

Coding for Duplication Channels with Applications to DNA Storage

Thesis submitted in partial fulfillment
of the requirements for the degree of
“DOCTOR OF PHILOSOPHY”

by

Yonatan Yehezkeally

Submitted to the Senate of Ben-Gurion University
of the Negev

March 18, 2020

Beer-Sheva

Coding for Duplication Channels with Applications to DNA Storage

Thesis submitted in partial fulfillment
of the requirements for the degree of
“DOCTOR OF PHILOSOPHY”

by

Yonatan Yehezkeally

Approved by the advisor



Approved by the Dean of the Kreitman School of Advanced Graduate Studies

Submitted to the Senate of Ben-Gurion University
of the Negev

March 18, 2020

Beer-Sheva

This work was carried out under the supervision of

Prof. Moshe Schwartz

In the School of Electrical and Computer Engineering

Faculty of Engineering Sciences


Research-Student's Affidavit when Submitting the Doctoral Thesis for Judgment

I Yonatan Yehezkeally, whose signature appears below, hereby declare that (Please mark the appropriate statements):

I have written this Thesis by myself, except for the help and guidance offered by my Thesis Advisors.

The scientific materials included in this Thesis are products of my own research, culled from the period during which I was a research student.

This Thesis incorporates research materials produced in cooperation with others, excluding the technical help commonly received during experimental work. Therefore, I am attaching another affidavit stating the contributions made by myself and the other participants in this research, which has been approved by them and submitted with their approval.

Date March 18, 2020 Student's name: Yonatan Yehezkeally Signature: 

I would like to gratefully acknowledge the following people and institutions, whose support and contributions – personal, scientific, and financial – were paramount in enabling the composition of the work enclosed herein (not to mention, in keeping its author’s sanity and well-being while composing it):

First and foremost, to my doctoral advisor Moshe Schwartz: your teaching, support, and guidance have carried me over this long (quite long) path. Your patience kept me from despairing. In a crossroad, I ask myself WHAT WOULD MOSHE DO (which is never an alternative to finding out from the source). Thank you so very much for being as available and approachable as you were inspirational throughout these years. None of it was ever taken for granted.

To Rutie Peled, Irit Natania, Avital Edri, and Debbie Israel-Feinaro: I have generated enough trouble over the years to keep things interesting; Thank you for guiding me through (and around) it all.

To professors Jehoshua Bruck and Antonia Wachter-Zeh, who each helped me take my first steps towards wherever it is that I’ll find myself going from here. Professor Bruck is also owed thanks for igniting my interest in the problem which is considered in Chapter 2 of this work.

Finally, I am also grateful for the support of BGU’s Kreitman School of Advanced Graduate Studies through its Negev Scholarship, of BGU’s department (now school) of electrical and computer engineering, of the Israeli Science Foundation (ISF), the U.S.-Israel Binational Science Foundation (BSF), and (of course!) that of Moshe, which kept the lights on as I was looking for my way.

Yonatan (Yoni) Yehezkeally
MARCH 2020

TO ORLI, WHOSE LIFE I'M GRATEFUL FOR
BEING A PART OF, AS SHE TRANSFORMS
MINE.

TO URI, WHO FILLS OUR HEARTS WITH
JOY AND OUR MINDS WITH HOPE.

AND TO THEY WHO ARE YET TO JOIN
OUR FAMILY.

Contents

List of Figures	xiii
List of Symbols and Abbreviations	xv
Abstract	xvii
Introduction	xix
Work structure	xxi
Publications	xxii
Notation and definitions	xxiii
1 Reconstruction Codes	1
1.1 Preliminaries	2
1.2 Structure of descendant cones	2
1.3 Size of reconstruction codes	3
1.4 Minimal distance of reconstruction codes	6
1.5 Capacity of reconstruction codes	7
1.6 Discussion of recent results	18
2 Uncertainty with List Decoding	21
2.1 Preliminaries and related works	21
2.2 Additional notation and definitions	23
2.3 Typical set	23
2.4 Uncertainty with underlying ECC	31
Appendix: Conclusion of proof of Lemma 2.2	37
3 Combined Substitutions Noise	41
3.1 Additional notation and definitions	41
3.2 Error-correction via constrained coding	43
3.3 Code Construction and Size	46
Discussion	49

Bibliography

51

List of Figures

- 1 The graph $G_q(k - 1)$ generating the $(0, k - 1)_q$ -RLL system. 9
- 2 Capacity $C(\gamma)$ in the cases (a) $q = k = 2$, $\theta = 0.7236$, and (b) $q = 4$, $k = 2$, $\theta = 0.8273$. The value at $\gamma = 1$ equals $\text{cap}(\text{Irr})$ 13

List of Symbols and Abbreviations

Σ	Finite alphabet
Σ^*	The set of all finite strings over Σ
$\Sigma^{\geq k}$	The set of all finite strings over Σ of length $\geq k$
$ x $	The length of the string x
\mathcal{T}_i	A tandem duplication at index i
$x \xrightarrow{t} y$	y is a t -descendant of x
$x \xrightarrow{*} y$	y is a descendant of x
$D^t(x)$	The set of all t -descendants of x
$D^*(x)$	The descendant cone of x
Irr	The set of all irreducible strings
Irr(n)	The set of all irreducible strings of length n
drt(x)	The duplication root of x
$x \sim_k y$	x, y have a common ancestor
C_x	$C \cap D^*(x)$, for some code $C \subseteq \Sigma^*$ and $x \in \text{Irr}$
wt(x)	The Hamming weight of the string x
$d(\cdot, \cdot)$	The duplication distance on $D^r(x)$
$\ \cdot\ _1$	The ℓ_1 norm on \mathbb{N}^{w+1}
$d_1(\cdot, \cdot)$	A metric on \mathbb{N}^{w+1} equal half the ℓ_1 norm of the difference
$\phi = \hat{\phi}\bar{\phi}$	The discrete derivative
$w(x)$	The Hamming weight of $\bar{\phi}(x)$
ψ_x	An isometric poset isomorphism from $D^*(x)$ to $\mathbb{N}^{w(x)+1}$
\vee and \wedge	The join and meet operations in a lattice, respectively
$r(x)$	The ℓ_1 norm of $\psi_{\text{drt}(x)}(x)$
$f(S)$	When f is a function, $S \subseteq \text{dom}(f)$: $\{f(s) : s \in S\}$
Δ_r^w	The simplex of dimension w and weight r

H_q	The q -ary entropy function; $H = H_2$
$A(\nu, 2\delta, \omega)$	The size of the largest length ν binary code with minimum Hamming distance 2δ and constant Hamming weight ω
$\text{red}(C)$	The redundancy of the code $C \subseteq \Sigma^n$
$R(C)$	The rate of the code $C \subseteq \Sigma^n$
$\text{cap}(\mathcal{C})$	The capacity of the family $\mathcal{C} \subseteq \Sigma^*$

Abstract

Motivated by applications to in-vivo DNA data storage, we study coding for error-control in channels with duplication noise.

We propose that schemes for in-vivo DNA data storage may benefit from an application of Levenshtein's reconstruction schema, which can be used whenever multiple reads of noisy data are available. This strategy is uniquely suited to the medium, which inherently replicates stored data in multiple distinct ways, caused by mutations.

We consider noise introduced solely by uniform tandem duplication, and utilize the relation to constant-weight integer codes in the Manhattan metric. By bounding the intersection of the cross-polytope with hyperplanes, we prove the existence of reconstruction codes with lower redundancy than error-correcting codes, as well as suggest an implicit construction for a family of reconstruction codes.

Next, we propose a list-decoding scheme for reconstruction codes in the context of uniform-tandem-duplication noise, which can be viewed as an application of the associative memory model to this setting. We find the uncertainty associated with $m > 2$ strings (where unique reconstruction corresponds to $m = 2$) in asymptotic terms, where codewords are taken from an error-correcting code. Thus, we find the trade-off between the designed minimum distance, the number of errors, the acceptable list size and the resulting uncertainty, which corresponds to the required number of distinct retrieved outputs for successful reconstruction. It is therefore seen that by accepting list decoding one may further decrease coding redundancy, or the required number of reads, or both.

A combination of duplication noise with substitution errors, which is also motivated by mutation processes occurring in in-vivo DNA data storage applications, is studied next. We focus on an unrestricted model, where substitution errors are permitted at any location in the string, and any time in the sequence of duplication events. By proposing a constrained coding approach, we develop error-correcting codes which can recover from any number of duplication events, together with a single substitution.

Introduction

With recent improvements in DNA sequencing and synthesis technologies, the case for DNA as a data-storage medium is now stronger than ever before. It offers a long lasting and high density alternative to current storage media, particularly for archival purposes [CGK12]. Moreover, due to medical necessities, the technology required for data retrieval from DNA is highly unlikely to become obsolete; as recent history shows, the same cannot be said of concurrent alternatives, such as flash memory, magnetic disks, and optical disks. Data storage in DNA may provide integral memory for synthetic-biology methods, where such is required, and offer a protected medium for long-period data storage [Bal13, WkWF03].

Several recent works have studied the inherent constraints of storing and retrieving data from DNA. While desired sequences (over quaternary alphabet) may be synthesized (albeit, while suffering from substitution noise), generally data can only be read by observation of its substrings, quite possibly an incomplete observation [KPM16]. Moreover, the nature of DNA and current technology results in asymmetric errors which depend upon the dataset [GKM17]. The medium itself also introduces other types of errors which are atypical in electronic storage, such as symbol/block deletion and adjacent transpositions (possibly complemented) [GYM18]. Finally, the purely combinatorial problem of recovering a sequence from the multiset of all its substrings (including their numbers of incidence), was also studied, e.g., [ADM⁺15, SCT16], as well as coding schemes involving only these multisets (or their profile vectors – describing the incidence frequency of each substring) [RSY19].

Other recent works were concerned with storing information in the DNA of living organisms (henceforth, *in-vivo* DNA storage) in particular; with the advent of CRISPR/Cas gene editing technique [SNMC16, SNMC17], it is now becoming more feasible. In-vivo DNA storage has somewhat lower data density than in-vitro storage (i.e., where synthesized DNA strands are kept cell-free), but it provides a reliable and cost-effective propagation via replication, in addition to some protection to stored data [SNMC17]. It also has applications including watermarking genetically modified organisms [AO04, HB07, LDB⁺12] or research material [WkWF03, JFS⁺10] and concealing sensitive information [CRB99]. However, muta-

tions introduce a diverse set of potential errors, including symbol- or burst- substitution/insertions/deletion, and duplication (including tandem- and interspersed-duplication). Naturally, therefore, data integrity in such storage schemes is of great interest.

In an effort to better understand these noise mechanisms, their potential to generate the diversity observed in nature was studied. [FSB16] classified the *capacity* and/or *expressiveness* of the systems of sequences over a finite alphabet generated by four distinct substring duplication rules: end duplication, tandem duplication, tandem-palindromic duplication, and interspersed duplication. [JFB17] fully characterized the expressiveness of bounded tandem-duplication systems, proved bounds on their capacity (and, in some cases, even exact values). [JFSB17b] later showed that when point mutations act together with tandem duplication as a sequence generation process, they may actually increase the capacity of the generated system. [ABFJ17] looked at the typical *duplication distance* of binary sequences; i.e., the number of tandem duplications generating a binary sequence from its root. It was proven that for all but an exponentially small number of sequences that number is proportional to the sequence length. Further, when tandem duplication is combined with point mutations (here, only within the duplicated string), it was shown that the frequency of substitutions governs whether that distance becomes logarithmic.

The generative properties of interspersed duplication were also studied from a probabilistic point of view. [FSB15, FSB19] showed (under assumption of uniformity) that the frequencies of incidence for each substring converge to the same limit achieved by an i.i.d. source, thus reinforcing the notion that interspersed duplication is – on its own – capable of generating diversity. [EFSB16] specifically looked at tandem- and end-duplication, and found exact capacities in the case of duplication length 1 by a generalization of the Pólya urn model that applies to strings. It also tightly bounded the capacity of complement tandem duplication, a process where the duplicated symbol is complemented (using binary alphabet).

Finally, the main focus of this work is tandem-duplication noise. In a tandem-duplication event, a substring of the DNA sequence, the *template*, is duplicated, and the resulting *copy* is inserted into the sequence next to the template [ZAM14]. Evidence of this process is found in the genomes of many organisms as patterns that are repeated multiple times [FSB19]. Error-correcting codes for data affected by tandem duplication have been studied in [JFSB17a], which presented a construction of optimal-size codes for correcting any number of errors under *uniform tandem duplication* (fixed duplication length), computing their (and thus, the optimal-) capacity. It also presented a framework for the construction of optimal codes for the correction of a fixed number of errors. Next, it studies bounded tandem duplications, where a characterization of the capacity of error-correcting

codes is made for small constants. In general, it characterized the cases where the process of tandem duplication can be traced back uniquely.

More recently, a flurry of activity in the subject includes works such as [LJWZ18, KT18a, LWY19] which provided some implicit constructions for uniform-tandem-duplication codes and presented bounds on minimal code redundancy for uniform tandem duplication and tandem-palindromic-duplication noise. Other related works include [SGSD17] which studied Levenshtein’s reconstruction problem in the context of insertions when messages are restricted to distance-deficient codes, [MV17] which developed explicit construction for sticky insertion noise in binary strings achieving asymptotically optimal redundancy, [Kov19] which studied optimal error-correcting codes for short tandem-duplication noise, and finally [TYSF19, TF19] which examined combined duplication and substitution noise.

Work structure

This work is organized as follows: At the end of this chapter, we present notation and definitions which will be used throughout.

Chapter 1 is devoted to a study of Levenshtein’s reconstruction problem in the context of uniform-tandem-duplication noise. In Section 1.1 the reconstruction problem is formalized in this context; it is then demonstrated in Sections 1.2 to 1.3 that reconstruction codes partition into error-correcting codes, and the requisite minimum-distance of each part is found in Section 1.4, as a function of the reconstruction parameters. We shall see that these parts can be isometrically embedded as constant-weight codes in the Manhattan metric. Section 1.5 shows that reconstruction codes exist with full capacity, and also suggest a construction for reconstruction codes. Lastly, recent results, published as this work was being compiled, are briefly reviewed in Section 1.6, and their implications discussed.

Chapter 2 extends the study of the previous chapter by allowing for list decoding in conjunction with reconstruction. Section 2.1 describes an overview of the results contained in this chapter and puts them in context of related works. Section 2.2 presents additional notations and definitions which we shall find useful, then in Section 2.3 the problem is solved assuming minimal constraints on the codebook, and an efficient decoding scheme is developed. Our analysis is then repeated and extended in Section 2.4 for subsets of the codebook which form error-correcting codes.

Chapter 3 is dedicated to a study of a different noise mechanism, involving substitution errors together with duplication errors. In Section 3.1, we provide additional notation as well as relevant background and known results. In Section 3.2, a constrained coding approach is proposed for combined error-correction, and in

Section 3.3 a construction based on that approach for an error-correcting code is presented, and the rate of the resulting codes is analyzed.

Publications

The work presented herein has been published in the following avenues:
Chapter 1 includes results which were presented in

[YS18] Yonatan Yehezkeally and Moshe Schwartz. Reconstruction codes for DNA sequences with uniform Tandem-Duplication errors. In *Proceedings of the 2018 IEEE International Symposium on Information Theory (ISIT'2018)*, Vail, CO, US, pages 2535–2539, June 2018.

Further results from that chapter were finally published in

[YS19] Yonatan Yehezkeally and Moshe Schwartz. Reconstruction codes for DNA sequences with uniform Tandem-Duplication errors. *IEEE Trans. on Inform. Theory*, 2019. accepted for publication.

Chapter 3 includes results which were presented, in a joint work based on equal contribution with additional co-authors, in

[TYSF19] Yuanyuan Tang, Yonatan Yehezkeally, Moshe Schwartz, and Farzad Farnoud (Hassanzadeh). Single-error detection and correction for duplication and substitution channels. In *Proceedings of the 2019 IEEE International Symposium on Information Theory (ISIT'2019)*, Paris, France, pages 300–304, July 2019.

Further results were submitted for publication in

[TYSFon] Yuanyuan Tang, Yonatan Yehezkeally, Moshe Schwartz, and Farzad Farnoud (Hassanzadeh). Single-error detection and correction for duplication and substitution channels. Submitted to the *IEEE Trans. on Inform. Theory* (revision submitted).
url: <https://arxiv.org/abs/1911.05413>

Finally, partial results from Chapter 2 were accepted for presentation at

[YSona] Yonatan Yehezkeally and Moshe Schwartz. Uncertainty of reconstruction with list-decoding from uniform-tandem-duplication noise. Accepted to appear in *Proceedings of the 2020 IEEE International Symposium on Information Theory (ISIT'2020)*.

Additional material was later submitted for publication in

[Y**Sonb**] Yonatan Yehezkeally and Moshe Schwartz. Uncertainty of reconstruction with list-decoding from uniform-tandem-duplication noise. Submitted to the *IEEE Trans. on Inform. Theory*.
url: <https://arxiv.org/abs/2001.07047>

Notation and definitions

The setting of this work is the set of finite strings Σ^* , over an alphabet Σ . Though DNA is composed of four nucleotide bases, the alphabet elements themselves are immaterial to our discussion, hence Σ assumed only to be a finite unital ring of size $q \geq 2$ (e.g., \mathbb{Z}_q , or when q is a prime power, $\text{GF}(q)$). Thus, addition (or subtraction) and multiplication of letters from the alphabet are well defined.

The length of a string $x \in \Sigma^*$ is denoted $|x|$. For any two strings $x, y \in \Sigma^*$, we denote their concatenation xy . x^k denotes concatenating k copies of x . To avoid confusion, the multiplication in the ring is denoted as $a \cdot b$. We say that $y \in \Sigma^*$ is a *substring* of $w \in \Sigma^*$ if there exist $x, z \in \Sigma^*$ such that $w = xyz$. If the need arises to refer to specific positions in strings, positions are numbered $1, 2, \dots$.

A *tandem duplication* (or *tandem repeat*) of fixed duplication-window length k (thus, *uniform* tandem-duplication noise) duplicates a substring of length k and inserts it in tandem into the string, namely, the copy immediately follows the template. As an example for $k = 3$ and alphabet $\Sigma = \mathbb{Z}_3$, consider

$$x = 1012121 \rightarrow y = 1012\underline{012}121,$$

where the underlined part is the copy. More formally, a tandem duplication at index i is defined as follows, for $x, y, z \in \Sigma^*$, $|x| = i$ and $|y| = k$:

$$\mathcal{T}_i(xyz) \triangleq xyyz.$$

Thus, uniform tandem-duplication noise with duplication-window length k acts only on strings of length $\geq k$, which we denote $\Sigma^{\geq k}$. Since throughout this work all duplications considered will be in tandem and of duplication-window-length k , we shall use the term “duplication” where convenient, to avoid cumbersome terminology.

If $y \in \Sigma^{\geq k}$ can be derived from $x \in \Sigma^{\geq k}$ by a sequence of tandem repeats, i.e., if there exist i_1, \dots, i_t such that

$$y = \mathcal{T}_{i_t}(\dots \mathcal{T}_{i_1}(x)),$$

then y is called a *t-descendant* (or simply *descendant*) of x (vice versa, x is an *ancestor* of y), and we denote $x \xrightarrow{t} y$. We observe that, in that case, $|y| = |x| + k$. We say that x is a 0-descendant of itself. If $t = 1$ we denote $x \implies y$. Where the

number of repeats is unknown or irrelevant, we may denote $x \xrightarrow{*} y$. We define the set of t -descendants of x as

$$D^t(x) \triangleq \{y \in \Sigma^* : x \xrightarrow{t} y\},$$

and the *descendant cone* of x as

$$D^*(x) \triangleq \{y \in \Sigma^* : x \xrightarrow{*} y\} = \bigcup_{t=0}^{\infty} D^t(x).$$

If there exists no $z \neq x$ such that $z \xrightarrow{*} x$, we say that x is *irreducible*. The sets of all irreducible strings, and those of length n , are denoted Irr and $\text{Irr}(n)$, respectively. It can be shown (see, e.g., [LMVM05, JFSB17a]) that for all $y \in \Sigma^{\geq k}$ there exists a unique irreducible x , called the *duplication root* of y and denoted $\text{drt}(y)$, such that $y \in D^*(x)$. This induces a partition of $\Sigma^{\geq k}$ into descendant cones; i.e., it induces an equivalence relation, denoted herein \sim_k . For a code $C \subseteq \Sigma^*$ and some $x \in \text{Irr}$, we denote $C_x \triangleq C \cap D^*(x)$, and observe that $\{C_x : x \in \text{Irr}\}$ is a partition of C .

