetition codes.

1.1 Combinatorial/existence results We start by showing that quite powerful linear insdel codes exist just below the $1/2$ rate upper bound of [1]. In particular, a random binary linear code with a rate just below $1/2$ is capable of correcting a linear number of errors, vastly more than the single error of the 2-fold repetition code of rate $1/2$.

THEOREM 1.1. *For any $\varepsilon > 0$, a random binary linear*

code of length n and rate $1/2 - \varepsilon$ has edit distance¹ $\Omega(\varepsilon \log^{-1} \frac{1}{\varepsilon})n$ with high probability and thus is capable of correcting a constant fraction of insertions and deletions.

We also extend the result of [1] into a distance-dependent rate upper bound showing that, independent of the alphabet used, no rate larger than $\frac{1}{2}(1 - \delta)$ is achievable by any linear code which can correct a δ -fraction of insdel errors. This shows that the asymptotics of Theorem 1.1 are optimal up to the $\Omega(\log^{-1} \frac{1}{\varepsilon})$ factor. It also shows that while linear codes for Hamming errors and non-linear insdel codes can get arbitrarily close to the Singleton bound, the best rate achievable by any linear insdel code (over a finite alphabet) is exactly half as large, for any error correcting capability $\delta > 0$:

THEOREM 1.2. (HALF-SINGLETON BOUND) *For any $\delta > 0$ and any $\varepsilon > 0$, for all large enough finite fields of size $q = 2^{\Theta(\varepsilon^{-1})}$, there exists a family of \mathbb{F}_q -linear insdel codes of rate $\frac{1}{2}(1 - \delta) - \varepsilon$ capable of correcting any δ -fraction of insdel errors. Furthermore, no such family of linear codes exists with a rate larger than $\frac{1}{2}(1 - \delta)$ over any alphabet.*

We also prove an alphabet-dependent rate upper bound, namely a “half-Plotkin” bound. This bound, which is rather subtle to establish, shows that the rate of a \mathbb{F}_q -linear code capable of correcting a fraction δ of insdel errors can be at most half as large as the $(1 - \frac{q}{q-1}\delta)$ Plotkin outer bound for q -ary codes of relative (Hamming) distance δ .

THEOREM 1.3. (HALF-PLOTKIN OUTER BOUND) *Any \mathbb{F}_q -linear code which can correct a δ fraction of deletions has rate at most $\frac{1}{2}(1 - \frac{q}{q-1}\delta) + o(1)$.*

1.2 Explicit constructions of linear/affine insdel codes We also give an explicit construction of asymptotically good linear insdel codes. Using the ideas from synchronization strings [14] we give a transformation which takes any explicit, efficiently decodable linear code C for Hamming errors and produces an explicit, efficiently decodable linear insdel code C' whose rate and error correcting capabilities are at most a constant factor smaller than those of the code C .

THEOREM 1.4. *Fix any finite field \mathbb{F}_q . There is an explicit construction of a \mathbb{F}_q -linear code family with*

¹Edit distance between two strings is the minimum number of insertions, deletions and replacements that can modify one string to be the other.

rate bounded away from zero, a linear time encoding algorithm, and a polynomial time decoding algorithm to correct a positive constant fraction of insertions and deletions. The generator matrix of codes in the family can be deterministically computed in polynomial time in the code block length.

The above result gives an analog of the classic result of Schulman and Zuckerman which gave an explicit construction of a *non-linear* binary code family with positive rate and an efficient algorithm to correct a constant fraction of insdel errors [20]. The construction in [20] starts with an outer code that can correct a constant fraction of Hamming errors (like Reed-Solomon codes), adds index information to its symbols, and then encodes the resulting indexed symbols by a suitable inner codes. The introduction of index information precludes getting a linear code via this approach.

Our linear code construction of Theorem 1.4 does not have rate approaching $1/2$, which is shown to be possible by the existence result of Theorem 1.1. Improving the rate of our construction remains a fascinating open question. An ultimate goal would be to construct linear codes approaching the half-Singleton bound over large finite fields.

A systematic code encodes a message x as x followed by some check symbols. For Hamming errors, one can make any linear code systematic by performing a basis change in the generator matrix and permuting the symbols if necessary. For insdel errors, permuting symbols can greatly alter the edit distance. We show our random coding result underlying Theorems 1.1 and 1.2 holds with the same parameters also for systematic linear codes. We can also get a version of Theorem 1.4 with systematic codes by simply appending the message at the beginning and losing constant factors in the rate and fraction of insdel errors corrected.

1.3 Affine insdel codes Our final result concerns affine codes, a concept introduced in [1] to overcome the limitations of linear and cyclic codes for deletion errors. Indeed, it was suggested in [1] to add a fixed symbol to the codewords of a cyclic/linear code to increase its edit distance. The codewords of such an encoding form an affine subspace. In general, an affine code is one with encoding function $\text{Enc} : \mathbb{F}_q^m \rightarrow \mathbb{F}_q^n$ given by an affine map $\text{Enc}(x) = xG + b$ for some generator matrix G and offset vector b . (In fact, the affine codes constructed in this paper are of the same simpler form as the codes in [1] where each position of a codeword is either a fixed symbol or a linear combination of input symbols.) While it is easy to see that affine codes and linear codes are equivalent in power for correcting Hamming errors,

the same is not true for insdel errors. Indeed, despite being structurally about as simple as linear codes the impossibility result of [1] and our outer bounds for linear insdel codes do *not* apply to affine codes. This opens up the possibility of obtaining affine codes with better rates.

We show that this is indeed the case by giving high-rate affine insdel codes to correct a constant fraction of insdel errors while achieving a rate arbitrarily close to 1. Our affine codes are binary, explicit, and feature efficient encoding and decoding algorithms. The construction is based on the binary codes in [18] which are a synchronization string based simplification of the codes of [11].

THEOREM 1.5. *For any $\varepsilon > 0$ there exists an explicit binary affine code family with rate $1 - \varepsilon$ that can be efficiently decoded from a fraction $O(\varepsilon^3)$ of insertions and deletions.*

We note that unlike linear codes, affine codes can allow the introduction of index information as in the original Schulman-Zuckerman construction [20]. However, their construction also requires inner codes with strong properties (such as a good density of 1's in every codeword, or good edit distance between all pairs of long enough subsequences of distinct codewords) which are difficult to ensure with affine codes, and in any case incompatible with obtaining high rate.

For general insdel codes, constructions with a better trade-off between rate and fraction of correctable insdel errors than the guarantee of Theorem 1.5 are known. The constructions in [11, 18] could correct a fraction $\tilde{\Omega}(\varepsilon^2)$ of insdel errors with rate $1 - \varepsilon$. Even more recently, via connections to the document exchange problem, codes with rate $1 - \varepsilon$ that can correct a fraction $\Omega(\varepsilon/\log^2(1/\varepsilon))$ of insertions and deletions have been constructed in [8, 12].

Our results expose the under-appreciated power of linear and affine codes in the context of correcting insdel errors and identify some limitations that are in contrast with general codes. Put together, they substantially improve our understanding of the landscape of possibilities for using linear/affine codes to tackle synchronization errors, and open up several intriguing questions that merit further study (see Section 10).

2 Overview of our approach

Our existential results follow from relatively straightforward applications of the probabilistic method: we take a random generator matrix and analyze its property. The half-Singleton bound also follows easily from a reduction to previous results: if the code can correct δn

errors then we can simply remove the first $\delta n - 1$ symbols and the code can still correct one deletion, which by the result of [1] limits the rate to $1/2$. The proof of the half-Plotkin bound is significantly more involved. From a q -ary linear code C that can correct δn errors, we drop a prefix consisting of the first $d \approx qn\delta/(q-1)$ symbols of codewords and prove that the remaining punctured code C' has dimension at most $\approx \frac{n}{2}(1 - q\delta/(q-1))$. For the latter, we give a clever reduction to the fact that in any linear code of large enough dimension, there is a non-zero codeword which has many 0's in *both* the first and the second half. We establish this latter fact using a carefully designed probabilistic argument combined with collapsing together linearly dependent coordinates. Now note that if instead of dropping the prefix consisting of the first $d \approx qn\delta/(q-1)$ symbols of all codewords, we just delete all symbols except d/q of the most frequently occurring element in the prefix, then after this operation any two different codewords in C are still distinct (since $d(1 - 1/q) \leq \delta n$). Note that the resulted codeword is of the form $\alpha^{d/q}$ followed by a codeword in C' , for some $\alpha \in \mathbb{F}_q$. This implies that the number of codewords in C is at most q times the number of codewords in C' , thus the original code has dimension at most $\dim(C') + 1$ which is also at most $\approx \frac{n}{2}(1 - q\delta/(q-1))$.

Now we give an informal overview of our constructions of linear codes and affine codes for insertions and deletions. As observed in many previous works, the main difficulty in handling edit errors is that insertions and deletions can shift the string, and thus cause the loss of index information. An elegant idea to handle this, introduced in the work of Haeupler and Shahrasbi [14], and later used in [16, 6] is to construct a fixed string that can be used to recover the index information. Such a string is called a *synchronization string* and one can then use it to transform a code for Hamming errors into a code for edit errors, just by appending the symbols of the synchronization string to every codeword.

However, even if we start with a linear code for Hamming errors, the resulting code after appending the synchronization string is not necessarily a linear code or an affine code. This is due to both the fact that the synchronization string is some fixed string, and the fact that when appending we are putting the symbol from the codeword and the symbol from the synchronization string together to form a symbol in a larger alphabet. If we restrict ourselves to binary codes, then the above approach has an additional problem of increasing the alphabet size. Instead, here we use similar ideas, but modify the synchronization string to obtain a linear or affine code without blowing up the alphabet size. Specifically, instead of appending

a synchronization string, we insert a fixed sequence of synchronization symbols into the codewords of a code for Hamming errors. This ensures that the resulting code is still a linear or affine code. We call such a sequence of symbols a *synchronization sequence*, and we further show how to explicitly design such sequences to allow proper decoding. We now give more details below.

2.1 Randomized construction of the linear code

We first start with a randomized construction of the synchronization sequence, which results in a randomized linear code. Later we describe how to derandomize the construction to give an explicit code. The construction of the randomized linear code is as follows.

To encode a given message $x \in \mathbb{F}_q^m$, we first encode it to be $y = C(x)$ using a linear code $C : \mathbb{F}_q^m \rightarrow \mathbb{F}_q^{n_C}$ which can correct k_C Hamming errors. Then before each symbol of y , we insert a short string of 0's, obtaining the final codeword

$$z = S[1] \circ y[1] \circ S[2] \circ y[2] \circ \dots \circ S[n_C] \circ y[n_C].$$

Here \circ stands for string concatenation. $S = (S[1], S[2], \dots, S[n_C])$ is a sequence of independent random variables s.t. each $S[i]$ is an all-0 string whose length is uniformly chosen from $\{1, 2, \dots, a\}$. We set $a = O(\frac{n_C}{\kappa_C})$ for a large enough constant hiding in $O(\cdot)$. Note that since we are only inserting a sequence of 0's into the codeword, the resulting code is still a linear code.

To see that the new code is good for insertions and deletions, let us focus on the case of a binary alphabet, where we start with an asymptotically good code for Hamming errors. We argue that for any two different new codewords after the transformation, with high probability their edit distance is large (i.e., $\Omega(n_C)$). Assume the two codewords from C are y_1 and y_2 , and the codewords after transformation are y'_1 and y'_2 . For simplicity, let us make the over simplified assumption that both y_1 and y_2 have Hamming weight βn_C , and their Hamming distance is also βn_C , for some constant $\beta > 0$. Note that in this case $a = O(\frac{n_C}{\kappa_C})$ is a constant and thus the length of the new codeword is $n = O(an_C) = O(n_C)$.

Consider a largest monotone matching between the bits of y'_1 and y'_2 , which corresponds to the longest common subsequence. We focus on the number of matched pairs of 1's in this matching, and we have two cases.

