

Fundamental Limits of Multiple Sequence Reconstruction from Substrings

Kel Levick

University of Illinois, Urbana-Champaign
Urbana, IL, USA
klevick2@illinois.edu

Ilan Shomorony

University of Illinois, Urbana-Champaign
Urbana, IL, USA
ilans@illinois.edu

Abstract—The problem of reconstructing a sequence from the set of its length- k substrings has received considerable attention due to its various applications in genomics. We study an uncoded version of this problem where multiple random sources are to be simultaneously reconstructed from the union of their k -mer sets. We consider an asymptotic regime where $m = n^\alpha$ i.i.d. source sequences of length n are to be reconstructed from the set of their substrings of length $k = \beta \log n$, and seek to characterize the (α, β) pairs for which reconstruction is information-theoretically feasible. We show that, as $n \rightarrow \infty$, the source sequences can be reconstructed if $\beta > \max(2\alpha + 1, \alpha + 2)$ and cannot be reconstructed if $\beta < \max(2\alpha + 1, \alpha + \frac{3}{2})$, characterizing the feasibility region almost completely. Interestingly, our result shows that there are feasible (α, β) pairs where repeats across the source strings abound, and non-trivial reconstruction algorithms are needed to achieve the fundamental limit.

I. INTRODUCTION

The problem of reconstructing a sequence from the set of its length- k substrings (or k -mers) has a long history, dating back to the development of the sequencing-by-hybridization (SBH) technology in the late 1980s [1]. Due to this technological advance, significant attention was originally devoted to understanding when the set of k -mers uniquely determines the underlying sequence, and to designing algorithms to perform the reconstruction efficiently [2–5].

Inspired by Ukkonen’s 1992 paper [2], which characterized how large k needs to be to allow reconstruction, a more recent line of work considered this problem from an information-theoretic perspective [6–10]. In particular, [7] considered a random source \mathbf{x} (modeling a genome sequence), and provided conditions on how large k needs to be (and how many length- k substrings need to be observed) in order for the reconstruction of a source sequence \mathbf{x} to be information-theoretically feasible.

In this work, we consider a multiple-source version of this problem. We assume that there are m random source sequences $\mathbf{x}_1, \dots, \mathbf{x}_m$, each of length n , and we wish to reconstruct them from the union of their sets of k -mers. This multiple-source setting can be motivated by modern sequencing assays such as metagenomics and RNA-seq, where one observes substrings from a mixture of distinct source sequences (different microbial genomes in the case of metagenomics and different mRNA molecules in the case of RNA-seq). It can also be motivated by the fact that many bioinformatics algorithms first convert a set of reads into the set of their constituent k -mers and then operate in k -mer space. It is thus relevant to

understand when such sets fully preserve the information in the original reads. Our setting is also connected with the problem of DNA-based data storage, where one wishes to recover a set of m synthesized DNA molecules by jointly sequencing them. Our work can thus be seen as an *uncoded* version of the recent work on multi-strand reconstruction from substrings [11], or a random coding approach to it (see also “Related work” below).

Our goal is to characterize how large k needs to be to allow perfect reconstruction of all m source sequences. We assume each source is i.i.d. $\text{Bern}(\frac{1}{2})$ and consider an asymptotic regime where $m = n^\alpha$ and $k = \beta \log n$, for some constants $\alpha, \beta > 0$. We aim to characterize for which (α, β) pairs the set of k -mers uniquely determines the set of source sequences $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ with vanishing error probability. Our main result, illustrated in Figure 1, characterizes the (α, β) feasibility region almost completely. We establish this feasibility region by carefully analyzing the (random) de Bruijn graph formed by the k -mer set. In this graph, the set of source sequences corresponds to a set of paths that cover all edges. Our analysis shows that there are feasible (α, β) pairs where the k -mer de Bruijn graph is fairly complex, due to the presence of repeats

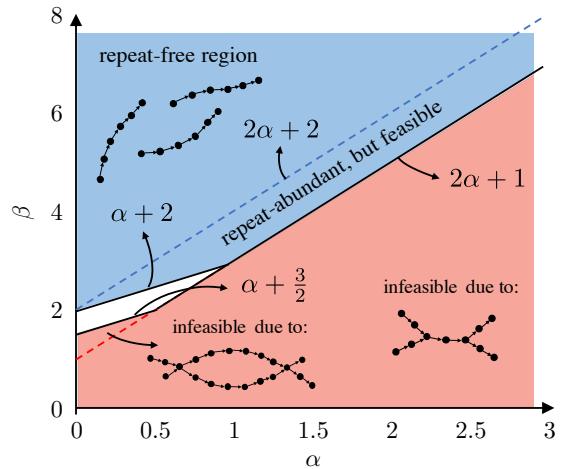


Fig. 1. The blue region is given by $\beta > \max(2\alpha + 1, \alpha + 2)$, and corresponds to (α, β) pairs for which $\mathbf{x}_1, \dots, \mathbf{x}_m$ are uniquely reconstructible from the set of k -mers with vanishing error probability. The red region is given by $\beta < \max(2\alpha + 1, \alpha + 2)$, and corresponds to (α, β) pairs where the solution is not unique. We illustrate the types of k -mer de Bruijn graphs that lead to the ambiguity. The white region corresponds to the unknown (α, β) pairs.

across the source sequences, but nevertheless there is a unique way to decompose the graph into m source paths.

