Optimal Mean-Based Algorithms for Trace Reconstruction*

Anindya De[†]
Northwestern University
Evanston, IL 60208 U.S.A.
anindya@eecs.northwestern.edu

Ryan O'Donnell[‡]
Carnegie Mellon University
Pittsburgh, PA 15213 U.S.A.
odonnell@cs.cmu.edu

Rocco A. Servedio[§]
Columbia University
New York, NY 10027 U.S.A.
rocco@cs.columbia.edu

ABSTRACT

In the (*deletion-channel*) *trace reconstruction problem*, there is an unknown n -bit *source string* x . An algorithm is given access to independent *traces* of x , where a trace is formed by deleting each bit of x independently with probability δ . The goal of the algorithm is to recover x exactly (with high probability), while minimizing samples (number of traces) and running time.

Previously, the best known algorithm for the trace reconstruction problem was due to Holenstein et al. [SODA 2008]; it uses $\exp(\tilde{O}(n^{1/2}))$ samples and running time for any fixed $0 < \delta < 1$. It is also what we call a “mean-based algorithm”, meaning that it only uses the empirical means of the individual bits of the traces. Holenstein et al. also gave a lower bound, showing that any mean-based algorithm must use at least $n^{\tilde{\Omega}(\log n)}$ samples.

In this paper we improve both of these results, obtaining matching upper and lower bounds for mean-based trace reconstruction. For any constant deletion rate $0 < \delta < 1$, we give a mean-based algorithm that uses $\exp(O(n^{1/3}))$ time and traces; we also prove that any mean-based algorithm must use at least $\exp(\Omega(n^{1/3}))$ traces. In fact, we obtain matching upper and lower bounds even for δ subconstant and $\rho := 1 - \delta$ subconstant: when $(\log^3 n)/n \ll \delta \leq 1/2$ the bound is $\exp(-\Theta(\delta n)^{1/3})$, and when $1/\sqrt{n} \ll \rho \leq 1/2$ the bound is $\exp(-\Theta(n/\rho)^{1/3})$.

Our proofs involve estimates for the maxima of Littlewood polynomials on complex disks. We show that these techniques can also be used to perform trace reconstruction with random insertions and bit-flips in addition to deletions. We also find a surprising result: for deletion probabilities $\delta > 1/2$, the presence of insertions can actually *help* with trace reconstruction.

CCS CONCEPTS

• **Theory of computation** $\rightarrow$ *Algorithm design techniques*;

KEYWORDS

Trace Reconstruction, deletion channel, Littlewood polynomials.

*A full version of the paper is available at [DOS16].

[†]Supported by start-up grant from Northwestern University.

[‡]Supported by NSF grant CCF-1618679.

[§]Supported by NSF grants CCF-1420349 and CCF-1563155.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC'17, Montreal, Canada

© 2017 ACM. 978-1-4503-4528-6/17/06...\$15.00

DOI: 10.1145/3055399.3055450

ACM Reference format:

Anindya De, Ryan O'Donnell, and Rocco A. Servedio. 2017. Optimal Mean-Based Algorithms for Trace Reconstruction. In *Proceedings of 49th Annual ACM SIGACT Symposium on the Theory of Computing, Montreal, Canada, June 2017 (STOC'17)*, 10 pages. DOI: 10.1145/3055399.3055450

1 INTRODUCTION

Consider a setting in which a string x of length n over an alphabet Σ is passed through a *deletion channel* that independently deletes each coordinate of x with probability δ . The resulting string, of length somewhere between 0 and n , is referred to as a *trace* of x , or as a *received string*; the original string x is referred to as the *source string*. The *trace reconstruction problem* is the task of reconstructing x (with high probability) given access to independent traces of x . This is a natural and well-studied problem, dating back to the early 2000's [Lev01b, Lev01a, BKKM04], with some combinatorial variants dating even to the early 1970's [Kal73]. However, perhaps surprisingly, much remains to be discovered both about the information-theoretic and algorithmic complexity of this problem. Indeed, in a 2009 survey [Mit09, Section 7], Mitzenmacher wrote that “the study of [trace reconstruction] is still in its infancy”.

Before discussing previous work, we briefly explain why one can assume a binary alphabet without loss of generality. In case of a general Σ , drawing $O(\frac{\log n}{1-\delta})$ traces will with high probability reveal the entire alphabet $\Sigma' \subseteq \Sigma$ of symbols that are present in x . For each symbol $\sigma \in \Sigma'$ we may consider the binary string $x|_{\sigma}$ whose i -th character is 1 iff $x_i = \sigma$; a trace of x is easily converted into a trace of $x|_{\sigma}$, so the trace reconstruction problem for x can be solved by solving the binary trace reconstruction problem for each $x|_{\sigma}$ and combining the results in the obvious way. For this reason, our work (and most previous work) focuses on the case of a binary alphabet.

1.1 Prior Work

As described in [Mit09], the trace reconstruction problem can arise in several natural domains, including sensor networks and biology. However, the apparent difficulty of the problem means that there is not too much published work, at least on the problem of “worst-case” trace reconstruction problem (“worst-case” in the sense that the source string may be any element of $\{0, 1\}^n$). Because of this, several prior authors have considered an “average-case” version of the problem in which the source string is assumed to be uniformly random over $\{0, 1\}^n$ and the algorithm is required to succeed with high probability over the random draw of the traces and over the uniform random choice of x . This average-case problem seems to have first been studied by Batu et al. [BKKM04], who showed

that a simple efficient algorithm which they call Bitwise Majority Alignment succeeds with high probability for sufficiently small deletion rates $\delta = O(1/\log n)$ using only $O(\log n)$ traces. Subsequent work of Kannan and McGregor [KM05] gave an algorithm for random x that can handle both deletions and insertions (both at rates $O(1/\log^2 n)$ as well as bit-flips (with constant probability bounded away from $1/2$) using $O(\log n)$ traces. Viswanathan and Swaminathan [VS08] sharpened this result by improving the deletion and insertion rates that can be handled to $O(1/\log n)$. Finally, [HMPW08] gave a $\text{poly}(n)$ -time, $\text{poly}(n)$ -trace algorithm for random x that succeeds with high probability for any deletion rate δ that is at most some sufficiently small absolute constant.

Several researchers have considered, from an information-theoretic rather than algorithmic perspective, various reconstruction problems that are closely related to the (worst-case) trace reconstruction problem. Kalashnik [Kal73] showed that any n -bit string is uniquely specified by its k -deck, which is the multiset of all its length- k subsequences, when $k = \lfloor n/2 \rfloor$; this result was later reproved by Manvel et al. [MMS⁺91]. Scott [Sco97] subsequently showed that $k = (1 + o(1))\sqrt{n \log n}$ suffices for reconstruction from the k -deck for any x , and simultaneously and independently Krasnikov and Roditty [KR97] showed that $k = \lfloor \frac{16}{7} \sqrt{n} \rfloor + 5$ suffices. (McGregor et al. observed in [MPV14] that the result of [Sco97] yields an information-theoretic algorithm using $\exp(\tilde{O}(n^{1/2}))$ traces for any deletion rate $\delta \leq 1 - O(\sqrt{\log(n)/n})$, but did not discuss the running time of such an algorithm.) On the other side, successively larger $\Omega(\log n)$ lower bounds on the value of k that suffices for reconstruction of an arbitrary $x \in \{0, 1\}^n$ from its k -deck were given by Manvel et al. [MMS⁺91] and Choffrut and Karhumäki [CK97], culminating in a lower bound of $2^{\Omega(\sqrt{\log n})}$ due to Dudík and Schulman [DS03].