Case 1: The number of matched pairs of 1's is at most $\frac{3\beta}{4}n_C$. In this case, there are at least $\frac{\beta}{2}n_C$ unmatched 1's, and thus the edit distance between y'_1 and y'_2 is at least $\frac{\beta}{2}n_C$.

Case 2: The number of matched pairs of 1's is at least $\frac{3\beta}{4}n_C$. In this case, since the Hamming distance between y_1 and y_2 is βn_C , at least $\frac{\beta}{4}n_C$ pairs of 1's are incorrectly matched (i.e., the positions of the matched 1's in y_1 and y_2 are not the same). Now consider the length of the interval between two adjacent matched 1's in y'_1 and y'_2 . Our main observation is that if at least one of the matches is incorrect, then the lengths of the two intervals in y'_1 and y'_2 are unlikely to be equal. Indeed, since the length of the interval involves the sum of a series of integers independently and uniformly chosen from $\{1, 2, \dots, a\}$, the expected difference between the lengths of these two intervals is $\Omega(a)$. Note that the edit distance between y'_1 and y'_2 is at least the sum of all the length differences of these intervals. Since we have at least $\frac{\beta}{4}n_C$ pairs of 1's that are incorrectly matched, by a Chernoff bound with high probability the edit distance is $\Omega(an_C) = \Omega(n_C)$.

We now turn to our decoding algorithm. Based on our observation above, our decoding algorithm will first compute some kind of optimal matching between the received word and a template of codeword to recover part of the codeword $y = C(x)$, and then use the decoding algorithm of C to recover x .

Specifically, to decode, let the received corrupted codeword be z' . We imagine that the correct codeword is in the following form

$$z' = S[1] \circ ? \circ S[2] \circ ? \circ \dots \circ S[n_C] \circ ? \in \{0, ?\}^n,$$

where each “?” mark is a special symbol indicating a corresponding symbol of y . To fill the “?” marks, we compute a monotone matching between z' and z , where we match the “?” marks of z to the non-zero-symbols in z' . The purpose of the matching algorithm is to achieve the following goals: 1) the size of the matching should be as large as possible, and 2) the number of symbols in z' which are matched to their correct positions in z should be as large as possible. After obtaining the matching we fill the unmatched “?” marks with 0 to get a string y' . The previous two conditions now guarantee that the Hamming distance between y and y' is not too large (as long as the number of edit errors is not too large), and we can use the decoding algorithm of C to get x .

However, the second condition above is hard to check, thus we again turn to our previous observation— if there is a pair of incorrectly matched symbols, then the lengths of the two intervals in z' and z involving this pair are unlikely to be equal. If there are many incorrectly matched symbols, then there are many corresponding intervals whose lengths are not equal. Thus,

we can simply count the number of such intervals. Inspired by this, we define the following cost function.

Given a ? to non-zero symbol matching $w = ((i_1, j_1), \dots, (i_{|w|}, j_{|w|}))$, where $(i_1, \dots, i_{|w|})$ stand for the relative indices of the matched ? marks in all ? marks of z' , and $(j_1, \dots, j_{|w|})$ stand for the relative indices of the matched non-zero symbols in all non-zero symbols of z' , we define the cost function $\text{cost}(w)$ to be

$$\text{cost}(w) = \sum_{k=1}^{|w|} \mathbb{1}(p_{i_k} - p_{i_{k-1}} \neq q_{j_k} - q_{j_{k-1}}).$$

Here p_{i_k} is the position of the i_k -th ? mark in z' , q_{j_k} is the position of the j_k -th non-zero symbol in z' , and $\mathbb{1}()$ is the indicator function. Our goal is now to maximize $\text{obj}(w) = |w| - \text{cost}(w)$, and we give a dynamic programming algorithm for this. In the following, let w be the optimal matching returned by our algorithm.

We now show why this approach ensures a small Hamming distance between y' and y , as long as the number of insertions and deletions κ is bounded, say $\kappa \leq 0.1\kappa_C$. We will actually prove the contrapositive, that is, if the Hamming distance between y' and y is large, then w will not be the optimal matching returned by our algorithm. Again, for simplicity let us assume that we are working with a binary alphabet.

First we bound the number of non-zero symbols in z' that are not matched, to be $O(\kappa)$, as in Lemma 6.5. This follows from the fact that the matching w returned by our algorithm maximizes the function $\text{obj}(w)$. Specifically, consider the ideal matching w^* that matches each non-inserted 1 in z' to its correct position. If there is no error then there are $\tau \geq 2\kappa_C + 1$ matched 1's since C has distance at least $2\kappa_C + 1$. Note that κ errors can only decrease $\text{obj}(w^*)$ by $O(\kappa)$ since each error is just one insertion or deletion. Hence $\text{obj}(w) \geq \text{obj}(w^*) \geq \tau - O(\kappa)$. This means that $|w| \geq \tau - O(\kappa)$. Further notice that the number of 1's in z' is at most $\tau + \kappa$, thus the number of non-zero symbols in z' that are not matched is bounded by $O(\kappa)$.

Next we focus on the matched symbols of z' and analyze how many of them are not matched to their correct positions in y . Towards that, we say a match (i, j) in w is bad if the j -th non-zero symbol of z' is not inserted and is originally the i -th symbol in y s.t. $i_j \neq i$.

Now suppose the Hamming distance between y and y' is larger than κ_C , and let $\Delta(w)$ be the number of bad matches in w . We claim that $\Delta(w) > (\kappa_C - 6\kappa)/2$, i.e., Lemma 6.7. To see this, recall that by Lemma 6.5, there are at most $O(\kappa)$ non-zeros in z' that are not matched. So a large fraction of the non-zero symbols

in z' are matched. Intuitively, since the Hamming distance between y and y' is larger than $\kappa_C \gg \kappa$, a certain fraction of the matches should be bad. We show that this is indeed true, although some effort and careful analysis are required here, since we need deal with several issues such as the non-zeros in z' may be inserted, and the Hamming errors may come from the unmatched “?” marks.

Since our definition of a bad match implies that the matched non-zero symbol in z' is not inserted, the above bad matches can be viewed as inducing a self matching \tilde{w} between the “?” marks of $z?$ and itself. Now we again use our crucial observation that if a match is bad, then the lengths of the corresponding two intervals involving this match are unlikely to be equal. Specifically, for any bad match (i, j) , we show that it will contribute 1 to cost with probability $1 - 1/a$, since the length of the 0's is uniform chosen from $[a]$. Given that there are $\Delta(w)$ bad matches, with high probability they will contribute $\Omega(\Delta(w))$ to cost . Now note that κ errors can decrease cost by at most κ , since each insertion or deletion can only affect at most one of the indicators in cost . As a result, we have $\text{cost}(w) \geq \Omega(\Delta(w)) - \kappa$ and thus $\text{obj}(w) = |w| - \text{cost}(w) \leq \tau + \kappa - (\Omega(\Delta(w)) - \kappa) = \tau + 2\kappa - \Omega(\Delta(w))$. By setting κ to be small enough (e.g., $\kappa \leq 0.1\kappa_C$) this contradicts the fact that $\text{obj}(w) \geq \text{obj}(w^*) \geq \tau - O(\kappa)$.

It is straightforward to see that the code is linear if C is linear, and is asymptotically good if C is.

2.2 Synchronization Separator Sequence and the explicit linear code From our discussion above, we can actually abstract out the property of the sequence S that we need. In short, we need the following property which holds with high probability from a random sequence: any self matching between the ? marks which has many bad matches must incur a large cost. Thus, to derandomize our construction, we just need to explicitly construct such a sequence S . Towards that, we give the following definition.

For a sequence $s[1]o? \circ s[2]o? \circ \dots \circ s[n]o?$, a ?-to-? self matching w is a monotone matching between the string and itself s.t. each match (i, j) matches the i -th question mark to the j -th question mark. A match (i, j) in w is called *undesired* if

- $i \neq j$ (i.e. the match is bad);
- $p_i - p_{i'} = p_j - p_{j'}$.

Here (i', j') is the immediate previous match of (i, j) in the matching. If there is no such (i', j') then the second bullet just checks $p_i = p_j$ instead. Also recall that p_i is the position of the i -th ?-mark in $z?$.

DEFINITION 2.1. $s[1], s[2], \dots, s[n]$ is called a (Λ, a) synchronization separator sequence, if for any self matching of $s[1]o? \circ s[2]o? \circ \dots \circ s[n]o?$, the number of undesired matches is at most Λ .

Here each $s[i], i \in [n]$ is an all 0 string whose length is in $\{1, \dots, a\}$.

It can be easily seen that a (Λ, a) synchronization separator sequence satisfies the property we require: any matching with Δ bad matches must incur a cost of at least $\Delta - \Lambda$. In this paper, for any $n \in \mathbb{N}, \Lambda \leq n$, we can construct an explicit synchronization separator sequence of length n , with $a = \Theta(\frac{n}{\Lambda})$. In particular, setting $\Lambda = \alpha n$ for some constant $0 < \alpha < 1$ results in $a = O(1)$ and an explicit asymptotically good linear code for insertions and deletions.

We now show how to construct this object explicitly. The high level idea is that we use a pseudorandom distribution, i.e., an ε almost k -wise independent distribution with small seed length, to generate the sequence S . To show this works, we reduce the condition of being a (Λ, a) synchronization separator sequence to some small scale checks, i.e., for every pair of short substrings (the sum of the lengths of the two substrings is $\Theta(\frac{n \log n}{\Lambda \log \frac{n}{\Lambda}})$), we check if there is no matching between them that has a large fraction of undesired matches. Note that these checks are local checks, so they can be fooled by an almost k -wise independent distribution with a proper setting of the parameter k . That is, under such a pseudorandom distribution, the probability that all checks passed is positive. Furthermore this distribution can be generated by a small number of random bits (i.e., $\Theta(\log n)$ random bits), and thus we can exhaustively search the entire sample space to find an explicit sequence that passes all checks. Note that for a specific sequence, whether it is a (Λ, a) synchronization separator sequence can be checked efficiently (e.g., by dynamic programming), so the whole construction runs in polynomial time.

To show that passing all small scale checks guarantees a good synchronization separator sequence, we use a matching cutting argument similar to that of [8]: For a pair of substrings, let γ be the ratio of the number of undesired matches to the total length of the pair of substrings. We show that we can cut the pair of strings into two pairs s.t. each pair contains half the number of undesired matches, and at least one pair has ratio at least γ . One can keep doing this in the pair with ratio at least γ , until the number of undesired matches is small enough (e.g., $\Theta(\frac{\log n}{\log \frac{n}{\Lambda}})$). This also implies that the total length of the current pair of substrings is $\Theta(\gamma^{-1} \frac{\log n}{\log \frac{n}{\Lambda}})$. As a result, if there is no pair of such short substrings

which has ratio γ then the original pair of strings cannot have ratio γ . If we start with $\gamma = \frac{\Lambda}{2n}$, then the total length of the pair of short substrings is $\Theta\left(\frac{n \log n}{\Lambda \log \frac{n}{\Lambda}}\right)$.

Systematic linear code. Our construction can be turned into a systematic linear code by simply concatenating the message with the corresponding codeword. To decode, we can simply ignore the first m symbols (m is the message length) and decode the rest of the received word. See Section 9 for the details.

2.3 The high rate affine code construction In order to construct an affine code C over the binary alphabet, our first step is to use the construction in [14] which works for a larger alphabet. Specifically, we take a linear code C_0 over the alphabet $\mathbb{F}_2^{l_0}$ for Hamming errors, and for each codeword we concatenate it coordinate-wise with a synchronization string over a constant size alphabet. This provides us with a code C' for insertions and deletions. However, due to the coordinate-wise concatenation, the resulted code is not linear or affine any more.