A useful tool in studying uniform tandem-duplication noise is the *k-discrete derivative* ϕ defined in [FSB16] for $x \in \Sigma^{\geq k}$:

$$\phi(x) \triangleq \hat{\phi}(x)\bar{\phi}(x),$$

where

$$\begin{aligned} \hat{\phi}(x) &\triangleq x(1), x(2), \dots, x(k), \\ \bar{\phi}(x) &\triangleq x(k+1) - x(1), \dots, x(|x|) - x(|x| - k). \end{aligned}$$

As seen, e.g., in [JFSB17a], ϕ is injective, and if $\bar{\phi}(x) = ab$ for $a, b \in \Sigma^*$, $|a| = i$, then $\bar{\phi}(\mathcal{T}_i(x)) = a0^k b$. For the example given above,

$$\begin{aligned} x = 1012121 &\rightarrow y = 1012\underline{0}12121, \\ \bar{\phi}(x) = 1112 &\rightarrow \bar{\phi}(y) = \underline{1000}112. \end{aligned}$$

The implications are better captured by the function $\psi_x : D^*(x) \rightarrow \mathbb{N}^{w+1}$, defined

$$\psi_x(y) \triangleq (\lfloor s(1)/k \rfloor, \dots, \lfloor s(w+1)/k \rfloor),$$

if

$$\bar{\phi}(y) = 0^{s(1)} a_1 0^{s(2)} \dots a_w 0^{s(w+1)},$$

where $w = w(x) \triangleq \text{wt}(\bar{\phi}(x))$ is the Hamming weight of $\bar{\phi}(x)$, and $a_1 \dots, a_w \in \Sigma \setminus \{0\}$. In particular, it is seen (e.g., in [JFSB17a]) that $x \in \text{Irr}$ if and only if

$\psi_x(x) = (0, 0, \dots, 0)$, that is, $\bar{\phi}(x)$ contains no zero-runs of length k ; such strings are called $(0, k - 1)_q$ -run-length-limited, or $(0, k - 1)_q$ -RLL.

We denote the 1-norm on \mathbb{N}^{w+1}

$$\|u\|_1 \triangleq \sum_{i=1}^{w+1} u(i),$$

and define a metric

$$d_1(u, v) \triangleq \frac{1}{2} \|u - v\|_1.$$

A metric can then be defined on $D^r(x)$ for each r (in particular, when x is irreducible) by

$$d(y_1, y_2) \triangleq \min\{t \in \mathbb{N} : D^t(y_1) \cap D^t(y_2) \neq \emptyset\},$$

and it is seen in [JFSB17a, Lem. 14] that this is well-defined, in the sense that there does exist such t , for $y_1, y_2 \in D^r(x)$, such that $D^t(y_1) \cap D^t(y_2) \neq \emptyset$ (in fact, such t exists only if $y_1 \sim_k y_2$.) Furthermore, [JFSB17a, Lem. 19] shows for $y_1, y_2 \in D^r(x)$ that

$$d(y_1, y_2) = d_1(\psi_x(y_1), \psi_x(y_2)).$$

As a consequence, $d(y_1, y_2)$ may be computed efficiently (note that foreknowledge of x is unnecessary).

In our analysis we shall use the following asymptotic notation: for two sequences a_n, b_n we say that $a_n \sim b_n$ if $a_n = b_n(1 + o(1))$.

Finally, we define the *redundancy* of a code $C \subseteq \Sigma^n$ as

$$\text{red}(C) \triangleq n - \log_q |C| = n - \log_{|\Sigma|} |C|,$$

and the code's *rate* as

$$R(C) \triangleq \frac{1}{n} \log_q |C| = \frac{1}{n} \log_{|\Sigma|} |C|.$$

For a system $\mathcal{C} \subseteq \Sigma^*$, we define its *capacity* as

$$\text{cap}(\mathcal{C}) \triangleq \limsup_{n \rightarrow \infty} R(\mathcal{C} \cap \Sigma^n).$$

Chapter 1

Reconstruction Codes

Classical error-correction coding, hence also relevant works which were listed in the last chapter, ignores some properties of the DNA storage channel; namely, stored information is expected to be replicated, even as it is mutated. This lends itself quite naturally to the reconstruction problem [Lev01], which assumes that data is simultaneously transmitted over several noisy channels, and a decoder must therefore estimate that data based on several (distinct) noisy versions of it.

Solutions to this problem have been studied in several contexts. It was solved in [Lev01] for sequence reconstruction over finite alphabets, where several error models were considered, such as substitutions, transpositions and deletions. Moreover, a framework was presented for solving the reconstruction problem in general cases of interest in coding theory, utilizing a graph representation of the error model, which was further developed in [LKKM08,LS09]. The problem was also studied in the context of permutation codes with transposition and reversal errors [Kon07,KLS07,Kon08], and partially solved therein. Later, applications were found in storage technologies [CB11,YBS16,CKV⁺18,YB19], since modern application might preclude the retrieval of a single data point, in favor of multiple-point requests. However, the problem hasn't been addressed yet for data storage in the DNA of living organisms, where it may be most applicable.

In this chapter, we shall study the reconstruction problem over DNA sequences, with uniform tandem-duplication errors. The main results we shall establish are the following: reconstruction codes in this setting necessarily partition into error-correcting codes, with appropriately chosen minimum distance, based on the uncertainty parameter. In two asymptotic regimes, we shall see that one can always obtain lower redundancy than error-correcting codes. These asymptotic regimes include what is believed to be the most interesting one, where the uncertainty is sub-linear, and the time (number of mutations) is bounded by a constant.

In the next chapter, we shall relax our assumptions, which will shed a new light on our findings here, as a private case of a more general setting.

1.1 Preliminaries

The reconstruction problem in the context of uniform tandem-duplication errors can be stated as follows: suppose data is encoded in $C \subseteq \Sigma^n$, and suppose we later are able to read distinct $x_0, x_1, \dots, x_N \in D^t(c)$ for some unknown $c \in C$ and a given $t \in \mathbb{N}$; can we uniquely identify c ?

It is apparent (see [Lev01]) that successful reconstruction is both assured by, and requires, the following property:

Definition 1.1 *Take $N, t, n > 0$. We say that $C \subseteq \Sigma^n$ is a uniform tandem-duplication reconstruction code, or simply a reconstruction code, for t duplications with uncertainty N , if*

$$\max\{|D^t(c) \cap D^t(c')| : c, c' \in C, c \neq c'\} \leq N.$$

The purpose of this section is to characterize reconstruction codes. By an evaluation of the size of intersection of descendant cones, we determine the achievable size of reconstruction codes. We shall state the solution to this problem in terms of error-correcting codes for the Manhattan metric, and devote the next section to a study of such codes.

1.2 Structure of descendant cones

The structure of \mathbb{N}^{w+1} as a partially ordered set with the product order (i.e., $u \leq v$ if for every coordinate $1 \leq i \leq w + 1$ it holds that $u(i) \leq v(i)$) is relatively easy to describe. We shall therefore find it more convenient to consider $D^*(x)$, for some $x \in \text{Irr}$, in these terms:

Lemma 1.2 *Take $x \in \text{Irr}$, and denote $w = w(x)$. Then ψ_x is a poset isomorphism from $(D^*(x), \xrightarrow{*})$ to (\mathbb{N}^{w+1}, \leq) . In particular,*

1. *For all $y, y' \in D^*(x)$ there exists $z \in D^*(y) \cap D^*(y')$ such that*

$$D^*(y) \cap D^*(y') = D^*(z);$$

2. *If in addition $|y| = |y'|$ then for all $t \in \mathbb{N}$*

$$|D^t(y) \cap D^t(y')| = \begin{cases} 0 & t < d(y, y'), \\ |D^{t-d(y, y')}(x)| & t \geq d(y, y'). \end{cases}$$

Proof We note that since $x \in \text{Irr}$ we know $\psi_x(x) = (0, 0, \dots, 0) \in \mathbb{N}^{w+1}$. Further, we note that in the image of $\bar{\phi}$, a tandem-duplication corresponds to increasing by one a single coordinate of $\psi_x(\cdot)$, i.e., an addition of a unit vector $e_j \in \mathbb{N}^{m+1}$.

Hence, ψ_x is indeed a poset isomorphism, and we see that $\xrightarrow{*}$ endows $D^*(x)$ with a lattice structure; We denote the *join* of $y, y' \in D^*(x)$ as $y \vee y'$, and their *meet* $y \wedge y'$. It follows that $z = y \vee y'$ satisfies Item 1.

Finally, if $|y| = |y'|$ then by definition of $d \triangleq d(y, y')$ we have $z = y \vee y' \in D^d(y) \cap D^d(y')$ (note, $d = \frac{|z|-|y|}{k}$), and it is now straightforward to prove from the poset-isomorphism that, when $t \geq d$, $D^t(y) \cap D^t(y') = D^{t-d}(z)$; since $z \in D^*(x)$, Item 2 also follows. \square

Given Lemma 1.2, we can now find the size of intersection of descendant cones for any $c, c' \in \Sigma^n$ ($n \geq k$), keeping in mind that $D^*(c) \cap D^*(c') \neq \emptyset$ if and only if $c \sim_k c'$.

Lemma 1.3 For $x \in \text{Irr}$ and $t \in \mathbb{N}$ we have $|D^t(x)| = \binom{t+w}{w}$, where $w \triangleq w(x)$.

Proof By Lemma 1.2 we know that

$$D^t(x) = \{y \in D^*(x) : \|\psi_x(y)\|_1 = t\}.$$

Since $\psi_x: D^*(x) \rightarrow \mathbb{N}^{w+1}$ is bijective, $|D^t(x)|$ equals the number of distinct integer solutions to $\sum_{i=1}^{w+1} u(i) = t$, where $u(1), \dots, u(w+1) \geq 0$ (equivalently, the number of distinct ways to distribute t identical balls into $w+1$ bins). \square

1.3 Size of reconstruction codes

In this section we aim to estimate the maximal size of reconstruction codes. We shall first need to make the following notation:

Definition 1.4 For $w, r > 0$ we denote the simplex of dimension w and weight r , or (w, r) -simplex

$$\Delta_r^w \triangleq \left\{ u \in \mathbb{N}^{w+1} : \sum_{i=1}^{w+1} u(i) = r \right\}.$$

Definition 1.5 For $x \in \Sigma^*$, denote $r(x) \triangleq \|\psi_{\text{drt}(x)}(x)\|_1$.

Corollary 1.6 For $x \in \Sigma^*$ it holds that $r(x) = \frac{|x| - |\text{drt}(x)|}{k} < \lfloor \frac{|x|}{k} \rfloor$.

Proof The claim follows from $x \in D^{r(x)}(\text{drt}(x))$. \square

Noting that $r(C_x)$, for some $C \subseteq \Sigma^N$ and $x \in \text{Irr}$, is a singleton (whenever $C_x \neq \emptyset$), we shall find it comfortable going forward to think of it as an integer, by abuse of notation.

Theorem 1.7 *We take positive integers N, t and $n > k$. Then $C \subseteq \Sigma^n$ is a reconstruction code for t duplications, with uncertainty N , if and only if for all $x \in \text{Irr}$ such that $C_x \neq \emptyset$, the image $\psi_x(C_x) \subseteq \Delta_{r(C_x)}^{w(x)}$ satisfies*

$$\min\{d_1(c, c') : c \neq c' \in \psi_x(C_x)\} \geq d_{N,t}(w(x)),$$

where we make the notation

$$d_{N,t}(w) \triangleq \min\left\{\delta \in \mathbb{N} : \binom{t - \delta + w}{w} \leq N\right\}.$$

Proof If $C \cap D^*(x) \neq \emptyset$ then it follows from the definitions that for some $r \in \mathbb{N}$ we have $|x| + rk = n$; since $|x| \geq k$, necessarily $r = r(C_x) < \lfloor \frac{n}{k} \rfloor$. Furthermore, $C \cap D^*(x) = C \cap D^r(x)$, hence we have seen in the proof of Lemma 1.2 that for all $y \in D^r(x)$ we have $\psi_x(y) = \sum_{u=1}^r e_{j_u} \in \Delta_r^{w(x)}$.

In addition, by Lemma 1.2 and Lemma 1.3, for all $x \in \text{Irr}$ and $y \neq y' \in C_x$ the size of intersection $D^t(y) \cap D^t(y')$ is $\binom{t - d(y, y') + w(x)}{w(x)}$. It follows that C_x is a reconstruction code with uncertainty N if and only if that size is no greater than N for all such $y, y' \in C_x$.

Recalling that ψ_x is bijective and distance-preserving, i.e., that $d(y, y') = d_1(\psi_x(y), \psi_x(y'))$, the claim follows for C_x .

To conclude the proof, we recall that for $x, x' \in \text{Irr}$ we have $D^*(x) \cap D^*(x') = \emptyset$, hence C is a reconstruction code if and only if the same is true for C_x , for all $x \in \text{Irr}$. \square

In other words, Theorem 1.7 states that the intersection of a uniform-tandem-duplication reconstruction code C with the descendant cone of any irreducible string $D^*(x)$ can be viewed as an error-correcting code with a suitable minimal distance. Further, we see that these error-correcting codes are equivalent to codes in the Manhattan metric over a simplex $\Delta_{r(C_x)}^{w(x)}$. We note here, however, that this does not hold for C in general: not only is each code's minimal distance dependent on x , but the dimension and weight of the simplex in which that code exists do, as well.

We therefore see that constructions and bounds on the size of error-correcting codes for uniform tandem-duplication depend on doing the same for error-correcting codes in the Manhattan metric over Δ_r^w . We start by notating the maximal size of such codes:

Definition 1.8 For $w, r > 0$ and $d \geq 0$ we define

$$M(w, r, d) \triangleq \max \left\{ |C| : C \subseteq \Delta_r^w, \min_{\substack{c, c' \in C \\ c \neq c'}} d_1(c, c') \geq d \right\}.$$

We now reiterate that if $C \subseteq \Sigma^n$, $x, x' \in \text{Irr}(n - rk)$ (i.e., $r(C_x) = r(C_{x'}) = r$, if indeed $C_x, C_{x'} \neq \emptyset$) and $w(x) = w(x')$, then $D^{n-rk}(x) \cong D^{n-rk}(x')$ (i.e., they are isomorphic, through, e.g., $\psi_{x'}^{-1} \circ \psi_x$). It is therefore practical to assume $|C_x| = |C_{x'}| = M(w, r, d_{N,t}(w))$ for all such x, x' . This results in the following corollary, which concludes this section:

Corollary 1.9 If $C \subseteq \Sigma^n$ is a reconstruction code, and for all $x \in \text{Irr}$ it holds that $|C_x| = M(w(x), r, d_{N,t}(w(x)))$, then

$$\begin{aligned} |C| &= \sum_{r=0}^{\lfloor n/k \rfloor - 1} \sum_w M(w, r, d_{N,t}(w)) \cdot |\{x \in \text{Irr}(n - rk) : w(x) = w\}| \\ &= \sum_{r=0}^{\lfloor n/k \rfloor - 1} \sum_w M(w, r, d_{N,t}(w)) \cdot q^k \cdot \left| \left\{ b \in \Sigma^{n-(r+1)k} : b \text{ is } (0, k-1)_q\text{-RLL} \right. \right. \\ &\quad \left. \left. \text{wt}(b) = w \right\} \right| \end{aligned}$$

Proof First, trivially, $|C| = \sum_{x \in \text{Irr}} |C_x|$. Observe that $x \in \text{Irr}$ satisfies $C_x \neq \emptyset$, $r(C_x) = r$, if and only if $x \in \Sigma^{n-rk}$ and $\bar{\phi}(x)$ is $(0, k-1)_q$ -RLL.

The rest now follows from Theorem 1.7. \square

Corollary 1.9 motivates us to estimate the optimal size of error-correcting codes in the Manhattan metric over the (w, r) -simplex. This topic was examined in some depth in [KT18b], which was published as this work was being compiled, where a construction based on Sidon sets (of particular interest for our application, see [KT17], and references therein) was proposed, leading to lower bounds tighter than the Gilbert-Varshamov bound. For our purposes, we cite an asymptotic result (we slightly rephrase):

Lemma 1.10 [KT18b, Eq. 36] Take $\omega \in (0, 1)$, $\rho > 0$ and integer sequences $(w_n)_{n>0}$, $(r_n)_{n>0}$ such that $\lim_{n \rightarrow \infty} \frac{w_n}{n} = \omega$ and $\lim_{n \rightarrow \infty} \frac{r_n}{n} = \rho$. Also take a fixed $d > 0$. Then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log_2 M(w_n, r_n, d) = (\omega + \rho) H \left(\frac{1}{1 + \frac{\rho}{\omega}} \right), \quad (1.1)$$

where H is the binary entropy function, defined by

$$H(p) \triangleq -p \log_2 p - (1 - p) \log_2 (1 - p).$$

1.4 Minimal distance of reconstruction codes

Next, before we can ascertain the sizes of error-correcting codes over simplices, we bound their requisite minimal distance. That is, given $N, t > 0$ and $w > 0$, we establish bounds on

$$d_{N,t}(w) \triangleq \min \left\{ \delta \in \mathbb{N} : \binom{t - \delta + w}{w} \leq N \right\}$$

seen in Theorem 1.7.

Lemma 1.11 *If $N \leq w$ then $d_{N,t}(w) = t$.*

Proof We may verify by substitution that $\delta = t$ satisfies $\binom{t - \delta + w}{w} \leq N$, while $\delta = t - 1$ does not. Using the strict monotonicity of $s \mapsto \binom{s+w}{w}$, we are done. \square

In order to find a practical bound for $d_{N,t}(w)$ when $N > w$, we first require the following three lemmas:

Lemma 1.12 *1. [MS78, Ch.10, Sec.11, Lem.7] For integers $0 < k < n$ it holds that*

$$\sqrt{\frac{n}{8k(n-k)}} 2^{nH\left(\frac{k}{n}\right)} \leq \binom{n}{k} \leq \sqrt{\frac{n}{2\pi k(n-k)}} 2^{nH\left(\frac{k}{n}\right)};$$

2.

$$nH\left(\frac{k}{n}\right) - \frac{1}{2} \log_2(2n) \leq \log_2 \binom{n}{k} < nH\left(\frac{k}{n}\right).$$

Proof For Item 2, we see that if $0 < k < n$ we have $n - 1 \leq k(n - k) \leq \frac{n^2}{4}$, hence

$$\begin{aligned} \frac{n}{2\pi k(n-k)} &\leq \frac{1}{2\pi} \left(1 + \frac{1}{n-1}\right) \leq \frac{1}{\pi} < 1, \\ \frac{n}{8k(n-k)} &\geq \frac{1}{2n}. \end{aligned}$$

Thus the claim trivially follows from Item 1. \square

For ease of notation in what follows, we make the notation, for $1 \leq x \in \mathbb{R}$:

$$\mathcal{H}(x) \triangleq xH\left(\frac{1}{x}\right).$$

Lemma 1.13 *For $N > w > 0$ and $t > 0$ it holds that*

$$d_{N,t}(w) \leq \min \left\{ \delta \in \mathbb{N} : \mathcal{H}\left(1 + \frac{t - \delta}{w}\right) \leq \frac{\log_2 N}{w} \right\}.$$

Proof Under the assumption, $\delta = t - 1$ satisfies the inequality $\binom{t-\delta+w}{w} \leq N$. Therefore we may restrict the minimum to $\delta < t$, giving $0 < w < (t - \delta) + w$. Now, Lemma 1.12 implies

$$\log_2 \binom{t - \delta + w}{w} \leq w \left(1 + \frac{t - \delta}{w} \right) H \left(\frac{1}{1 + \frac{t - \delta}{w}} \right),$$

which completes the proof. \square

Lemma 1.14 For $x \geq 1$ it holds that $\mathcal{H}(x) \leq 2\sqrt{x - 1}$.

Proof The claim can be restated by the substitution $p = \frac{1}{x}$ as the known inequality $H(p)^2 \leq 4p(1 - p)$ (its proof follows elementary calculus, and is omitted here). \square

Finally,

Theorem 1.15 Take $N > w > 0$. Then

$$d_{N,t}(w) \leq \max \left\{ 1, t - \left\lfloor \frac{(\log_2 N)^2}{4w} \right\rfloor \right\}.$$

Proof Using Lemma 1.14 we may bound $\mathcal{H}\left(1 + \frac{t-\delta}{w}\right) \leq 2\sqrt{\frac{t-\delta}{w}}$. Lemma 1.13 therefore implies that it suffices to require $2\sqrt{\frac{t-\delta}{w}} \leq \frac{\log_2 N}{w}$, and reordering the inequality we get $\delta \geq t - \frac{(\log_2 N)^2}{4w}$, yielding the claim. \square

1.5 Capacity of reconstruction codes

We are interested in $\sup\{\text{cap}(\mathcal{C})\}$, where \mathcal{C} is any family of reconstruction codes (i.e., $\mathcal{C} \cap \Sigma^n$ is a reconstruction code for t_n duplications, with uncertainty N_n , for all n).