Related work: An important line of related work deals with the *coded* version of the problem we consider. Since a sequence \mathbf{x} containing no repeats of length $k-1$ can be reconstructed from the set of its k -mers [2], [12] studied how large a *repeat-free code* can be. They showed that, as long as $k = \lceil \beta \log n \rceil$, for $\beta > 1$, one can asymptotically build $(k-1)$ -repeat-free codes with rate 1. The setup where one seeks to reconstruct a codeword \mathbf{x} from a noisy version of its set of k -mers have also received considerable attention [13–16].

A multi-sequence version of the coding problem in [12] was considered in [13]. Similar to our paper, they provide conditions on how k needs to scale with n and m to guarantee that the asymptotic rate of all k -mer-reconstructible codes approaches 1. Our work can be seen as an uncoded counterpart to [13], where we characterize how k needs to scale with n and m to allow the reconstruction of random source sequences with vanishing failure probability. As an important distinction, we notice that the code constructions in [13] guarantee that there are no repeats across the m sequences to be reconstructed, while in our setting, m source sequences with repeats may still be reconstructible for some (α, β) .

II. PROBLEM SETTING AND MAIN RESULTS

We consider m source sequences $\mathbf{x}_1, \dots, \mathbf{x}_m$, each of which is generated as an i.i.d. $\text{Ber}(\frac{1}{2})$ length- n sequence. The goal is to reconstruct the set of source sequences $\mathcal{X} := \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ from the set of all of their $(k+1)$ -mers, i.e.,

$$\mathcal{Y}(\mathcal{X}) := \{\mathbf{x}_i[j : j+k] : 1 \leq i \leq m, 1 \leq j \leq n-k\},$$

where $\mathbf{x}_i[a : b]$ is the substring of \mathbf{x}_i between (and including) the a th and b th symbols. Here, we use $(k+1)$ -mers rather than k -mers for notational convenience, and the final results are not affected. When clear from context, we write \mathcal{Y} for $\mathcal{Y}(\mathcal{X})$. We will also use the notation $\mathbf{x}_i(a)$ to refer to the a th k -mer of \mathbf{x}_i , for $1 \leq a \leq n-k+1$.

We consider an asymptotic regime where $m = n^\alpha$ and $k = \beta \log n$, for positive constants α and β . As it turns out, this is the asymptotic regime of interest, where the fundamental limits are non-trivial. We say that \mathcal{Y} reconstructs \mathcal{X} if \mathcal{X} is the unique set of m length- n sequences (up to relabeling) that could have generated \mathcal{Y} . It is then natural to define a notion of feasibility as follows.

Definition 1: We say that (α, β) is a feasible pair if for $m = n^\alpha$ and $k = \beta \log n$, $\lim_{n \rightarrow \infty} \Pr(\mathcal{Y} \text{ reconstructs } \mathcal{X}) = 1$.

The goal is to characterize the feasibility region $\mathcal{F} \subset \mathbb{R}_+^2$, defined as the set of all feasible (α, β) pairs.

An equivalent and useful representation of the k -mer set \mathcal{Y} is its de Bruijn graph $G(\mathcal{Y})$. The de Bruijn graph $G(\mathcal{Y})$ can be obtained by taking the set of k -mers in \mathcal{Y} (i.e., length- k prefixes and suffixes of the $(k+1)$ -mers in \mathcal{Y}) as the node set, and connecting two k -mers with a *directed* edge if they are the prefix and suffix of a $(k+1)$ -mer in \mathcal{Y} . Notice that the true set of sequences \mathcal{X} can be seen as a set of paths on $G(\mathcal{Y})$ that cover all the edges (edges can be used multiple times).

Our main result is to characterize, almost completely, the feasibility region \mathcal{F} as follows:

Theorem 1: The following bounds on \mathcal{F} hold:

- $(\alpha, \beta) \in \mathcal{F}$ if $\beta > \max(2\alpha + 1, \alpha + 2)$,
- $(\alpha, \beta) \notin \mathcal{F}$ if $\beta < \max(2\alpha + 1, \alpha + \frac{3}{2})$.

As illustrated in Figure 1, Theorem 1 nearly completely characterizes \mathcal{F} . In particular, for $\alpha > 1$, $(\alpha, \beta) \in \mathcal{F}$ if $\beta > 2\alpha + 1$, and $(\alpha, \beta) \notin \mathcal{F}$ if $\beta < 2\alpha + 1$. For $\alpha < 1$ there is a small uncharacterized region of (α, β) pairs.

The rest of the paper is organized as follows. In Section III, as a warm-up, we characterize the repeat-free region, a (strict) subset of \mathcal{F} where there are no length- k repeats across the source sequences (and the source strings are thus reconstructible). In Section IV, we prove the inner bound part of Theorem 1, and in Section V, we prove the outer bound part of Theorem 1.

III. REPEAT-FREE REGION

We start by considering a simple scenario where \mathcal{X} can be recovered correctly: when each k -mer is unique across all source sequences. In this case, it is well known that the de Bruijn graph $G(\mathcal{Y})$ will contain m disjoint paths, each of which corresponds to one of the \mathbf{x}_i s, and \mathcal{X} is thus reconstructible. We prove the following result.

Lemma 1: If $\beta > 2\alpha + 2$, with probability tending to 1 as $n \rightarrow \infty$, there are no repeated k -mers across or within \mathbf{x}_i s. This region is marked in Figure 1. Notice that it does not include all pairs of feasible (α, β) pairs.