We can convert the code C' back into an affine code by directly expressing each symbol in binary. Since we start with a linear code C_0 over the alphabet $\mathbb{F}_2^{l_0}$, the conversion of C_0 into the binary alphabet keeps its linearity. The symbols of the synchronization string now become fixed strings, and hence now the code becomes affine. However, simply doing this will be problematic, since now in a corrupted codeword it will be hard to tell the boundaries between different symbols.

To solve this problem, we create special boundaries in the codewords by taking a special string B which consists of $t+1$ 1's, for some parameter t . We insert the string B before each symbol in a codeword y' of C' . To distinguish the boundaries from the bits in a codeword, we also modify the codeword in the following way: when we express each symbol of y' in binary, we insert a 0 after every t bits. We refer to the binary representation of a symbol in y' , and the inserted 0's as the content. We refer to the content together with the boundary before it as a block. This finishes our construction.

The decoding strategy and correctness are as follows. Assume there are κ insertions and deletions. Note that each insertion or deletion can only affect at most two blocks, since the worst situation is that it creates a new boundary and also modifies the content of a block. Therefore there are only $O(\kappa)$ corrupted blocks. Now we can just locate every correct boundary and view the substrings between two adjacent boundaries as contents. By doing this, every uncorrupted block of y' will be recovered correctly. Now we can use the decoding algorithm in [14] to recover the message, as long as there are at

most κ insertions/deletions.

To get a high information rate, we start with a linear code C_0 that has codeword length n_0 , decoding radius $\kappa_0 = \varepsilon n_0$ and information rate $1 - \Theta(\varepsilon)$. This code has $l_0 = O(\varepsilon^{-2})$, and can be constructed using algebraic geometry codes. We set $t = O(\frac{1}{\varepsilon})$. Note that the synchronization string in [14] has constant alphabet size, thus it follows from our construction that the new code C has codeword length $O(l_0 n_0 + t n_0) = O(\varepsilon^{-2} n_0)$. From our previous discussion, the number of insertions and deletions C can correct is $\kappa = \Theta(\kappa_0) = \Theta(\varepsilon^3 n)$. For the information rate, notice that l_0 is much larger than t , and the synchronization string has constant alphabet size, thus the information rate of C is still $1 - \Theta(\varepsilon)$.

Since the only modification we make to the codewords is inserting some fixed strings in fixed positions, the code is still an affine code.

3 Preliminaries

Notation. Let Σ be an alphabet. For a string $x \in \Sigma^*$,

1. $|x|$ denotes the length of the string.
2. $x[i, j]$ denotes the substring of x from position i to position j (both endpoints included).
3. $x[i]$ denotes the i -th symbol of x .
4. $x \circ x'$ denotes the concatenation of x and some other string $x' \in \Sigma^*$.
5. x^N the concatenation of N copies of the string x .

3.1 Edit distance and longest common subsequence

DEFINITION 3.1. (EDIT DISTANCE) For any two strings $x, y \in \Sigma^n$, the edit distance $\text{ED}(x, y)$ is the minimum number of edit operations (insertions and deletions) required to transform x into y .²

DEFINITION 3.2. (LONGEST COMMON SUBSEQUENCE) For any two strings x, y over an alphabet Σ , the longest common subsequence of x and y is the longest pair of subsequences of x and y that are equal as strings. We use $\text{LCS}(x, y)$ to denote the length of the longest common subsequence between x and y .

$$\text{Note that } \text{ED}(x, y) = |x| + |y| - 2 \cdot \text{LCS}(x, y).$$

²The standard definition of edit distance also allows substitution, but for simplicity we only consider insertions and deletions here, as a substitution can be replaced by a deletion followed by an insertion.

3.2 Error correcting codes

DEFINITION 3.3. An (n, m, d) -code C is an error-correcting code (for Hamming errors) with codeword length n , message length m , such that the Hamming distance between every pair of codewords in C is at least d .

Next we recall the definition of error-correcting codes (ECC) for edit errors.

DEFINITION 3.4. Fix an alphabet Σ , an error-correcting code $C \subseteq \Sigma^n$ for edit errors with message length m and codeword length n consists of an encoding function $Enc : \Sigma^m \rightarrow \Sigma^n$ and a decoding function $Dec : \Sigma^n \rightarrow \Sigma^m$. The code can correct k edit errors if for every y, s, t . $ED(y, Enc(x)) \leq k$, we have $Dec(y) = x$. The rate of the code is defined as $\frac{m}{n}$. The alphabet size is $|\Sigma|$.

The code family C is explicit (or has an explicit construction) if both encoding and decoding can be done in polynomial time. We say C is a linear code if the alphabet Σ is a finite field \mathbb{F}_q and the encoding function $Enc : \mathbb{F}_q^m \rightarrow \mathbb{F}_q^n$ is a \mathbb{F}_q -linear map.

Let C_1, C_2 be two linear codes over \mathbb{F}_q^n . We use $C_1 \cap C_2$ to denote the intersection of the two linear spaces, which is also a linear space. Also we use $C_1 \cup C_2$ to denote the linear space which is spanned by all basis vectors of C_1, C_2 .

3.3 Pseudorandom generator We use U_n to denote the uniform distribution on $\{0, 1\}^n$, and $\mathbb{1}(\cdot)$ to denote the indicator function.

DEFINITION 3.5. ([2]) *Random variables* X_1, X_2, \dots, X_n over $\{0, 1\}$ are ε -almost k -wise independent in max norm if $\forall i_1, i_2, \dots, i_k \in [n]$, $\forall x \in \{0, 1\}^k$, $|\Pr[X_{i_1} \circ X_{i_2} \circ \dots \circ X_{i_k} = x] - 2^{-k}| \leq \varepsilon$.

A function $g : \{0, 1\}^d \rightarrow \{0, 1\}^n$ is an ε -almost k -wise independence generator in max norm if $g(U) = X = X_1 \circ \dots \circ X_n$ are ε -almost k -wise independent in max norm.

In the following passage, unless specified, when we say ε -almost k -wise independence, we mean in max norm.

THEOREM 3.1. ([2]) *There exists an explicit construction s.t. for every $n, k \in \mathbb{N}$, $\varepsilon > 0$, it computes an ε -almost k -wise independence generator $g : \{0, 1\}^d \rightarrow \{0, 1\}^n$, where $d = O(\log \frac{k \log n}{\varepsilon})$.*

The construction is highly explicit in the sense that, $\forall i \in [n]$, the i -th output bit can be computed in time $\text{poly}(k, \log n, \frac{1}{\varepsilon})$ given the seed and i .

4 Existential Bounds of Linear Codes

In this section we show existential bounds of linear codes for insdel errors. We have the following theorem.

THEOREM 4.1. Let \mathbb{F}_q denote the finite field of size q . For any $\delta > 0$ there exists a linear code family over \mathbb{F}_q that can correct up to δn insertions and deletions, with rate $(1 - \delta)/2 - H(\delta)/\log_2 q$. Here $H(\delta) = -\delta \log_2 \delta - (1 - \delta) \log_2(1 - \delta)$ is the binary entropy function.

Proof. Let n be a large enough integer and let m be an integer so that $m/n \leq (1 - \delta)/2 - H(\delta)/\log_2 q$. Consider the random linear code with encoding function $y = xG$, where $x \in \mathbb{F}_q^m$ is the message, $y \in \mathbb{F}_q^n$ is the codeword, and $G \in \mathbb{F}_q^{m \times n}$ is an m by n matrix where each entry is chosen independently and uniformly from \mathbb{F}_q .

To ensure that the code can correct up to $\delta n - 1$ insertions and deletions, we just need to make sure that every two different codewords have edit distance larger than $2\delta n$, or equivalently the length of the longest common subsequence is less than $(1 - \delta)n$. We now bound the probability that this does not happen. First we have the following claim the proof of which is deferred to the full version.

CLAIM 4.1. For any two different messages x^i, x^j and codewords $C^i = x^i G, C^j = x^j G$, consider any fixed common subsequence between C^i and C^j of length t , where the corresponding indices in C^i are $\{s_1, \dots, s_t\}$ and the corresponding indices in C^j are $\{r_1, \dots, r_t\}$. Then

$$\Pr[\forall k \in [t], C_{s_k}^i = C_{r_k}^j] \leq q^{-t}.$$

Using the claim, and noticing that the total number of possible cases where two strings of length n have a common subsequence of length $(1 - \delta)n$ is at most $\binom{n}{(1 - \delta)n}^2$, we have

$$\begin{aligned} & \Pr[\text{LCS}(C^i, C^j) \geq (1 - \delta)n] \\ & \leq \binom{n}{(1 - \delta)n}^2 q^{-((1 - \delta)n)} \\ & = \binom{n}{\delta n}^2 q^{-((1 - \delta)n)} \\ & \leq 2^{2H(\delta)n} q^{(\delta - 1)n}. \end{aligned}$$

Thus by a union bound

$$\Pr[\exists C^i \neq C^j, \text{LCS}(C^i, C^j) \geq (1 - \delta)n] < q^{2m} 2^{2H(\delta)n} q^{(\delta - 1)n}.$$

The right hand side is at most 1 as long as $m/n \leq (1 - \delta)/2 - H(\delta)/\log_2 q$. Thus the rate of the code family can be made to approach $(1 - \delta)/2 - H(\delta)/\log_2 q$. \square

In particular, for the case of binary and large finite alphabets we have the following corollaries.

COROLLARY 4.1. *For any $\delta > 0$ there exists a binary linear code family that can correct up to a fraction δ of insertions and deletions and which has rate at least $(1 - \delta)/2 - H(\delta)$.*

Note that this corollary implies Theorem 1.1.

COROLLARY 4.2. *For any $\delta > 0$ and $\epsilon > 0$ there exists a finite alphabet size $q = 2^{O(\epsilon^{-1})}$ and a q -ary linear code family that can correct a fraction δ of insertions and deletions, and which has rate at least $(1 - \delta)/2 - \epsilon$.*

This implies the existential result part of Theorem 1.2.

A systematic code is one which encodes a message x as the codeword $x \circ y$ where y are some check symbols computed from x . That is, the message appears in raw form at the beginning of the codeword. For Hamming errors, any linear code can be transformed into a systematic code by doing a basis change of the generator matrix, followed by a possible permutation of the symbols. This transformation does not change the parameters of code. For insertion and deletion errors, however, a permutation of the symbols may change the edit distance of the code, and hence it is not a priori clear that one can get a systematic linear code for insertion and deletion errors with the same parameters. Nevertheless, we now show that there exist such systematic linear codes with almost the same parameters. The proof appears in the full version.

THEOREM 4.2. *Let \mathbb{F}_q denote the finite field of size q . For any $\delta > 0$ there exists a systematic \mathbb{F}_q -linear code family that can correct up to δn insertions and deletions, with rate $(1 - \delta)/2 - H(\delta)/\log_2 q - o(1)$. Here $H(\delta) = -\delta \log_2 \delta - (1 - \delta) \log_2 (1 - \delta)$ is the binary entropy function.*

We can thus claim the analog of Corollaries 4.1 and 5.1 for systematic linear codes.

5 Outer Bound for Rate of Linear Codes

To complement our existence results for linear insdel codes from Section 4 we prove here a novel upper bound on the rate of any linear code that can correct a given fraction of edit errors:

Here is the restate of Theorem 1.3

THEOREM 5.1. (HALF-PLOTKIN BOUND) *Fix a finite field \mathbb{F}_q . Every \mathbb{F}_q -linear insdel code which is capable of correcting a $\delta > 0$ fraction of deletions has rate at most $\frac{1}{2}(1 - \frac{q}{q-1}\delta) + o(1)$.*

For binary codes this is a rate upper bound of $\frac{1}{2}(1 - 2\delta) + o(1)$. Furthermore since $\frac{q}{q-1} \geq 1$ for any q the following Half-Singleton bound follows directly for any linear codes for edit errors, independent of alphabet size q :

COROLLARY 5.1. (HALF-SINGLETON BOUND) *Every linear insdel code which is capable of correcting a $\delta > 0$ fraction of deletions has rate at most $\frac{1}{2}(1 - \delta) + o(1)$.*

This implies the negative result part of Theorem 1.2.