Surprisingly few algorithms have been given for the worst-case trace reconstruction problem as defined in the first paragraph of this paper. Batu et al. [BKKM04] showed that a variation of their Bitwise Majority Alignment algorithm succeeds efficiently using $O(n \log n)$ traces if the deletion rate δ is quite low, at most $O(1/n^{1/2+\epsilon})$. Holenstein et al. [HMPW08] gave a “mean-based” algorithm (we explain precisely what is meant by such an algorithm later) that runs in time $\exp(\tilde{O}(\sqrt{n}))$ and uses $\exp(\tilde{O}(\sqrt{n}))$ traces for any deletion rate δ that is bounded away from 1 by a constant; this is the prior work that is most relevant to our main positive result. [HMPW08] also gave a lower bound showing that for any δ bounded away from 0 by a constant, at least $n^{\Omega(\frac{\log n}{\log \log n})}$ traces are required for any mean-based algorithm. Since the result of [HMPW08], several researchers (such as [Mos13]) have raised the question of finding (potentially inefficient) algorithms which have a better sample complexity; however, no progress had been made until this work.

One may also ask (as was done in the “open questions” of [Mit09, Section 7]) for trace reconstruction for more general channels, such as those that allow deletions, insertions, and bit-flips. The only work we are aware of along these lines is that of Andoni et al. [ADHR12], which gives results for trace reconstruction for *average-case* words in the presence of insertions, deletions, and substitutions on a tree.

1.2 Our Results

THEOREM 1.1 (DELETION CHANNEL POSITIVE RESULT). *There is an algorithm for the trace reconstruction problem which, for any constant $0 < \delta < 1$, uses $\exp(O(n^{1/3}))$ traces and running time.*

Theorem 1.1 significantly improves the running time and sample complexity of the [HMPW08] algorithm, which is $\exp(\tilde{O}(n^{1/2}))$ for fixed constant δ . Furthermore, we can actually extend Theorem 1.1 to the case of $\delta = o(1)$ or $\delta = 1 - o(1)$; see Theorem 1.3 below.

The algorithm of Theorem 1.1 is a “mean-based” algorithm, meaning that it uses only the empirical mean of the trace vectors it receives. We prove an essentially matching lower bound for such algorithms:

THEOREM 1.2 (DELETION CHANNEL NEGATIVE RESULT). *For any constant $0 < \delta < 1$, every mean-based algorithm must use at least $\exp(\Omega(n^{1/3}))$ traces.*

As mentioned, we can also treat $\delta = o(1)$ and $\delta = 1 - o(1)$:

THEOREM 1.3 (DELETION CHANNEL GENERAL MATCHING BOUNDS). *The matching bounds in Theorems 1.1 and 1.2 extend as follows: For $O(\log^3 n)/n \leq \delta \leq 1/2$, the matching bound is $\exp(\Theta(\delta n)^{1/3})$ (and for any smaller δ we have a $\text{poly}(n)$ upper bound). Writing $\rho = 1 - \delta$ for the “retention” probability, for $O(1/n^{1/2}) \leq \rho \leq 1/2$ the matching bound is $\exp(\Theta(n/\rho)^{1/3})$.*

For simplicity in the main portion of the paper we consider only the deletion channel and prove the above results. In Appendix A we consider a more general channel that allows for deletions, insertions, and bit-flips, and prove the following result, which extends Theorem 1.1 to that more general channel and includes Theorem 1.1 as a special case.

THEOREM 1.4 (GENERAL CHANNEL POSITIVE RESULT). *Let C be the general channel described in Section A.1 with deletion probability $\delta = 1 - \rho$, insertion probability σ , and bit-flip probability $\gamma/2$. Define*

$$r := \frac{\rho + \delta\sigma}{1 + \sigma}.$$

Then there is an algorithm for C -channel trace reconstruction using samples and running time bounded by

$$\text{poly}\left(\frac{1}{1-\delta}, \frac{1}{1-\sigma}, \frac{1}{1-\gamma}\right) \cdot \begin{cases} \exp(O(n/r)^{1/3}) & \text{if } \frac{C}{n^{1/2}} \leq r \leq \frac{1}{2}, \\ \exp(O((1-r)n)^{1/3}) & \text{if } \frac{O(\log^3 n)}{n} \leq 1-r \leq \frac{1}{2}. \end{cases}$$

Since some slight technical and notational unwieldiness is incurred by dealing with the more general channel, we defer the proof of Theorem 1.4 to Appendix A; however, we note here that the main core of the proof is unchanged from the deletion-only case. We additionally note that, as discussed in Appendix A, a curious aspect of the upper bound given by Theorem 1.4 is that having a constant insertion rate can make it possible to perform trace reconstruction in time $\exp(O(n^{1/3}))$ even when the deletion rate is much higher than Theorem 1.3 could handle in the absence of insertions. A possible intuitive explanation for this is that having random insertions could serve to “smooth out” worst-case instances that are problematic for a deletion-only model.

1.3 Independent and Concurrent Work

Fedor Nazarov and Yuval Peres have independently obtained results that are substantially similar to Theorems 1.1 and 1.2 (see these proceedings). Also, Elchanan Mossel has informed us [Mos] that around 2008, Mark Braverman, Avinatan Hassidim and Elchanan Mossel had independently proven (unpublished) superpolynomial lower bounds for mean-based algorithms.

1.4 Our Techniques

For simplicity of discussion, we restrict our focus in this section to the question of upper bounding the sample complexity of trace reconstruction for the deletion channel, where every bit gets deleted independently with probability δ . (As discussed above, generalizing the results to channels which also allow for insertions and flips is essentially a technical exercise that does not require substantially new ideas.) As we discuss in Section 3.2, an efficient *algorithm* follows easily from a sample complexity upper bound via the observation that the minimization problem whose solution yields a sample complexity upper bound, extends to a slightly larger *convex* set, and thus one can use convex (in fact, linear) programming to get an algorithmic result. Hence the technical meat of the argument lies in upper bounding the sample complexity.

The key enabling idea for our work is to take an analytic view on the combinatorial process defined by the deletion channel. More precisely, consider two distinct strings $x, x' \in \{-1, 1\}^n$. A necessary (and sufficient) condition to upper bound the sample complexity of trace reconstruction is to lower bound the statistical distance between the two distributions of traces of x versus x' (let us write $C(x)$ and $C(x')$ to denote these two distributions). Since analyzing the statistical distance $d_{TV}(C(x), C(x'))$ between the distributions $C(x)$ and $C(x')$ turns out to be a difficult task, we approach it by considering a limited class of statistical tests.