To prove Lemma 1, we note that, if $i \neq j$,

$$\Pr(\mathbf{x}_i(a) = \mathbf{x}_j(b)) = 2^{-k} = 2^{-\beta \log n} = n^{-\beta}.$$

Hence, the probability that any two strings \mathbf{x}_i and \mathbf{x}_j , for $i \neq j$ have at least one k -mer in common is upper-bounded using the union bound by

$$\binom{m}{2} n^2 n^{-\beta} \leq m^2 n^{2-\beta} \leq n^{2\alpha+2-\beta}.$$

Additionally, we need to consider the case where a single string contains a repeat k -mer. In this case, not all of the k -mers are independent due to potential overlaps. However, a careful analysis reveals that, even if $|b-a| < k$ and $\mathbf{x}_i(a)$ and $\mathbf{x}_i(b)$ overlap, we still have $\Pr(\mathbf{x}_i(a) = \mathbf{x}_i(b)) = n^{-\beta}$ (see [17] for a similar analysis). We conclude that the probability that there is repeated k -mer within some \mathbf{x}_i is at most $mn^2 n^{-\beta} = n^{\alpha+2-\beta}$, and the probability that any k -mer is repeated across or within $\mathbf{x}_1, \dots, \mathbf{x}_m$ is upper-bounded by

$$n^{2\alpha+2-\beta} + n^{\alpha+2-\beta},$$

which tends to zero for $\beta > 2\alpha + 2$. A converse result can also be shown; i.e., for $\beta < 2\alpha + 2$, there will be repeats with probability tending to 1 (but we do not describe that). In the next section, we show that there are feasible (α, β) pairs that are in this repeat-abundant regime.

IV. INNER BOUND TO FEASIBILITY REGION

In order to characterize \mathcal{F} , we define the error event $\mathcal{E} = \{\mathcal{Y} \text{ does not reconstruct } \mathcal{X}\}$, and seek to characterize the asymptotic behavior of $\Pr(\mathcal{E})$ for a given choice of (α, β) . To do so, we will need to define several events related to repeats:

$$\begin{aligned}\mathcal{A} &= \{\exists i, a, b : \mathbf{x}_i(a) = \mathbf{x}_i(b)\} \\ \mathcal{B} &= \{\exists i, j, l, a, b, c, d : \mathbf{x}_i(a) = \mathbf{x}_j(c), \mathbf{x}_i(b) = \mathbf{x}_l(d), \\ &\quad 0 \leq b - a < k, \text{ and either } \{j \neq l\} \text{ or} \\ &\quad \{j = l \text{ and } d - c \neq b - a\}\} \\ \mathcal{C} &= \{\exists i, j, a : \mathbf{x}_i(1) = \mathbf{x}_j(a)\} \cup \{\exists i, j, a : \mathbf{x}_i(n') = \mathbf{x}_j(a)\} \\ \mathcal{D} &= \{\exists i, j, a : \mathbf{x}_i(a) = \mathbf{x}_j(a)\}\end{aligned}$$

where we let $n' = n - k + 1$ for conciseness.

Event \mathcal{A} is the event that there is at least one intra-sequence repeat. Event \mathcal{B} is the event that there are two pairs of repeats with an overlap (i.e., $\mathbf{x}_i(a)$ and $\mathbf{x}_i(b)$ overlap in \mathbf{x}_i), but they are not simply a longer length- t repeat for $t > k$. Notice that \mathcal{B} includes the event that a triple repeat $\mathbf{x}_i(a) = \mathbf{x}_j(b) = \mathbf{x}_l(c)$ occurs. Event \mathcal{C} is the event that the first or last k -mer of some \mathbf{x}_i appears somewhere else. Event \mathcal{D} is the event that the a th k -mer of two sequences \mathbf{x}_i and \mathbf{x}_j is the same. Notice that events \mathcal{A} , \mathcal{B} and \mathcal{C} are not error events per se (they each do not imply \mathcal{E}), but they will be useful in the analysis. On the other hand, $\mathcal{D} \Rightarrow \mathcal{E}$, since given \mathcal{D} it is possible to replace \mathbf{x}_i and \mathbf{x}_j with two incorrect sequences

$$\begin{aligned}\tilde{\mathbf{x}}_i &= \mathbf{x}_i[1 : a - 1]\mathbf{x}_j[a : n] \\ \tilde{\mathbf{x}}_j &= \mathbf{x}_j[1 : a - 1]\mathbf{x}_i[a : n].\end{aligned}$$