The naming for these two bounds stems from the fact that they prove outer bounds on the rate of linear codes which are exactly half as large as the Plotkin bound of $(1 - \frac{q}{q-1}\delta) + o(1)$ and the alphabet-independent Singleton bound of $(1 - \delta) + o(1)$ which hold for codes (not necessarily linear) with relative Hamming distance δ . Note that since two strings at fractional Hamming distance δ can lead to the same subsequence by deleting the δ fraction of symbols where they differ, the Plotkin and Singleton bounds also apply for general insdel codes.

Note that our Half-Singleton bound matches the existential bound for linear insdel codes given by Corollary 4.2 proving that the best limiting rate achievable by a linear code over a finite alphabet that can correct up to a δ fraction of insdel errors is exactly $\frac{1}{2}(1 - \delta)$.

The proof for our Half-Plotkin bound is significantly more involved than in the Hamming case. The Half-Singleton bound of Corollary 5.1 has a much simpler direct proof which we first present as a warm-up. Both our proofs (implicitly) build on the $\frac{1}{2}$ rate upper bound of [1] recorded below.

LEMMA 5.1. ([1]) *Let C be a linear code over any finite field \mathbb{F}_q with message length m and codeword length n . If C can correct even a single deletion, then $m/n \leq 1/2$.*

Proof. [Proof of Corollary 5.1 independent from Theorem 5.1] Let C be a linear code with message length m and codeword length n which can correct any δ fraction of deletions. For any two different codewords $c_1, c_2 \in C$, we must have $\text{ED}(c_1, c_2) \geq 2\delta n + 1$. Now consider the new code C' obtained by removing the first $\delta n - 1$ symbols of every codeword in C . For any two different codewords $c'_1, c'_2 \in C'$, we must have $\text{ED}(c'_1, c'_2) \geq 3$. This means that C' is a linear code capable to correct a single deletion. Thus Lemma 5.1 implies that $m/(n - \delta n + 1) \leq 1/2$. Therefore $m/n \leq \frac{1}{2}(1 - \delta) + \frac{1}{2n}$, giving our desired upper bound on rate. \square

We now return to the proof of the Half-Plotkin bound of Theorem 5.1. The proof will make crucial use of the following technical lemma.

LEMMA 5.2. *Suppose $A \subseteq \mathbb{F}_q^{2d}$ is a subspace of dimension $t > 16q$. Then A contains a non-zero vector $v = (v_1, v_2)$ with both $v_1, v_2 \in \mathbb{F}_q^d$ containing at least $\frac{d}{q}(1 - 4\sqrt{\frac{q}{t}})$ zeros.*

We present the proof of the above lemma in the full version but elaborate here the main non-trivial aspect tackled by it. Note that a random vector in A contains in expectation at least $\frac{2d}{q}$ zeros and each of the halves v_1 and v_2 contains in expectation at least $\frac{d}{q}$ zeros. The crux of Lemma 5.2 therefore lies in proving that there cannot exist correlations between the expectations for the two halves that are strong enough to prevent both expectation bounds to be satisfied simultaneously. Note that the $\frac{2d}{q}$ quantity of total number of zeros in a non-zero vector is tight (up to lower order terms) whenever A is a good low-rate Hamming code. Indeed the minimum-Hamming distance of a linear Hamming code is exactly equal to the weight of the non-zero codeword with the lowest Hamming weight and codes matching (up to lower order terms) the $2d(1 - 1/q)$ Hamming-Distance given by the Plotkin bound are known. For such a good code, Lemma 5.2 therefore guarantees a codeword with (up to lower order terms) minimal Hamming weight in which the $\frac{2d}{q} \pm o(1)$ zeroes are (up to lower order terms) exactly equally distributed between the two halves of the codeword.

Proof. [Proof of Theorem 5.1] Suppose that C is a q -ary linear insdel code of block-length n which is capable of correcting any δn deletions and suppose, for sake of contradiction that C has a rate of at least $(1 - \frac{q}{q-1}\delta)/2 + \epsilon$ for some $\epsilon > 0$.

Let C' be the linear code which encodes any input like C but then deletes the first $d = \frac{q}{q-1}\delta n - \frac{\epsilon n}{5}$ symbols.

Let C'_1 be the linear code which encodes any input like C' but then deletes the first symbol.

Let C'_2 be the linear code which encodes any input like C' but then deletes the last symbol.

For sake of contradiction, assume that $\dim(C'_1 \cap C'_2) > \frac{\epsilon n}{5}$ and let $C'' \subseteq (C'_1 \cap C'_2) \subseteq \mathbb{F}_q^{n-d-1}$ be some linear sub-space of \mathbb{F}_q^{n-d-1} with dimension $\frac{\epsilon n}{5}$ which does not contain the all-ones vector $(1, \dots, 1)$.

Note that by definition of C'_1, C'_2 and C'' , every non-zero $c'' \in C''$ can be completed to two vectors $c = (v, \sigma, c'')$ and $c' = (v', c'', \sigma')$ in C , where $v, v' \in \mathbb{F}_q^d$ and $\sigma, \sigma' \in \mathbb{F}_q$. Note that c and c' cannot be identical because they each contain a copy of c'' albeit shifted by one symbol and c'' does not consist of only identical symbols since such vectors are excluded from C'' .

Consider a basis $c''_1, c''_2, \dots \in C''$ of C'' . For each such basis vector c''_i fix one pair of completions $v_i, v'_i \in$

\mathbb{F}_q^d and let $A \subseteq \mathbb{F}_q^{2d}$ be the subspace spanned by all such vectors, i.e., $A = \langle \{(v_i, v'_i) \mid i\} \rangle$. Note that each non-zero linear combination of basis vectors indeed gives rise to the same linear combination of its completion vectors, justifying the definition of A as a subspace. Furthermore we claim that $\dim(A) = \dim(C'') = \frac{\epsilon n}{5}$, i.e., that the set of vectors $\{(v_i, v'_i) \mid 1 \leq i \leq \epsilon n/5\}$ is linearly independent. Indeed suppose that there is a non-trivial linear combination $c'' \neq 0$ of basis vectors in C'' that completes to $(v, v') = 0$. In this case C would contain two codewords starting with d zeros followed by $n - d$ symbols containing a (shifted) copy of c'' plus one extra symbol (at the very end or after the d zeros respectively). The edit distance between these two different codewords in C is merely 2, contradicting the assumption that C can correct a much larger number of deletions. By Lemma 5.2, this means that there exists a $(v, v') \in A$ in which both v and v' contain a $1/q - o(1)$ fraction of 0's. The above implies the existence of codewords $c, c' \in C$ starting with v and v' respectively, again contradicting the distance property of C because v and v' can be transformed into an all-zero string using $(1 - 1/q + O_{q,\epsilon}(\frac{1}{n}))d = (1 - 1/q + o(1))d < \delta n - 2$ deletions while the latter parts of c and c' can be made equal using a single deletion. This completes the proof that $\dim(C'_1 \cap C'_2) \leq \frac{\epsilon n}{5}$.

We now have that

$$\begin{aligned} n - d - 1 &\geq \dim(C'_1 \cup C'_2) \\ &= \dim(C'_1) + \dim(C'_2) - \dim(C'_1 \cap C'_2) \\ &\geq 2(\dim(C') - 1) - \frac{\epsilon n}{5}, \end{aligned}$$

and therefore

$$\dim(C') \leq \frac{1}{2}(n - d + \frac{3\epsilon n}{5}) = \frac{1}{2}(1 - \frac{q}{q-1}\delta)n + \frac{2\epsilon n}{5}.$$

Note that it is possible for the channel to take any sent codeword of C , determine the most frequent symbol σ in the the first d coordinates, and delete all symbols in these d coordinates but $\frac{d}{q}$ occurrences of σ , since $d(1 - \frac{1}{q}) \leq \delta n$ is below the budget of allowed deletions. This means that after this operation all codewords in C are still distinct. Furthermore, the span of the corrupted codewords received through this channel would form a subspace of dimension at most $\dim(C') + 1$, since any corrupted codeword is of the form $\alpha^{d/q}$ followed by a codeword in C' , for some $\alpha \in \mathbb{F}_q$. This contradicts the assumption that C has a rate of at least $\frac{1}{2}(1 - \frac{q}{q-1}\delta) + \epsilon$. \square

6 Monte Carlo construction of linear insdel codes

In this section we give a randomized Monte Carlo construction of linear codes s.t. with high probability over the randomness, we get an asymptotically good linear code for edit distance, with an efficient decoding algorithm for a constant fraction of insdel errors.

The construction uses a random sequence where each element in the sequence is a string of 0-symbols having length uniformly chosen from $\{1, 2, \dots, a\}$ where a is an integer specified later. This sequence serves as a special kind of synchronization strings first introduced in [14], then explicitly constructed in [16, 6].

The high level idea is to use this sequence to cut a codeword y for Hamming distance into blocks, and argue that a carefully designed matching algorithm, which is matching the edited codeword and a “templet”, can recover most of the symbols in y as long as the total number of edit errors is bounded.

6.1 Construction Let C be an $(n_C, m, 2\kappa_C + 1)$ linear ECC over $\mathbb{F}_q = \{0, 1, \dots, q - 1\}$ that can correct κ_C Hamming errors.

Let $S = (S[1], S[2], \dots, S[n_C])$ be a sequence of independent random variables s.t. $S[i]$ is an all-0 string with length uniformly chosen from $\{1, 2, \dots, a\}$, where $a = O(\frac{n_C}{\kappa_C})$ and the constant hiding in big-O is large enough. Once chosen, both the encoder and the decoder will use the same S .

We construct an $(n, m, 2\kappa + 1)$ linear code for edit distance over \mathbb{F}_q , with $n = O(n_C)$, correcting $\kappa = 0.01\kappa_C$ insdel errors.

CONSTRUCTION 6.1. (ENCODING) *The encoding operates as follows:*

1. Input message $x \in \mathbb{F}_q^m$;
2. Compute $y = C(x)$;
3. Let

$$z = S[1] \circ y[1] \circ S[2] \circ y[2] \circ \dots \circ S[n_C] \circ y[n_C];$$

4. Output codeword $z \in \mathbb{F}_q^n$.

CONSTRUCTION 6.2. (DECODING) *The decoding algorithm is as follows:*

1. Input $z' \in \mathbb{F}_q^{n'}$, $n' = \Theta(n)$;
2. Let $z_\? = S[1] \circ ? \circ S[2] \circ ? \circ \dots \circ S[n_C] \circ ? \in \{0, ?\}^n$, where “?” is a special symbol different from elements in \mathbb{F}_q , indicating a blank to be filled, and 0-symbol is the 0-element in \mathbb{F}_q ;

3. Compute the ?-to-non-zero matching w between $z_\?$ and z' , by Construction 6.3;
4. Fill the blanks by using the matching w , to get y' which has a certain hamming distance from y ; (Each blank is filled with the matched non-zero symbol, if there is no such a match, fill it with a 0-symbol.)
5. Apply the decoding of C on y' to get x .
6. Output x .

Here we define the ?-to-non-zero matching between two strings $z_\? \in \{0, ?\}^*$ and $z' \in \{0, 1, \dots, q, ?\}^*$, to be a monotone matching, and further, every match in such a matching is a pair of indices (i, j) that matches an i -th ?-symbol of $z_\?$, to a j -th non-zero symbol of z' .

Given a ?-to-non-zero matching $w = ((i_1, j_1), \dots, (i_{|w|}, j_{|w|}))$, we define the cost function $\text{cost}(w)$ to be

$$\text{cost}(w) = \sum_{k=1}^{|w|} \mathbb{1}(p_{i_k} - p_{i_{k-1}} \neq q_{j_k} - q_{j_{k-1}}).$$

Here p_{i_k} is the position of the i_k -th ?-symbol in $z_\?$, q_{j_k} is the position of the j_k -th non-zero-symbol of z' . $\mathbb{1}$ is the indicator function.