In [HMPW08] the authors consider “mean-based” algorithms; such algorithms correspond to statistical tests that only use 1-bit marginals of the distribution of the received string. More precisely, for any $1 \leq j \leq n$, consider the quantities $\Pr_{\mathbf{y} \leftarrow C(x)}[\mathbf{y}_j = 1]$ and $\Pr_{\mathbf{y}' \leftarrow C(x')}[\mathbf{y}'_j = 1]$. The difference $|\Pr_{\mathbf{y} \leftarrow C(x)}[\mathbf{y}_j = 1] - \Pr_{\mathbf{y}' \leftarrow C(x')}[\mathbf{y}'_j = 1]|$ is a lower bound on $d_{TV}(C(x), C(x'))$.

Let us define the vector $\beta_{x,x'} = (\beta_{x,x'}(1), \dots, \beta_{x,x'}(n)) \in [-1, 1]^n$ by

$$\beta_{x,x'}(j) = \Pr_{\mathbf{y} \leftarrow C(x)}[\mathbf{y}'_j = 1] - \Pr_{\mathbf{y}' \leftarrow C(x')}[\mathbf{y}_j = 1].$$

In this terminology, giving a sample complexity upper bound on mean-based algorithms correspond to showing a lower bound on $\min_{x \neq x' \in \{-1, 1\}^n} \|\beta_{x,x'}\|_1$. A central idea in this paper is to analyze $\|\beta_{x,x'}\|_1$ by studying the Z -transform of the vector $\beta_{x,x'}$. More precisely, for $z \in \mathbb{C}$, we consider $\widehat{\beta}_{x,x'}(z) := \sum_{j=1}^n \beta_{x,x'}(j) \cdot z^{j-1}$. Elementary complex analysis can be used to show that

$$\sup_{|z|=1} |\widehat{\beta}_{x,x'}(z)| \leq \|\beta_{x,x'}\|_1 \leq \sqrt{n} \cdot \sup_{|z|=1} |\widehat{\beta}_{x,x'}(z)|.$$

Thus, for our purposes, it suffices to study $\sup_{|z|=1} |\widehat{\beta}_{x,x'}(z)|$. By analyzing the deletion channel and observing that $\widehat{\beta}_{x,x'}(z)$ is a polynomial in z , we are able to characterize this supremum as the supremum of a certain polynomial (induced by x and x') on a certain disk in the complex plane. Thus giving a sample complexity

upper bound amounts to lower bounding $\sup_{|z|=1} |\widehat{\beta}_{x,x'}(z)|$ across all polynomials $\widehat{\beta}_{x,x'}$ induced by distinct $x, x' \in \{-1, 1\}^n$ (essentially, across a class of polynomials closely related to *Littlewood polynomials*: those polynomials with all coefficients in $\{-1, 0, 1\}$). The technical heart of our sample complexity upper bound is in establishing such a lower bound. Finally, similar ideas and arguments are used to lower bound the sample complexity of mean-based algorithms, by upper bounding $\sup_{|z|=1} |\widehat{\beta}_{x,x'}(z)|$ across all polynomials $\widehat{\beta}_{x,x'}$ induced by distinct $x, x' \in \{-1, 1\}^n$.

2 PRELIMINARIES AND TERMINOLOGY

Throughout this paper we will use two slightly nonstandard notational conventions. Bits will be written as $\{-1, 1\}$ rather than $\{0, 1\}$, and strings will be indexed starting from 0 rather than 1. Thus the *source string* will be denoted $x = (x_0, x_1, \dots, x_{n-1}) \in \{-1, 1\}^n$; this is the unknown string that the reconstruction algorithm is trying to recover.

We will write C for the *channel* through which x is transmitted. In the main body of the paper our main focus will be on the *deletion channel* $C = \text{Del}_\delta$, in which each bit of x is independently deleted with probability $\delta < 1$. We will also often consider $\rho = 1 - \delta > 0$, the retention probability of each coordinate. In Appendix A we will see that a more general channel that also involves *insertions* and *bit-flips* can be handled in a similar way.

We will use boldface to denote random variables. We typically write $\mathbf{y} \leftarrow C(x)$ to denote that $\mathbf{y} = (\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{n-1})$ is a random *trace* (or *received string* or *sample*), obtained by passing x through the channel C . Notice the slight inconvenience that the length of $\mathbf{y}$ is a random variable (for the deletion channel this length is always between 0 and n); we denote this length by n .

We define a *trace reconstruction algorithm* for channel C to be an algorithm with the following property: for any unknown source string $x \in \{-1, 1\}^n$, when given access to independent strings $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots$ each distributed according to $C(x)$, it outputs x with probability at least (say) 99%. The *sample complexity* of the trace reconstruction algorithm is the number of draws from $C(x)$ that it uses (in the worst case across all $x \in \{-1, 1\}^n$ and all draws from $C(x)$). We are also interested in the algorithm’s (worst-case) running time.

As mentioned earlier we will use basic complex analysis. The following notation will be useful:

Notation 2.1. We write $D_r(c)$ for the closed complex disk of radius r centered at c ; i.e., $\{z \in \mathbb{C} : |z - c| \leq r\}$. We write $\partial D_r(c)$ for the boundary of this disk; thus, e.g., $\partial D_1(0) = \{z \in \mathbb{C} : |z| = 1\}$ is the complex unit circle.

3 MEAN TRACES

We now come to a key definition, that of the *mean trace*. For now we restrict our focus to C being the deletion channel Del_δ (we consider a more general channel in Appendix A).

Although a random trace $\mathbf{y} \leftarrow \text{Del}_\delta(x)$ does not have a fixed length, we can simply define the mean trace of a source string $x \in \{-1, 1\}^n$ to be

$$\mu_{\text{Del}_\delta}(x) = \mathbb{E}_{\mathbf{y} \leftarrow \text{Del}_\delta(x)}[\mathbf{y}'] \in [-1, 1]^n, \quad (1)$$

where y' is y padded with zeros so as to be of length exactly n . Here “0” has a natural interpretation as a “uniformly random bit” (indeed, a trace reconstruction algorithm could always pad deletion-channel traces with random bits by itself, and this would not change the definition of the mean trace $\mu_{\text{Del}_\delta}(x)$).