We will upper bound the error probability as

$$\begin{aligned}\Pr(\mathcal{E}) &\leq \Pr(\mathcal{A} \cup \mathcal{B} \cup \mathcal{C} \cup \mathcal{D} \cup \mathcal{E}) \\ &\leq \Pr(\mathcal{A}) + \Pr(\mathcal{B}) + \Pr(\mathcal{C}) + \Pr(\mathcal{D}) + \\ &\quad + \Pr(\mathcal{E} \cap \bar{\mathcal{A}} \cap \bar{\mathcal{B}} \cap \bar{\mathcal{C}} \cap \bar{\mathcal{D}})\end{aligned}\tag{1}$$

and analyze the conditions for each term to go to zero. First we notice that, by the union bound,

$$\Pr(\mathcal{A}) \leq mn^2n^{-\beta} = n^{\alpha+2-\beta},\tag{2}$$

from which we see that $\Pr(\mathcal{A}) \rightarrow 0$ if $\beta > \alpha + 2$. Similarly,

$$\begin{aligned}\Pr(\mathcal{B}) &\leq m^3n^32kn^{-2\beta} + m^2n^2(2k)^2n^{-2\beta} \\ &\leq n^{3\alpha+3-2\beta+o(1)},\end{aligned}\tag{3}$$

$$\Pr(\mathcal{C}) \leq 2m^2nn^{-\beta} = 2n^{2\alpha+1-\beta},\tag{4}$$

$$\Pr(\mathcal{D}) \leq m^2nn^{-\beta} = 4n^{2\alpha+1-\beta}.\tag{5}$$

We point out that to establish (3), we need a careful analysis of the case where $\mathbf{x}_j(c)$ and $\mathbf{x}_j(d)$ also overlap, presented in [17].

As a result we see that, if $\beta > \max(\alpha + 2, \frac{3}{2}\alpha + \frac{3}{2}, 2\alpha + 1)$, $\Pr(\mathcal{A}) + \Pr(\mathcal{B}) + \Pr(\mathcal{C}) + \Pr(\mathcal{D}) \rightarrow 0$. What remains is to bound $\Pr(\mathcal{E} \cap \bar{\mathcal{A}} \cap \bar{\mathcal{B}} \cap \bar{\mathcal{C}} \cap \bar{\mathcal{D}})$.

Now suppose $\bar{\mathcal{A}} \cap \bar{\mathcal{B}} \cap \bar{\mathcal{C}} \cap \bar{\mathcal{D}}$ holds. Let the multiplicity $\mu_{\mathcal{X}}(\mathbf{v})$ of a node \mathbf{v} in $G(\mathcal{Y})$ be the number of times it is traversed by paths in \mathcal{X} . Notice that $\bar{\mathcal{B}}$ guarantees that no node can be

traversed three or more times, so given $\bar{\mathcal{B}}$, node multiplicities can only be 1 or 2.

The proof of the following lemma is provided in [17].

Lemma 2: Suppose $\bar{\mathcal{A}}$, $\bar{\mathcal{B}}$, and $\bar{\mathcal{C}}$ hold. Then the multiplicity of every node in $G(\mathcal{Y})$ is fully determined. In other words, every valid solution of $G(\mathcal{Y})$ traverses a given node in $G(\mathcal{Y})$ the same number of times.

Now suppose \mathcal{E} occurs. By definition, there is at least one $\tilde{\mathcal{X}} \neq \mathcal{X}$ such that $\mathcal{Y}(\tilde{\mathcal{X}}) = \mathcal{Y}(\mathcal{X})$. Let $\tilde{\mathcal{X}}$ be one such choice with minimum set difference $|\mathcal{X} - \tilde{\mathcal{X}}|$. Let $|\mathcal{X} - \tilde{\mathcal{X}}| = c$. By Lemma 2, we know that for all nodes \mathbf{v} ,

$$\mu_{\mathcal{X}}(\mathbf{v}) = \mu_{\tilde{\mathcal{X}}}(\mathbf{v}) := \mu_G(\mathbf{v}).\tag{6}$$

Additionally,

$$\mu_{\mathcal{X} \cap \tilde{\mathcal{X}}}(\mathbf{v}) + \mu_{\mathcal{X} - \tilde{\mathcal{X}}}(\mathbf{v}) = \mu_{\mathcal{X}}(\mathbf{v}) = \mu_G(\mathbf{v})\tag{7}$$

$$\mu_{\mathcal{X} \cap \tilde{\mathcal{X}}}(\mathbf{v}) + \mu_{\tilde{\mathcal{X}} - \mathcal{X}}(\mathbf{v}) = \mu_{\tilde{\mathcal{X}}}(\mathbf{v}) = \mu_G(\mathbf{v}),\tag{8}$$

implying that $\mu_{\mathcal{X} - \tilde{\mathcal{X}}}(\mathbf{v}) = \mu_{\tilde{\mathcal{X}} - \mathcal{X}}(\mathbf{v})$.

We now build a “difference de Bruijn graph” $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$ on the $(k + 1)$ -mers in \mathcal{Y} that are present in $\mathcal{X} - \tilde{\mathcal{X}}$. It is easy to see that $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$ is the same as the subgraph of $G(\mathcal{Y})$ induced by the edges corresponding to $(k + 1)$ -mers in $\mathcal{X} - \tilde{\mathcal{X}}$ with nodes $\{\mathbf{v} : \mu_{\mathcal{X} - \tilde{\mathcal{X}}}(\mathbf{v}) > 0\}$. Note that the nodes and their multiplicities are equivalent to those of the set $\{\mathbf{v} : \mu_{\tilde{\mathcal{X}} - \mathcal{X}}(\mathbf{v}) > 0\}$, and the set of edges in $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$ must also correspond to all of the $(k + 1)$ -mers in \mathcal{Y} that are present in $\tilde{\mathcal{X}} - \mathcal{X}$. Therefore $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$ is the same as the de Bruijn graph on the $(k + 1)$ -mers in \mathcal{Y} that are present in $\tilde{\mathcal{X}} - \mathcal{X}$; i.e., $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}}) = \tilde{G}(\mathcal{Y}, \tilde{\mathcal{X}}, \mathcal{X})$. Notice that $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$ is a de Bruijn graph for both the set of c missing true paths $\mathcal{X} - \tilde{\mathcal{X}} = \{\mathbf{x}_1, \dots, \mathbf{x}_c\}$ and for the set of incorrectly reconstructed paths $\tilde{\mathcal{X}} - \mathcal{X} = \{\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_c\}$.