Next we show our matching algorithm which returns a matching w that maximizes $\text{obj}(w) = |w| - \text{cost}(w)$.

CONSTRUCTION 6.3. *On input strings $z_\?$ and z' , the algorithm returns a ?-to-non-zero matching w .*

- Let n_1 denote the number of non-zero-symbols in z' .
- Let p_i denote the position of the i -th question mark in $z_\?$, and q_j denote the position of the j -th non-zero in z' .
- We want to use $f[i][j], i \in [n_C], j \in [n_1]$ to record the maximum $\text{obj}(w)$, among all w , each being a ?-to-non-zero matching between $z_\?[1, p_i]$ and $z'[1, q_j]$ with the last match being (i, j) .

1. Initialization: $f[i][j] = 0, i \in \{0, 1, \dots, n_C\}, j \in \{0, 1, \dots, n_1\}$;
2. Transition step: For $i = 1$ to n_C , $j = 1$ to n_1 ,
 - if (i, j) can be a match, then

$$f[i][j] = \max_{i' < i, j' < j} \{f[i'][j'] + \mathbb{1}(p_i - p_{i'} = q_j - q_{j'})\},$$

where if there is no such (i', j') , then $p_{i'} = q_{j'} = 0$;

- Else $f[i][j] = 0$;

3. Output $\max_{i,j} \{f[i][j]\}$.

To output the corresponding matching, we only need to store the corresponding matching along with each time we compute an $f[i][j]$. For the initialization step, those matchings are all empty. For the transition step,

- if $p_i - p_{i'} = q_j - q_{j'}$, then the corresponding matching to $f[i][j]$ is the concatenation of the matching for $f[i'][j']$ and the match (i, j) , where i', j' are the indices we pick when taking the maximum;
- else the corresponding matching to $f[i][j]$ is an empty matching.

6.2 Analysis

LEMMA 6.4. *The ?-to-non-zero matching algorithm returns a matching which maximizes obj.*

Proof. We use induction to show that $f[i][j]$ records the maximum $\text{obj}(w)$ among all matchings between $z?[1, p_i]$ and $z'[1, q_j]$ whose last match is (i, j) .

The base case is $f[1][j], j \in [n_1], f[i][1], i \in [n_C]$. $f[1][j]$ is one if $p_1 = q_j$, otherwise it's 0. $f[i][1]$ is one if $p_i = q_1$, otherwise it's 0. Notice that the transition function indeed achieves this.

For the induction case, in order to compute $f[i][j], i \in [n_1], j \in [n_C]$, we assume every $f[i'][j'], i' < i, j' < j$ and its corresponding matching with maximum target function value has already been computed.

Let w be the matching attained when $f[i][j]$ is computed. So $f[i][j] = |w| - \text{cost}(w)$. Notice that w can be divided into two parts. The first part w_1 includes every match of w except for the last one. So it is a matching between $z?[0, p_{i'}]$ and $z'[0, q_{j'}]$, where (i', j') is the second last match in w and also the last match in w_1 . The second part is the last match (i, j) of w .

We claim that $\text{obj}(w_1)$ has to be $f[i'][j']$. Suppose it is not. Then it cannot be larger because $f[i'][j']$ is the maximum by the induction hypothesis. So it can only be smaller. But then we can use the matching w'_1 corresponding to $f[i'][j']$ to replace the first part and get another matching $w' = w'_1 \circ (i, j)$ s.t.

$$\begin{aligned}
 & |w'| - \text{cost}(w') \\
 (6.1) \quad & = |w'_1| - \text{cost}(w'_1) + \mathbb{1}(p_i - p_{i'} = q_j - q_{j'}) \\
 & > |w_1| - \text{cost}(w_1) + \mathbb{1}(p_i - p_{i'} = q_j - q_{j'}) \\
 & = |w| - \text{cost}(w).
 \end{aligned}$$

The first equation is because of the definition of cost and the structure of w' . The inequality is because our

picking of w'_1 . The last equation is because the definition of cost . However (6.1) contradicts the definition of $f[i][j]$. This shows the claim.

As a result $f[i][j]$ has to be in the form $\max_{i' < i, j' < j} \{f[i'][j'] + \mathbb{1}(p_i - p_{i'} = q_j - q_{j'})\}$. This is exactly what the algorithm computes. So it can compute $f[i][j]$ and get the corresponding matching correctly. This shows the induction step. \square

LEMMA 6.5. *The number of non-zeros in z' that are not matched, is at most 3κ .*

Proof. Consider the intuitive matching w^* induced by matching those non-inserted non-zero-symbols in z' to their original positions in z . The original position of a symbol means the index of this symbol in z . Since $z?$ and z has the same length, these original positions are also one-on-one correspond to positions in $z?$. When there are no edits, $\text{obj}(w^*)$ is equal to τ , which is the number of non-zero-symbols in z . Now we analyze the effects of κ insdel errors on obj .

Each insertion can

- insert a 0-symbol and thus increase the number of consecutive 0-symbols immediately before a non-zero-symbol by one;
- insert a non-zero-symbol.

Both bullets can decrease obj by at most one. Because they can increase cost by at most one, as each insertion may change the distance between the two matched symbols immediately before and after that inserted symbol.

Each deletion can

- delete a 0-symbol and thus decrease the number of consecutive 0-symbols immediately before a non-zero-symbol by one;
- delete a non-zero-symbol.

The first bullet will decrease obj by one as it changes the distance between the two matched symbols immediately before and after that deleted symbol, and thus increasing cost by at most one. The second bullet will decrease obj by ≤ 2 since it can cancel a possible match and also changes the distance between the two matched symbols immediately before and after that deleted symbol.

Thus κ insdel errors can decrease the target function value by at most 2κ .

Let the number of non-zero symbols in y be τ . Notice that $\tau \geq 2\kappa_C + 1$ number of non-zeros, since C

is a linear code with distance $2\kappa_C + 1$. By the previous analysis,

$$\text{obj}(w^*) = |w^*| - \text{cost}(w^*) \geq \tau - 2\kappa.$$

For any matching w' whose number of unmatched non-zeros is more than 3κ , we have

$$|w'| - \text{cost}(w') < \tau - 2\kappa \leq |w^*| - \text{cost}(w^*).$$

Here the first inequality is because there can be at most κ inserted non-zeros and so there are more than $3\kappa - \kappa = 2\kappa$ unmatched non-zero-symbols which are from the original z . Hence by the definition of the target function, its value must be less than $\tau - 2\kappa$.

Therefore w' cannot be the maximum matching computed by the algorithm. Hence the number of unmatched non-zeros in z' cannot be more than 3κ . \square

Next we show that the decoding can correct κ insdel errors.

DEFINITION 6.6. *A match (i, j) is bad, if the j -th non-zero symbol of z' is the i_j -th symbol of z , but $i_j \neq i$.*

Note that there are some non-zero-symbols in z' which are inserted, we will not consider them here i.e. the badness are only defined over matches whose second entry points to a symbol of z' , which is already in z before insertions and deletions. We denote $\Delta(w)$ as the number of bad matches in w . When w is clear in the context, we simply use Δ .

LEMMA 6.7. *Suppose $\text{HD}(y, y') = t > \kappa_C$, then $\Delta(w) > (\kappa_C - 6\kappa)/2$.*

Proof. The t Hamming errors come from two different types of matches, which are

1. an i -th question mark is not matched, but $y[i] \neq 0$;
2. an i -th question mark is matched, but $y[i] = 0$.

We denote the number of first type errors to be t_1 , and the number of second type errors to be t_2 . So

$$(6.2) \quad t_1 + t_2 = t > \kappa_C.$$

For t_2 , the errors are only from two sources, i.e. the non-zero-symbol in such a match is from (1) inserted non-zero-symbols or (2) non-zero-symbols in y . Each inserted non-zero-symbol may be in one of such matches, contributing to one such error. Each non-zero symbol in y may be matched to a question mark which

should a 0, contributing to one such error, noticing that this can only happen when the match is bad. Hence,

$$(6.3) \quad t_2 \leq \kappa + \Delta.$$

For t_1 , we are actually counting the number of non-zero symbols in z that are not matched, because y is a subsequence of z which contains all the non-zeros in z .

We claim that t_1 is upper bounded by the summation of

1. the number of non-zeros that are deleted from z ,
2. the number of non-zeros in z' that are not matched,
3. the number of non-zeros in z' that are matched to some i -th question mark but $y[i] = 0$.

This is because t_1 is the number of not matched non-zero symbols in z . So it pluses the number of matched non-zero-symbols in z , is equal to the total number of non-zero-symbols in z .

$$(6.4) \quad t_1 + \text{the number of matched non-zero-symbols in } z = \text{the total number of non-zero-symbols in } z.$$

On the other hand, the total number of non-zero-symbols in z is at most the number of non-zero-symbols in z' pluses the number of deleted non-zero-symbols (term 1). And the number of non-zero-symbols in z' is the summation of term 2, term 3 and the number of matched non-zero-symbols in z . This is because the number of non-zeros in z' that are matched to some i -th question mark and $y[i] \neq 0$ is equal to the number of matched non-zero-symbols in z . Hence

$$(6.5) \quad \begin{aligned} & \text{the total number of non-zero-symbols in } z \\ &= \text{term 1} + \text{term 2} + \text{term 3} \\ &+ \text{the number of matched non-zero-symbols in } z. \end{aligned}$$

By equality 6.4 and 6.5, t_1 is bounded as desired.

The number of non-zero-symbols deleted can be at most κ , since there are at most κ deletions. By Lemma 6.5 the 2nd term is at most 3κ . The third item is at most t_2 by their definitions. So

$$(6.6) \quad t_1 \leq 4\kappa + t_2.$$

By equality (6.2), (6.3), (6.6), $\Delta > (\kappa_C - 6\kappa)/2$. \square

For a bad match (i, j) , recall that i_j be the original position in z of the j -th non-zero-symbol of z' . Note that since we only define bad match on some j which

points to a non-zero-symbol that is not deleted from z , thus there is always an i_j for j .

An interesting property of those matches in w that are not involved with inserted symbols, is that they induce a self-matching between $z_?$ and $z_?$. We denote this matching as \tilde{w} , i.e. each match (i, i_j) of \tilde{w} is a match (i, j) in w s.t. j points to a non-zero symbol which is not deleted from z . Since bad matches are only defined for matches whose second entries are non-zero symbols of z , $\Delta(\tilde{w}) = \Delta(w)$.

LEMMA 6.8. *With probability $1 - 2^{O(n_C)} a^{-\Theta(\Delta)}$ over S , for any self-matching \tilde{w} on $z_?$ having Δ bad matches, there are at most 0.1Δ bad matches (i, j) s.t. $p_i - p_{i'} = p_j - p_{j'}$, and hence $\text{cost}(\tilde{w}) \geq 0.9\Delta$.*

Notice that the probability is meaningful when both a and Δ are large enough and hence $a^{-\Theta(\Delta)}$ is the dominating term. We will see later that this is indeed the case in our setting of parameters.

Proof. Consider a specific sequence of Δ bad matches. We first claim that with probability $\geq 1 - a^{-\Theta(\Delta)}$, for any matching \tilde{w} containing this specific sequence of bad matches, there are $\Delta' < 0.1\Delta$ bad matches (i, j) s.t. $p_i - p_{i'} = p_j - p_{j'}$.