The following is immediate:

PROPOSITION 3.1. *Viewing the domain of μ_{Del_δ} as the real vector space $\mathbb{R}^n$, $\mu_{\text{Del}_\delta}(x)$ is a (real-)linear function of x ; that is, each $\mu_{\text{Del}_\delta}(x)_j$ can be written as $\sum_i a_{i,j} x_i$ for some constants $a_{i,j} \in \mathbb{R}$.*

3.1 The Mean-Based (Deletion-Channel) Trace Reconstruction Model

One of the most basic things that a trace reconstruction algorithm can do is calculate an empirical estimate of the mean trace. A simple Chernoff/union bound shows that, with $\text{poly}(n/\epsilon)$ samples and time, an algorithm can compute an estimator $\hat{\mu}_{\text{Del}_\delta}(x) \in [-1, 1]^n$ satisfying $\|\hat{\mu}_{\text{Del}_\delta}(x) - \mu_{\text{Del}_\delta}(x)\|_1 \leq \epsilon$ with very high probability. The algorithm might then proceed to base its reconstruction solely on $\hat{\mu}_{\text{Del}_\delta}(x)$, without relying on further traces. We call such algorithms “mean-based trace reconstruction algorithms” (Holenstein et al. [HMPW08] called them algorithms based on “summary statistics”). We give a formal definition:

Definition 3.2. An algorithm in the *mean-based (deletion-channel) trace reconstruction model* works as follows. Given an unknown source string $x \in \{-1, 1\}^n$, the algorithm first specifies a parameter $T \in \mathbb{N}$. The algorithm is then given an estimate $\hat{\mu}_{\text{Del}_\delta}(x) \in [-1, +1]^n$ of the mean trace satisfying

$$\|\hat{\mu}_{\text{Del}_\delta}(x) - \mu_{\text{Del}_\delta}(x)\|_1 \leq 1/T. \quad (2)$$

We define the “cost” of this portion of the algorithm to be T . Having been given $\hat{\mu}_{\text{Del}_\delta}(x)$, the algorithm has no further access to x , but may do further “postprocessing” computation involving $\hat{\mu}_{\text{Del}_\delta}(x)$. The algorithm should end by outputting x .

From the above discussion, we see that an algorithm in the mean-based trace reconstruction model with cost T_1 and postprocessing time T_2 may be converted into a normal trace reconstruction algorithm using $\text{poly}(n, T_1)$ samples and $\text{poly}(n, T_1) + T_2$ time.

3.2 The Complexity of Mean-Based (Deletion-Channel) Trace Reconstruction

As discussed in [HMPW08], the sample complexity of mean-based trace reconstruction is essentially determined by the minimum distance between the mean traces $\mu_{\text{Del}_\delta}(x)$ and $\mu_{\text{Del}_\delta}(x')$ of two distinct source strings $x, x' \in \{-1, 1\}^n$. Furthermore, one can get an upper bound on the *time* complexity of mean-based trace reconstruction if a certain “fractional relaxation” of this minimum mean trace distance is large. We state these observations from [HMPW08] here, using slightly different notation.

Definition 3.3. Given n and $0 \leq \delta < 1$, we define:

$$\begin{aligned} \epsilon_{\text{Del}_\delta}(n) &:= \min_{\substack{x, x' \in \{-1, 1\}^n \\ x \neq x'}} \|\mu_{\text{Del}_\delta}(x) - \mu_{\text{Del}_\delta}(x')\|_1 \\ &= 2 \min_{\substack{b \in \{-1, 0, +1\}^n \\ b \neq 0}} \|\mu_{\text{Del}_\delta}(b)\|_1; \\ \epsilon_{\text{Del}_\delta}^{\text{frac}}(n) &:= \min_{0 \leq i < n} \min_{\substack{x, x' \in \{-1, +1\}^n \\ x_j = x'_j \in \{-1, 1\} \forall j < i \\ x_i = -x'_i \in \{-1, 1\}}} \|\mu_{\text{Del}_\delta}(x) - \mu_{\text{Del}_\delta}(x')\|_1 \\ &= 2 \min_{d \in [n]} \min_{b \in \{0\}^{d-1} \times \{1\} \times [-1, +1]^{n-d}} \|\mu_{\text{Del}_\delta}(b)\|_1. \end{aligned}$$

In both cases, the equality on the right uses Proposition 3.1.

It's easy to see that in the mean-based trace reconstruction model, it is information-theoretically possible for an algorithm to succeed if and only if its cost T exceeds $2/\epsilon_{\text{Del}_\delta}(n)$. Thus characterizing the sample complexity of mean-based trace reconstruction essentially amounts to analyzing $\epsilon_{\text{Del}_\delta}(n)$. For example, to establish our lower bound Theorem 1.2, it suffices to prove that the $\epsilon_{\text{Del}_\delta}(n) \leq \exp(-\Omega(n^{1/3}))$ for constant $0 < \delta < 1$.

Furthermore, as observed in [HMPW08], given an $\epsilon_{\text{Del}_\delta}^{\text{frac}}(n)/4$ -accurate estimate of $\mu_{\text{Del}_\delta}(x)$, as well as the ability to compute the linear function $\mu_{\text{Del}_\delta}(x')$ for any $x' \in [-1, +1]^n$ (or even estimate it to $\epsilon_{\text{Del}_\delta}^{\text{frac}}(n)/4$ -accuracy), one can recover x exactly in $\text{poly}(n, \log(1/\epsilon_{\text{Del}_\delta}^{\text{frac}}(n)))$ time by solving a sequence of n linear programs.¹ Thus to establish our Theorem 1.1, it suffices to prove that $\epsilon_{\text{Del}_\delta}^{\text{frac}}(n) \geq \exp(-O(n^{1/3}))$ for constant $0 < \delta < 1$.

3.3 Reduction to Complex Analysis

Our next important definition is of a polynomial that encodes the components of $\mu_C(x)$ in its coefficients — kind of a generating function for the channel. We think of its parameter z as a complex number.

Definition 3.4. Given $x \in \{-1, 1\}^n$ and $0 \leq \delta < 1$, we define the *deletion-channel polynomial*

$$P_{\text{Del}_\delta, x}(z) = \sum_{j < n} \mu_{\text{Del}_\delta}(x)_j \cdot z^j,$$

a polynomial of degree less than n . We extend this definition to $x \in [-1, +1]^n$ using the linearity of μ_{Del_δ} .

We now make the step to elementary complex analysis, by relating the size of a mean trace difference $\mu_{\text{Del}_\delta}(b)$ to the maximum modulus of $P_{\text{Del}_\delta, b}(z)$ on the unit complex circle (or equivalently, the unit complex disk, by the Maximum Modulus Principle):

PROPOSITION 3.5. *For any $b \in [-1, 1]^n$, we have*

$$\max_{z \in \partial D_1(0)} |P_{\text{Del}_\delta, b}(z)| \leq \|\mu_{\text{Del}_\delta}(b)\|_1 \leq \sqrt{n} \max_{z \in \partial D_1(0)} |P_{\text{Del}_\delta, b}(z)|.$$

¹If the algorithm “knows” δ it can efficiently compute $\mu_{\text{Del}_\delta}(x')$ exactly. But even if it doesn't “know” δ , it can estimate δ to sufficient accuracy so that $\mu_{\text{Del}_\delta}(x')$ can be estimated to the necessary accuracy, with no significant algorithmic slowdown.