The purpose of this section is to determine that optimal capacity in two asymptotic regimes:

Regime I When $N_n = o(n)$ and $t_n = t$ is fixed.

Regime II When $N_n = 2^{\alpha n}$ and $t_n = \beta n$ for constants $\alpha, \beta > 0$ (such that $N_n, t_n \in \mathbb{N}$ for some, hence infinitely many, indices).

In practical applications, Regime I is likely to apply, since we may indeed expect the number of duplications t , which is dependent on the period of time before data is read, to be fixed w.r.t. n . The allowed uncertainty N_n will also likely be

bounded. Regime II requires Theorem 1.15 (and some restrictions over the values of α, β), but allows us to calculate capacity in much the same way, which we do after presenting the former.

Note, since [JFSB17a] showed that $\text{Irr}(n)$ can correct any number of tandem-duplication errors, they are trivially reconstruction codes for all N, t (more precisely, they are reconstruction codes with uncertainty $N = 0$ for all t). In comparison, in the setting we consider only t tandem-duplications are assumed to have occurred, therefore the codes we seek are less restrictive. Nevertheless, at the time of this work's compilation no bounds on the size of error-correcting codes for a fixed number of tandem-duplications were known; It is our purpose, then, to demonstrate that reconstruction codes exist which have strictly higher capacity than Irr , and suggest constructions for families of such codes. In the next chapter, we shall take a more general approach, and see that these results may indeed be strengthened.

First, we denote for any $n, r \in \mathbb{N}$ such that $n \geq k$ and $r < \lfloor \frac{n}{k} \rfloor$, and any $N, t \in \mathbb{N}$

$$\mathcal{M}_{N,t}(n, r) \triangleq \sum_w M(w, r, d_{N,t}(w)) \cdot \left| \left\{ b \in \Sigma^{n-(r+1)k} : \begin{array}{l} b \text{ is } (0, k-1)_q\text{-RLL} \\ \text{wt}(b)=w \end{array} \right\} \right|.$$

We recall for all n , if $r_n = \arg \max_r \mathcal{M}_{N,t}(n, r)$, that by Corollary 1.9 we have a reconstruction code $C \subseteq \Sigma^n$ with $|C| \geq q^k \mathcal{M}_{N,t}(n, r_n)$. Corollary 1.9 also implies that for all $C \subseteq \Sigma^n$ it holds that $|C| \leq \frac{n}{k} q^k \mathcal{M}_{N,t}(n, r_n)$. We therefore focus on maximizing $\limsup_{n \rightarrow \infty} \frac{1}{n} \log_q \mathcal{M}_{N,t}(n, r_n)$ by choice of r_n .

In what follows, we take $\gamma \in (0, 1)$ and set $r_n = \frac{1-\gamma}{k}n - 1$ for any $n \in \mathbb{N}$ for which $r_n \in \mathbb{N}$; we shall assume that such n exist (hence, infinitely many exist), and refer only to such indices.

For all $x \in \text{Irr}(n - r_n k) = \text{Irr}(k + \gamma n)$, recall that $\bar{\phi}(x) \in \Sigma^{\gamma n}$ is $(0, k-1)_q$ -RLL. We shall build a reconstruction code in the descendant cones of only such x (note, $D^r(x) \subseteq \Sigma^n$), which we denote C_γ .

Lemma 1.16 *There exists a system $\mathcal{S} \subseteq (0, k-1)_q$ -RLL and $\theta \in (\frac{1}{2}, 1)$ such that*

$$\text{cap}(\mathcal{S}) = \text{cap}((0, k-1)_q\text{-RLL})$$

and for all $b \in \mathcal{S}$ it holds that $\text{wt}(b) \geq \theta|b|$.

Proof Let $G_q(k-1)$ be the strongly connected deterministic digraph representing the $(0, k-1)_q$ -RLL system, seen in Figure 1, whose adjacency matrix is

$$T_q(k-1) = \begin{pmatrix} q-1 & 1 & 0 & \cdots & 0 \\ q-1 & 0 & 1 & & \vdots \\ \vdots & \vdots & & \ddots & 0 \\ q-1 & 0 & \cdots & 0 & 1 \\ q-1 & 0 & \cdots & \cdots & 0 \end{pmatrix}$$




Figure 1: The graph $G_q(k-1)$ generating the $(0, k-1)_q$ -RLL system.

As is well known for the case of $q = 2$ (see, e.g., [ZW88, How89]), its characteristic polynomial is

$$p_q^{(k-1)}(x) = x^k - (q-1) \sum_{j=0}^{k-1} x^j = \frac{x^{k+1} - qx^k + (q-1)}{x-1},$$

hence the Perron eigenvalue λ of $T_q(k-1)$ is the unique positive root of $\hat{p}_q^{(k-1)}(x) = x^{k+1} - qx^k + (q-1)$ greater than 1 (in fact, $\lambda \in (q-1, q)$, which can readily be confirmed either using elementary calculus or by information-theoretic methods, since $(\Sigma \setminus \{0\})^* \subseteq (0, k-1)_q\text{-RLL} \subseteq \Sigma^*$).

Further, $T_q(k-1)$ has positive right- and left-eigenvectors associated with λ , which we denote \bar{v}, \bar{w} respectively; specifically,

$$\bar{v} = \left(1, \lambda - (q-1), \dots, \lambda^{j-1} - (q-1) \sum_{i=0}^{j-2} \lambda^i, \dots, \lambda^{k-1} - (q-1) \sum_{i=0}^{k-2} \lambda^i \right),$$

$$\bar{w} = (\lambda^{k-1}, \lambda^{k-2}, \dots, \lambda^{k-j}, \dots, 1).$$

and we may verify that

$$v_k = \lambda^{k-1} - (q-1) \sum_{i=0}^{k-2} \lambda^i = \frac{1}{\lambda} \left[\lambda^k - (q-1) \sum_{j=1}^{k-1} \lambda^j \right]$$

$$= \frac{q-1}{\lambda} > 0$$

and $v_j = \frac{v_{j+1} + (q-1)}{\lambda}$, hence every entry of \bar{v} is indeed positive.

Denoting $q_{i,j} = (T_q(k-1))_{i,j} \cdot \frac{v_j}{\lambda v_i}$, it follows (see, e.g., [MRS01][Sec. 3.5]) that $Q = (q_{i,j})_{1 \leq i,j \leq k}$ is stochastic, and represents a transition matrix of a stationary

Markov chain \mathcal{P} on $G_q(k-1)$ (a probability measure on its edges set $E_q(k-1)$) satisfying $H(\mathcal{P}) = \log_q \lambda = \text{cap}((0, k-1)_q\text{-RLL})$. Further, the stationary distribution of the Markov chain, i.e., a positive $\bar{\pi} = (\pi_1, \dots, \pi_k)$ such that $\sum_{j=1}^k \pi_j = 1$ and $\bar{\pi}^T Q = \bar{\pi}^T$, is given by $\pi_j = \frac{\hat{\pi}_j}{\sum_{i=1}^k \hat{\pi}_i}$, where $\hat{\pi}$ is defined by $\hat{\pi}_j = w_j v_j$. It holds for all j that π_j is the sum of probabilities $\sum \mathcal{P}(e)$ of edges terminating at the j 'th node.

Note, then, that

$$\begin{aligned} \sum_{i=1}^k \hat{\pi}_i &= \lambda^{k-1} + \sum_{i=2}^k \left[\lambda^{k-1} - (q-1) \frac{\lambda^{k-1} - \lambda^{k-i}}{\lambda-1} \right] \\ &= \lambda^{k-1} \left[1 + (k-1) \left(1 - \frac{q-1}{\lambda-1} \right) \right] + \frac{q-1}{\lambda-1} \sum_{i=2}^k \lambda^{k-i} \\ &= \lambda^{k-1} \left[k - (k-1) \frac{q-1}{\lambda-1} \right] + \frac{q-1}{\lambda-1} \sum_{j=0}^{k-2} \lambda^j \\ &= \lambda^{k-1} \left[k - (k-1) \frac{q-1}{\lambda-1} \right] + \frac{\lambda^k - (q-1)\lambda^{k-1}}{\lambda-1} \\ &= \frac{\lambda^{k-1}}{\lambda-1} [\lambda - k(q-\lambda)] \end{aligned}$$

and in particular $\pi_1 = \frac{\lambda-1}{\lambda-k(q-\lambda)}$. (Incidentally, it follows from $\pi_1 \in (0, 1)$ that $1 < k(q-\lambda) < \lambda$, that is, $q - \frac{q}{k+1} < \lambda < q - \frac{1}{k}$.)

Next, recall that for a given $\epsilon > 0$, a (\mathcal{P}, ϵ) -strongly-typical path in G is a path $\gamma = (e_1, e_2, \dots, e_l)$ (denoted by its edges $\{e_1, e_2, \dots, e_l\} \subseteq E_q(k-1)$) such that each $e \in E_q(k-1)$ appears in the path $l \cdot \tau$ times, for some τ satisfying $|\tau - \mathcal{P}(e)| \leq \epsilon$. If we let $\mathcal{S}_\epsilon \subseteq \Sigma^*$ be the system induced by $(\mathcal{P}, \frac{\epsilon}{k(q-1)})$ -strongly-typical paths, then it is well known that $\text{cap}(\mathcal{S}_\epsilon) = \text{cap}((0, k-1)_q\text{-RLL})$. Note, for $b \in \mathcal{S}_\epsilon$ of length $|b| = l$, which is generated by the path $\gamma = (e_1, \dots, e_l)$, $\text{wt}(b)$ is precisely the number of edges which terminate at the first node; since γ is $(\mathcal{P}, \frac{\epsilon}{k(q-1)})$ -strongly-typical,

$$\text{wt}(b) \geq \sum_{\substack{e \text{ terminates} \\ \text{at first node}}} l \cdot \left(\mathcal{P}(e) - \frac{\epsilon}{k(q-1)} \right) = l(\pi_1 - \epsilon)$$

To conclude the proof, note

$$\begin{aligned} \lambda + k(q-\lambda) &= q + (k-1)(q-\lambda) > q \geq 2 \\ \implies \lambda &> 2 - k(q-\lambda) \\ \implies 2(\lambda-1) &> \lambda - k(q-\lambda) \implies \pi_1 > \frac{1}{2} \end{aligned}$$

Hence we can take any $0 < \epsilon < \pi_1 - \frac{1}{2}$, and observe that $\mathcal{S} = \mathcal{S}_\epsilon$, $\theta = \pi_1 - \epsilon$ satisfy the proposition. \square

Lemma 1.16 implies the existence of a subset $S_k \subseteq \text{Irr}$ such that $\text{cap}(S_k) = \text{cap}(\text{Irr})$, and for every $x \in S_k$ of length $|x| = k + \gamma n$ we have $w(x) \geq \lceil \theta \cdot \gamma n \rceil$. For the rest of this section we only build codes \mathcal{C}_γ^n in the descendant cones of roots in S_k . Note, then, that if we denote $w_n = \lceil \theta \cdot \gamma n \rceil$ and $\mathcal{C}_\gamma \triangleq \bigcup \mathcal{C}_\gamma^n$, then

$$\text{cap}(\mathcal{C}_\gamma) = \limsup_{n \rightarrow \infty} \frac{1}{n} \log_q \sum_{w \geq w_n} M(w, r_n, d_{N_n, t}(w)) \cdot \left| \left\{ x \in S_k : \begin{array}{l} |x| = n - (r+1)k \\ w(x) = w \end{array} \right\} \right|.$$

We evaluate this quantity in the following theorem:

Theorem 1.17 *Take \mathcal{C}_γ as defined above. Then*

$$\text{cap}(\mathcal{C}_\gamma) \geq \gamma \text{cap}(\text{Irr}) + \frac{\theta \gamma}{\log_2 q} \cdot \mathcal{H} \left(1 + \frac{1 - \gamma}{k \theta \gamma} \right)$$

in both of the aforementioned two regimes:

1. *Regime I: when $N_n = o(n)$ and $t_n = t$ is fixed.*
2. *Regime II: when $N_n = 2^{\alpha n}$ and $t_n = \beta n$, if we additionally require $\alpha^2 > 4\beta$.*

Proof 1. Note, for sufficiently large n , that $N_n < \theta \cdot \gamma n \leq w_n$, resulting by Lemma 1.11 in $d_{N_n, t}(w) = t$ for all $w \geq w_n$. It follows that

$$\begin{aligned} \text{cap}(\mathcal{C}_\gamma) &= \limsup_{n \rightarrow \infty} \frac{1}{n} \log_q \sum_{w \geq w_n} M(w, r_n, t) \cdot \left| \left\{ x \in S_k : \begin{array}{l} |x| = n - (r+1)k \\ w(x) = w \end{array} \right\} \right| \\ &\geq \limsup_{n \rightarrow \infty} \frac{1}{n} \log_q |S_k \cap \Sigma^{n - (r+1)k}| \cdot M(w_n, r_n, t) \\ &= \gamma \text{cap}(\text{Irr}) + \limsup_{n \rightarrow \infty} \frac{1}{n} \log_q M(w_n, r_n, t). \end{aligned}$$

We note that $\lim_{n \rightarrow \infty} \frac{r_n}{n} = \frac{1 - \gamma}{k}$ and $\lim_{n \rightarrow \infty} \frac{w_n}{n} = \theta \gamma$, hence the claim is proven by Lemma 1.10 when t is fixed.

2. By Theorem 1.15:

$$\begin{aligned} d_{N_n, t_n}(w) &\leq \max \left\{ 1, \beta n - \left\lfloor n \frac{\alpha^2 n}{4 \lceil \theta \cdot \gamma n \rceil} \right\rfloor \right\} \\ &\leq \max \left\{ 1, \left\lfloor \left(\beta - \frac{\alpha^2}{4} \right) n \right\rfloor \right\}. \end{aligned}$$

If $\alpha^2 > 4\beta$ then for sufficiently large n we have $d_{N_n, t_n}(w) = 1$ for all w . Since it is fixed, we may now apply the same argument used in the previous part. \square

Going forward, we shall view the lower bound to $\text{cap}(\mathcal{C}_\gamma)$,

$$C(\gamma) \triangleq \gamma \text{cap}(\text{Irr}) + \frac{\theta\gamma}{\log_2 q} \cdot \mathcal{H}\left(1 + \frac{1-\gamma}{k\theta\gamma}\right),$$

as a function of γ . Before moving on to show that it may be made to exceed $\text{cap}(\text{Irr})$ by a careful choice of γ , we look at the following example.

Example 1.18 *Set $q = k = 2$. Then the Perron eigenvalue of $T_2(1)$ is $\lambda = \frac{1+\sqrt{5}}{2}$, and*

$$\text{cap}(\text{Irr}_2) = \log_2(\lambda) = \log_2\left(\frac{1+\sqrt{5}}{2}\right) \approx 0.6942.$$

In addition, any θ which is less than $\pi_1 = \frac{1}{2}\left(1 + \frac{1}{\sqrt{5}}\right) \approx 0.7236$ satisfies Lemma 1.16.

Alternatively, we may set $q = 4$ (for the special case of DNA) and duplication-length $k = 2$. Now the Perron eigenvalue of $T_4(1)$ is given by $\lambda = \frac{3+\sqrt{21}}{2}$, hence

$$\text{cap}(\text{Irr}_2) = \log_4(\lambda) = \log_4\left(\frac{3+\sqrt{21}}{2}\right) \approx 0.9613.$$

Further, we may choose any θ which is less than $\pi_1 = \frac{1}{2}\left(1 + \sqrt{\frac{3}{7}}\right) \approx 0.8273$.

$C(\gamma)$ is shown for both cases in Figure 2, under the assumptions of asymptotic regime made in Theorem 1.17. The figure demonstrates that the capacity of reconstruction codes (bounded from below by the maximum of the curve) is greater than $\text{cap}(\text{Irr})$. \blacksquare

We now attempt to maximize $C(\gamma)$ by a proper choice of $\gamma \in (0, 1)$. Analysis of $C(\gamma)$ is simpler using the following change of variable:

Definition 1.19 *Define $x: (0, 1) \rightarrow (0, \infty)$ by $x(\gamma) \triangleq \frac{1-\gamma}{\gamma}$.*

We observe that $x(\gamma)$ is a decreasing diffeomorphism, and $\gamma = \frac{1}{1+x(\gamma)}$.

Lemma 1.20 *One has*

$$C(\gamma) = \gamma \text{cap}(\text{Irr}) + \theta\gamma \left[\left(1 + \frac{x(\gamma)}{k\theta}\right) \log_q \left(1 + \frac{x(\gamma)}{k\theta}\right) - \frac{x(\gamma)}{k\theta} \log_q \left(\frac{x(\gamma)}{k\theta}\right) \right]$$




Figure 2: Capacity $C(\gamma)$ in the cases (a) $q = k = 2$, $\theta = 0.7236$, and (b) $q = 4$, $k = 2$, $\theta = 0.8273$. The value at $\gamma = 1$ equals $\text{cap}(\text{Irr})$.

Proof We observe that for all $x > 0$, $\log\left(1 + \frac{1}{x}\right) = \log\left(\frac{x+1}{x}\right) = \log(x+1) - \log x$; in particular

$$\log_q\left(1 + \frac{k\theta\gamma}{1-\gamma}\right) = \log_q\left(1 + \frac{1-\gamma}{k\theta\gamma}\right) - \log_q\left(\frac{1-\gamma}{k\theta\gamma}\right)$$

Hence,

$$\begin{aligned} C(\gamma) &= \gamma \text{cap}(\text{Irr}) + \frac{\theta\gamma}{\log_2 q} \cdot \mathcal{H}\left(1 + \frac{1-\gamma}{k\theta\gamma}\right) \\ &= \gamma \text{cap}(\text{Irr}) + \theta\gamma \log_q\left(1 + \frac{1-\gamma}{k\theta\gamma}\right) + \frac{1-\gamma}{k} \log_q\left(1 + \frac{k\theta\gamma}{1-\gamma}\right) \\ &= \gamma \text{cap}(\text{Irr}) + \left(\theta\gamma + \frac{1-\gamma}{k}\right) \log_q\left(1 + \frac{1-\gamma}{k\theta\gamma}\right) - \frac{1-\gamma}{k} \log_q\left(\frac{1-\gamma}{k\theta\gamma}\right) \\ &= \gamma \text{cap}(\text{Irr}) + \theta\gamma \left[\left(1 + \frac{1-\gamma}{k\theta\gamma}\right) \log_q\left(1 + \frac{1-\gamma}{k\theta\gamma}\right) \right. \\ &\quad \left. - \frac{1-\gamma}{k\theta\gamma} \log_q\left(\frac{1-\gamma}{k\theta\gamma}\right) \right] \end{aligned}$$

$$\begin{aligned}
C'(\gamma) &= \text{cap}(\text{Irr}) + \frac{dx}{d\gamma} \cdot \frac{d}{dx} \left[\frac{\theta}{1+x} \left(\left(1 + \frac{x}{k\theta}\right) \log_q \left(1 + \frac{x}{k\theta}\right) - \frac{x}{k\theta} \log_q \left(\frac{x}{k\theta}\right) \right) \right]_{x=x(\gamma)} \\
&= \text{cap}(\text{Irr}) - \frac{1}{\gamma^2} \left[\frac{-\theta}{(1+x)^2} \left(\left(1 + \frac{x}{k\theta}\right) \log_q \left(1 + \frac{x}{k\theta}\right) - \frac{x}{k\theta} \log_q \left(\frac{x}{k\theta}\right) \right) \right. \\
&\quad \left. + \frac{\theta}{(1+x)} \cdot \left(\frac{1}{k\theta} \log_q \left(1 + \frac{x}{k\theta}\right) - \frac{1}{k\theta} \log_q \left(\frac{x}{k\theta}\right) \right) \right]_{x=x(\gamma)} \\
&= \text{cap}(\text{Irr}) + \frac{1}{k} \left[(k\theta - 1) \log_q \left(1 + \frac{x(\gamma)}{k\theta}\right) + \log_q \left(\frac{x(\gamma)}{k\theta}\right) \right] \tag{1.2}
\end{aligned}$$

□

We can now show that there always exists a choice of γ for which we get $R(C_\gamma^n) > \text{cap}(\text{Irr})$:

Theorem 1.21 $\max_{\gamma \in (0,1)} C(\gamma) > \text{cap}(\text{Irr})$.

Proof Observe that $C(\gamma)$ is continuously differentiable and satisfies $C(0) = 0$, $C(1) = \text{cap}(\text{Irr})$ (when extended continuously). We find $C'(\gamma)$ in Equation (1.2); Thus, We can show that $C'(\gamma) = 0$ if and only if

$$q^{-k \text{cap}(\text{Irr})} = \left(1 + \frac{x(\gamma)}{k\theta}\right)^{k\theta-1} \cdot \frac{x(\gamma)}{k\theta} \tag{1.3}$$

This equation has a unique solution $x_0 = x(\gamma_0)$, since the RHS is a monotonic increasing function of x , vanishing at $x = 0$ and unbounded as x grows. Moreover, $0 < x_0 < k\theta$, since $k\theta > 1$, hence the RHS is greater than 1 at $x = k\theta$. Thus $C(\gamma)$ has a unique local extremum in $(0, 1)$.