For a bad match (i, j) in the sequence, if it is the first match in the matching, then $p_i \neq p_j$ since $|S[i]|, |S[j]| > 0$. If it has previous match (i', j') , then (i', j') is either a bad match or a good (not bad) match. If it is a good match, then $p_i - p_{i'} \neq p_j - p_{j'}$, since $|S[k]| > 0, \forall k$. If it is a bad match, then conditioned on $S[1], \dots, S[\max(i, j) - 1]$ being arbitrary fixed values, $p_i - p_{i'} = p_j - p_{j'}$ happens with probability at most $1/a$. Because, in this condition, $p_i - p_{i'} = p_j - p_{j'}$ happens only if $S[\max(i, j)]$ is equal to one specific value in $[a]$.

So for any matching \tilde{w} including this specific sequence of bad matches, the probability that there are $\Delta' \geq 0.1\Delta$ bad matches (i, j) s.t. $p_i - p_{i'} = p_j - p_{j'}$, is at most $a^{-\Delta'}$.

The number of sequences of Δ bad matches is at most

$$\binom{n_C}{\Delta} \binom{n_C}{\Delta} \leq \left(\frac{en_C}{\Delta}\right)^{2\Delta}.$$

So by a union bound, with probability $1 - \left(\frac{en_C}{\Delta}\right)^{2\Delta} a^{-\Theta(\Delta)}$, there are at most 0.1Δ bad matches s.t. (i, j) s.t. $p_i - p_{i'} = p_j - p_{j'}$, and hence $\text{cost}(\tilde{w}) \geq (1 - 0.1)\Delta = 0.9\Delta$. \square

LEMMA 6.9. *For S being any string in $\text{supp}(S)$, for any κ insdels, $\text{cost}(w) \geq \text{cost}(\tilde{w}) - \kappa$.*

Proof. Recall that \tilde{w} consists of all (i, i_j) s.t. (i, j) is a match in w and j points to a non-inserted symbol. Each insdel will only decrease cost function by 1, since for each match (i, j) , $\mathbb{1}(p_i - p_{i'} = q_j - q_{j'})$ only involves (i, j) and its immediate previous match. As there are at most κ insdels, the lemma holds. \square

LEMMA 6.10. *With probability $1 - 2^{-\Omega(\kappa_C)}$ over S , the encoding and decoding give a code which can correct κ insdels.*

Proof. Suppose the decoding cannot compute x correctly. Then it must be the case that $\text{HD}(y, y') > \kappa_C$. Hence by Lemma 6.7, $\Delta(w) > (\kappa_C - 6\kappa)/2$. Then by definition of \tilde{w} , $\Delta(\tilde{w}) = \Delta(w) > (\kappa_C - 6\kappa)/2$. By Lemma 6.8, and since $\Delta > (\kappa_C - 6\kappa)/2, a = O(n_C/\kappa_C)$, with probability $1 - 2^{O(n_C)} a^{-\Theta(\Delta)} = 1 - a^{-\Theta(\Delta)}$ over S , for any \tilde{w}' having Δ bad matches, we should have $\text{cost}(\tilde{w}') \geq 0.9\Delta$. So for this \tilde{w} , we should also have $\text{cost}(\tilde{w}) \geq 0.9\Delta$. and hence by Lemma 6.9, $\text{cost}(w) \geq 0.9\Delta - \kappa \geq 0.9(\kappa_C - 6\kappa)/2 - \kappa > 2\kappa$, recalling that $\kappa = 0.01\kappa_C$. So $\text{obj}(w) = |w| - \text{cost}(w) \leq |w| - 2\kappa$. However, by Lemma 6.4, $\text{obj}(w)$ is maximized, which is at least $\tau - 2\kappa$. By definitions of τ and w , this is at least $|w| - 2\kappa$. A contradiction is reached. \square

Next we show that the codeword length is $O(n_C)$.

LEMMA 6.11. *For any $\gamma > 0$, with probability at most $e^{-\Omega(\gamma n_C)}$,*

$$n \in [(1 - \gamma)n_C(a + 1)/2 + n_C, (1 + \gamma)n_C(a + 1)/2 + n_C].$$

Proof. Since the length of $S[i]$ is uniformly chosen to be $1, \dots, a$, the expectation of their total length is $n_C(a + 1)/2$. By a Chernoff bound, with probability at least $1 - e^{-\Omega(\gamma n_C)}$, $\sum_{i \in [n_C]} |S_i| \in [(1 - \gamma)n_C(a + 1)/2, (1 + \gamma)n_C(a + 1)/2]$. So $n = n_C + \sum_{i \in [n_C]} |S_i| \in [(1 - \gamma)n_C(a + 1)/2 + n_C, (1 + \gamma)n_C(a + 1)/2 + n_C]$. \square

The efficiency of the encoding and decoding follows directly from the construction.

LEMMA 6.12. *The encoding is in linear time of the encoding time of C ; The decoding is in time $O(n^4)$ plus the decoding time of C .*

We remark that the $O(n^4)$ time in decoding is because of the dynamic programming matching algorithm.

THEOREM 6.1. *If there is an explicit q -ary linear $(n_C, m, 2\kappa_C + 1)$ code for Hamming distance, then for any $\alpha > 0$, there is a q -ary linear $(n, m, 2\kappa + 1)$ code for edit distance, with $n \in$*

$[(1 - \alpha)n_C(a + 1)/2 + n_C, (1 + \alpha)n_C(a + 1)/2 + n_C], \kappa = 0.01\kappa_C, a = O(\frac{n_C}{\kappa_C})$. The encoding is randomized with failure probability at most $2^{-\Omega(\kappa_C)} + 2^{-\Omega(\alpha a n_C)}$ over the random choice of S . The decoding is explicit when given S , correcting κ insdels.

Proof. Construction 6.1 6.2 gives such a code. Message length and the unique decoding radius directly follows from the construction. Codeword length follows from 6.11. Correctness follows Lemma 6.10. Computing efficiency follow from 6.12.

Note that the code constructed is linear. Since by the construction, C is a linear code and the encoding is just inserting some number of zero-symbols to some fixed positions of $y = C(x)$. So taking arbitrary two codewords z_1, z_2 corresponding to $y_1, y_2 \in C$, their summation is still a codeword z which corresponds to $y_1 + y_2 \in C$. Also, taking any codeword z , multiplying every coordinate with σ in \mathbb{F}_q , the result is still a codeword, because the multiplication does not change the zero-symbols and for y , multiplying σ with every coordinate results a codeword in C since C is a linear code. So the constructed code is linear. \square

We state the following theorem to explicitly describe the error rate and information rate. It immediately follows from Theorem 6.1, by taking α to be a constant.

THEOREM 6.2. *If there is an explicit linear code for Hamming distance with block length n_C , error rate δ_C , information rate γ_C , then there is a linear code for insdel errors, with block length $O(n_C/\delta_C)$, error rate $\delta = O(\delta_C^2)$, information rate $O(\delta_C \cdot \gamma_C)$. The encoding is randomized with failure probability at most $2^{-\Omega(\delta_C n_C)}$ over the random choice of S . The decoding is explicit when given S , correcting δ fraction of insdel errors.*

Notice that this directly implies Theorem 1.4, since there are various explicit constructions of asymptotically good linear code for Hamming distance.

7 Derandomization and explicit construction of linear insdel codes

In this section, we derandomize the encoding of the previous section. The key idea is to develop a special string s to replace the random sequence S .

7.1 Synchronization Separator Sequence We use the following terminologies about self-matchings for strings consists of 0s and question marks.

For a sequence $S[1] \circ S[2] \circ \dots \circ S[n] \circ$, a $?$ -to- $?$ self matching w is a monotone matching between

two substrings of this sequence s.t. each match (i, j) matches the i -th question mark to the j -th question mark.

DEFINITION 7.1. *A match (i, j) in such a matching is called undesired if*

- $i \neq j$ (i.e. bad);
- $p_i - p_{i'} = p_j - p_{j'}$.

Here (i', j') is the immediate previous match of (i, j) in the matching. If there is no such (i', j') then the second bullet just checks $p_i = p_j$ instead.

Next we show the string defined as follows can be used to replace the random sequence $S[1], S[2], \dots, S[n]$.

DEFINITION 7.2. $s[1], s[2], \dots, s[n]$ is called a (Λ, a) synchronization separator sequence, if for any self-matching of $s[1] \circ s[2] \circ \dots \circ s[n] \circ$, the number of undesired matches is at most Λ .

Here each $s[i], i \in [n]$ is an all 0 string with length in $\{1, \dots, a\}$.

We show that synchronization separator sequences can be constructed explicitly.

LEMMA 7.3. *There is an explicit construction of (Λ, a) synchronization separator sequences of length n , for every $n \in \mathbb{N}, \Lambda \in \mathbb{N}, a = \Theta(\frac{n}{\Lambda})$ being a power of 2 and the constant hiding in Θ large enough.*

We give the following construction.

CONSTRUCTION 7.4. 1. Let $g : \{0, 1\}^{d_g} \rightarrow \{0, 1\}^{n_g = n \log a}$ be an ε_g almost n_g -wise independence generator, with $d_g = O(\log \frac{n_g \log n_g}{\varepsilon_g^2})$, $\varepsilon_g = 1/\text{poly}(n)$, from Theorem 3.1;

2. For every $r \in \{0, 1\}^{d_g}$,

(a) compute $g(r)$, and then partition it into a sequence of length $\log a$ blocks, where the i -th block, $i \in [n]$, corresponds to a binary number $a_i \in \{1, \dots, a\}$;

(b) Generate sequence $s \in \{0^*\}^n$ s.t. $\forall i \in [n], s[i] = \underbrace{0 \circ 0 \circ \dots \circ 0}_{a_i}$;

(c) For some large enough constant c , for every pair of substrings u, v of s , s.t. $|u| + |v| \leq c \frac{2n \log n}{\Lambda \log \frac{n}{\Lambda}}$, compute a matching which only consists of bad matches, having size $\Delta' = 2c \frac{\log n}{\log \frac{n}{\Lambda}}$ and the maximum number of undesired matches;

- (d) Let Λ_0 be the maximum number of undesired matches, among all the matchings computed in the above step;
- (e) If $\Lambda_0 < c \frac{\log n}{\log \frac{n}{\Lambda}}$, then halt and output s .

3. Abort.

LEMMA 7.5. For a string $s \in \Sigma^*$, if there is a self-matching w between two substrings u, v of s , having Λ undesired matches, then for every $i \in \mathbb{N}$, there is a self-matching w' , having at most $\Lambda' = \Lambda/2^i$ undesired matches, between two substrings u', v' , where u' is a substring of u , v' is a substring of v , and $\frac{\Lambda'}{|u'|+|v'|} \geq \frac{\Lambda}{|u|+|v|}$.

We use the matching cutting strategy similar to that of [8]. The difference is that in [8] the object of interest was the matching size. But here we care about the number of undesired matches in the matching. We cannot just pick out these undesired matches to form a new matching to consider, because the “undesired” property depends on other matches which may not be an undesired matches.

Proof. We repeatedly use the following cutting technique on the current matching w . We denote ratio $\gamma = \frac{\Lambda}{l}$ and length $l = |u| + |v|$.

We find the match (i, j) s.t. the number of undesired matches from the first match to (i, j) (include), is $\Lambda_1 = \lfloor \Lambda/2 \rfloor$. Thus the number of undesired matches from (i, j) (not include) to the end of w is $\Lambda_2 = \Lambda - \lfloor \Lambda/2 \rfloor$. Here (i, j) also divides w into two halves. The first half matching w_1 is from the first match to (i, j) (include). The second half from (i, j) (not include) to the last match. Note that u, v are also divided by (i, j) . Let l_1 be the length of the summation of the lengths of first halves of u, v , including the symbols that i, j pointed to. For the second halves it's l_2 , not including the symbols that i, j pointed to. We compare Λ_1/l_1 and Λ_2/l_2 . Pick the larger one, and then consider the corresponding halves of w and u, v .

The procedure ends until we get a pair of strings u, v and a matching w between them s.t. the number of undesired matches is at most Λ' .