PROOF. Recall that $\mu_{\text{Del}_\delta}(b)$ is the length- n vector of coefficients for the polynomial $P_{\text{Del}_\delta, b}(z)$. The lower bound above is immediate from the triangle inequality. For the upper bound, we use

$$\begin{aligned} \|\mu_{\text{Del}_\delta}(b)\|_1^2 &\leq n \|\mu_{\text{Del}_\delta}(b)\|_2^2 \\ &= n \cdot \text{avg}_{z \in \partial D_1(0)} |P_{\text{Del}_\delta, b}(z)|^2 \\ &\leq n \left(\max_{z \in \partial D_1(0)} |P_{\text{Del}_\delta, b}(z)| \right)^2. \end{aligned}$$

Here the first inequality is Cauchy–Schwarz, the equality is an elementary fact about complex polynomials (or Fourier series), and the final inequality is obvious. $\square$

Let us reconsider Definition 3.3. As a factor of $\sqrt{n}$ is negligible compared to the bounds we will prove (which are of the shape $\exp(-\Theta(n^{1/3}))$), we may as well analyze $\max_{z \in \partial D_1(0)} |P_{\text{Del}_\delta, b}(z)|$ rather than $\|\mu_{\text{Del}_\delta}(b)\|_1$ in the definition of $\epsilon_{\text{Del}_\delta}(n)$ and $\epsilon_{\text{Del}_\delta}^{\text{frac}}(n)$. We therefore take a closer look at the deletion-channel polynomial.

4 THE DELETION-CHANNEL POLYNOMIAL

In this section we compute the deletion-channel polynomial. When the deletion channel is applied to some source string x , each bit x_i is either deleted with probability δ or else is transmitted at some position $j \leq i$ in the received string y . Let us introduce (non-independent) random variables $J_0, \dots, J_{n-1}$, where $J_i = \perp$ if x_i is deleted and otherwise J_i is the position in y at which x_i is transmitted. We thus have

$$\begin{aligned} P_{\text{Del}_\delta, x}(z) &= \sum_{j < n} \mathbf{E}_{y \leftarrow C(x)} [y_j] \cdot z^j \\ &= \sum_{j < n} z^j \cdot \sum_{i < n} \Pr[J_i = j] x_i \\ &= \sum_{i < n} x_i \cdot \sum_{j < n} \Pr[J_i = j] z^j \\ &= \sum_{i < n} x_i \cdot \text{"E"}[z^{J_i}]. \end{aligned}$$

Here we put the expectation $\mathbf{E}$ in quotation marks because the expression should count 0 whenever $J_i = \perp$. Observing that $\Pr[J_i \neq \perp]$ equals the retention probability $\rho = 1 - \delta$, if we define the conditional random variable

$$\tilde{J}_i = (J_i \mid J_i \neq \perp)$$

(so $\tilde{J}_i$ is an $\mathbb{N}$ -valued random variable), then we have

$$P_{\text{Del}_\delta, x}(z) = \rho \sum_{i < n} x_i \cdot \mathbf{E}[z^{\tilde{J}_i}]. \quad (3)$$

Observing that $\tilde{J}_i$ is distributed as Binomial(i, ρ), and letting $B_1, \dots, B_i$ denote independent Bernoulli random variables with “success” probability ρ , we easily compute

$$\mathbf{E}[z^{\tilde{J}_i}] = \mathbf{E}[z^{B_1 + \dots + B_i}] = \mathbf{E}[z^{B_1}]^i = ((1 - \rho) + \rho z)^i.$$

Denoting

$$w = 1 - \rho + \rho z,$$

we conclude that

$$P_{\text{Del}_\delta, x}(z) = \rho \sum_{i < n} x_i w^i.$$

As z ranges over the unit circle $\partial D_1(0)$, w ranges over the radius- ρ circle $\partial D_\rho(1 - \rho)$. Recalling Definition 3.3 and Proposition 3.5, we are led to consider the following two quantities for $0 < \rho < 1$ (note that by the Maximum Modulus Principle, these quantities are unchanged whether the max is over $D_\rho(1 - \rho)$ or $\partial D_\rho(1 - \rho)$):

$$\kappa_{\text{Littlewood}}(\rho, n) = \min \left\{ \max_{w \in D_\rho(1 - \rho)} |P(w)| : P(w) = b_0 + b_1 w + \dots + b_{n-1} w^{n-1}, b_i \in \{0, \pm 1\} \text{ not all } 0 \right\},$$

$$\kappa_{\text{bounded}}^{\text{frac}}(\rho, d) = \min \left\{ \max_{w \in D_\rho(1 - \rho)} |P(w)| : P(w) = w^d + b_{d+1} w^{d+1} + \dots + b_N w^N, N \geq d, b_i \in D_1(0) \right\}.$$

By the Maximum Modulus Principle, both $\kappa_{\text{Littlewood}}(\rho, n)$ and $\kappa_{\text{bounded}}^{\text{frac}}(\rho, d)$ are nondecreasing functions of $0 < \rho < 1$. It's also easy to see that both are nonincreasing functions of their second argument for all $0 < \rho < 1$ (for $\kappa_{\text{bounded}}^{\text{frac}}(\rho, d)$, consider replacing $P(w)$ by $wP(w)$ and observe that $|wP(w)| \leq |P(w)|$ for all $w \in D_\rho(1 - \rho)$). It thus follows that

$$\kappa_{\text{bounded}}^{\text{frac}}(\rho, d) \leq \kappa_{\text{Littlewood}}(\rho, d).$$

Our main technical theorems are the following:

THEOREM 4.1. *There is a universal constant $C \geq 1$ such that:*

$$\begin{aligned} &\text{for } 1/d \leq \delta \leq 1/2, \quad \kappa_{\text{bounded}}^{\text{frac}}(1 - \delta, d) \geq \exp(-C(\delta d)^{1/3}); \\ &\text{for } 1/d^{1/2} \leq \rho \leq 1/2, \quad \kappa_{\text{bounded}}^{\text{frac}}(\rho, d) \geq \exp(-C(d/\rho)^{1/3}). \end{aligned}$$

THEOREM 4.2. *There is a universal constant $C \geq 1$ such that:*

$$\begin{aligned} &\text{for } C(\log^3 n)/n \leq \delta \leq 1/2, \\ &\kappa_{\text{Littlewood}}(1 - \delta, n) \leq \exp(-\Omega(\delta n)^{1/3}); \\ &\text{for } C/n^{1/2} \leq \rho \leq 1/2, \\ &\kappa_{\text{Littlewood}}(\rho, n) \leq \exp(-\Omega(n/\rho)^{1/3}). \end{aligned}$$

By Definition 3.3, Proposition 3.5, and the discussion at the end of Section 3.2, we have that Theorem 4.2 implies both Theorem 1.2 and the more general sample complexity lower bound in Theorem 1.3. Regarding the algorithmic upper bounds in Theorems 1.1 and 1.3, again from Definition 3.3 and Proposition 3.5 we get that

$$\begin{aligned} \epsilon_{\text{Del}_\delta}^{\text{frac}}(n) &\geq 2\rho \cdot \min_{0 \leq d < n} \left\{ \max_{w \in D_\rho(1 - \rho)} |P(w)| : P(w) = w^d + b_{d+1} w^{d+1} + \dots + b_{n-1} w^{n-1}, b_i \in [-1, +1] \right\} \\ &\geq 2\rho \cdot \min_{0 \leq d < n} \kappa_{\text{bounded}}^{\text{frac}}(\rho, d) \geq 2\rho \cdot \kappa_{\text{bounded}}^{\text{frac}}(\rho, n). \end{aligned}$$

Thus the upper bounds Theorems 1.1 and 1.3 likewise follow from Theorem 4.1 and the discussion at the end of Section 3.2. (Note that

if $\delta \leq O(\log^3 n)/n$, we can always pay the bound for the larger value $\delta = \Theta(\log^3 n)/n$, which is $\text{poly}(n)$.)