Define a *maximal shared subpath* in $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$ as a directed path \mathbf{p} of nodes with maximal length such that $\mu_{\mathcal{X} - \tilde{\mathcal{X}}}(\mathbf{v}) = 2$ for all nodes $\mathbf{v} \in \mathbf{p}$. As seen in Figure 2(a), a single node can be a maximal shared subpath; the two nodes with in-degree and out-degree 2 each have multiplicity 2. We have the following lemma.

Lemma 3: Given $\bar{\mathcal{A}}, \bar{\mathcal{B}}, \bar{\mathcal{C}}, \bar{\mathcal{D}}$, there are at least c maximal shared subpaths in $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$.

Proof: Given $\bar{\mathcal{A}}, \bar{\mathcal{B}}, \bar{\mathcal{C}}$, and Lemma 2, the multiplicities of all nodes in $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$ are fully determined and must be either 1 or 2. Now consider the paths corresponding to $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_c$ in $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$. We claim that, given $\bar{\mathcal{C}}$ and $\bar{\mathcal{D}}$, each $\tilde{\mathbf{x}}_i$ must traverse at least two maximal shared subpaths. Otherwise, suppose $\tilde{\mathbf{x}}_i$ only traverses a single maximal shared subpath, say $\mathbf{x}_1[s : t]$, which is shared between \mathbf{x}_1 and \mathbf{x}_2 . Given $\bar{\mathcal{C}}$, $\tilde{\mathbf{x}}_i$ must start and end at nodes with zero in-degree and out-degree respectively. Hence we must be able to write $\tilde{\mathbf{x}}_i$, without loss of generality, as

$$\tilde{\mathbf{x}}_i = \mathbf{x}_1[1 : s - 1]\mathbf{x}_1[s : t]\mathbf{x}_2[t + 1 : n]$$

But this implies that $\mathbf{x}_1(s) = \mathbf{x}_2(s)$, which contradicts $\bar{\mathcal{D}}$.

Finally, since each $\tilde{\mathbf{x}}_i$, $i = 1, \dots, c$, must traverse at least two maximal shared subpaths and, given $\bar{\mathcal{B}}$, every maximal

shared subpath can be shared by at most two true paths \mathbf{x}_i , we conclude that there must be at least c maximal shared subpaths in $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$. \blacksquare

Figure 2 shows examples of the graph $\tilde{G}(\mathcal{Y}, \mathcal{X}, \tilde{\mathcal{X}})$ given $\bar{\mathcal{A}}, \bar{\mathcal{B}}, \bar{\mathcal{C}}, \bar{\mathcal{D}}$, for $c = 2$ and $c = 4$. Notice that we have c paths and c maximal shared subpaths. Also notice that, in both examples, there exist two choices for the set of c paths.

Lemma 3 allows us to analyze $\Pr(\mathcal{E})$ by partitioning it into events $\mathcal{E}_c := \mathcal{E} \cap \{|\tilde{\mathcal{X}} - \mathcal{X}| = c\}$. Notice that a maximal shared subpath between \mathbf{x}_i and \mathbf{x}_j occurs with probability at most $n^{-\beta}$. Moreover, by the minimality of $\tilde{\mathcal{X}}$, each $\tilde{\mathbf{x}}_i$ must contain at least one maximal shared subpath. By the union bound, we conclude that

$$\begin{aligned} \Pr(\mathcal{E}_c \cap \bar{\mathcal{A}} \cap \bar{\mathcal{B}} \cap \bar{\mathcal{C}} \cap \bar{\mathcal{D}}) \\ \leq \binom{m}{c} n^c (cn)^c n^{-\beta c} \leq m^c n^{2c} n^{-\beta c} = n^{c(\alpha+2-\beta)}, \end{aligned}$$

where we used the fact that $\binom{m}{c} \leq (m/c)^c$. Summing over all values of c , we obtain

$$\begin{aligned} \Pr(\mathcal{E} \cap \bar{\mathcal{A}} \cap \bar{\mathcal{B}} \cap \bar{\mathcal{C}} \cap \bar{\mathcal{D}}) &\leq \sum_{c=2}^m n^{c(\alpha+2-\beta)} \\ &= \frac{n^{2(\alpha+2-\beta)} (1 - n^{(m-1)(\alpha+2-\beta)})}{1 - n^{\alpha+2-\beta}}, \end{aligned}$$

which tends to zero as $n \rightarrow \infty$ provided that $\beta > \alpha + 2$. Plugging this back into (1), we conclude that any (α, β) pair with $\beta > \max(\alpha + 2, \frac{3}{2}\alpha + \frac{3}{2}, 2\alpha + 1) = \max(\alpha + 2, 2\alpha + 1)$ is feasible. In other words, we showed that

$$\{(\alpha, \beta) : \beta > \max(\alpha + 2, 2\alpha + 1)\} \subseteq \mathcal{F},$$

concluding the proof of the achievability part of Theorem 1.