Notice that $\Lambda_1 + \Lambda_2 = \Lambda$, $l_1 + l_2 = |u| + |v|$. Let $\gamma_1 = \Lambda_1/l_1$, $\gamma_2 = \Lambda_2/l_2$. Then one of γ_1 and γ_2 has to be at least γ . Otherwise, $\Lambda/l = \frac{\Lambda_1 + \Lambda_2}{l_1 + l_2}$ has to be smaller than γ , causing a contradiction. So the one we pick has ratio at least γ . Also the number of undesired matches is divided into two halves each time we apply a cut.

As a result, after doing cutting for i times, we get the desired matching and substrings. \square

Next we prove Lemma 7.3.

Proof. [Proof of 7.3]

We use Construction 7.4.

We show that there is a seed $r \in \{0, 1\}^{d_g}$ which can let Construction 7.4 output the $s = s(r)$, which is a synchronization separator sequence we want.

The proof strategy is to consider a uniform random seed r and show that with some probability the construction can output a desired s .

If there is a self-matching w having Λ undesired matches, then by Lemma 7.5, there is a self-matching w' between some u' and v' which are substrings of s , s.t. there are $\Lambda' = \Lambda/2^i = c \frac{\log n}{\log \frac{n}{\Lambda}}$ undesired matches in w' , for some large enough i and constant c , where $|u'| + |v'| = l' \leq c \frac{2n}{\Lambda} \frac{\log n}{\log \frac{n}{\Lambda}}$. Also we can think of w' as only consisting of bad matches since good matches cannot be undesired by definition and for any undesired match, its immediate previous match has to be a bad match due to $s[k] > 0, \forall k$. Also we only need to consider w' of size $\Delta' = 2\Lambda'$ since there is always a size w' matching which can contain the Λ' undesired bad matches.

On the other hand, for arbitrary u, v in the algorithm, we first consider that $s[1], \dots, s[n]$ are at uniform random. The probability that a undesired match (i, j) in w' happens, conditioned on fixing $s[1], \dots, s[\max(i, j) - 1]$ is at most $1/a$, since we need $s[\max(i, j)]$ to take a specific value to make $p_i - p_{i'} = p_j - p_{j'}$. So the probability that a specific matching with Λ' undesired matches, happens with probability $\rho \leq a^{-\Lambda'}$.

Next we consider using ε_g almost n -wise independence. Each sample point instead happens with probability at most $\rho + \varepsilon_g \leq 1.1\rho$ when $\varepsilon_g = 1/\text{poly}(n)$ is small enough. Hence the probability that a specific matching with Λ' undesired matches, happens with probability at most $1.1a^{-\Lambda'}$.

There are at most n^4 such pair of s'_1, s'_2 . For each pair the number of matchings w' between s'_1, s'_2 is at most $\binom{l'}{\Delta'} = \text{poly}(n)$ since we only consider w' with size Δ' . So by a union bound, with probability $1 - \text{poly}(n)a^{-\Lambda'} = 1 - 1/\text{poly}(n)$, there is no size Δ' matching between any pair of substrings s'_1, s'_2 with $l' \leq c \frac{2n}{\Lambda} \frac{\log n}{\log \frac{n}{\Lambda}}$ having Λ' undesired matches, and hence there is no self-matching with Λ undesired matches for the string s .

We can find a seed r which gives us such a string $s(r)$, in polynomial time. This is because $|r| = O(\log n)$ and in the above computation the number of testings is $\text{poly}(n)$. The testings are also efficient. Thus an

exhaustive search for a (Λ, a) synchronization separator sequence can be done in polynomial time. \square

7.2 Deterministic Encoding

THEOREM 7.1. *If there is an explicit linear q -ary $(n_C, m, 2\kappa_C + 1)$ code for Hamming distance, then there is an explicit linear q -ary $(n, m, 2\kappa + 1)$ code for edit distance, where $n = O(an_C)$, for $a = \Theta(\frac{n_C}{\kappa_C})$, $\kappa = 0.01\kappa_C$.*

Proof. We replace the random sequence $S[1], S[2], \dots, S[n_C]$ in Construction 6.1 by a (Λ', a) synchronization separator sequence from Lemma 7.3. Here we let $\Lambda' = 0.2\kappa_C$, $a = \Theta(\frac{n_C}{\Lambda'})$.

Suppose $\text{HD}(y, y') > \kappa_C$. By Lemma 6.7 the number of bad matches in w is at least $(\kappa_C - 6\kappa)/2 - \kappa \geq 0.4\kappa_C$. Recall that w is the matching returned by the matching algorithm during the decoding.

Consider the induced self matching \tilde{w} on z from w . There are also $\geq 0.4\kappa_C$ bad matches in \tilde{w} by the definition of bad matches. As s is a $(\Lambda' = 0.2\kappa_C, a)$ synchronization separator sequence, there are only at most Λ' matches are undesired. The remaining ones, has number at least $0.4\kappa_C - 0.2\kappa_C = 0.2\kappa_C$, are desired which means each of them will contribute one to $\text{cost}(\tilde{w})$. Hence by Lemma 6.9, $\text{cost}(w) = \text{cost}(\tilde{w}) \geq 0.2\kappa_C$.

As a result, $\text{obj}(w)$ should be at most $\tau + \kappa - 0.2\kappa_C = \tau - 0.19\kappa_C$, where τ is the number of non-zero-symbols in y . This is because there are at most κ inserted non-zero-symbols contributed to the number of matches and the cost function value is at least $0.2\kappa_C$.

However, the matching algorithm will return a matching w with maximum obj by Lemma 6.4. Let w^* be the natural ideal matching which matches every non-deleted non-zero-symbols to itself. Then the target function value of w has to be at least that of w^* , which is $|w^*| - \text{cost}(w^*) \geq \tau - \kappa - \kappa = \tau - 2\kappa = \tau - 0.02\kappa_C$. The inequality is because there can be at most κ non-zero-symbols being deleted and thus $|w^*| \geq \tau - \kappa$; Also each insdel can increase the cost function by at most one, so κ insdels make the cost function of w^* to be at most κ .

Therefore we reach a contradiction. Thus $\text{HD}(y, y') \leq \kappa_C$. So the decoding can get the correct x by the definition of code C .

The efficiency of encoding and decoding follows from Lemma 6.12 and Lemma 7.3. The parameters follow from the construction. The linearity follows from the same argument as that of Theorem 6.1. \square

We state the following theorem to explicitly describe the error rate and information rate. It immedi-

tely follows from Theorem 7.1.

THEOREM 7.2. *If there is an explicit linear code for Hamming distance with block length n_C , error rate δ_C , information rate γ_C , then there is an explicit linear code for insdel errors, with block length $O(n_C/\delta_C)$, insdel error rate $\delta = O(\delta_C^2)$, information rate $O(\delta_C \cdot \gamma_C)$.*

We remark that if we pick a very good code C , with error rate δ_C , information rate $\gamma_C = 1 - \Theta(\delta_C)$. then if $\delta = \kappa/n$ is the error rate for the constructed code, then the information rate for it is $O(\delta\gamma_C/\delta_C) = \Theta(\sqrt{\delta}(1 - \Theta(\sqrt{\delta})))$.

By using an explicit asymptotically good code with linear time encoding and decoding [23], we can immediately get the following corollary. The overall decoding time is due to Lemma 6.12.

COROLLARY 7.1. *There exists an explicit asymptotically good binary linear code for edit distance, with linear encoding time and decoding time $O(n^4)$, and a polynomial time pre-computation for the encoding matrix.*

We remark that the pre-computation is mainly used to compute the synchronization separator sequence.

8 Explicit Affine Codes

In this section, we give an explicit construction of binary affine codes for insdel errors. We first recall the code construction using synchronization strings [14].

THEOREM 8.1. ([14]) *Assume $s \in \Sigma^n$ is an η -synchronization string. C is an $(n, m, d = 2\epsilon n + 1)$ linear code for Hamming distance with efficient decoding.*

Then attaching s to every codeword of C , symbol by symbol, gives a code which can efficiently correct $\frac{2\epsilon n}{1-\eta}$ insertion/deletion errors.

Here is our construction.

CONSTRUCTION 8.1. *Let $s \in \Sigma^{n_0}$ be an η -synchronization string, where $\eta = 0.01$, alphabet Σ_s has size $O(\frac{1}{\eta^2})$ from [6], where each symbol's binary representation has length l_s .*

Let C_0 be an $(n_0, m_0, d_0 = 2\epsilon n_0 + 1) \mathbb{F}_{2^{l_0}}$ -linear code for Hamming distance, where $l_0 = O(\frac{1}{\epsilon^2})$, $m_0 = (1 - \Theta(\epsilon))n_0$. One can use suitable off-the-shelf algebraic-geometric codes for this purpose [21].

Let $l = l_s + l_0$, $t = \Theta(\frac{1}{\epsilon})$. Let string $p = 0 \circ \underbrace{1 \dots 1}_{t+1 \text{ number of } 1\text{'s}}$. Also call this pattern the boundary string.

The encoding $C : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ is as follows, where $m = m_0l$, $n = 2n_0l(1 + \frac{1}{t}) + n_0(t + 1)$.

On input $x \in \mathbb{F}_2^m$:

- View x as in $(\mathbb{F}_2^{l_0})^{m_0}$ by partitioning every l_0 bits as an element in $\mathbb{F}_2^{l_0}$.
- Compute $C_0(x)$, and let $y_i, i \in [n_0]$ be the l_0 bits corresponding to the i -th coordinate, i.e. think of the i -th coordinate as a degree $\leq l_0$ polynomial with coefficients in \mathbb{F}_2 , taking the binary coefficients of its monomials to be y_i . Let $y = y_1 \circ y_2 \circ \dots \circ y_{n_0}$.
- Let $y'_i = s_i \circ y_i$ where s_i is the binary representation of the i -th symbol of s ;
- After every t bits of y'_i , insert a 0 to attain y''_i ;
- Let $z_i = p \circ y''_i$;
- The codeword $z = z_1 \circ \dots \circ z_{n_0}$.

The decoding is as follows.

On input \tilde{z} ;

- Start from the beginning of \tilde{z} ;
- Locate every appearance of the boundary p ;
For the i -th appearance of p ,
 - Take the bits between the i -th boundary and the $i + 1$ -th boundary to be the i -th block;
 - Then for this i -th block, eliminate every $t + 1$ -th symbol (suppose to be the inserted 0) to get \tilde{y}'_i ;
- View \tilde{y}' as s.t. each symbol \tilde{y}'_i is a symbol which is the concatenation of a symbol from s and an element from $\mathbb{F}_2^{l_0}$;
- Use the decoding algorithm from Theorem 8.1 on \tilde{y}' to get x .

THEOREM 8.2. (RESTATE OF THEOREM 1.5) For any $\epsilon > 0$, there exists an explicit affine code over \mathbb{F}_2 , with rate $1 - \Theta(\epsilon)$, which can be efficiently decoded from any $O(\epsilon^3)$ fraction of insertions and deletions.

Proof. We show that Construction 8.1 gives such a code.

First we prove the correctness of the decoding. Assume there are k insdel errors. Since we insert a 0-symbol after every t bits of y'_i to attain y''_i , there are no appearance of p in y''_i . Consider the adversary's insdels on z . We claim that each insdel, in the worst case, can corrupt one of the blocks and at the same

time delete or insert a block. This is because for each insertion or deletion, it can corrupt a boundary or the b bits following it to modify a block. Also at the same time, it may create a new boundary to insert a new block, or may corrupt an existing boundary to delete a block. Note that this will not affect any other block which has no insdels and is not on the left of corrupted boundary. Therefore, for \tilde{y}' , it can be viewed as having $2k$ insdels or modifications from y' . Also recall that \tilde{y} has the structure of the code described in Theorem 8.1. Thus by Theorem 8.1, we can correct $\frac{2\epsilon n_0}{1-\eta} = \Theta(k)$ insdel errors if $\epsilon n_0 = \Theta(k)$.