5 PROOF OF THEOREM 4.1

We will need the following:

THEOREM 5.1. ([BE97], Corollary 3.2, $M = 1$ case.) *Let $Q(w)$ be a polynomial with constant coefficient 1 and all other coefficients bounded by 1 in modulus. Fix any $0 < \theta \leq \pi$, and let A be the arc $\{e^{it} : -\theta \leq t \leq \theta\}$. Then $\sup_{w \in A} |Q(w)| \geq \exp(-C_1/\theta)$ for some universal constant C_1 .*

We remark that for any $0 < r < 1$, Theorem 5.1 holds for the arc $A = \{re^{it} : -\theta \leq t \leq \theta\}$ with no change in the constant C_1 . This is immediate by applying the theorem to $\tilde{Q}(w) = Q(rw)$.

PROOF OF THEOREM 4.1. Fix $d \geq 2$ (else the hypotheses are vacuous) and $\delta + \rho = 1$. We call Case I when $1/d \leq \delta < 1/2$, and we call Case II when $1/d^{1/2} \leq \rho \leq 1/2$. Select

$$\theta = \begin{cases} \frac{1}{2(\delta d)^{1/3}} & \text{in Case I,} \\ \left(\frac{\rho}{d}\right)^{1/3} & \text{in Case II.} \end{cases}$$

In Case I we have $\theta \leq 1/2$, and in Case II we have $\theta \leq \rho \leq 1/2$.

Let $P(w) = w^d \cdot Q(w)$, where $Q(w)$ is a polynomial with constant coefficient 1 and all other coefficients bounded by 1 in modulus. We need to show

$$\max_{w \in D_\rho(\delta)} |P(w)| \geq \begin{cases} \exp(-C(\delta d)^{1/3}) & \text{in Case I,} \\ \exp(-C(d/\rho)^{1/3}) & \text{in Case II.} \end{cases} \quad (4)$$

In Case I, the ray $\{re^{i\theta} : r > 0\}$ intersects $\partial D_\rho(\delta)$ at a unique point, call it w_0 . In Case II, the same ray intersects $D_\rho(\delta)$ twice (this uses $\theta \leq \rho$); call the point of larger modulus w_0 . In either case, consider the triangle formed in the complex plane by the points 0, δ , and w_0 ; it has some acute angle α at w_0 and an angle of θ at 0. By the Law of Sines,

$$\begin{aligned} \frac{\rho}{\sin \theta} &= \frac{\delta}{\sin \alpha} \\ &= \frac{|w_0|}{\sin(\pi - \theta - \alpha)} \\ &= \frac{|w_0|}{\sin(\theta + \alpha)} \\ &= \frac{|w_0|}{\sin \theta \cos \alpha + \sin \alpha \cos \theta}, \end{aligned}$$

which implies that

$$\begin{aligned} |w_0| &= \delta \cos \theta + \rho \cos \alpha \\ &= \delta \cos \theta + \rho \sqrt{1 - \left(\frac{\delta}{\rho}\right)^2 \sin^2 \theta} \\ &\geq \delta(1 - \theta^2) + \rho(1 - \left(\frac{\delta}{\rho}\right)^2 \theta^2) \\ &= 1 - \frac{\delta}{\rho} \theta^2. \end{aligned}$$

(The last inequality used $\theta \leq \rho$ in Case II.) Writing $r_0 = |w_0|$, Theorem 5.1 (and the subsequent remark) implies that

$$\max_{w \in A} |Q(w)| \geq \exp(-C_1/\theta) \quad \text{for } A = \{r_0 e^{it} : -\theta \leq t \leq \theta\} \subset D_\rho(\delta). \quad (5)$$

As a consequence, we have

$$\begin{aligned} \max_{w \in D_\rho(\delta)} |P(w)| &\geq \max_{w \in A} |P(w)| \\ &\geq r_0^d \cdot \exp(-C_1/\theta) \\ &\geq (1 - (\delta/\rho)\theta^2)^d \cdot \exp(-C_1/\theta) \\ &\geq \exp(-2(\delta/\rho)\theta^2 d - C_1/\theta) \end{aligned}$$

(the last inequality again using $\theta \leq \rho$ in Case II). Substituting in the value of θ yields (4). $\square$

5.1 An Improved Version

Although we don't need it for our application, we can actually provide a stronger version of the results in the previous section that is also self-contained — i.e., it does not rely on Borwein and Erdélyi's Theorem 5.1. We used that theorem to establish (5); but more strongly than (5), we can show there exists an arc $A \subset D_\rho(\delta)$ such that

$$\text{GM}_{w \in A} |Q(w)| \geq \exp(-O(1/\theta)),$$

where the left-hand side here denotes the *geometric mean* of $|Q|$ along A . (Of course, this is at most the max of $|Q|$ along A .) To keep the parameters simpler, we will assume $\rho \leq 1/3$ (this is the more interesting parameter regime anyway, and it is sufficient to yield our Theorem 1.1). Our alternate arc A will be

$$A = \{1/3 + re^{it} : -\theta \leq t \leq \theta\},$$

where $0 < r < 2/3$ is the larger real radius such that $1/3 + re^{\pm i\theta} \in \partial D_\rho(\delta)$. We remark that still $A \subset D_\rho(\delta)$, by virtue of $\theta \leq \rho \leq 1/3$, and it is not hard to show that the endpoint of A , call it $w' = 1/3 + re^{i\theta} \in \partial D_\rho(\delta)$, again satisfies $|w'| \geq 1 - \Omega(\frac{\delta}{\rho}\theta^2)$. Thus instead of using Theorem 5.1 as a black box, we could have completed our proof of Theorem 4.1 using the following:

THEOREM 5.2. *Let $Q(w)$ be a polynomial with constant coefficient 1 and all other coefficients in $D_1(0)$. Fix any $0 < \theta \leq \pi$, $0 \leq r \leq 2/3$, and let A be the arc $\{1/3 + re^{it} : -\theta \leq t \leq \theta\}$. Then $\text{GM}_{w \in A} (|Q(w)|) \geq 9/18\pi^{1/\theta}$.*