V. OUTER BOUND TO FEASIBILITY REGION

In this section, we will prove the outer bound part of Theorem 1. First, we will show that $\Pr(\mathcal{E}) \not\rightarrow 0$ as $n \rightarrow \infty$, if $\beta < 2\alpha + 1$. Second, we will show that $\Pr(\mathcal{E}) \not\rightarrow 0$ if $\beta < \alpha + \frac{3}{2}$. To do so, we will show that in those regimes, with positive probability, there will be multiple reconstructions.

Define $Z_{i,j,a} = \mathbb{1}_{\mathbf{x}_i(a) = \mathbf{x}_j(a)}$ for $i < j$; i.e., $Z_{i,j,a}$ indicates whether \mathbf{x}_i and \mathbf{x}_j share a k -mer at a common position a . Let $V = \sum_{i=1}^{m-1} \sum_{j=i+1}^m \sum_{a=1}^{n-k+1} Z_{i,j,a}$, and note that $\mathcal{D} = \{V > 0\}$. In particular, since any $Z_{i,j,a} = 1$ implies that there are multiple valid reconstructions, we have that $\Pr(\mathcal{E}) \geq \Pr(V > 0)$. By the Paley-Zygmund inequality, we then have

$$\Pr(\mathcal{E}) \geq \Pr(V > 0) \geq \frac{\mathbb{E}[V]^2}{\mathbb{E}[V^2]}. \quad (9)$$

It is straightforward to see $\mathbb{E}[Z_{i,j,a}] = n^{-\beta}$, and by the linearity of expectation,

$$\mathbb{E}[V] = \binom{m}{2} (n - k + 1) n^{-\beta}.$$

The computation of $\mathbb{E}[V^2]$ requires more care in handling $\Pr(\mathbf{x}_i(a) = \mathbf{x}_j(a), \mathbf{x}_s(b) = \mathbf{x}_t(b))$. First, we note that when

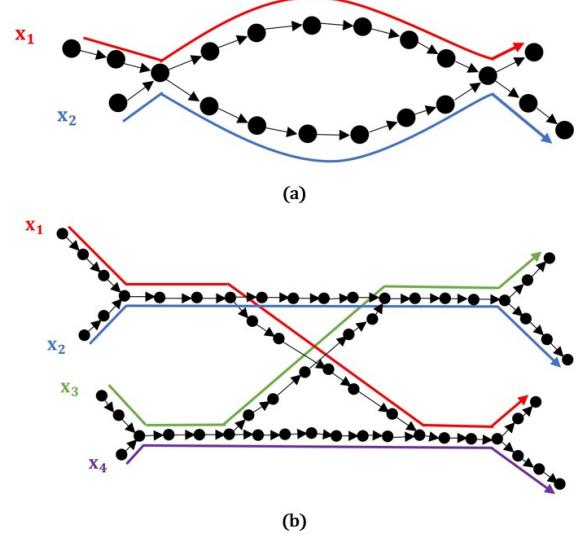


Fig. 2. Two examples of difference de Bruijn graphs \tilde{G} . (a) Example \tilde{G} for $c = 2$. The red and blue paths represent the correct strings, while the other reconstruction follows the paths red-blue-red and blue-red-blue. (b) Example \tilde{G} for $c = 4$. The red, blue, green, and purple paths represent the correct strings, while a possible incorrect reconstruction follows the paths red-blue-green, blue-red-purple, green-purple-red, purple-green-blue.

the pairs (i, j) and (s, t) are disjoint, the events $\mathbf{x}_i(a) = \mathbf{x}_j(a)$ and $\mathbf{x}_s(b) = \mathbf{x}_t(b)$ are independent:

$$\Pr(\mathbf{x}_i(a) = \mathbf{x}_j(a), \mathbf{x}_s(b) = \mathbf{x}_t(b)) = n^{-2\beta}.$$

The events are also independent when $i = s$ but $j \neq t$, or $j = t$ but $i \neq s$,

$$\Pr(\mathbf{x}_i(a) = \mathbf{x}_j(a), \mathbf{x}_i(b) = \mathbf{x}_t(b)) = n^{-2\beta},$$

and when $i = s$, $j = t$, and $|a - b| \geq k$. However, when $i = s$, $j = t$, and $|a - b| < k$,

$$\Pr(\mathbf{x}_i(a) = \mathbf{x}_j(a), \mathbf{x}_i(b) = \mathbf{x}_j(b)) = 2^{-(k+|a-b|)} \leq n^{-\beta}$$

as there are $|a - b|$ symbols of overlap between $\mathbf{x}_i(a)$ and $\mathbf{x}_j(b)$. Therefore,

$$\begin{aligned} \mathbb{E}[V^2] &= \mathbb{E} \left[\sum_{i=1}^{m-1} \sum_{j=i+1}^m \sum_{a=1}^{n-k+1} Z_{i,j,a} \sum_{s=1}^{m-1} \sum_{t=s+1}^m \sum_{b=1}^{n-k+1} Z_{s,t,b} \right] \\ &\leq m^4 n^{2-2\beta} + 2m^2 n^{1-\beta} k. \end{aligned} \quad (10)$$

Combining $\mathbb{E}[V]$ and (10), we have that

$$\Pr(\mathcal{E}) \geq \frac{\binom{m}{2}^2 (n - k + 1)^2 n^{-2\beta}}{m^4 n^{2-2\beta} + 2m^2 n^{1-\beta} k} \quad (11)$$

$$\sim \frac{m^4 n^2 n^{-2\beta}}{4m^4 n^{2-2\beta} + 8m^2 n^{1-\beta} k} \quad (12)$$

$$= \frac{1}{4 + \frac{8\beta \log n}{n^{2\alpha+1-\beta}}} \quad (13)$$

where (12) is asymptotically equivalent as $n \rightarrow \infty$. Finally, we notice that if $\beta < 2\alpha + 1$, (13) converges to $\frac{1}{4}$, implying that $\Pr(\mathcal{E})$ is bounded away from zero and $(\alpha, \beta) \notin \mathcal{F}$.