It now suffices to show that $C(\gamma)$ is concave, hence the extremum is a maximum. Indeed,

$$\begin{aligned}
C''(\gamma) &= \frac{1}{k} \frac{dx}{d\gamma} \cdot \frac{d}{dx} \left[(k\theta - 1) \log_q \left(1 + \frac{x}{k\theta}\right) + \log_q \left(\frac{x}{k\theta}\right) \right]_{x=x(\gamma)} \\
&= \frac{-1}{k \ln(q) \gamma^2} \left[\frac{k\theta - 1}{k\theta + x(\gamma)} + \frac{1}{x(\gamma)} \right] < 0
\end{aligned}$$

It follows that $C(\gamma_0) > \lim_{\gamma \rightarrow 1} C(\gamma) = \text{cap}(\text{Irr})$. □

Thus, the main result of this paper is established. In what remains of this section we show that we can bound γ_0 which maximizes $C(\gamma)$, in practice, to any desired level of accuracy. We begin by establishing bounds in the following lemma.

Lemma 1.22 *Let $\gamma_0 \in (0, 1)$ be the unique maximum of $C(\gamma)$, and denote $x_0 = x(\gamma_0)$. Then*

$$x_0 \geq \frac{k\theta}{(2^\theta q^{\text{cap}(\text{Irr})k})^k - 1}$$

and

$$\begin{aligned} x_0 &\leq \frac{1}{2} \left[\sqrt{(1 - q^{-\text{cap}(\text{Irr})k})^2 + k\theta q^{2-\text{cap}(\text{Irr})k}} - (1 - q^{-\text{cap}(\text{Irr})k}) \right] \\ &\leq \frac{k\theta q^2}{4(q^{\text{cap}(\text{Irr})k} - 1)}. \end{aligned}$$

Proof For fixed $x \in [0, \infty)$ define $g_x: (0, \infty) \rightarrow \mathbb{R}$ by $g_x(y) = y \ln\left(1 + \frac{x}{y}\right)$. Then

$$\begin{aligned} g'_x(y) &= \ln\left(1 + \frac{x}{y}\right) + \frac{y}{1 + \frac{x}{y}} \cdot \frac{-x}{y^2} = \ln\left(1 + \frac{x}{y}\right) - \frac{x}{y+x} \\ &= -\ln\left(1 - \frac{x}{x+y}\right) - \frac{x}{y+x} \\ &\geq -\left(-\frac{x}{x+y}\right) - \frac{x}{y+x} \geq 0. \end{aligned}$$

Therefore, $f_x(y) = e^{g_x(y)} = \left(1 + \frac{x}{y}\right)^y$ satisfies $1 + x = f_x(1) \leq f_x(y) = \left(1 + \frac{x}{y}\right)^y$ for all $y \geq 1$. In our case $k\theta > 1$ and x_0 satisfies Equation (1.3), hence

$$q^{-\text{cap}(\text{Irr})k} = \left(1 + \frac{x_0}{k\theta}\right)^{k\theta-1} \frac{x_0}{k\theta} \geq \frac{1+x_0}{1+\frac{x_0}{k\theta}} \cdot \frac{x_0}{k\theta} = \frac{x_0 + x_0^2}{k\theta + x_0}$$

which we simplify to $0 \geq x_0^2 + (1 - q^{-\text{cap}(\text{Irr})k})x_0 - k\theta q^{-\text{cap}(\text{Irr})k}$. Thus, the first upper bound is proven. For the second, we require only that for $a, b > 0$ it holds that $\sqrt{a+b^2} - b \leq \frac{a}{2b}$, which is readily shown by differentiation.

On the other hand, Equation (1.3) implies that $x_0 \leq k\theta$. Therefore

$$\begin{aligned} q^{-\text{cap}(\text{Irr})k} &= \left(1 + \frac{x_0}{k\theta}\right)^{k\theta-1} \frac{x_0}{k\theta} \leq \frac{2^{k\theta}}{1 + \frac{x_0}{k\theta}} \cdot \frac{x_0}{k\theta} \\ &\iff k\theta q^{-\text{cap}(\text{Irr})k} \leq (2^{k\theta} - q^{-\text{cap}(\text{Irr})k})x_0 \end{aligned}$$

which proves the lower bound. \square

Next, we show that we may tighten the bounds we derived in the previous lemma.

Lemma 1.23 *Let $x_0 > 0$ be the unique solution to Equation (1.3), and denote $z_0 = \frac{x_0}{k\theta}$. If $\underline{z} \leq z_0 \leq \bar{z}$ then $F(\underline{z}) \leq z_0 \leq F(\bar{z})$, where*

$$F(z) \triangleq \frac{q^{-\text{cap}(\text{Irr})k}}{\left(1 + \frac{q^{-\text{cap}(\text{Irr})k}}{(1+z)^{k\theta-1}}\right)^{k\theta-1}}.$$

Proof By assumption we have

$$q^{-\text{cap}(\text{Irr})k} = (1 + z_0)^{k\theta-1} \cdot z_0,$$

hence $q^{-\text{cap}(\text{Irr})k} \leq (1 + \bar{z})^{k\theta-1} \cdot z_0$, implying that $z_0 \geq G(\bar{z})$ where $G(z) = \frac{q^{-\text{cap}(\text{Irr})k}}{(1+z)^{k\theta-1}}$. Similarly, $z_0 \leq G(\underline{z})$. The proposition now trivially follows for $F(z) = G(G(z))$. \square

Finally, we can show that x_0 may be found by the following limiting process:

Theorem 1.24 *The unique solution to Equation (1.3) is given by $x_0 = k\theta \lim_{n \rightarrow \infty} F^n(z_1)$, for all $z_1 \in [0, 1]$.*

Proof As before, we denote the unique solution $x_0 > 0$, and take $z_0 = \frac{x_0}{k\theta}$.

Note that Lemma 1.23 implies that $z_0 = F(z_0)$. We will prove that $F: [0, 1] \rightarrow [0, 1]$ is a contraction; that is, for all $z_1, z_2 \in [0, 1]$ we have $|F(z_1) - F(z_2)| \leq c|z_1 - z_2|$ for some $c < 1$. Indeed, recalling $k\theta > 1$ we find

$$\begin{aligned} F'(z) &= \frac{2^{-2\text{cap}(\text{Irr})k}(k\theta - 1)^2}{(1+z)^{k\theta} \left(1 + \frac{q^{-\text{cap}(\text{Irr})k}}{(1+z)^{k\theta-1}}\right)^{k\theta}} \\ &\leq \frac{(k\theta - 1)^2}{(2^{2\text{cap}(\text{Irr})k})^k} \leq \frac{(k - 1)^2}{2^k} \leq \frac{9}{16} < 1, \end{aligned}$$

where the next to last inequality may be directly verified for all small k .

Having done so, we utilize Banach's fixed-point theorem to deduce that F has a unique fixed point (necessarily z_0), and for all $z_1 \in [0, 1]$, defining $z_{n+1} = F(z_n)$ we get $\lim_{n \rightarrow \infty} z_n = z_0$. \square

We can now suggest a construction for reconstruction codes achieving better capacity than the error-correcting codes $\text{Irr}(n)$ suggested in [JFSB17a] (provided that one is willing to consider reconstruction codes over unambiguous decoding of any single output).

Construction 1.A We set the alphabet size q , duplication length k . In the case that our application falls within Regime I, we also set a fixed decoding-delay t , and restrict the ambiguity N_n to be sub-linear in n . (with the necessary adjustments, this construction also applies for Regime II.)

- Start by finding the Perron eigenvalue λ of $T_q(k-1)$, and $\pi_1 = \frac{\lambda-1}{\lambda-k(q-\lambda)}$, as in the proof of Lemma 1.16. Set some $\theta < \pi_1$.
- The upper and lower bounds on x_0 from Lemma 1.22 can be made tighter by a repetitive application of $F(\cdot)$ from Lemma 1.23; Theorem 1.24 guarantees that the bounds—hence the acceptable error—can be made as tight as desired for our application.
- With $\gamma_0 = \frac{1}{1+x_0}$ we may find $r_n = \frac{1-\gamma_0}{k}n - 1$, and we note that a capacity-achieving subset of $\text{Irr}(n - r_n k) = \text{Irr}(k + \gamma n)$ has weight $w(x) \geq w_n = \lceil \theta \cdot \gamma n \rceil$.
- Within $D^{r_n}(x)$ of just such irreducible strings x we may utilize any construction of codes for the Manhattan metric over $\Delta_{r_n}^{m_n}$ with minimal distance t , if it produces codes of size sufficiently close to $M(w_n, r_n, t)$. For practical applications, [KT18b, Sec. IV-A] showed that if w_n is a prime power, then by [BC62] there exist such codes of size $|\Delta_{r_n}^{w_n}| / \frac{w_n^t - 1}{w_n - 1}$ (which improves on the Gilbert-Varshamov bound, and is sufficiently tight to achieve the same result as in Theorem 1.17).

■

Note that we do not establish that Construction 1.A produces a system of codes of capacity 1, rather only greater than $\text{cap}(\text{Irr})$. To conclude this section, we also present a non-constructive argument proving the existence of a system of reconstruction codes with capacity 1 by an application of the Gilbert-Varshamov bound.

Recall that in the proof of Theorem 1.17 we have shown that the minimal distance, $d_{N_n, t_n}(w_n)$ was bounded. In particular, in the case of interest Regime I, we used the fact that $w_n = \Theta(n)$; This does not, in general, hold for $w(\text{drt}(y))$ for all $y \in \Sigma^n$.

However, if we show that to be the case for a sufficiently large subset $S^n \subseteq \Sigma^n$, then we may note the following: by [KT18b, Lem. 1] the size of ball in the $d(\cdot, \cdot)$ metric of radius d in the descendant cone of $x \in \text{Irr}$, where $w(x) \geq d$, is

$$\begin{aligned}
& \sum_{j=0}^d \binom{w(x)}{j} \binom{d}{j} \binom{d+w(x)-j}{d} \\
& \leq (d+1) \cdot \binom{w(x)}{d} \binom{d}{\lfloor d/2 \rfloor} \binom{d+w(x)}{d} \\
& = O(w(x)^d) = O(n^d)
\end{aligned}$$

It would follow that a code of size $\frac{|S^n|}{O(n^d)}$ exists (and, again, the rates of these codes will be $R(S^n)$).

It now suffices to show that except for a vanishingly small portion of $y \in \Sigma^n$, it holds that $w(\text{drt}(y)) = \Theta(n)$. Indeed, recall that $w(\text{drt}(y)) = w(y) = \text{wt}(\bar{\phi}(y))$, where $\bar{\phi}(y) \in \Sigma^{n-k}$. Then, for any real $0 < \xi < 1 - \frac{1}{q}$,

$$\frac{|\{b \in \Sigma^{n-k} : \text{wt}(b) \leq \xi(n-k)\}|}{q^{n-k}} \leq q^{(n-k)(H_q(\xi)-1)},$$

where $H_q(\cdot)$ is the q -ary entropy function,

$$H_q(\xi) \triangleq -\xi \log_q \xi - (1-\xi) \log_q(1-\xi) + \xi \log_q(q-1),$$

and where we used a standard bounding of the size on the Hamming ball, e.g., see [Rot06, Lemma 4.7].

1.6 Discussion of recent results

Before finishing, we shall note here that the last argument also shows via the GV bound that error-correcting codes for a fixed number of tandem-duplications achieve capacity 1. Indeed, during the compilation of this work [KT18a, LJWZ18] were made available, wherein bounds on the optimal size of such error-correcting codes were presented; these bounds show that the redundancy required to correct a fixed number of tandem-duplications is logarithmic in n .

More specifically, both works showed (see [KT18a, Thm. 4], [LJWZ18, Lem. 6]) that there exist codes $C^n \subseteq \Sigma^n$ that correct up to t tandem-duplications, for a fixed $t \in \mathbb{N}$, satisfying

$$|C^n| \geq \frac{(1+o(1))q^t}{(q-1)^t} \cdot \frac{q^n}{n^t}$$

They also showed that the optimal size was $O\left(\frac{q^n}{n^t}\right)$. Finally, [KT18a, Lem. 3] demonstrated that C^n can be assumed w.l.o.g. to only contain strings which roots satisfy $w(x) = \Theta(n)$.

We note that error-correcting codes for t tandem-duplications have minimal $d(\cdot, \cdot)$ distance $t+1$; In comparison, then, we have showed that reconstruction codes, where t is fixed and $N = o(n)$, have minimal distance t (when restricted to descendant cones of irreducible strings with $w(x) = \Theta(n)$). The observations above imply that codes designed in the aforementioned works for correcting $t-1$ tandem-duplications, of size $\geq \frac{(1+o(1))q^{t-1}}{(q-1)^{t-1}} \cdot \frac{q^n}{n^{t-1}}$, are reconstruction codes. Importantly, this validates the hypothesis that reconstruction codes for data storage in the DNA of living organisms offer greater data-density than error-correcting codes.

Namely, in comparison to the $t \log(n) + O(1)$ redundancy achieved by optimal error-correcting codes in [KT18a, LJWZ18], reconstruction codes for t duplications achieve redundancy $(t - 1) \log(n) + O(1)$. This result is further generalized in the next chapter, where we shall take advantage of another observation which allows us to disregard some of the specific parameters of any particular element of our codebook.

We also note for completeness that our results in Regime II, albeit less applicable in practice, are unique to this work.

Finally, these findings give rise to a natural question, whether an explicit encoder can be found for reconstruction codes, which achieves optimal (or close to optimal) redundancy. We finish by demonstrating that, with a few amendments to previous works, the answer to that question is in the affirmative. For reference, the encoding procedure of such an encoder is briefly outlined below. We describe an encoder for a code which can correct t duplications for simplicity; as we've seen, a reconstruction code requires protecting against only $t - 1$ duplications (and then some vanishing fraction of the codebook might need to be eliminated, to achieve $w(x) = \Omega(|x|)$ for every codeword x).

1. Take $x \in \Sigma^n$, and denote $w \triangleq w(x) \leq n$; Protecting against tandem repeats is equivalent to protecting $\psi_{\text{drt}(x)}(x)$, as we've seen, given the (in fact, any) descendant of x .
2. If at most t tandem repeats occur, then [MV17] showed that it suffices to know the $w + 1$ elements of $\psi_{\text{drt}(x)}(x)$ modulo a prime $p > t$ in order to fully recover data. [MV17] studied sticky insertions of 1's in binary strings, and utilized results from [RS94] on the l_1 error-correcting capability of BCH codes to demonstrate that a systematic encoder of such codes allows one to only store $t \log_p(n)$ redundancy symbols (modulo p); the same approach is applicable to protecting $\psi_{\text{drt}(x)}(x)$ from t additions of 1 (recall, this is equivalent to protecting x from t tandem repeats).
3. We define a q -ary balanced x sequence as such with a typical number of zeroes (i.e., $\lfloor \frac{q-1}{q}|x| \rfloor$). Building on Knuth's original work on balanced binary sequences [Knu86], it can be seen that every symbol modulo p may be encoded to a q -ary balanced sequence of length $N = \log_q(p) + O(\log(\log_q(p)))$. An explicit encoding for this correspondence is achieved by extending Knuth's method; we flip occurrences of 0 and one chosen other q -ary symbol, up to a certain position, to generate a balanced sequence, and that index then encoded iteratively in similar fashion. This process can be efficiently performed (also, reversed).

4. Following a similar approach to [MV17], the parity symbols themselves are then encoded to protect against t deletions, e.g., by using similar BCH codes, and every resulting symbol encoded as a balanced q -ary sequence; two consecutive symbols are delimited with a single element of $\Sigma \setminus \{0\}$, which is chosen as the q -ary symbol which was used in the last part. Duplication noise is then easily identified at the redundancy symbols, and interpreted as deletions, and successful recovery of the original redundancy symbols is assured. This suffices for decoding.

We only need note that if we allow p to increase arbitrarily (e.g., $p = \log(n)$), then

$$t \log_p(n) [\log_q(p) + O(\log(\log_q(p)))] = t \log_q(n) + o(\log(n))$$

Recall that the minimum required redundancy is $t \log_q(n) + O(1)$; Asymptotically, the redundancy we achieve in this fashion is therefore optimal.

Chapter 2

Uncertainty with List Decoding

In Chapter 1 we have seen that a classical error-correction coding approach is sub-optimal for the application of in-vivo DNA data storage, as it does not take advantage of the cost-effective data replication offered inherently by the medium; instead, it was shown that reframing the problem as a *reconstruction* scheme reduces the redundancy required for any fixed number of duplication errors. Of works considering applications of the reconstruction problem to storage technologies, [YB19] in particular extended the reconstruction model to *associative memory*, where one retrieves the set of all entries (or codewords) *associated* with every element of a given set. For a given size of entry set, the maximal number of entries being possibly associated with all of them was dubbed the *uncertainty* of the memory.

Study of this extended model for in-vivo DNA data storage is motivated by a list-decoding reconstruction scheme, whereby tolerance for decoding a list of possible inputs, given multiple channel outputs, enables coding with a lower minimum distance, thereby reducing the redundancy of the code. Alternatively, given the same code, it allows reducing the number of required outputs for reconstruction.

This chapter again focuses on uniform tandem-duplication noise. Our main goal shall be to analyze the uncertainty associated with codes which are subsets of a typical set of strings (consisting of most strings in Σ^n , a definition which is made precise in Lemma 2.2) as a function of the acceptable list size m and code minimum distance d , where the number of tandem repeats t which channel outputs undergo is fixed.

2.1 Preliminaries and related works

Associative memory was discussed in [YB19], where items are retrieved by association with other items; the human mind seems to operate in this fashion,

one concept bringing up memories of other, related, concepts or events. The more items one considers together, the smaller the set of items associated with all of them. Giving a precise definition to that notion, one defines the uncertainty of an associative memory as the largest possible size of set $N(m)$ whose members are associated with all elements of an m -subset of the memory codebook.

This model is a generalization of the reconstruction problem posed by Levenshtein in [Lev01], wherein a transmission model is assumed with the decoder receiving multiple channel outputs of the same input. N is then the largest size of intersection of balls of radius t about two distinct codewords, where at most t errors are assumed to have occurred in each transmission; if $N + 1$ outputs are available to the decoder, the correct input can be deduced.

This can be viewed as a reduction of the associative memory model to the case of $m = 2$, allowing a precise reconstruction of the unique ($m - 1 = 1$) input. When $m > 2$, the decoder seeing $N(m) + 1$ channel outputs can only unambiguously infer which list of $l < m$ codewords contains the correct input; thus, a list-decoding model is suggested.

Chapter 1 studied the reconstruction problem for uniform-tandem-duplication noise, which is applicable to in-vivo DNA data storage. An uncertainty which is sub-linear in the message length was assumed (as it represents the number of distinct reads required for decoding), and it was shown that the redundancy required for unique reconstruction was $(t - 1) \log_q(n) + O(1)$ (compared to the $t \log_q(n) + O(1)$ redundancy required for unique decoding from a single output [KT18a, LJWZ18]), where n is the message length, t the number of errors, and q the alphabet size.

In this chapter, we apply the associative memory model from [YB19] (where binary vectors with the Hamming distance were considered) to the setting of uniform-tandem-duplication noise in finite strings, i.e., we consider list-decoding instead of a unique reconstruction. We shall restrict our attention to codebooks contained in a typical subspace, asymptotically achieving the full space size.

Our goal is to find the trade-off between the code redundancy, the number of tandem-duplication errors, the uncertainty, and the decoded list size. We find the asymptotic behavior, as the message length n grows, of the uncertainty, or required number of reads (more precisely, that number minus one) N , where it is viewed as a function of the list size (plus one) m , the designed minimum distance d , and the number of tandem-duplication errors t . The main result of this chapter (see Corollary 2.24) can informally be summarized in

$$\log_n N + \lceil \log_n(m) \rceil + d = t + \epsilon + o(1),$$

where $\epsilon \in \{0, 1\}$ is a non-increasing function of m , which we find. Thus, such a trade-off is established.

This can be seen as an extension to the results in Chapter 1, where unique reconstruction ($m = 2$) was required, and it was seen that coding with minimum distance $d = t$ enables sub-linear uncertainty (i.e., $\log_n(N) = o(1)$).

In conclusion, we shall see that list-decoding is not only theoretically feasible, but may be efficiently performed. This is done using an isometric transform to integer vectors, and by utilizing combination generators; efficient list-decoding algorithms are developed, given a sufficient number of distinct channel outputs. If the codebook is restricted, then this task is reduced to that of decoding an error-correcting code.