Our proof will require one standard fact from the theory of “Mahler measures”:

FACT 5.3. *Let Q be a complex polynomial and let O be a circle in the complex plane with center c . Then $\text{GM}_{w \in O} (|Q(w)|) \geq |Q(c)|$.*

PROOF. By a linear transformation we may assume O is the unit circle $\partial D_1(0)$. Express $Q(w) = a_0 \prod_i (w - \alpha_i)$, where the α_i 's are the roots of Q . Then $\text{GM}_{w \in O} (|Q(w)|)$ — known as Q 's *Mahler measure*, see e.g. [Smy08] — is exactly equal to $|a_0| \prod_{i \in I} |\alpha_i|$, where $I = \{i : |\alpha_i| \geq 1\}$. (Since $\text{GM}_{w \in O} (|\cdot|)$ is multiplicative, this statement follows immediately from the elementary fact that $\text{GM}_{w \in O} (|w - \alpha|) = \max\{|\alpha|, 1\}$.) But clearly we have $|a_0| \prod_{i \in I} |\alpha_i| \geq |a_0| \prod_i |\alpha_i| = |Q(0)|$. $\square$

We can now establish Theorem 5.2:

PROOF OF THEOREM 5.2. Using the bounds on Q 's coefficients we have:

$$|Q(w)| \leq 1 + |w| + |w|^2 + \dots = \frac{1}{1 - |w|} \text{ for } w \in D_1(0); \quad (6)$$

$$|Q(1/3)| \geq 1 - |1/3| - |1/3|^2 - \dots = 1/2. \quad (7)$$

Let us apply Fact 5.3 with $O = \partial D_r(1/3) \supset A$, writing A' for the complementary arc to A in O . We get

$$\begin{aligned} 1/2 &\leq \text{GM}_{w \in O}(|Q(w)|) \\ &= \text{GM}_{w \in A}(|Q(w)|)^{\theta/\pi} \cdot \text{GM}_{w \in A'}(|Q(w)|)^{1-\theta/\pi}. \end{aligned} \quad (8)$$

And by (6) we have

$$\begin{aligned} \text{GM}_{w \in A'}(|Q(w)|) &\leq \text{GM}_{w \in A'}\left(\frac{1}{1-|w|}\right) \\ &\leq \text{GM}_{w \in O}\left(\frac{1}{1-|w|}\right) \\ &\leq \text{GM}_{w \in \partial D_{2/3}(1/3)}\left(\frac{1}{1-|w|}\right), \end{aligned} \quad (9)$$

where the second inequality is because the points $w \in A$ only have larger $\frac{1}{1-|w|}$ than the points in A' , and the third inequality is because increasing the radius of O from r to $2/3$ only increases the value of $\frac{1}{1-|w|}$ for points on O . But now for $-\pi < t \leq \pi$, the point $w = 1/3 + (2/3)e^{it} \in D_{2/3}(1/3)$ has $|w|^2 = 1 - \frac{4}{9}(1 - \cos t)$ and hence

$$\frac{1}{1 - |w|} = \frac{1}{1 - \sqrt{1 - \frac{4}{9}(1 - \cos t)}} \leq \frac{9}{2(1 - \cos t)}.$$

Thus

$$\begin{aligned} \text{GM}_{w \in \partial D_{2/3}(1/3)}\left(\frac{1}{1-|w|}\right) &\leq \exp\left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \ln\left(\frac{9}{2(1-\cos t)}\right) dt\right) \\ &= \frac{9}{2} \exp\left(-\frac{1}{2\pi} \int_{-\pi}^{\pi} \ln(1 - \cos t) dt\right) \\ &= 9, \end{aligned} \quad (10)$$

the last integral being known. (One can get a much easier integral, with a slightly worse constant, by lower-bounding $1 - \cos t \geq (2/\pi^2)t^2$.) Combining (8), (9), (10) yields the theorem. $\square$

6 PROOF OF THEOREM 4.2

The key ingredient is the following theorem from [BEK99]. (Recall that a *Littlewood polynomial* has all nonzero coefficients either -1 or 1 .)

THEOREM 6.1 ([BEK99], THEOREM 3.3). *For all $k \geq 2$ there is a nonzero Littlewood polynomial Q_k of degree at most k satisfying $|Q_k(t)| \leq \exp(-c_0\sqrt{k})$ for all real $0 \leq t \leq 1$. Here $c_0 > 0$ is a universal constant.*

By a simple use of the Hadamard Three-Circle Theorem and Maximum Modulus Principle, Borwein and Erdélyi proved in [BE97] that the polynomials in Theorem 6.1 establish tightness of their Theorem 5.1 (up to the constant C_1). We quote a result that appears within their proof:

THEOREM 6.2 ([BE97], IN THE FIRST PROOF OF THEOREM 3.3 IN THE "SPECIAL CASE", P. 11). *There are universal constants $c_1, c_2, c_3 > 0$ such that the following holds: For all $0 < a \leq c_1$ there exists an integer*

$2 \leq k \leq c_2/a^2$ such that $\max_{w \in D_{6a}(1)} |Q_k(w)| \leq \exp(-c_3/a)$, where Q_k is the nonzero Littlewood polynomial from Theorem 6.1.

Remark 6.3. Actually, Borwein and Erdélyi proved this with an elliptical disk $\mathcal{E}_a$ in place of $D_{6a}(1)$, where $\mathcal{E}_a$ has foci at $1 - 8a$ and 1 and major axis $[1 - 14a, 1 + 6a]$. It is easy to see that $D_{6a}(1) \subset \mathcal{E}_a \subset D_{14a}(1)$, so we wrote $D_{6a}(1)$ in Theorem 6.2 for simplicity and because it loses almost nothing.

We can now prove Theorem 4.2. We state here a slightly more precise version:

THEOREM 6.4. *Using the notation $\delta = 1 - \rho$, and the notation $\text{Exp}(t) = \exp(c \cdot t)$ for an unspecified universal constant $c > 0$, we have*

$$\kappa_{\text{Littlewood}}(\rho, n) \leq \begin{cases} \text{Exp}(-(\delta n)^{1/3}) & \text{in Case I:} \\ C(\log^3 n)/n \leq \delta \leq 1/2, \\ \text{Exp}(-(n/\rho)^{1/3}) & \text{in Case II:} \\ C/n^{1/2} \leq \rho \leq 1/2, \end{cases}$$

provided $n \geq n_0$. Here $n_0, C \geq 1$ are universal constants.