Next, we show that $\Pr(\mathcal{E}) \not\rightarrow 0$ when $\beta < \alpha + \frac{3}{2}$. For this regime, we will define the event \mathcal{H} that there are two strings $\mathbf{x}_i, \mathbf{x}_j$ and indices a, b, c such that

$$\mathbf{x}_i(a) = \mathbf{x}_j(c), \mathbf{x}_i(b) = \mathbf{x}_j(c + b - a).$$

This is illustrated in Figure 2(a). Essentially, we have two repeated k -mers across sequences \mathbf{x}_i and \mathbf{x}_j , and the gap between them is the same in both sequences. Given \mathcal{H} , an incorrect reconstruction exists because we can swap the middle parts of \mathbf{x}_i and \mathbf{x}_j , creating

$$\begin{aligned}\tilde{\mathbf{x}}_i &= \mathbf{x}_i[1 : a - 1]\mathbf{x}_j[c : c + b - a - 1]\mathbf{x}_i[b : n] \\ \tilde{\mathbf{x}}_j &= \mathbf{x}_j[1 : c - 1]\mathbf{x}_i[a : b - 1]\mathbf{x}_j[c + b - a : n],\end{aligned}$$

which can be verified to be two length- n sequences. We will show that $\Pr(\mathcal{H}) \not\rightarrow 0$ as $n \rightarrow \infty$.

Define $W_{i,j,\mathbf{a}} = \mathbb{1}_{\mathbf{x}_i(a_1) = \mathbf{x}_j(a_3), \mathbf{x}_i(a_2) = \mathbf{x}_j(a_3 + a_2 - a_1)}$ for $\mathbf{a} = (a_1, a_2, a_3)$ and let $U = \sum_{i,j} \sum_{\mathbf{a}} W_{i,j,\mathbf{a}}$. We can again use the Paley-Zygmund inequality to find a lower bound on $\Pr(U > 0)$. The second moment $\mathbb{E}[U^2]$ can be computed similarly to $\mathbb{E}[V^2]$, considering separately the terms of the form $\mathbb{E}[W_{ija}W_{ijb}]$ where the k -mers indicated by \mathbf{a} overlap with those indicated by \mathbf{b} :

$$\begin{aligned}\mathbb{E}[U^2] &= \mathbb{E} \left[\sum_{i=1}^{m-1} \sum_{j=i+1}^m \sum_{\mathbf{a}} W_{i,j,\mathbf{a}} \sum_{s=1}^{m-1} \sum_{t=s+1}^m \sum_{\mathbf{b}} W_{s,t,\mathbf{b}} \right] \\ &\leq (m^2 n^3 n^{-2\beta})^2 + m^2 n^3 (2k)^3 n^{-2\beta}\end{aligned}\quad (14)$$

where the first term in (14) covers all k -mers that are independent of each other and the second term covers the case where $i = s, j = t$, and a_i overlaps with b_i for $i = 1, 2, 3$.

To compute the first moment, $\mathbb{E}[U]$, we must account for which assignments of \mathbf{a} are possible. More precisely, we need to count the vectors $(a_1, a_2, a_3) \in [1 : n']^3$ such that $a_1 < a_2$ and $a_2 + a_3 - a_1 \in [1 : n']$, where $n' = n - k + 1$. This is cumbersome, so we consider a lower bound by considering $a_1 \in [1 : n'/4]$, $a_2 \in [n'/4 + 1 : n'/2]$, and $a_3 \in [n'/4 + 1 : n'/2]$. Notice that for any of the $(n'/4)^3$ choices, $a_2 + a_3 - a_1 \in [1 : n']$. We obtain the (rather loose) lower bound

$$\mathbb{E}[U] \geq \binom{m}{2} \left(\frac{n'}{4} \right)^3 n^{-2\beta}. \quad (15)$$

Therefore, we can lower bound the error probability as

$$\begin{aligned}\Pr(\mathcal{E}) &\geq \frac{\mathbb{E}[U]^2}{\mathbb{E}[U^2]} \geq \frac{\binom{m}{2}^2 (n'/4)^6 n^{-4\beta}}{m^4 n^6 n^{-6\beta} + m^2 n^3 (2k)^3 n^{-2\beta}} \\ &\sim \frac{4^{-7} m^4 n^{6-4\beta}}{m^4 n^{6-4\beta} + 8m^2 k^3 n^{3-2\beta}} \\ &= \frac{4^{-7}}{1 + \frac{8k^3}{n^{2\alpha+3-2\beta}}},\end{aligned}$$

which tends to 4^{-7} as $n \rightarrow \infty$, as long as $\beta < \alpha + \frac{3}{2}$. We conclude that for any (α, β) such that $\beta < \alpha + \frac{3}{2}$ or $\beta < 2\alpha + 1$, $\Pr(\mathcal{E}) \not\rightarrow 0$ as $n \rightarrow \infty$. This concludes the proof of the outer bound in Theorem 1.