2.2 Additional notation and definitions

In order to simplify our analysis, we assume throughout this chapter that $k \geq 2$.

The focus of this chapter is to find the uncertainty, after t tandem repeats, as a function of the acceptable list size m . This is made precise by the following definition.

Definition 2.1 *Given $n, t \in \mathbb{N}$ and $x_1, \dots, x_m \in \Sigma^n$, we define*

$$S_t(x_1, \dots, x_m) \triangleq \bigcap_{i=1}^m D^t(x_i).$$

Then, the uncertainty associated with a code $C \subseteq \Sigma^n$ is

$$N_t(m, C) \triangleq \max_{\substack{x_1, \dots, x_m \in C \\ x_i \neq x_j}} |S_t(x_1, \dots, x_m)|.$$

Correspondingly, for $w, r \in \mathbb{N}$ and $u_1, \dots, u_m \in \Delta_r^w$ we define

$$\begin{aligned} \bar{S}_t(u_1, \dots, u_m) &\triangleq \bigcap_{i=1}^m \{v \in \mathbb{N}^{w+1} : v \geq u_i, \|v - u_i\|_1 = t\}; \\ \bar{N}_t(m, w, r) &\triangleq \max_{u_1, \dots, u_m \in \Delta_r^w} |\bar{S}_t(u_1, \dots, u_m)|. \end{aligned}$$

In the next section we describe a typical set of strings in Σ^n , then by ascertaining $\bar{N}_t(m, w, r)$ for that set we find an asymptotic expression (in the string length n) for the uncertainty associated with that set, as a function of m .

2.3 Typical set

We observe that the sets introduced in the previous section have many parameters. A complete combinatorial analysis of those would be riddled with pathological

extreme cases, tedious, and not enlightening; this is particularly so since these extreme cases occur in a vanishingly small fraction of the space. Since our main goal is an asymptotic analysis, we proceed by eliminating those rare pathological cases, and focus on the common typical ones. In particular, we would like to limit our attention to strings $x \in \Sigma^n$ for which the Hamming weight of $\bar{\phi}(x)$ and the 1-norm of $\psi_{\text{drt}(x)}(x)$, as well as the difference between them, are asymptotically linearly proportional to the string length n . Those strings would form the code which we study. Thus, we start by presenting in the following lemma the code C for which it shall be our goal to find $N_t(m, C)$.

Lemma 2.2 *Define the family of codes*

$$\text{Typ}^n \triangleq \left\{ x \in \Sigma^n : \begin{array}{l} |w(x) - \frac{q-1}{q}(n-k)| < n^{3/4} \\ |r(x) - \frac{q-1}{q(q^k-1)}(n-k)| < 2n^{3/4} \end{array} \right\},$$

where $w(x) \triangleq \text{wt}(\bar{\phi}(x))$ and $r(x) \triangleq \|\psi_{\text{drt}(x)}(x)\|_1$. Then for sufficiently large n :

$$\frac{|\text{Typ}^n|}{|\Sigma^n|} \geq 1 - 4e^{-\sqrt{n}/2} \xrightarrow{n \rightarrow \infty} 1.$$

Proof We note that if $x, y \in \Sigma^n$ differ only in a single coordinate, then $|w(x) - w(y)|, |r(x) - r(y)| \leq 2$. If the $x(i)$'s are thought of as independent and uniformly distributed random variables on Σ , then by McDiarmid's inequality [Doo40] we have

$$\begin{aligned} \frac{1}{|\Sigma^n|} \left| \left\{ x \in \Sigma^n : |w(x) - \mathbb{E}[w(x)]| \geq n^{3/4} \right\} \right| &\leq 2e^{-\sqrt{n}/2}, \\ \frac{1}{|\Sigma^n|} \left| \left\{ x \in \Sigma^n : |r(x) - \mathbb{E}[r(x)]| \geq n^{3/4} \right\} \right| &\leq 2e^{-\sqrt{n}/2}. \end{aligned}$$

Further note that if $\mathbb{E}[r(x)] = \alpha(n-k) + o(n^{3/4})$ then for large enough n we also have

$$\frac{1}{|\Sigma^n|} \left| \left\{ x \in \Sigma^n : |r(x) - \alpha(n-k)| \geq 2n^{3/4} \right\} \right| \leq 2e^{-\sqrt{n}/2},$$

and hence

$$\frac{1}{|\Sigma^n|} \left| \left\{ x \in \Sigma^n : \begin{array}{l} |w(x) - \mathbb{E}[w(x)]| < n^{3/4} \\ |r(x) - \alpha(n-k)| < 2n^{3/4} \end{array} \right\} \right| \geq 1 - 4e^{-\sqrt{n}/2}.$$

Next, note that $u(i) \triangleq (\bar{\phi}(x))(i)$ are also independent and uniformly distributed. Define the indicator functions $a(i) \triangleq \mathbb{1}_{\{u(i) \neq 0\}}$. Clearly

$$\mathbb{E}[w(x)] = \sum_{i=1}^{n-k} \mathbb{E}[a(i)] = \sum_{i=1}^{n-k} \Pr(u(i) = 1) = \frac{q-1}{q}(n-k).$$

See the Appendix for proof that $\mathbb{E}[r(x)] = \frac{q-1}{q(q^k-1)}(n-k) + O(1)$, which concludes the proof. \square

We remark that a similar concentration result (for $w(x)$ and $\text{wt}(\psi_{\text{drt}(x)}(x))$ instead of $r(x)$) was derived in [KT18a, Lem. 3] using a different approach.

Next, for Typ^n we show that the uncertainty can be calculated by \bar{N}_t , which provides an expression we may more easily analyze.

Lemma 2.3 *For $C \subseteq \Sigma^n$, there exist $x \in \text{Irr}$ and $u_1, \dots, u_m \in \psi_x(C_x)$ such that*

$$N_t(m, C) = |\bar{S}_t(u_1, \dots, u_m)|.$$

Proof Take $x_1, \dots, x_m \in C$ such that $|S_t(x_1, \dots, x_m)| = N_t(m, C)$, and note that if there exist $x_i \not\sim_k x_j$, then $S_t(x_1, \dots, x_m) = \emptyset$, in contradiction. Hence there exists $x = \text{drt}(\{x_1, \dots, x_m\})$. The claim now follows from the isometry ψ_x . \square

Corollary 2.4 *For $k \geq 2$ and sufficiently large n ,*

$$N_t(m, \text{Typ}^n) = \max \left\{ \bar{N}_t(m, w, r) : \begin{array}{l} |w - \frac{q-1}{q}(n-k)| < n^{3/4} \\ |r - \frac{q-1}{q(q^k-1)}(n-k)| < 2n^{3/4} \end{array} \right\}.$$

Proof Lemma 2.3 proves the inequality from left to right. The other direction follows from the observation that for every pair w, r satisfying

$$\left| w - \frac{q-1}{q}(n-k) \right| < n^{3/4} \quad \text{and} \quad \left| r - \frac{q-1}{q(q^k-1)}(n-k) \right| < 2n^{3/4},$$

there exists $x \in \text{Irr}(n - kr)$ (so that $D^r(x) \subseteq \text{Typ}^n$) for which $w(x) = w$. This follows from counting the required number of zeros in $\bar{\phi}(x)$ for such x , which is $n - (1+r)k - w$; for large enough n this number is positive and no greater than $(k-1)(w+1)$, so that any choice of w non-zero elements can be padded with runs of no more than $k-1$ zeros to achieve a total length of $n - (1+r)k$. Thus, a set maximizing \bar{S}_t necessarily has pre-images in Typ^n . \square

Hence, the quantity one needs to assess is $\bar{N}_t(m, w, r)$. We do that next by exploiting the lattice structure of \mathbb{N}^{w+1} , and introducing the connection to supremum height and lower-bound-set size in that lattice.

Lemma 2.5 *Given $u_1, \dots, u_m \in \Delta_r^w$, denote $u \triangleq \bigvee_{i=1}^m u_i$. Then,*

$$|\bar{S}_t(u_1, \dots, u_m)| = \begin{cases} 0 & \|u\|_1 > r + t, \\ \binom{w+t+r-\|u\|_1}{w} & \text{otherwise.} \end{cases}$$

Proof The proposition follows from the lattice structure of \mathbb{N}^{w+1} , i.e.,

$$\bar{S}_t(u_1, \dots, u_m) = \left\{ v \in \mathbb{N}^{w+1} : v \geq \bigvee_{i=1}^m u_i, \|v - u_1\|_1 = t \right\}$$

\square

Definition 2.6 Denote the minimum supremum height

$$\sigma(m, w, r) \triangleq \min_{u_1, \dots, u_m \in \Delta_r^w} \left\| \bigvee_{i=1}^m u_i \right\|_1 - r.$$

Conversely, for $w, r, s \in \mathbb{N}$ and $u \in \Delta_{r+s}^w$, denote the lower-bounds set $A_r(u) \triangleq \{v \in \Delta_r^w : v \leq u\}$ and the maximal lower-bounds-set size

$$\mu(w, r, s) \triangleq \max\{|A_r(u)| : u \in \Delta_{r+s}^w\}.$$

Corollary 2.7 $\bar{N}_t(m, w, r) = \binom{w+t-\sigma(m, w, r)}{w}$.

Proof The proposition follows from Lemma 2.5. \square

It is therefore seen that the main task is to find or estimate the minimum supremum height. We next show the duality between $\sigma(m, w, r)$ and $\mu(w, r, s)$, which we shall use to calculate the former.

Lemma 2.8 Take $w, r, s \in \mathbb{N}$. If $s \geq wr$ then

$$\mu(w, r, s) = |\Delta_r^w| = \binom{r+w}{r} \quad \text{and} \quad \sigma(|\Delta_r^w|, w, r) = wr.$$

For $s < wr$ we have

$$\sigma(\mu(w, r, s), w, r) = s.$$

Proof The first part of the proposition is justified by $(r, r, \dots, r) \in \Delta_{(w+1)r}^w$.

For the second, take $u \in \Delta_{r+s}^w$ satisfying $|A_r(u)| = \mu(w, r, s)$. Since $\bigvee A_r(u) \leq u$ we have

$$\sigma(\mu(w, r, s), w, r) \leq s.$$

However, if in contradiction $\sigma(\mu(w, r, s), w, r) < s$ then we may find $v = \bigvee A_r(v)$ satisfying $|A_r(v)| \geq \mu(w, r, s)$ and $\|v\|_1 < r + s < (w + 1)r$. Therefore, we know that $A_r(v) \neq \Delta_r^w$, hence there exist $v', v'' \in \Delta_r^w$, $v' \notin A_r(v)$ (thus $v' \not\leq v$) and $v'' \in A_r(v)$, satisfying $d_1(v', v'') = 1$. It follows that $\|v \vee v'\|_1 = \|v\|_1 + 1 \leq r + s$, in contradiction to $|A_r(u)| = \mu(w, r, s)$. It follows that $\sigma(\mu(w, r, s), w, r) = s$. \square

Corollary 2.9 If $\mu(w, r, s) < m \leq \mu(w, r, s + 1)$ then

$$\sigma(m, w, r) = s + 1.$$

Proof Firstly, since $m \mapsto \sigma(m, w, r)$ is non-decreasing by definition,

$$s = \sigma(\mu(w, r, s), w, r) \leq \sigma(m, w, r) \leq \sigma(\mu(w, r, s + 1), w, r) = s + 1.$$

However, if $\sigma(m, w, r) = s$, by finding $u_1, \dots, u_m \in \Delta_r^w$ with $\|\bigvee_{i=1}^m u_i\|_1 = r + s$ we deduce $\mu(w, r, s) \geq m$, in contradiction. \square

Since we now know that calculating $\mu(w, r, s)$ is sufficient for our purposes, we turn to that task; since our focus is Typ^n , we may do so for the relevant ranges of w, r , where that is simpler.

Lemma 2.10 *For $w, r, s \in \mathbb{N}$ there exists $u \in \Delta_{r+s}^w$ such that $|A_r(u)| = \mu(w, r, s)$ and for all $1 \leq i < j \leq w + 1$ it holds that $|u(i) - u(j)| < 2$.*

Proof Take $u \in \Delta_{r+s}^w$ satisfying $|A_r(u)| = \mu(w, r, s)$, and assume to the contrary that there exist i, j such that, w.l.o.g., $u(j) \geq u(i) + 2$. Denote by u' the vector which agrees on u on all coordinates except $u'(j) = u(j) - 1$ and $u'(i) = u(i) + 1$.

Further, partition $A_r(u)$ and $A_r(u')$ by the projection on all other coordinates. For any matching classes $C, C' \subseteq \Delta_r^w$ in the corresponding partitions, denote by $t(C) = t(C')$ the difference between r and the sum of all coordinates other than i, j ; Note that $|C|$ is the number of ways to distribute $t(C)$ balls into two bins with capacities $u(i), u(j)$ (and correspondingly $u'(i), u'(j)$ for $|C'|$), hence

$$\begin{aligned} |C| &= \min\{t(C), u(i)\} - \max\{t(C) - u(j), 0\} + 1 \\ &\leq \min\{t(C), u(i) + 1\} - \max\{t(C) - u(j) + 1, 0\} + 1 \\ &= \min\{t(C'), u'(i)\} - \max\{t(C) - u'(j), 0\} + 1 = |C'|, \end{aligned}$$

where the inequality is justified by cases for $t(C)$, and is strict only if $u(i) < t(C) < u(j)$. Thus, the proof is concluded. \square

Lemma 2.10 allows us to find $\mu(w, r, s)$ with relative ease; perhaps the most straightforward example of that is a precise calculation for the cases $s = 1, 2$, which we present next; following the examples we conduct a more extensive evaluation, for $s > 2$ and the relevant ranges of w, r .

Example 2.11 *Any vector $u \in \Delta_{r+1}^w$ having $1 + \min\{w, r\}$ positive coordinates has precisely*

$$|A_r(u)| = 1 + \min\{w, r\}.$$

By Lemma 2.10 one such vector satisfies $\mu(w, r, 1) = |A_r(u)|$, therefore

$$\mu(w, r, 1) = 1 + \min\{w, r\}.$$

■

Example 2.12 *We define an injection*

$$\xi: \{v \in \mathbb{N}^{w+1} : v \leq u\} \rightarrow \mathbb{N}^{w+1}$$

by $\xi(v) \triangleq u - v$; then clearly, ξ is distance preserving, and in particular injective. Hence,

$$\mu(w, r, 2) \leq |\Delta_2^w| = \binom{w+2}{2}.$$

This is achieved with equality when $r + 2 \geq 2(w + 1)$, as evidenced by any vector greater than $(2, 2, \dots, 2)$. The inequality is strict, however, when $r < 2w$.

To examine the remaining cases, note first that increasing any coordinate of u above 2 has no effect on $|A_r(u)|$. Further, we again know by Lemma 2.10 that $\mu(w, r, 2)$ is achieved when u has the greatest number of positive coordinates, and among such vectors, the greatest number greater than or equal to 2. Now, by counting the number of lower bounds for any such $u \in \Delta_{r+2}^w$ we see that

$$\mu(w, r, 2) = \begin{cases} \binom{w+2}{2}, & r \geq 2w; \\ \binom{w+1}{2} + (r - w + 1), & w - 1 \leq r < 2w; \\ \binom{r+2}{2}, & r < w - 1. \end{cases}$$

■

As can now be seen, a complete evaluation of $\mu(w, r, s)$ for $s > 2$ is possible using Lemma 2.10, but it involves application of the inclusion-exclusion principle and its results are not illuminating. We shall see instead that an asymptotic evaluation of $\mu(w, r, s)$ for typical ranges of w, r will suffice. To do so, we note the following proposition.

Lemma 2.13 Fix t , and take w, r such that $r + t \leq w + 1$. For all $s \leq t$ it holds that

$$\mu(w, r, s) = \binom{r + s}{s}.$$

Proof By Lemma 2.10 we know that $u \in \Delta_{r+s}^w$ achieving $|A_r(u)| = \mu(w, r, s)$ is such that $r + s$ of its coordinates equal 1, and the remaining $w + 1 - r - s$ equal 0. The proposition follows. □

We can use what we now know about maximal size of lower-bounds sets to establish the main result of this section, in the following theorem. Before doing so, we note a consequence of, e.g., Lemma 2.13, namely that for any string $x \in \text{Typ}^n$, and any $y \in D^t(x)$, it holds that

$$|\{x' \in \text{Typ}^n : y \in D^t(x')\}| = O(n^t).$$

Hence, we have for $m_n = \omega(n^t)$ that $N_t(m_n, \text{Typ}^n) = o(1)$; it is therefore only interesting to find an asymptotic expression for $N_t(m_n, \text{Typ}^n)$ when $m_n = O(n^t)$.

Theorem 2.14 Fix t and a sequence $m_n = O(n^t)$. Then

$$N_t(m_n, \text{Typ}^n) \sim \frac{1}{(e_t(m_n, n))!} \left(\frac{q-1}{q}n\right)^{e_t(m_n, n)},$$

where $e_t(m_n, n) = t - \lceil \log_n(m_n) \rceil - \delta(m_n, n)$ and $\delta(m, n) \in \{0, 1\}$ is a non-decreasing function in m .

Proof Let $s \triangleq \lceil \log_n(m_n) \rceil$.

Recall from Lemma 2.13 that for $w \geq r + t - 1$

$$\mu(w, r, s - 1) = \binom{r + s - 1}{r} < \frac{(r + s - 1)^{s-1}}{(s - 1)!},$$

hence for r satisfying $\left| r - \frac{q-1}{q(q^k-1)}(n-k) \right| < 2n^{3/4}$ and sufficiently large n

$$\log_n \mu(w, r, s - 1) < s - 1.$$

On the other hand we have

$$\mu(w, r, s + 1) = \binom{r + s + 1}{r} > \frac{r^{s+1}}{(s + 1)!},$$

and therefore, for such r ,

$$\begin{aligned} \log_n \mu(w, r, s + 1) &> \log_n \left(\frac{1 + o(1)}{(s + 1)!} \left(\frac{q-1}{q(q^k-1)} n \right)^{s+1} \right) \\ &= s + 1 + o(1). \end{aligned}$$

Since $s - 1 < \log_n(m_n) \leq s$ it now follows from Corollary 2.9, for sufficiently large n (which does not depend on s , i.e., on m_n), and w, r satisfying

$$\left| w - \frac{q-1}{q}(n-k) \right| < n^{3/4} \quad \text{and} \quad \left| r - \frac{q-1}{q(q^k-1)}(n-k) \right| < 2n^{3/4},$$

that

$$\sigma(m_n, w, r) = s + \delta(m_n, n, r),$$

where

$$\delta(m_n, n, r) = \begin{cases} 1, & m_n > \binom{r + \lceil \log_n(m_n) \rceil}{r}; \\ 0, & \text{otherwise.} \end{cases}$$

Next, for such n, w, r we have

$$\binom{w + t - \sigma(m_n, w, r)}{w} = \frac{1 + o(1)}{(t - (s + \delta(m_n, n, r)))!} \left(\frac{q-1}{q} n \right)^{t - (s + \delta(m_n, n, r))}.$$

It therefore follows from Corollary 2.4 and Corollary 2.7 that

$$\begin{aligned} N_t(m_n, \text{Typ}^n) &= \frac{1 + o(1)}{(t - (s + \delta(m_n, n, r)))!} \left(\frac{q-1}{q} n \right)^{t - (s + \delta(m_n, n, r))} \\ &= \frac{1 + o(1)}{e_t(m_n, n)!} \left(\frac{q-1}{q} n \right)^{e_t(m_n, n)}, \end{aligned}$$

where $\delta(m_n, n) = 1$ if and only if $\delta(m_n, n, r) = 1$ for all r satisfying the above requirement, and $e_t(m_n, n)$ is as defined in the theorem's statement. \square

Finally, we note that the process of list-decoding given $N_t(m, \text{Typ}^n) + 1$ distinct strings in Σ^{n+kt} , i.e., finding $x_1, \dots, x_l \in \text{Typ}^n$, $l < m$, such that these strings lie in $S_t(x_1, \dots, x_l) \setminus \bigcup_{x \in \text{Typ}^n \setminus \{x_1, \dots, x_l\}} D^t(x)$, is straightforward:

Algorithm 2.A Denote $N \triangleq N_t(m, \text{Typ}^n)$ and assume as input distinct $y_1, \dots, y_{N+1} \in \Sigma^{n+kt}$ such that there exists $x \in \text{Typ}^n$ satisfying $y_1, \dots, y_{N+1} \in D^t(x)$.

1. Apply $\psi_{\text{drt}(y_1)}$ to map them to $v_1, \dots, v_{N+1} \in \Delta_{r+t}^w$ where

$$w = \text{wt}(\bar{\phi}(\text{drt}(y_1))) \quad \text{and} \quad r = \|\psi_{\text{drt}(y_1)}(y_1)\|_1;$$

note that prior computation of $\text{drt}(y_1)$ is not required to perform this mapping, and that it may be found as a byproduct of finding any v_i .