Next we show the error rate and information rate are as stated. We take $\eta = 0.01$, $l_0 = O(\epsilon^{-2})$, $t = O(\epsilon^{-1})$. So $l_s = \Theta(\eta^{-2}) = \Theta(1)$. Recall that

$$n = n_0l(1 + \frac{1}{t}) + n_0|p| = n_0(l_0 + l_s)(1 + \frac{1}{t}) + n_0(t + 2).$$

So combining with these parameters we get $k = \Theta(\epsilon n_0) = O(\epsilon^3 n)$. The information rate is

$$\begin{aligned} \frac{m_0 l_0}{n} &= \frac{(1 - \Theta(\epsilon))l_0}{(l_0 + l_s)(1 + 1/t) + (t + 2)} \\ &= \frac{1 - \Theta(\epsilon)}{(1 + \Theta(\epsilon^2))(1 + \Theta(\epsilon)) + \Theta(\epsilon)} \\ &\geq 1 - \Theta(\epsilon). \end{aligned}$$

Now we prove that the code C we constructed is affine. By definition of affine code, we need to show that the space of the code is a linear subspace $\subseteq \mathbb{F}_2^n$, plus a shift w in \mathbb{F}_2^n . We let w be the codeword in C , which is the codeword for the zero element of \mathbb{F}_2^m . Recall that the encoding will first compute $C_0(x)$ and then partition it into binary bits, i.e. attaining $y(x)$. After that the encoding just inserting some boundaries and symbols to some specific positions of $y(x)$. This is doing the same as putting y 's coordinates to some specific entries of a vector in $\{0, 1\}^n$ setting other entries to be 0 and then add w to this vector. So we can regard w as the shift, and then to show C is affine, we only need to show that that $A = \{y(x) \mid x \in \mathbb{F}_2^m\} \subseteq \{0, 1\}^{n_0 l_0}$ is a linear space. Since coordinates of A are over \mathbb{F}_2 , any vector in A times an element in \mathbb{F}_2 is still in A . On the other hand, we claim that for two arbitrary vectors $u, v \in A$ corresponding to messages x_u, x_v , it is true that $u \oplus v \in A$. To see this, recall that C_0 is a linear code and thus $C_0(x_u) + C_0(x_v) = C_0(x_u + x_v) \in C_0$. Also for each $i \in [n_0]$ the coordinates addition is done over $\mathbb{F}_2^{l_0}$. So it is adding polynomials in $\mathbb{F}_2[x]$ whose degrees are at most l_0 . Thus the binary representation of the i -th coordinate of $C_0(x_u + x_v) = C_0(x_u) + C_0(x_v)$ is the bit-wise xor over the coordinates' binary representations i.e. $u_i \oplus v_i, i \in [n_0]$, where $u_i, v_i \in \{0, 1\}^{l_0}$ are the binary

representation of the i -th coordinates of $C_0(x_u), C_0(x_v)$. Therefore $u \oplus v = y(x_u + x_v)$. So $u \oplus v \in A$. So A is linear and hence C is affine. \square

9 Explicit Systematic Linear Codes

In this section we construct explicit systematic linear codes for insertions and deletions. Our construction is quite simple, and can be described as follows.

CONSTRUCTION 9.1. *Let $C : \mathbb{F}_q^m \rightarrow \mathbb{F}_q^n$ be a q -ary linear code for insertions and deletions. The systematic linear code is defined below.*

Encoding: *For any message $x \in \mathbb{F}_q^m$, the encoding function is $\text{Enc}(x) = (x, C(x))$. That is, we simply concatenate x with the codeword $C(x)$.*

Decoding: *For any received word y , simply remove the first m symbols to get a substring y' , and run the decoding function of C on y' .*

We have the following lemma.

LEMMA 9.2. *If C is a code that can correct up to k deletions and insertions, then Construction 9.1 gives a code \tilde{C} that can also correct up to k deletions and insertions, with codeword length $n + m$.*

Proof. Let $\text{Enc}(x) = (x, C(x))$ and y be a string obtained from $\text{Enc}(x)$ by at most k deletions and insertions. We only need to show that y' can also be obtained from $C(x)$ by at most k deletions and insertions. To see this, suppose that there are r insertions and t deletions in the x part of $\text{Enc}(x)$, that changes x into x' , where $r + t \leq k$. Thus the total number of insertions and deletions in the $C(x)$ part is at most $k - r - t$. Suppose this changes $C(x)$ into $C'(x)$ and thus $y = (x', C'(x))$.

Note that after r insertions and t deletions, the length of x' is $m + r - t$. Thus to change $C(x)$ into y' , we can first use the same at most $k - r - t$ insertions and deletions in the $C(x)$ part, and then either insert or delete $|r - t|$ symbols. The total number of insertions and deletions needed is at most

$$k - r - t + |r - t| \leq k.$$

\square

This gives the following corollary.

COROLLARY 9.1. *Suppose there exists an explicit q -ary linear code with rate R that can correct up to δ fraction of insertions and deletions, then there exists an explicit q -ary systematic linear code with rate $R/(1 + R)$ that can correct up to $\delta/(1 + R)$ fraction of insertions and deletions.*

Combined with Corollary 7.1 this gives the following theorem.

THEOREM 9.1. *For any n there exists an explicit systematic linear code C_n with constant rate which can be efficiently encoded in linear time and decoded from $\Theta(n)$ insertions and deletions in time $O(n^4)$. The generator matrix of C_n can be deterministically computed in polynomial time.*

10 Open Questions

There are numerous intriguing questions in the subject that remain open making this work the beginning of a new line of research aimed toward painting a more complete picture of the power and limitations of linear and affine codes for synchronization errors. Concretely, the following are some of the several interesting questions brought to the fore by our work:

- Can one find a better distribution of run lengths of 0's that we intersperse in our construction that improves the rate of our construction? Could we possibly approach a rate of 1/2 with this method, or are there any inherent limitations to this approach?
- Can one explicitly construct efficient linear insdel codes over large alphabets which get arbitrarily close to the Half-Singleton bound in Theorem 1.2?
- Can one explicitly construct efficient affine codes over large alphabets that approach the Singleton bound similar to the non-affine insdel codes in [14]?
- Is there any separation between affine insdel codes and unrestricted general insdel codes?
- Obtain better bounds (even non-constructively) on the asymptotic rate-distance trade-offs in the high-rate regime for linear insdel codes over the binary (or fixed q -ary) alphabet.
- What can be said about the zero-rate regime? That is, what fraction of deletions can be corrected by linear (or affine) codes of rate bounded away from 0? Can one do much better than the random coding bound, as was done in [5] for general codes?

References

- [1] Khaled A.S. Abdel-Ghaffar, Hendrik C. Ferreira, and Ling Cheng. On linear and cyclic codes for correcting deletions. In *2007 IEEE International Symposium on Information Theory (ISIT)*, pages 851–855, 2007. 0, 1, 2, 7

- [2] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k -wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992. 6
- [3] Joshua Brakensiek, Venkatesan Guruswami, and Samuel Zbarsky. Efficient low-redundancy codes for correcting multiple deletions. *CoRR*, abs/1507.06175, 2015. 0
- [4] Joshua Brakensiek, Venkatesan Guruswami, and Samuel Zbarsky. Efficient low-redundancy codes for correcting multiple deletions. *IEEE Transactions on Information Theory*, 64(5):3403–3410, 2018. Preliminary version in SODA 2016. -1, 0
- [5] Boris Bukh, Venkatesan Guruswami, and Johan Håstad. An improved bound on the fraction of correctable deletions. *IEEE Trans. Information Theory*, 63(1):93–103, 2017. Preliminary version in SODA 2016. -1, 17
- [6] K. Cheng, B. Haeupler, X. Li, A. Shahrasi, and K. Wu. Synchronization Strings: Efficient and Fast Deterministic Constructions over Small Alphabets. In *2019 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2019. -1, 2, 9, 15
- [7] Kuan Cheng, Venkatesan Guruswami, Bernhard Haeupler, and Xin Li. Efficient linear and affine codes for correcting insertions/deletions. *arXiv preprint arXiv:2007.09075*, 2020. -1
- [8] Kuan Cheng, Zhengzhong Jin, Xin Li, and Ke Wu. Deterministic document exchange protocols, and almost optimal binary codes for edit errors. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 200–211. IEEE Computer Society, 2018. -1, 1, 4, 14
- [9] Kuan Cheng, Zhengzhong Jin, Xin Li, and Ke Wu. Block Edit Errors with Transpositions: Deterministic Document Exchange Protocols and Almost Optimal Binary Codes. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 37:1–37:15, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. -1
- [10] Venkatesan Guruswami, Bernhard Haeupler, and Amirbehshad Shahrasi. Optimally resilient codes for list-decoding from insertions and deletions. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing*, pages 524–537, 2020. -1
- [11] Venkatesan Guruswami and Carol Wang. Deletion codes in the high-noise and high-rate regimes. *IEEE Trans. Information Theory*, 63(4):1961–1970, 2017. -1, 1
- [12] Bernhard Haeupler. Optimal document exchange and new codes for insertions and deletions. In *60th IEEE Annual Symposium on Foundations of Computer Science*, pages 334–347, 2019. -1, 1
- [13] Bernhard Haeupler, Aviad Rubinfeld, and Amirbehshad Shahrasi. Near-Linear Time Insertion-Deletion Codes and $(1+\epsilon)$ -Approximating Edit Distance via Indexing. *Proceeding of the ACM Symposium on Theory of Computing (STOC)*, pages 697–708, 2019. -1
- [14] Bernhard Haeupler and Amirbehshad Shahrasi. Synchronization strings: codes for insertions and deletions approaching the singleton bound. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*, pages 33–46. ACM, 2017. -1, 0, 2, 5, 9, 15, 17
- [15] Bernhard Haeupler and Amirbehshad Shahrasi. Synchronization strings: explicit constructions, local decoding, and applications. In *STOC’18—Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 841–854. ACM, New York, 2018. -1
- [16] Bernhard Haeupler and Amirbehshad Shahrasi. Synchronization strings: Explicit constructions, local decoding, and applications. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing*, pages 841–854, 2018. 2, 9
- [17] Bernhard Haeupler, Amirbehshad Shahrasi, and Madhu Sudan. Synchronization strings: List decoding for insertions and deletions. *Proceeding of the International Colloquium on Automata, Languages and Programming (ICALP)*, pages 76:1–76:14, 2018. -1
- [18] Bernhard Haeupler, Amirbehshad Shahrasi, and Ellen Vitercik. Synchronization strings: Channel simulations and interactive coding for insertions and deletions. *Proceeding of the International Colloquium on Automata, Languages and Programming (ICALP)*, pages 75:1–75:14, 2018. -1, 1
- [19] Tomohiro Hayashi and Kenji Yasunaga. On the list decodability of insertions and deletions. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 86–90. IEEE, 2018. -1
- [20] Leonard J. Schulman and David Zuckerman. Asymptotically good codes correcting insertions, deletions, and transpositions. *IEEE Trans. Inf. Theory*, 45(7):2552–2557, 1999. Preliminary version in SODA 1997. 1
- [21] Kenneth W. Shum, Ilia Aleshnikov, P. Vijay Kumar, Henning Stichtenoth, and Vinay Deolalikar. A low-complexity algorithm for the construction of algebraic-geometric codes better than the Gilbert-Varshamov bound. *IEEE Trans. Inf. Theory*, 47(6):2225–2241, 2001. 15
- [22] Jin Sima and Jehoshua Bruck. Optimal k -deletion correcting codes. In *IEEE International Symposium on Information Theory*, pages 847–851, 2019. -1
- [23] Daniel A. Spielman. Linear-time encodable and decodable error-correcting codes. *IEEE Trans. Inf. Theory*, 42(6):1723–1731, 1996. 15
- [24] Antonia Wachter-Zeh. List decoding of insertions and deletions. *IEEE Transactions on Information Theory*, 64(9):6297–6304, 2017. -1