VI. CONCLUDING REMARKS

In this paper, we have characterized the feasibility of a multiple sequence reconstruction problem using k -mers for a large region of (α, β) parameters. We identified a region where the k -mers of each string are disjoint with high probability and correct recovery of the source sequences is trivial, as well as a region where, while repeats exist, the only possible reconstruction of the source sequences is the correct one. In the latter repeat-abundant region, careful examination of the de Bruijn graph $G(\mathcal{Y})$ reveals the unique set of paths covering the graph. Finally, we have determined a region where incorrect reconstructions exist with positive probability.

The above findings have characterized nearly all (α, β) pairs for $\alpha > 0, \beta > 0$, except for a small region between $\beta = \max(2\alpha + 1, \alpha + \frac{3}{2})$ and $\beta = \alpha + 2$. We believe that part of the remaining region may be where the solution is not unique, and close analysis of additional difference de Bruijn graph topologies for a larger number of strings and shared k -mers might result in a larger red region in Figure 1.

This work invites future study in several directions. One of these is generalization to other source distributions; some results here may be immediately trivially extended to a general i.i.d. source. Additionally, motivated by real-world DNA sequencing technologies, future work may analyze source reconstruction from noisy k -mers or a random number of copies of each k -mer.

ACKNOWLEDGEMENTS

The work of I.S. was supported in part by the National Science Foundation under CCF grants 2007597 and 2046991. The authors would like to thank the anonymous reviewers of this manuscript for their detailed feedback and suggestions for improvement.

REFERENCES

- [1] R. Dramanac, I. Labat, I. Brulkner, and R. Crkvenjakov, “Sequencing of megabase plus DNA by hybridization: theory of the method,” *Genomics*, vol. 4, no. 2, pp. 114–128, 1989.
- [2] E. Ukkonen, “Approximate string-matching with q-grams and maximal matches,” *Theoretical computer science*, vol. 92, no. 1, pp. 191–211, 1992.
- [3] P. A. Pevzner, “l-tuple DNA sequencing: computer analysis,” *Journal of Biomolecular structure and dynamics*, vol. 7, no. 1, pp. 63–73, 1989.
- [4] P. A. Pevzner, H. Tang, and M. S. Waterman, “An eulerian path approach to DNA fragment assembly,” *Proceedings of the national academy of sciences*, vol. 98, no. 17, pp. 9748–9753, 2001.
- [5] S. S. Skiena and G. Sundaram, “Reconstructing strings from substrings,” *Journal of Computational Biology*, vol. 2, no. 2, pp. 333–353, 1995.
- [6] F. P. Preparata and E. Upfal, “Sequencing-by-hybridization at the information-theory bound: an optimal algorithm,” in *Proceedings of the fourth annual international conference on Computational molecular biology*, 2000, pp. 245–253.
- [7] A. S. Motahari, G. Bresler, and N. David, “Information theory of DNA shotgun sequencing,” *IEEE Transactions on Information Theory*, vol. 59, no. 10, pp. 6273–6289, 2013.
- [8] S. Mohajer, A. Motahari, and D. Tse, “Reference-based DNA shotgun sequencing: Information theoretic limits,” in *2013 IEEE International Symposium on Information Theory*. IEEE, 2013, pp. 1635–1639.
- [9] A. Motahari, K. Ramchandran, D. Tse, and N. Ma, “Optimal DNA shotgun sequencing: Noisy reads are as good as noiseless reads,” in *2013 IEEE International Symposium on Information Theory*. IEEE, 2013, pp. 1640–1644.

- [10] I. Shomorony, T. Courtade, and D. Tse, “Do read errors matter for genome assembly?” in *2015 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2015, pp. 919–923.
- [11] Y. Yehezkeally, S. Marcovich, and E. Yaakobi, “Multi-strand reconstruction from substrings,” in *2021 IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–6.
- [12] O. Elishco, R. Gabrys, E. Yaakobi, and M. Médard, “Repeat-free codes,” *IEEE Transactions on Information Theory*, vol. 67, no. 9, pp. 5749–5764, 2021.
- [13] Y. Yehezkeally and N. Polyanskii, “On codes for the noisy substring channel,” in *2021 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2021, pp. 1700–1705.
- [14] S. Marcovich and E. Yaakobi, “Reconstruction of strings from their substrings spectrum,” *IEEE Transactions on Information Theory*, vol. 67, no. 7, pp. 4369–4384, 2021.
- [15] Z. Chang, J. Chrisnata, M. F. Ezerman, and H. M. Kiah, “Rates of DNA sequence profiles for practical values of read lengths,” *IEEE Transactions on Information Theory*, vol. 63, no. 11, pp. 7166–7177, 2017.
- [16] R. Gabrys and O. Milenkovic, “Unique reconstruction of coded sequences from multiset substring spectra,” in *2018 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2018, pp. 2540–2544.
- [17] K. Levick and I. Shomorony, “Fundamental limits of multiple sequence reconstruction from substrings,” 2023, *arXiv:2305.05820*. [Online]. Available: <http://arxiv.org/abs/2305.05820>.