2. Find $u \triangleq \bigwedge_{i=1}^{N+1} v_i \in \Delta_{r'}^w$ by calculating the minimum over each coordinate.
3. Calculate $A_r(u)$.
4. Return $\psi_{\text{drt}(y_1)}^{-1}(A_r(u))$ as a list. ■

Theorem 2.15 *Algorithm 2.A operates in $O(n^t) = \text{poly}(N)$ steps, and produces $x_1, \dots, x_l \in \text{Typ}^n$, $l < m$, such that*

$$y_1, \dots, y_{N+1} \in S_t(x_1, \dots, x_l) \setminus \bigcup_{x \in \text{Typ}^n \setminus \{x_1, \dots, x_l\}} D^t(x).$$

Proof First, note that the existence of an ancestor for all y_1, \dots, y_{N+1} implies that $y_i \in D^*(\text{drt}(y_1))$ for all i . Moreover, note that finding any v_i may be done in $O(n)$ steps (by calculating $\bar{\phi}(y_i)$ and recording lengths of runs of zeros in the process). Any one of these can also produce $\text{drt}(y_1)$. Hence Step 1 concludes in $O(Nn)$ steps.

Step 2 can also be performed in $O(Nw) = O(Nn)$ steps.

Now, note that since an ancestor of all y_i 's exists in Σ^n , $r' \geq r$. It is hence possible to compute $A_r(u)$. This may be achieved by finding all ways of distributing $r' - r < t$ balls into $w + 1$ bins with capacities $u(j)$, e.g., by utilizing combination generators for all $\binom{w+r'-r}{w}$ combinations, then discarding combination which violate the bin-capacity restriction. Combination generating algorithms exist which generate all combinations in $O\left(\binom{w+r'-r}{w}\right) = O(n^{t-1})$ steps (e.g., see [RSW12]), and pruning illegal combinations can be done in $O(w)$ steps each. Step 3 can therefore be performed in $O(n^t)$ steps.

Finally, the pre-image $\psi_{\text{drt}(y_1)}^{-1}(A_r(u))$ is a set of ancestors of y_1, \dots, y_{N+1} , which is a subset Typ^n , and no other element of Typ^n is an ancestor of y_1, \dots, y_{N+1} . We also know that $|A_r(u)| < m$, otherwise a contradiction is reached to the definition of N . Computing $\psi_{\text{drt}(y_1)}^{-1}(A_r(u))$ given $\text{drt}(y_1)$ requires $O(|A_r(u)|w) \leq O(mn)$ steps. □

2.4 Uncertainty with underlying ECC

In the previous section, a reconstruction problem with a list-decoding algorithm was considered, when the underlying message space was unconstrained (more precisely, constrained only to a typical set). However, one is naturally interested in a more general setting, in which the message space may be a code with a given minimum distance. Thus, in this section, we shall consider the uncertainty associated with codes $C \subseteq \text{Typ}^n$ such that for all distinct $c, c' \in C$, $d(c, c') \geq d$, for some $d > 0$. We start with a definition of a typical set with a minimum distance.

Definition 2.16 *Given $m, n, t, d \in \mathbb{N}$, the uncertainty associated with the minimum distance d (in the typical sense) is defined as*

$$N_t^{\text{Typ}}(m, n, d) \triangleq \max_{\substack{x_1, \dots, x_m \in \text{Typ}^n \\ d(x_i, x_j) \geq d}} |S_t(x_1, \dots, x_m)|.$$

We again define correspondingly, for $w, r \in \mathbb{N}$,

$$\begin{aligned} \bar{N}_t(m, w, r, d) &\triangleq \max_{\substack{u_1, \dots, u_m \in \Delta_r^w \\ d_1(u_i, u_j) \geq d}} |\bar{S}_t(u_1, \dots, u_m)|, \\ \mu(w, r, s, d) &\triangleq \max_{u \in \Delta_{r+s}^w} \max \left\{ |C| : \begin{array}{l} C \subseteq A_r(u) \\ \forall v \neq v' \in C: d_1(v, v') \geq d \end{array} \right\} \\ \sigma(m, w, r, d) &\triangleq \min_{\substack{u_1, \dots, u_m \in \Delta_r^w \\ d_1(u_i, u_j) \geq d}} \left\| \bigvee_{i=1}^m u_i \right\|_1 - r. \end{aligned}$$

It should be noted that if $d > t$ then $N^{\text{Typ}}(2, n, d) = 0$, meaning that unique decoding from a single noisy output is possible. As was seen in Chapter 1, $d = t$ suffices for unique reconstruction ($m = 2$) with sub-linear uncertainty (in fact, $N = 1$, which corresponds to receiving two distinct noisy outputs, suffices). We shall incidentally see that again while considering $d \leq t$.

Corollary 2.17 *For all sufficiently large n ,*

$$N_t^{\text{Typ}}(m, n, d) = \max \left\{ \bar{N}_t(m, w, r, d) : \begin{array}{l} |w - \frac{q-1}{q}(n-k)| < n^{3/4} \\ |r - \frac{q-1}{q(q^k-1)}(n-k)| < 2n^{3/4} \end{array} \right\}.$$

Proof Similarly to the proof of Lemma 2.3, a choice of $x_1, \dots, x_m \in \text{Typ}^n$ satisfying $d(x_i, x_j) \geq d$ and $|S_t(x_1, \dots, x_m)| = N_t^{\text{Typ}}(m, n, d)$ must also satisfy $x_i \sim_k x_j$ (otherwise $S_t(x_1, \dots, x_m) = \emptyset$), hence we may find $x \triangleq \text{drt}(x_1) = \dots = \text{drt}(x_m)$. In addition

$$|\bar{S}_t(\psi_x(x_1), \dots, \psi_x(x_m))| = |S_t(x_1, \dots, x_m)|$$

and $d_1(\psi_x(x_i), \psi_x(x_j)) \geq d$. The other direction follows as in the proof of Corollary 2.4. \square

We shall continue using an analogous approach to that of the previous section, in finding $\bar{N}_t(m, w, r, d)$ in order to estimate $N_t^{\text{Typ}}(m, n, d)$.

Corollary 2.18 $\bar{N}_t(m, w, r, d) = \binom{w+t-\sigma(m, w, r, d)}{w}$.

Proof This proposition follows from Lemma 2.5 as well. \square

Lemma 2.19 Take some $m, w, r, s, d \in \mathbb{N}$. If

$$\mu(w, r, s, d) < m \leq \mu(w, r, s + 1, d)$$

then

$$\sigma(m, w, r, d) = s + 1.$$

Proof The proof follows the same arguments as in the proofs of Lemma 2.8 and Corollary 2.9. \square

Lemma 2.20 For $w, r, s, d \in \mathbb{N}$ there exist $u \in \Delta_{r+s}^w$, and $C \subseteq A_r(u)$ with minimum d_1 distance d , satisfying $|C| = \mu(w, r, s, d)$, such that for no pair $1 \leq i, j \leq w + 1$, $i \neq j$, it holds that $u(i) \geq 2$ and $u(j) = 0$.

Proof Take $u \in \Delta_{r+s}^w$ and $C \subseteq A_r(u)$ satisfying $|C| = \mu(w, r, s, d)$, and assume to the contrary that there exist such i, j ; denote by u' the vector which agrees on u on all coordinates except $u'(j) = 1$ and $u'(i) = u(i) - 1$. The proposition is justified by finding any isometric injection $\rho: A_r(u) \rightarrow A_r(u')$.

Indeed, define $\rho(v) \triangleq v$ if $v(i) < u(i)$, otherwise

$$(\rho(v))(l) \triangleq \begin{cases} u(i) - 1, & l = i; \\ 1, & l = j; \\ v(l), & \text{otherwise.} \end{cases}$$

Then ρ is well defined. Moreover, take any $v_1, v_2 \in A_r(u)$. If $v_1(i), v_2(i) < u(i)$ then clearly $d_1(\rho(v_1), \rho(v_2)) = d_1(v_1, v_2)$. The same trivially holds when $v_1(i) = v_2(i) = u(i)$. If, w.l.o.g. $v_1(i) < v_2(i) = u(i)$, then

$$\begin{aligned} |(\rho(v_1))(i) - (\rho(v_2))(i)| &= |v_1(i) - (\rho(v_2))(i)| \\ &= |v_1(i) - v_2(i)| - 1 \end{aligned}$$

but

$$\begin{aligned} |(\rho(v_1))(j) - (\rho(v_2))(j)| &= |v_1(j) - (\rho(v_2))(j)| = |0 - 1| \\ &= 1 = |v_1(j) - v_2(j)| + 1, \end{aligned}$$

hence, once again, $d_1(\rho(v_1), \rho(v_2)) = d_1(v_1, v_2)$. \square

As in Section 2.3, Lemma 2.20 allows us to find $\mu(w_n, r_n, s)$ for typical ranges of w_n, r_n , using binary constant-weight codes. This is given precise meaning in the following definition and lemma.

Definition 2.21 Denote the field with two elements $\text{GF}(2)$, and the Hamming metric d_H . Denote by $A(\nu, 2\delta, \omega)$ the size of the largest length ν binary code with minimum Hamming distance 2δ and constant Hamming weight ω .

Lemma 2.22 Fix t , and take w, r such that $r + t \leq w + 1$. For all $s \leq t$ it holds that

$$\mu(w, r, s, d) = A(r + s, 2d, s).$$

Proof By Lemma 2.20 we know that there exist $u \in \Delta_{r+s}^w$ and $C \subseteq A_r(u)$ satisfying

- $|C| = \mu(w, r, s, d)$.
- For all $v_1, v_2 \in C$, $v_1 \neq v_2$, it holds that $d_1(v_1, v_2) \geq d$.
- u has $r + s$ of its coordinates equal 1, and the remaining $w + 1 - r - s$ equal 0.

Define $\rho: A_r(u) \rightarrow \text{GF}(2)^{r+s}$ by restricting $u - v$ to the support of u (and identifying $\text{GF}(2)$ with $\{0, 1\} \subseteq \mathbb{N}$). Then ρ is a bijection onto constant-Hamming-weight s elements of $\text{GF}(2)^{r+s}$. Further, for all $v_1, v_2 \in A_r(u)$ it holds that

$$d_H(\rho(v_1), \rho(v_2)) = 2d_1(v_1, v_2).$$

Hence, there's a size-preserving one-to-one correspondence between codes $C' \subseteq A_r(u)$ with minimum d_1 distance d , and codes in $\text{GF}(2)^{r+s}$ with minimum Hamming distance $2d$ and constant Hamming weight s . The proposition follows. \square

We can now summarize our observations in the following theorem.

Theorem 2.23 Fix $d \leq t$ and a sequence $m_n = O(n^{t-d+1})$. Then

$$N_t^{\text{Typ}}(m_n, n, d) \sim \frac{1}{(e_t(m_n, n, d))!} \left(\frac{q-1}{q} n \right)^{e_t(m_n, n, d)},$$

where $e_t(m_n, n, d) = t - \lceil \log_n(m_n) \rceil - d + \epsilon(m_n, n, d)$ and $\epsilon(m, n, d) \in \{0, 1\}$ is a non-increasing function of m .

Proof The proof follows the same lines as that of Theorem 2.14. Let $s \triangleq \lceil \log_n(m_n) \rceil + d - 1$.

Recall from the first Johnson bound [Joh62, Th. 2] that

$$\begin{aligned} A(r + s - 1, 2d, s - 1) &\leq \binom{r + s - 1}{s - d} \Big/ \binom{s - 1}{s - d} \\ &< \frac{(d - 1)!}{(s - 1)!} (r + s - 1)^{s - d}, \end{aligned}$$

hence for r satisfying $\left| r - \frac{q-1}{q(q^k-1)}(n - k) \right| < 2n^{3/4}$ and sufficiently large n

$$\log_n A(r + s - 1, 2d, s - 1) < s - d.$$

On the other hand, by [GS80, Th. 6] we have

$$A(r + s + 1, 2d, s + 1) \geq \frac{1}{p^{d-1}} \binom{r + s + 1}{s + 1}$$

for any prime power p , $p > r + s$. By the prime number theorem (a weaker version, or even Bertrand's postulate, suffices. See, e.g., [Che52]) there exists in fact such prime number p satisfying $r + s < p \leq n$ for sufficiently large n and r satisfying $\left| r - \frac{q-1}{q(q^k-1)}(n - k) \right| < 2n^{3/4}$, hence in particular

$$\begin{aligned} A(r + s + 1, 2d, s + 1) &\geq \frac{1}{n^{d-1}} \binom{r + s + 1}{s + 1} \\ &> \frac{r^{s+1}}{n^{d-1}(s + 1)!}, \end{aligned}$$

and therefore

$$\begin{aligned} \log_n A(r + s + 1, 2d, s + 1) &> \log_n \left(\frac{1 + o(1)}{n^{d-1}(s + 1)!} \left(\frac{q - 1}{q(q^k - 1)} n \right)^{s+1} \right) \\ &= s - d + 2 + o(1). \end{aligned}$$

Since $s - d < \log_n(m_n) \leq s - d + 1$ it now follows from Lemma 2.19 and Lemma 2.22, for sufficiently large n (which does not depend on s , i.e., on m_n), and w, r satisfying

$$\left| w - \frac{q-1}{q}(n - k) \right| < n^{3/4} \quad \text{and} \quad \left| r - \frac{q-1}{q(q^k-1)}(n - k) \right| < 2n^{3/4},$$

that

$$\sigma(m_n, w, r, d) = s + \delta(m_n, n, r, d),$$

where

$$\delta(m_n, n, r, d) = \begin{cases} 1, & m_n > A(r + s, 2d, s); \\ 0, & \text{otherwise.} \end{cases}$$

(Note that that s is a function of m_n, n .)

Next, for such n, w, r we have

$$\binom{w + t - \sigma(m_n, w, r, d)}{w} = \frac{1+o(1)}{(t-(s+\delta(m_n, n, r, d)))!} \left(\frac{q-1}{q}n\right)^{t-(s+\delta(m_n, n, r, d))}.$$

It therefore follows from Corollary 2.17 and Corollary 2.18 that

$$\begin{aligned} N_t^{\text{Typ}}(m_n, n, d) &= \frac{1+o(1)}{(t-(s+\delta(m_n, n, d)))!} \left(\frac{q-1}{q}n\right)^{t-(s+\delta(m_n, n, d))} \\ &= \frac{1+o(1)}{e_t(m_n, n, d)!} \left(\frac{q-1}{q}n\right)^{e_t(m_n, n, d)}, \end{aligned}$$

where $\delta(m_n, n, d) = 1$ if and only if $\delta(m_n, n, r, d) = 1$ for all r satisfying the above requirement, $\epsilon(m_n, n, d) \triangleq 1 - \delta(m_n, n, d)$, and $e_t(m_n, n, d)$ is as defined in the theorem's statement. \square

It is again remarked here that in the case that coding is performed with $d = t$, we observe that unique reconstruction ($m = 2$) is possible with just two reads ($N = 1$); To see that, note that $\delta(2, r, d) = 0$ for all $r \geq d$, hence for sufficiently large n we have $\epsilon(2, d) = 1$ and therefore $e_t(2, n, d) = 0$. This result, as mentioned above, was already observed in Chapter 1.

The trade-off established in Theorem 2.23 between the code minimum distance d (equivalently, its redundancy, since as seen in [KT18a, LJWZ18] and mentioned above, a code with minimum distance d has optimal redundancy $(d - 1) \log_q(n) + O(1)$), the number of tandem-duplication errors t , the decoded list size m_n , and the resulting uncertainty $N_t^{\text{Typ}}(m_n, n, d)$, is perhaps better visualized in the following corollary.

Corollary 2.24 *Fix $d \leq t$ and a sequence $m_n = O(n^{t-d+1})$. Then*

$$\log_n N_t^{\text{Typ}}(m_n, n, d) + \lceil \log_n(m_n) \rceil + d = t + \epsilon(m_n, d) + o(1),$$

where $\epsilon(m, d) \in \{0, 1\}$ is a non-increasing function of m .

Before concluding, we present a list-decoding scheme given sufficiently many $(N_t^{\text{Typ}}(m, n, d) + 1)$ distinct strings in

$$D^t(C) \triangleq \bigcup_{c \in C} D^t(c),$$

for some given code $C \subseteq \text{Typ}^n$ with minimum distance d . We shall assume that a decoding scheme for recovering from at most $d - 1$ errors is known for C , which we denote $\mathcal{D}: \Sigma^{n+k(d-1)} \rightarrow C$.

Algorithm 2.B Fix n, m and $d \leq t$; take $C \subseteq \text{Typ}^n$ with minimum $d(\cdot, \cdot)$ distance d , and assume a decoding scheme for recovering up to $d - 1$ tandem-duplication errors is provided. Denote $N \triangleq N_t^{\text{Typ}}(m, n, d)$ and assume as input distinct $y_1, \dots, y_{N+1} \in \Sigma^{n+kt}$ such that there exists $x \in C$ satisfying $y_1, \dots, y_{N+1} \in D^t(x)$.

1. Apply Algorithm 2.A to obtain $z_1, \dots, z_l \in \Sigma^{n+k(d-1)}$ such that

$$y_1, \dots, y_{N+1} \in S_{t-d+1}(z_1, \dots, z_l) \setminus \bigcup_{\substack{z \in \text{Typ}^{n+k(d-1)} \\ z \notin \{z_1, \dots, z_l\}}} D^{t-d+1}(z).$$

2. Decode each z_i with the provided algorithm to produce $x_i \triangleq \mathcal{D}(z_i) \in C$; if $z_i \notin D^{d-1}(x_i)$, discard x_i .
3. Return every x_i that was not discarded in the last step, as a list. ■

Before proving correctness for Algorithm 2.B, we would like to estimate l from Step 1. To that end, note that the number of t -ancestors of $y \in \Sigma^{n+kt}$ is bound from above by $\mu(w, r, t)$ where $w = \text{wt}(\bar{\phi}(y)) \leq n - k$ and $r = \|\psi_{\text{drt}(y)}(y)\|_1 - t$. As in Example 2.12, using ξ we note that

$$\begin{aligned} \mu(w, r, t) &\leq |\Delta_t^w| = \binom{w+t}{w} < \frac{1}{t!}(w+t)^t \\ &\leq \frac{1}{t!}(n-k+t)^t < (n+t)^t. \end{aligned}$$

Hence $N_t((n+t)^t, \text{Typ}^n) = 0$; this in particular implies that for $\hat{m} \triangleq (n+t + (k-1)(d-1))^{t-d+1}$ we have

$$N_{t-d+1}(\hat{m}, \text{Typ}^{n+k(d-1)}) = 0 \leq N_t^{\text{Typ}}(m, n, d).$$

Note, then, that $l < \hat{m}$. This result can be considerably improved by noting that for all m' satisfying

$$N_{t-d+1}(m', \text{Typ}^{n+k(d-1)}) \leq N_t^{\text{Typ}}(m, n, d)$$

it holds that $l < m'$, but for our purposes \hat{m} will suffice.

Theorem 2.25 *Algorithm 2.B produces $x_1, \dots, x_{l'} \in C$, $l' < m$, such that*

$$y_1, \dots, y_{N+1} \in S_t(x_1, \dots, x_{l'}) \setminus \bigcup_{\substack{x \in \text{Typ}^n \\ x \notin \{x_1, \dots, x_{l'}\}}} D^t(x).$$

Further, it operates in $O(n^t + n^{t-d+1}\mathcal{C})$ steps, where \mathcal{C} is the run-time complexity of \mathcal{D} .

Proof There is one assumption to Algorithm 2.A and Theorem 2.15 which may now not be satisfied, that indeed there exists $z \in \Sigma^{n+k(d-1)}$ such that $y_1, \dots, y_{N+1} \in D^t(z)$. If there does not, then Step 1 might fail because Algorithm 2.A finds $\hat{z} \triangleq \bigwedge_{i=1}^N y_i$ with $|\hat{z}| = n + ks$ and $s < (d - 1)$. If that is the case, however, such \hat{z} may still be passed on to the next step, since we may still decode it to a unique $x \in C$ for which $z \in D^s(x)$ (since C has minimum distance d , there cannot exist two distinct ancestors of z in C), which justifies the claim. Otherwise, Theorem 2.15 proves that the first step produces what is claimed, and we may assume w.l.o.g. that $s \geq d - 1$.

This assumption now implies that for each $x \in C$ such that $y_1, \dots, y_{N+1} \in D^t(x)$ there exists $z \in D^{d-1}(x)$ such that $y_1, \dots, y_{N+1} \in D^t(z)$, hence $z \in \{z_1, \dots, z_l\}$; this is because one may arbitrarily choose such $x \leq z \leq \hat{z}$. On the other hand, each $z \in \Sigma^{n+k(d-1)}$ can be decoded to at most a single $x \in C$ for which $z \in D^{d-1}(x)$ (again, due to the code's minimum distance), and that x satisfies $y_1, \dots, y_{N+1} \in D^t(x)$. We remark that it is possible that the first step produces $z_i \notin D^{d-1}(C)$, hence $x_i = \mathcal{D}(z_i)$ may be erroneous (as the decoder receives invalid input); however, as $y_1, \dots, y_{N+1} \in D^{t-d+1}(z_i)$, such results can indeed be discarded by testing if $z_i \in D^{d-1}(x_i)$.

Note that if distinct $x_1, \dots, x_m \in C$ are produced by Step 2, we have $|S_t(x_1, \dots, x_m)| \geq |\{y_1, \dots, y_{N+1}\}| = N + 1$ and therefore a contradiction. Hence, $l' < m$.

Finally, we know that Step 1 operates in $O(n^t) = \text{poly}(N)$ steps. Step 2 clearly operates in $O(n^{t-d+1}C)$ steps, which concludes the proof. \square

Appendix: Conclusion of proof of Lemma 2.2

As in the proof of Lemma 2.2, we define $u(i) \triangleq (\bar{\phi}(x))(i)$. Further define for all $1 \leq i \leq n - k$ and $1 \leq j < n - k - i + 1$ the indicator $I_i(j)$ of the event of a run of precisely j zeros starting in u at index i . Then

$$\begin{aligned} \mathbb{E}[r(x)] &= \sum_{i=1}^{n-k} \sum_{j=1}^{n-k-i+1} \left\lfloor \frac{j}{k} \right\rfloor \Pr(I_i(j) = 1) \\ &= \left\lfloor \frac{n-k}{k} \right\rfloor \Pr(I_1(n-k) = 1) + \sum_{j=1}^{n-k-1} \left\lfloor \frac{j}{k} \right\rfloor \Pr(I_1(j) = 1) \\ &\quad + \sum_{i=2}^{n-k} \left\lfloor \frac{n-k-i+1}{k} \right\rfloor \Pr(I_i(n-k-i+1) = 1) \end{aligned}$$

$$\begin{aligned}
& + \sum_{i=2}^{n-k-1} \sum_{j=1}^{n-k-i} \lfloor \frac{j}{k} \rfloor \Pr(I_i(j) = 1) \\
& = \left\lfloor \frac{n-k}{k} \right\rfloor \frac{1}{q^{n-k}} + \sum_{j=1}^{n-k-1} \left\lfloor \frac{j}{k} \right\rfloor \frac{q-1}{q^{j+1}} \\
& \quad + \sum_{i=2}^{n-k} \left\lfloor \frac{n-k-i+1}{k} \right\rfloor \frac{q-1}{q^{n-k-i+2}} \\
& \quad + \sum_{i=2}^{n-k-1} \sum_{j=1}^{n-k-i} \left\lfloor \frac{j}{k} \right\rfloor \frac{(q-1)^2}{q^{j+2}} \\
& = \frac{\lfloor n/k \rfloor - 1}{q^{n-k}} + 2 \frac{q-1}{q} \sum_{j=1}^{n-k-1} \frac{\lfloor j/k \rfloor}{q^j} \\
& \quad + \frac{(q-1)^2}{q^2} \sum_{i=2}^{n-k-1} \sum_{j=1}^{n-k-i} \frac{\lfloor j/k \rfloor}{q^j}
\end{aligned}$$

We note that

$$\begin{aligned}
\sum_{j=1}^p \frac{\lfloor j/k \rfloor}{q^j} & = \sum_{j=k}^p \frac{\lfloor j/k \rfloor}{q^j} \\
& = \sum_{j=k \lfloor p/k \rfloor}^p \frac{\lfloor p/k \rfloor}{q^j} + \sum_{i=1}^{\lfloor p/k \rfloor - 1} \sum_{j=0}^{k-1} \frac{i}{q^{ik+j}} \\
& = \frac{q}{q-1} \left[\lfloor p/k \rfloor \left(\frac{1}{q^{k \lfloor p/k \rfloor}} - \frac{1}{q^{p+1}} \right) \right. \\
& \quad \left. + \left(1 - \frac{1}{q^k} \right) \sum_{i=1}^{\lfloor p/k \rfloor - 1} \frac{i}{q^{ik}} \right] \\
& = \frac{q}{q-1} \left[\lfloor p/k \rfloor \left(\frac{1}{q^{k \lfloor p/k \rfloor}} - \frac{1}{q^{p+1}} \right) \right. \\
& \quad \left. + \frac{1}{q^k - 1} \left(1 - \frac{1}{q^{k(\lfloor p/k \rfloor - 1)}} \right) \right. \\
& \quad \left. - \frac{\lfloor p/k \rfloor - 1}{q^{k \lfloor p/k \rfloor}} \right] \\
& = \frac{q}{q-1} \left[\frac{1}{q^k - 1} \left(1 - \frac{1}{q^{k(\lfloor p/k \rfloor - 1)}} \right) \right. \\
& \quad \left. + \frac{1}{q^{k \lfloor p/k \rfloor}} - \frac{\lfloor p/k \rfloor}{q^{p+1}} \right]
\end{aligned}$$

Now

$$\frac{\lfloor n/k \rfloor - 1}{q^{n-k}} + 2\frac{q-1}{q} \sum_{j=1}^{n-k-1} \frac{\lfloor j/k \rfloor}{q^j} = O(1).$$

Hence, it suffices to find

$$\begin{aligned} \frac{(q-1)^2}{q^2} \sum_{i=2}^{n-k-1} \sum_{j=1}^{n-k-i} \frac{\lfloor j/k \rfloor}{q^j} &= \frac{(q-1)^2}{q^2} \sum_{p=1}^{n-k-2} \sum_{j=1}^p \frac{\lfloor j/k \rfloor}{q^j} \\ &= \frac{q-1}{q(q^k-1)} \sum_{p=1}^{n-k-2} \left(1 - \frac{1}{q^{k(\lfloor p/k \rfloor - 1)}} \right) \\ &\quad - \frac{q-1}{q} \sum_{p=1}^{n-k-2} \left[\frac{\lfloor p/k \rfloor}{q^{p+1}} - \frac{1}{q^{k\lfloor p/k \rfloor}} \right]. \end{aligned}$$

Again, note that $\sum_{p=1}^{n-k-2} \frac{\lfloor p/k \rfloor}{q^{p+1}} = O(1)$; in addition, we note that $\sum_{p=1}^{n-k-2} \frac{1}{q^{k\lfloor p/k \rfloor}} = O(1)$.

We therefore find $\mathbb{E}[r(x)] = \frac{q-1}{q(q^k-1)}(n-k) + O(1)$.

Chapter 3

Combined Substitutions Noise

In this chapter, we shall focus on error-control when combining substitution errors with duplication noise. In a substitution event, a symbol in the sequence is changed to another symbol of the alphabet. It has been observed that point mutations such as substitutions are more common in tandem repeat regions of the genomes [POB08].

We shall consider the *unrestricted-substitution model* for combined duplication and substitution errors, in which substitutions are allowed at any position in the affected string. In comparison, one may also consider a *noisy-duplication model*, in which the copy is a noisy version of the template. I.e., noisy duplications in this model can be viewed as exact duplications followed by substitutions that are restricted to the newly added copy. See [TYSF19, TF19] for a study of both these models.

We shall construct, in this chapter, error-correcting codes that are capable of correctly handling any number of tandem duplications of a fixed length k , and at most a single substitution error. The main approach is to reverse the duplication process while accounting for the single substitution (which may spuriously create the appearance of a duplication that never happened, or eliminate one that did).

3.1 Additional notation and definitions

In this section, we observe the case of many tandem-duplications and a single substitution, occurring at any point during the duplication sequence, and in any place in the affected string (i.e., not necessarily in a duplicated substring). To illustrate the effects of such a substitution, we continue the example given in the Introduction with a substitution event:

$$\begin{aligned}x' &= 1012\underline{0}12121 \rightarrow x'' = 1012\underline{1}12121, \\ \phi(x') &= 101, 1\underline{0}00\underline{1}12 \rightarrow \phi(x'') = 101, 1\underline{1}00\underline{0}12.\end{aligned}$$

Formally, a substitution may be considered as the mapping $x \rightarrow x + ae_i$, where $e_i \in \Sigma^n$ is a standard unit vector at index i , and $a \in \Sigma$, $a \neq 0$. Since ϕ is linear over Σ (i.e., $\phi(x + ae_i) = \phi(x) + a\phi(e_i)$), we denote the transform of e_i as $\epsilon_i \triangleq \phi(e_i)$, and observe that $\epsilon_i = e_i - e_{i+k}$ for $i \leq n - k$ and $\epsilon_i = e_i$ for $n - k < i \leq n$. We note that substitutions might affect two positions in the ϕ -transform domain.

We define $D^{t,p}(x)$ to be the set of strings obtained from x through t tandem duplications and p substitutions, where substitutions can occur in any position. We further define

$$D^{*,p}(x) \triangleq \bigcup_{t=0}^{\infty} D^{t,p}(x), \quad D^{*(P)}(x) \triangleq \bigcup_{p \in P} D^{*,p}(x),$$

where P is a set of non-negative integers. We shall denote $P = \{0, 1\}$ as ≤ 1 .

We note that a code $C \subseteq \Sigma^n$ can correct any number of tandem-duplication errors if and only if no two distinct codewords $c_1, c_2 \in C$ have a common descendant, namely,

$$D^{*,0}(c_1) \cap D^{*,0}(c_2) = \emptyset.$$

As mentioned before, it was proved in [JFSB17a] that this condition is equivalent to all codewords having distinct roots; It was suggested in [JFSB17a] that error-correcting codes that protect against any number of duplications may be obtained simply by taking irreducible strings as codewords. Up to a minor tweaking, this strategy was shown in [JFSB17a] to produce optimal codes. Our approach herein shall be similar, while accounting for the effects of the substitution.

We shall refer to codes able to correct any number of tandem-duplication error, and a single substitution error occurring at any point in the duplication sequence, as a *single-substitution correcting* (1S-correcting) code. Obviously, a code C is 1S-correcting if and only if for any two distinct codewords $c_1, c_2 \in C$, we have

$$D^{*,\leq 1}(c_1) \cap D^{*,\leq 1}(c_2) = \emptyset.$$

In this context, we will find it easier to consider strings in the ϕ -transform domain. We also define the *substitution distance* $\sigma(u, v)$ to measure the number of substitutions required to transform one string into the other, when u, v are assumed to be in the transform domain. More precisely, if $u, v \in \Sigma^n$ and $v - u = \sum_{i=1}^n a_i \cdot \epsilon_i$, then

$$\sigma(u, v) \triangleq |\{i \in [n] : a_i \neq 0\}|.$$

Put differently, $\sigma(u, v) = d_H(\phi^{-1}(u), \phi^{-1}(v))$, where again $d_H(\cdot, \cdot)$ is the Hamming metric.

3.2 Error-correction via constrained coding

When considering the combination of duplications with even a single substitution in the transform domain, we come across the following example:

Example 3.1 *Set $\Sigma = \mathbb{Z}_2$ and $k = 3$, and observe the following two sequences of duplication and substitution, as seen in the ϕ -transform domain:*

$$\begin{aligned} u &\triangleq 111010111 \rightarrow 111010111\underline{000} \rightarrow 1110\underline{00}1\underline{0}1000 \\ v &\triangleq 111101010 \rightarrow 111\underline{000}101010 \rightarrow 1110001010\underline{00} \end{aligned}$$

It is clear that if $C \subseteq \Sigma^{\geq k}$ is a code correcting even a single duplication and a single substitution, even given the order in which they occur, then $\phi^{-1}(u) = 111101010$ and $\phi^{-1}(v) = 111010000$ cannot both belong to C . Observing that u, v are $(0, k-1)_q$ -RLL, and $\sigma(u, v) = 4$, however, we find that $C \triangleq \{\phi^{-1}(u), \phi^{-1}(v)\}$ can correct any number of duplications, or correct a single substitution. Why it cannot do both at once, then, is not immediately apparent. ■

In what follows, we propose a constrained-coding approach which resolves the issue demonstrated in the last example. It relies on the following observation: substitution noise might create a 0^k substring in the transform domain—that is not due to a duplication—as well as break a run of zeros. However, a constrained system exists which allows us to de-couple the effects of duplication and substitution noise.

More precisely, we denote

$$\mathcal{W} \triangleq \{u \in \Sigma^{\geq k} : \forall \text{ substring } v \text{ of } u, |v| = k : \text{wt}(v) > 1\}.$$

We aim to show that restricting codewords to be taken from \mathcal{W} (in the transform domain), the following holds.

Lemma 3.2 *Take an irreducible $x \in \Sigma^{\geq k}$, and $y \in D^{*, \leq 1}(x)$. If $v \triangleq \bar{\phi}(y)$ contains a 0^k substring, and \bar{v} is derived from v by removing that substring, and if $\bar{\phi}(x) \in \mathcal{W}$, then $\bar{v} \in \bar{\phi}(D^{*, \leq 1}(x))$.*

Proof We denote

$$v = \alpha c 0^k \beta$$

for $0 \neq c \in \Sigma$ and $\alpha, \beta \in \Sigma^*$, and by abuse of notation assume $|\alpha c| \geq 0$ is the shortest with the properties stated above (allowing $v = 0^k \beta$ as a private case).

We also take $y' \in D^{*, 0}(x)$ to be the descendant of x derived by the same sequence of duplications as y , where a substitution never occurs, and

$$v' = \bar{\phi}(y') = \alpha' c' 0^j a 0^{k-j-1} \beta',$$

for $0 \leq j < k$, $c', a \in \Sigma$, $\alpha', \beta' \in \Sigma^*$, where $|\alpha'c'| = |\alpha c|$. (We know v' can be represented in this fashion since y suffered a single substitution.)

If $a = 0$ then the claim is trivial. Assume, therefore, $a \neq 0$. Note that $\bar{\phi}(x) \in \mathcal{W}$ and $\text{wt}(0^j a 0^{k-j-1}) = 1$, implying that $0^{k-j-1}\beta'$ begins with a k -tuple of zeros. I.e., $\beta' = 0^{j+1}\beta''$, for some $\beta'' \in \Sigma^*$. Thus, a descendant of x is also z' , where $\bar{\phi}(z') = \alpha'c'0^j a \beta''$.

We now reexamine v', v :

$$\begin{aligned} v' &= \alpha' c' 0^j a 0^{k-j-1} \beta' \\ v &= \alpha c 0^j 0 0^{k-j-1} \beta \end{aligned}$$

and since y is derived from x by the same sequence of tandem-duplications as y' , with a single substitution, we may deduce that α, β and α', β' differ, respectively, in precisely one of the following manners:

- There exist $b \in \Sigma$ and $\alpha_1, \alpha_2 \in \Sigma^*$, with $|\alpha_2 c| = k - j - 1$, such that

$$\begin{aligned} v' &= \alpha_1 (b - a) \alpha_2 c' 0^j a 0^{k-j-1} \beta \\ v &= \alpha_1 b \alpha_2 c 0^j 0 0^{k-j-1} \beta \end{aligned}$$

and, again, by abuse of notation, including the case of $|\alpha_2 c| = 0$, meaning $b = c$ and $b - a = c'$; in all other cases $c' = c$.

In this case

$$\begin{aligned} \bar{v} &= \alpha c \beta = \alpha_1 b \alpha_2 c 0^{j+1} \beta'' \\ &= \alpha_1 (b - a) \alpha_2 c' 0^j a \beta'' + a \cdot \epsilon_{|xc|+j-k} \\ &= \bar{\phi}(z) + a \cdot \epsilon_{|xc|+j-k}. \end{aligned}$$

- $\beta = 0^j a \beta''$, implying $\alpha'c' = \alpha c$ and

$$\bar{v} = \alpha c \beta = \alpha c 0^j a \beta'' = \bar{\phi}(z).$$

- There exist $s \geq 0$, $b \in \Sigma$, $\gamma \in \Sigma^{k-1}$ and $\beta''' \in \Sigma^*$ such that $\beta'' = 0^{sk} \gamma b \beta'''$, and

$$\begin{aligned} v' &= \alpha c 0^j a 0^{k-j-1} 0^{j+1+sk} \gamma b \beta''' \\ v &= \alpha c 0^j 0 0^{k-j-1} 0^{j+1+sk} \gamma (b + a) \beta''' \end{aligned}$$

Let z'' be the ancestor of z' (thus descendant of x) satisfying

$$\bar{\phi}(z'') = \alpha c 0^j a \gamma b \beta'''$$

and note that

$$\begin{aligned} \bar{v} &= \alpha c 0^{j+sk} 0 \gamma (b + a) \beta''' \\ &= \alpha c 0^{j+sk} a \gamma b \beta''' + (-a) \cdot \epsilon_{|\alpha c|+sk+j} \\ &\in \bar{\phi}(D^{*,0}(z'' + (-a) \cdot e_{|\alpha c|+j})) \end{aligned}$$

□

Recall from [JFSB17a] that a decoder for correcting an unbounded number of duplications simply has to remove incidents of 0^k from the $\bar{\phi}$ -part of the noisy string. Lemma 3.2 shows that the same approach can be taken with the addition of a single substitution—without increasing the substitution distance—provided that coding is done in \mathcal{W} .

Next, we consider the case where a substitution breaks a run of zeros (in the transform domain). The following lemma allows us to remove appearances of $0^j a 0^{k-1-j}$ from the $\bar{\phi}$ -part of a noisy string (by applying an appropriate substitution) without increasing the substitution distance.

Lemma 3.3 *Suppose $u \in \Sigma^{\geq k}$ contains a substring 0^k starting at index i , and suppose $v = u + a \cdot \epsilon_\ell$ for some $i \leq j < i + k$, $0 \neq a \in \Sigma$, and $\ell \in \{j, j - k\}$ (so that $v_j \neq 0$). Note that $v' \triangleq v - v_j \cdot \epsilon_j$ has a 0^k substring at index i (like u); We remove that substring from both u, v' to produce \bar{u}, \bar{v} , respectively. Then, irrespective of what value ℓ takes, $\sigma(\bar{u}, \bar{v}) \leq 1$.*

Proof The lemma is straightforward to prove by case for ℓ . If $\ell = j$ then $v' = u$, and consequently $\bar{v} = \bar{u}$.

Otherwise, $\ell = j - k$ and $v_j = -a$, hence

$$v' = u + a \cdot (\epsilon_{j-k} + \epsilon_j)$$

and $\bar{v} = \bar{u} + a \cdot \epsilon_{j-k}$, which concludes the proof. □

It is therefore seen that a restriction to \mathcal{W} allows the correction of the substitution error without encountering the issue demonstrated in Example 3.1. This fact is more precisely stated in the following theorem:

Theorem 3.4 *If $C \subseteq \Sigma^n$, $n \geq k$, is an error-correcting code for a single substitution, and $\bar{\phi}(C) \subseteq \mathcal{W}$, then C is a 1S-correcting code.*

Proof Take $x \in C$, $y \in D^{*, \leq 1}(x)$, and define $u \triangleq \hat{\phi}(x)$, $v \triangleq \bar{\phi}(y)$. We first remove 0^k substrings from v , stopping if we reach length $n - k$. By Lemma 3.2, every removal of 0^k does not increase the substitution distance of the received string from a duplication descendant of x ; if indeed it is possible to arrive at \hat{v} of length $n - k$, then the error-correcting capabilities of C now suffice to deduce x from $\phi^{-1}(u\hat{v})$.

The only other possible case is that we ultimately arrive at \hat{v} of length n which contains a substring of length k of weight 1. We remove that substring to obtain \hat{v}' , and reverse the ϕ -transform, namely, $y' \triangleq \phi^{-1}(u\hat{v}')$. By Lemma 3.3, this produces y' of the same length as x and differing from it by at most a single substitution, which we may once more correct in the standard fashion. □

3.3 Code Construction and Size

In this section we construct a family of codes satisfying Theorem 3.4. We also study the redundancy and rate of the proposed construction. We start by bounding the rate loss of using constrained coding by restricting codes to \mathcal{W} :

Lemma 3.5 *For every integers $q \geq 2$ and $n \geq k \geq 1$,*

$$\frac{\text{red}(\mathcal{W} \cap \Sigma^n)}{n} \leq \frac{2}{k} \log_q \frac{q}{q-1}.$$

Proof We note that $C_n \subseteq \mathcal{W} \cap \Sigma^n$, where C_n is the set of length- n strings in which, divided into blocks of length k , every block ends with two non-zero elements. Hence,

$$\begin{aligned} \frac{\text{red}(\mathcal{W} \cap \Sigma^n)}{n} &\leq \frac{\text{red}(C_n)}{n} = \frac{1}{n} \left(\left\lfloor \frac{n}{k} \right\rfloor + \left\lfloor \frac{n+1}{k} \right\rfloor \right) \\ &\leq \frac{2}{k} \log_q \frac{q}{q-1}. \end{aligned}$$

□

Theorem 3.6 *If q is a prime power, $r \geq 2$, and $n = \frac{q^r-1}{q-1} + \lceil \frac{2r}{k} \rceil$, then a 1S-correcting code $C \subseteq \mathcal{W} \cap \mathbb{F}_q^n$ exists, with*

$$R(C) \geq 1 - \frac{2}{k} \log_q \frac{q}{q-1} - o(1).$$

Proof We begin by encoding data into $\mathcal{W} \cap \mathbb{F}_q^{\frac{q^r-1}{q-1}-r}$, incurring by Lemma 3.5 redundancy

$$\text{red}\left(\mathcal{W} \cap \mathbb{F}_q^{\frac{q^r-1}{q-1}-r}\right) \leq \left(\frac{q^r-1}{q-1} - r\right) \frac{2}{k} \log_q \frac{q}{q-1}.$$

Next, a systematic encoder for the $\left[\frac{q^r-1}{q-1}, r, 3\right]$ Hamming code (under the change of basis to $\{\epsilon_i\}$) can encode $\mathcal{W} \cap \mathbb{F}_q^{\frac{q^r-1}{q-1}-r} \rightarrow \mathbb{F}_q^{\frac{q^r-1}{q-1}}$, incurring r additional symbols of redundancy, and resulting in a code which can correct a single substitution.

Note, due to the systematic encoding, that the projection of this code onto the first $\frac{q^r-1}{q-1} - r$ coordinates is contained in \mathcal{W} . We may simply cushion the last r symbols with $\lceil \frac{2r}{k} \rceil$ interleaved 1's (two per k data symbols) to achieve a code $C \subseteq \mathcal{W} \cap \mathbb{F}_q^n$ which may still correct a single substitution. □

Taking $n \rightarrow \infty$, we can compare the rate obtained by the code in Theorem 3.6 to a simple upper bound of the best codes correcting only tandem duplications of length k (see [JFSB17a]),

$$R(C) \leq 1 - \frac{(q-1) \log_q e}{q^{k+2}} + o(1).$$

While both the upper bound and lower bound approach 1 as $k \rightarrow \infty$, the lower bound does so as $\Theta(k^{-1})$ whereas the upper bound is much faster as $\Theta(q^{-k})$, implying a gap yet to be resolved.

Discussion

We have proposed that reconstruction codes can be used for in-vivo DNA data-storage, due to the channel's inherent property of data replication. We have showed, under the assumption of uniform tandem-duplication noise, that any reconstruction code is partitioned into error-correcting codes for the Manhattan metric over a simplex, with minimal distances dependent on the reconstruction parameters. We then proved the existence of reconstruction codes with lower redundancy than optimal classical error-correcting codes in that context.

Next, we established a trade-off when allowing a list-decoding scheme, between the acceptable list size, the reconstruction uncertainty (i.e., number of reads), and the designed minimum distance (corresponding to code redundancy).

We have also studied a combination of duplication noise with a different error model, namely, substitutions. We have suggested a construction for error-correcting codes capable of recovering from any number of tandem-duplication errors, and a single duplication.

In the future, we believe that a study of reconstruction schemes, with or without list-decoding, is of interest with other error models which affect in-vivo DNA data storage; related models to uniform tandem-duplication noise, which have recently been studied on their own and may now be easier to analyze in that setting, and therefore are a logical first step in this direction, may be bounded tandem-duplication (see, e.g., [JFSB17a, JFB17, Kov19]) or combined uniform-tandem-duplication and substitution noise [TYSF19, TF19]. Other remaining open questions include error-control for different noise models such as interspersed-duplication (perhaps complemented- or palindromic-duplication), as well as combinations of other, or multiple, error models.

Bibliography

- [ABFJ17] Noga Alon, Jehoshua Bruck, Farzad Farnoud, and Siddharth Jain. Duplication distance to the root for binary sequences. *IEEE Trans. on Inform. Theory*, 63(12):7793–7803, December 2017.
- [ADM⁺15] Jayadev Acharya, Hirakendu Das, Olgica Milenkovic, Alon Orlitsky, and Shengjun Pan. String reconstruction from substring compositions. *SIAM J. Discrete Math.*, 29(3):1340–1371, 2015.
- [AO04] M. Arita and Y. Ohashi. Secret signatures inside genomic DNA. *Biotechnology Progress*, 20(5):1605–1607, 2004.
- [Bal13] F. Balado. Capacity of DNA data embedding under substitution mutations. *IEEE Trans. on Inform. Theory*, 59(2):928–941, February 2013.
- [BC62] Raj Chandra Bose and Sarvadaman Chowla. Theorems in the additive theory of numbers. *Commentarii Mathematici Helvetici*, 37(1):141–147, December 1962.
- [CB11] Yuval Cassuto and Mario Blaum. Codes for symbol-pair read channels. *IEEE Trans. on Inform. Theory*, 57(12):8011–8020, December 2011.
- [CGK12] George M. Church, Yuan Gao, and Sriram Kosuri. Next-generation digital information storage in DNA. *Science*, 337(6102):1628–1628, 2012.
- [Che52] Pafnuty Lvovich Chebyshev. Mémoire sur les nombres premiers. *J. Math. Pures Appl.*, 17:366–390, 1852.
- [CKV⁺18] Yeow Meng Chee, Han Mao Kiah, Alexander Vardy, Van Khu Vu, and Eitan Yaakobi. Coding for racetrack memories. *IEEE Trans. on Inform. Theory*, 64(11):7094–7112, November 2018.
- [CRB99] C. T. Clelland, V. Risca, and C. Bancroft. Hiding messages in DNA microdots. *Nature*, 399(6736):533–534, 1999.

- [Doo40] Joseph Leo Doob. Regularity properties of certain families of chance variables. *Transactions of the American Mathematical Society*, 47(3):455–486, 1940.
- [EFSB16] Ohad Elishco, Farzad Farnoud, Moshe Schwartz, and Jehoshua Bruck. The capacity of some Pólya string models. In *Proceedings of the 2016 IEEE International Symposium on Information Theory (ISIT2016), Barcelona, Spain*, pages 270–274, July 2016.
- [FSB15] Farzad Farnoud, Moshe Schwartz, and Jehoshua Bruck. A stochastic model for genomic interspersed duplication. In *Proceedings of the 2015 IEEE International Symposium on Information Theory (ISIT2015), Hong Kong, China*, pages 904–908, June 2015.
- [FSB16] Farzad Farnoud, Moshe Schwartz, and Jehoshua Bruck. The capacity of string-duplication systems. *IEEE Trans. on Inform. Theory*, 62(2):811–824, February 2016.
- [FSB19] Farzad Farnoud, Moshe Schwartz, and Jehoshua Bruck. Estimation of duplication history under a stochastic model for tandem repeats. *BMC Bioinformatics*, 20(1):64–74, February 2019.
- [GKM17] Ryan Gabrys, Han Mao Kiah, and Olgica Milenkovic. Asymmetric Lee distance codes for DNA-based storage. *IEEE Trans. on Inform. Theory*, 63(8):4982–4995, August 2017.
- [GS80] Ronald Lewis Graham and Neil James Alexander Sloane. Lower bounds for constant weight codes. *IEEE Trans. on Inform. Theory*, 26(1):37–43, January 1980.
- [GYM18] Ryan Gabrys, Eitan Yaakobi, and Olgica Milenkovic. Codes in the Damerau distance for deletion and adjacent transposition correction. *IEEE Trans. on Inform. Theory*, 64(4):2550–2570, April 2018.
- [HB07] Dominik Heider and Angelika Barnekow. DNA-based watermarks using the DNA-Crypt algorithm. *BMC Bioinformatics*, 8(1):176–185, May 2007.
- [How89] T. D. Howell. Statistical properties of selected recording codes. *IBM Journal of Research and Development*, 33(1):60–73, January 1989.
- [JFB17] Siddharth Jain, Farzad Farnoud, and Jehoshua Bruck. Capacity and expressiveness of genomic tandem duplication. *IEEE Trans. on Inform. Theory*, 63(10):6129–6138, October 2017.

- [JFS⁺10] Daniel C. Jupiter, Thomas A. Ficht, James Samuel, Qing-Ming Qin, and Paul de Figueiredo. DNA watermarking of infectious agents: Progress and prospects. *PLoS pathogens*, 6(6):e1000950, 2010.
- [JFSB17a] Siddharth Jain, Farzad Farnoud, Moshe Schwartz, and Jehoshua Bruck. Duplication-correcting codes for data storage in the DNA of living organisms. *IEEE Trans. on Inform. Theory*, 63(8):4996–5010, August 2017.
- [JFSB17b] Siddharth Jain, Farzad Farnoud, Moshe Schwartz, and Jehoshua Bruck. Noise and uncertainty in string-duplication systems. In *Proceedings of the 2017 IEEE International Symposium on Information Theory (ISIT2017), Aachen, Germany*, pages 3120–3124, June 2017.
- [Joh62] Selmer Martin Johnson. A new upper bound for error-correcting codes. *IRE Trans. on Inform. Theory*, 8(3):203–207, April 1962.
- [KLS07] Elena Konstantinova, Vladimir Levenshtein, and Johannes Siemons. Reconstruction of permutations distorted by single transposition errors. *arXiv preprint arXiv:math/0702191*, 2007.
- [Knu86] Donald E. Knuth. Efficient balanced codes. *IEEE Trans. on Inform. Theory*, 32(1):51–53, January 1986.
- [Kon07] E. Konstantinova. Reconstruction of permutations distorted by single reversal errors. *Discrete Appl. Math.*, 155:2426–2434, 2007.
- [Kon08] E. Konstantinova. On reconstruction of signed permutations distorted by reversal errors. *Discrete Math.*, 308:974–984, 2008.
- [Kov19] Mladen Kovačević. Codes correcting all patterns of tandem-duplication errors of maximum length 3. *arXiv preprint arXiv:1911.06561*, 2019.
- [KPM16] Han Mao Kiah, Gregory J. Puleo, and Olgica Milenkovic. Codes for DNA sequence profiles. *IEEE Trans. on Inform. Theory*, 62(6):3125–3146, June 2016.
- [KT17] Mladen Kovačević and Vincent Yan Fu Tan. Improved bounds on Sidon sets via lattice packings of simplices. *SIAM J. Discrete Math.*, 31(3):2269–2278, 2017.

- [KT18a] Mladen Kovačević and Vincent Yan Fu Tan. Asymptotically optimal codes correcting fixed-length duplication errors in DNA storage systems. *IEEE Communications Letters*, 22(11):2194–2197, November 2018.
- [KT18b] Mladen Kovačević and Vincent Yan Fu Tan. Codes in the space of multisets—coding for permutation channels with impairments. *IEEE Trans. on Inform. Theory*, 64(7):5156–5169, July 2018.
- [LDB⁺12] M. Liss, D. Daubert, K. Brunner, K. Kliche, U. Hammes, A. Leiberer, and R. Wagner. Embedding permanent watermarks in synthetic genes. *PLoS ONE*, 7(8):e42465, 2012.
- [Lev01] Vladimir Iosifovich Levenshtein. Efficient reconstruction of sequences. *IEEE Trans. on Inform. Theory*, 47(1):2–22, January 2001.
- [LJWZ18] Andreas Lenz, Niklas Jünger, and Antonia Wachter-Zeh. Bounds and constructions for multi-symbol duplication error correcting codes. *arXiv preprint arXiv:1807.02874*, 2018.
- [LKKM08] V. I. Levenshtein, E. Konstantinova, E. Konstantinov, and S. Molodtsov. Reconstruction of a graph from 2-neighborhoods of its vertices. *Discrete Appl. Math.*, 156:1399–1406, 2008.
- [LMVM05] Peter Leupold, Carlos Martín-Vide, and Victor Mitrana. Uniformly bounded duplication languages. *Discrete Appl. Math.*, 146(3):301–310, 2005.
- [LS09] V. I. Levenshtein and J. Siemons. Error graphs and the reconstruction of elements in graphs. *J. Combin. Theory Ser. A*, 116:795–815, 2009.
- [LWY19] Andreas Lenz, Antonia Wachter-Zeh, and Eitan Yaakobi. Duplication-correcting codes. *Designs, Codes and Cryptography*, 87(2):277–298, March 2019.
- [MRS01] Brian H. Marcus, Ron M. Roth, and Paul H. Siegel. An introduction to coding for constrained systems. Unpublished Lecture Notes, October 2001.
- [MS78] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes*. North-Holland, 1978.
- [MV17] Hessam MahdaviFar and Alexander Vardy. Asymptotically optimal sticky-insertion-correcting codes with efficient encoding and decoding.

- In *Proceedings of the 2017 IEEE International Symposium on Information Theory (ISIT2017)*, Aachen, Germany, pages 2683–2687, June 2017.
- [POB08] Danilo Pumpernik, Borut Oblak, and Branko Borštnik. Replication slippage versus point mutation rates in short tandem repeats of the human genome. *Molecular Genetics and Genomics*, 279:53–61, 2008.
- [Rot06] Ron M. Roth. *Introduction to Coding Theory*. Cambridge Univ. Press, 2006.
- [RS94] Ron M. Roth and Paul H. Siegel. Lee-metric BCH codes and their application to constrained and partial-response channels. *IEEE Trans. on Inform. Theory*, 40(4):1083–1096, July 1994.
- [RSW12] Frank Ruskey, Joe Sawada, and Aaron Williams. De bruijn sequences for fixed-weight binary strings. *SIAM J. Discrete Math.*, 26(2):605–617, 2012.
- [RSY19] Netanel Raviv, Moshe Schwartz, and Eitan Yaakobi. Rank-modulation codes for DNA storage with shotgun sequencing. *IEEE Trans. on Inform. Theory*, 65(1):50–64, January 2019.
- [SCT16] I. Shomorony, T. A. Courtade, and D. Tse. Fundamental limits of genome assembly under an adversarial erasure model. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, 2(2):199–208, December 2016.
- [SGSD17] Frederic Sala, Ryan Gabrys, Clayton Schoeny, and Lara Dolecek. Exact reconstruction from insertions in synchronization codes. *IEEE Trans. on Inform. Theory*, 63(4):2428–2445, April 2017.
- [SNMC16] Seth L. Shipman, Jeff Nivala, Jeffrey D. Macklis, and George M. Church. Molecular recordings by directed CRISPR spacer acquisition. *Science*, 353(6298):aaf1175, June 2016.
- [SNMC17] Seth L. Shipman, Jeff Nivala, Jeffrey D. Macklis, and George M. Church. CRISPR-Cas encoding of a digital movie into the genomes of a population of living bacteria. *Nature*, 547:345, July 2017.
- [TF19] Yuanyuan Tang and Farzad Farnoud (Hassanzadeh). Error-correcting codes for noisy duplication channels. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 140–146, September 2019.

- [TYSF19] Yuanyuan Tang, Yonatan Yehezkeally, Moshe Schwartz, and Farzad Farnoud (Hassanzadeh). Single-error detection and correction for duplication and substitution channels. In *Proceedings of the 2019 IEEE International Symposium on Information Theory (ISIT) (ISIT'2019), Paris, France*, pages 300–304, July 2019.
- [WkWF03] Pak Chung Wong, Kwong kwok Wong, and Harlan Foote. Organic data memory using the DNA approach. *Communications of the ACM*, 46(1):95–98, January 2003.
- [YB19] Eitan Yaakobi and Jehoshua Bruck. On the uncertainty of information retrieval in associative memories. *IEEE Trans. on Inform. Theory*, 65(4):2155–2165, April 2019.
- [YBS16] Eitan Yaakobi, Jehoshua Bruck, and Paul H. Siegel. Constructions and decoding of cyclic codes over b -symbol read channels. *IEEE Trans. on Inform. Theory*, 62(4):1541–1551, April 2016.
- [ZAM14] Kai Zhou, Abram Aertsen, and Chris W. Michiels. The role of variable DNA tandem repeats in bacterial adaptation. *FEMS Microbiology Reviews*, 38(1):119–141, January 2014.
- [ZW88] E. Zehavi and J. K. Wolf. On runlength codes. *IEEE Trans. on Inform. Theory*, 34(1):45–54, January 1988.

תקציר

אנו חוקרים קידוד לבקרת שגיאות בערוצים עם רעש שכפול, עבור שימושים לשם אחסון מידע ב-DNA. אנו מציעים כי שימוש בסכימת השחזור של לווינשטיין, אשר רלוונטי כאשר מספר עותקים מורעשים של המידע זמינים למפענח, יכול לתרום לשימושים אלו. זאת שכן סכימות אחסון מידע בתאים חיים תואמות מאליהן לסכימה, כיוון שהן מציעות שכפול באופן טבעי של המידע בתהליך ההתרבות, היוצר עותקים מורעשים רבים עקב מוטציות. אנו חוקרים רעש הנוצר עקב שכפולים עוקבים בלבד, עם אורך קבוע לחלון השכפול. אנו משתמשים בהקבלה לקודים קבועי-משקל במרחב וקטורים של מספרים שלמים, עם מטריקת מנהטן. ע"י חסימת גודל החיתוך של היפר-מישורים עם הרבעון החיובי אנו מוכיחים כי קודים המיועדים לסכימת השחזור משיגים יתירות נמוכה מזו הדרושה לקודים מתקני-שגיאות. כמו כן, אנו מציעים בניה למשפחת קודי שחזור.

בנוסף, אנו חוקרים את בעיית השחזור כאשר אנו מרשים שחזור של רשימה, במקום מחרוזת בודדת. בעיה זו ניתנת להצגה כסכימת זכרון אסוציאטיבי בהקשר הנוכחי; אנו מוצאים את חוסר-הודאות השייך ל- $2 < m$ מחרוזות (כאשר שחזור מחרוזת בודדת מתאים ל- $m = 2$) במובנים אסימפטוטיים, כאשר מילות-קוד מוגבלות לקוד מתקן שגיאות עם מרחק מזערי נתון. כך אנו מוצאים את שקלול התמורות בין המרחק המזערי של הקוד הנבחר, הסיבולת למספר נתון של שגיאות שחזור, גודל הרשימה המורשה וחוסר הודאות המתאים, אשר שקול למספר העותקים המורעשים שעלינו לקבל לצורך שחזור מוצלח. אנו מדגימים בדרך זו כי כאשר מרשים שחזור של רשימה, ניתן להפחית אפילו יותר את יתירות הקידוד, או את מספר הקריאות הדרוש, או שניהם. לסיום, אנו חוקרים שילוב של רעש שכפול עם שגיאות נקודתיות, אשר נצפה גם הוא בשימושי אחסון מידע ב-DNA בתאים חיים. אנו מתמקדים במודל לא מוגבל, בו השגיאות הנקודתיות יכולות להתרחש בכל מקום במחרוזת, כמו גם בכל שלב ברצף שגיאות השכפול. ע"י שימוש בגישת קידוד מאולץ, אנו מפתחים קודים מתקני שגיאות המסוגלים לתקן כמות כלשהי של שגיאות שכפול, יחדיו עם שגיאה נקודתית בודדת.

הצהרת תלמיד המחקר עם הגשת עבודת הדוקטור לשיפוט

אני החתום מטה מצהיר/ה בזאת: (אנא סמן):

X חיברתי את חיבורי בעצמי, להוציא עזרת ההדרכה שקיבלתי מאת מנחה/ים.

X החומר המדעי הנכלל בעבודה זו הינו פרי מחקרי מתקופת היותי תלמיד/ת מחקר.

___ בעבודה נכלל חומר מחקרי שהוא פרי שיתוף עם אחרים, למעט עזרה טכנית הנהוגה בעבודה ניסיונית. לפיכך מצורפת בזאת הצהרה על תרומתי ותרומת שותפי למחקר, שאושרה על ידם ומוגשת בהסכמתם.

תאריך 18 במרץ 2020 שם התלמיד/ה יונתן יחזקאלי חתימה יונתן

העבודה נעשתה בהדרכת

פרופ' משה שורץ

בבית הספר להנדסת חשמל ומחשבים

בפקולטה למדעי ההנדסה

קידוד לערוצי שכפול עם שימושים לאחסון מידע ב-DNA

מחקר לשם מילוי חלקי של הדרישות לקבלת תואר
"דוקטור לפילוסופיה"

מאת

יונתן יחזקאלי



אישור המנחה

אישור דיקן בית הספר ללימודי מחקר מתקדמים ע"ש קרייטמן

הוגש לסינאט אוניברסיטת בן-גוריון בנגב

18 במרץ 2020

כב' אדר התש"פ

באר שבע

קידוד לערוצי שכפול עם שימושים לאחסון מידע ב-DNA

מחקר לשם מילוי חלקי של הדרישות לקבלת תואר
"דוקטור לפילוסופיה"

מאת

יונתן יחזקאלי

הוגש לסינאט אוניברסיטת בן-גוריון בנגב

18 במרץ 2020

כב' אדר התש"פ

באר שבע