PROOF OF THEOREM 4.2. With $C \geq 1$ to be specified later, select

$$a = \begin{cases} C_1/(\delta n)^{1/3} & \text{in Case I: } C(\log^3 n)/n < \delta \leq 1/2, \\ C_1(\rho/n)^{1/3} & \text{in Case II: } 1/n^{1/2} < \rho < 1/2, \end{cases}$$

where $C_1 \geq 1$ is a universal constant to be specified later. Assuming $n_0 = n_0(C_1)$ is sufficiently large we get that $a \leq c_1$, where c_1 is as in Theorem 6.2. Applying that theorem, we obtain

$$\max_{w \in A} |Q_k(w)| \leq \text{Exp}(-1/a), \quad (11)$$

where

$$A := D_{6a}(1), \quad k \leq c_2/a^2 < n/2.$$

Here the inequality $c_2/a^2 < n/2$ holds in Case I by assuming $n_0 = n_0(C_1, c_2)$ large enough, and in Case II by taking $C_1 = C_1(c_2)$ large enough. Now define

$$P(w) = w^{\lfloor n/2 \rfloor} \cdot Q_k(w),$$

a nonzero Littlewood polynomial of degree less than n . We wish to bound

$$\max_{w \in R} |P(w)|, \quad R := D_\rho(\delta)$$

by the expression in the theorem statement. For the points $w \in R \cap A$, we are done by (11) (and the fact that $|w^{\lfloor n/2 \rfloor}| \leq 1$). For the points in $w \in R \setminus A$, we claim that

$$|w|^2 \leq 1 - 36\frac{\delta}{\rho}a^2 \leq \exp(-36\frac{\delta}{\rho}a^2) \quad \forall w \in R \setminus A. \quad (12)$$

Assuming (12), we get

$$\begin{aligned} \max_{w \in R \setminus A} |P(w)| &\leq \max_{w \in R \setminus A} |w|^{\lfloor n/2 \rfloor} \cdot \max_{w \in R \setminus A} |Q_k(w)| \\ &\leq \exp(-18\frac{\delta}{\rho}a^2)^{\lfloor n/2 \rfloor} \cdot (n/2 + 1) \\ &\leq \text{Exp}(-n\frac{\delta}{\rho}a^2) \cdot (n/2 + 1), \end{aligned}$$

where the factor $n/2 + 1$ is an upper bound on $|Q_k(w)|$ over all of $D_1(0)$ (recall that Q_k is a Littlewood polynomial of degree less than $n/2$). By inspection, this is sufficient to complete the proof in both Case I and Case II (in Case I we need to assume C large enough to absorb the factor of $(n/2 + 1)$).

It remains to establish (12). For this we first note that $\rho > 3a$ in both Case I and Case II (Case I is easier to check; for Case II we need to use that $C = C(C_1)$ is sufficiently large). This in particular means that $R \setminus A \neq \emptyset$. Writing w_0 for either of the intersection points of ∂R and ∂A , we have $\max_{w \in R \setminus A} |w| \leq |w_0|$. Thus it suffices to upper-bound $|w_0|^2$.

In the complex plane, consider the triangle formed by δ , 1, and w_0 . Note that w_0 has distance ρ from δ and distance $6a$ from 1. Let θ denote the triangle's angle at δ . By the Cosine Law, $(6a)^2 = \rho^2 + \rho^2 - 2\rho^2 \cos \theta$ and hence $\cos \theta = 1 - 18a^2/\rho^2$. Now consider the triangle formed by δ , 0, and w_0 . Its angle at δ is $\pi - \theta$ and the adjacent sides have length δ , ρ . Thus by the Cosine Law,

$$\begin{aligned} |w_0|^2 &= \delta^2 + \rho^2 - 2\delta\rho \cos(\pi - \theta) \\ &= \delta^2 + \rho^2 + 2\delta\rho \cos \theta \\ &= (\delta + \rho)^2 - 36\delta\rho a^2/\rho^2 \\ &= 1 - 36\frac{\delta}{\rho}a^2, \end{aligned}$$

as needed for (12). $\square$

7 CONCLUSIONS

A natural direction for future work is to go beyond mean-based algorithms. For example, an efficient algorithm can estimate the covariances of all *pairs* of trace bits. If different sources strings lead to sufficiently different trace-covariances, one could potentially get a more efficient trace reconstruction algorithm. Analyzing this strategy is equivalent to analyzing a certain problem concerning the maxima of Littlewood-like polynomials on $\mathbb{C}^2$; however we could not make any progress on this problem. It would also be interesting to develop lower bound techniques that apply to a broader class of algorithms than just mean-based algorithms.

Finally, we mention that the authors have applied the techniques in this paper (specifically, the technique used in Section 5.1) to several aspects of the population recovery problem. Details will appear in a forthcoming work.

A RESULTS ON CHANNELS THAT ALLOW INSERTIONS, DELETIONS AND FLIPS

A.1 Defining the General Channel

We now describe the most general channel C that we analyze, which we subsequently refer to as “the general channel”. As stated earlier, this channel allows for three different types of corruptions: deletions with probability δ , insertions with probability σ , and bit-flips with probability $\gamma/2$. We comment that for mean-based algorithms, the presence of bit-flips makes hardly any difference; thus the reader may focus just on the combination of deletions and insertions.

Our definition of this general channel is essentially the same as that of Kannan and McGregor [KM05]. More precisely, for parameters $\delta, \sigma, \gamma \in [0, 1]$, we define how the channel acts on a single source bit $b \in \{-1, 1\}$:

- (1) First, the channel performs “insertions”; i.e., it repeatedly does the operation “with probability σ , transmit a uniformly random bit; with probability $1 - \sigma$, stop”.

- (2) Having stopped, the channel “deletes” (completes transmission without sending b or $-b$) with probability δ .
- (3) Otherwise (with probability $1 - \delta$), the channel transmits one more bit: namely, b with probability $1 - \gamma/2$, or $-b$ with probability $\gamma/2$.

As usual, the channel operates on an entire source string $x \in \{-1, 1\}^n$ by operating on its individual bits independently, concatenating the results. That is,

$$C(x) = C(x_0)C(x_1) \cdots C(x_{n-1}) \in \{-1, 1\}^*.$$

Of course, if we set $\sigma = \gamma = 0$, we get the deletion channel Del_δ that was analyzed in the main body of the paper.

An alternative description of the channel's operation on a single bit x_i is as follows:

$$C(x_i) = \begin{cases} \mathbf{w} & \text{with probability } \delta, \\ (\mathbf{w}, \mathbf{a}) & \text{with probability } (1 - \delta) \cdot \gamma, \\ (\mathbf{w}, x_i) & \text{with probability } (1 - \delta) \cdot (1 - \gamma), \end{cases} \quad (13)$$

where $\mathbf{a} \in \{-1, 1\}$ is a uniformly random bit, and where $\mathbf{w} \in \{-1, 1\}^G$ is a uniformly random string of G bits, with G in turn being a Geometric random variable of parameter $1 - \sigma$.² From this description one can see that in a received word $\mathbf{y} \leftarrow C(x)$, each received bit either “comes from a properly transmitted source bit x_i ”, or else is uniformly random. (The probability each x_i comes through is $(1 - \delta)(1 - \gamma)$.) As a consequence, we have that Proposition 3.1 continues to hold for C : for every $j \in \mathbb{N}$, the mean value $\mathbb{E}_{\mathbf{y} \leftarrow C(x)}[\mathbf{y}_j]$ is a (real-)linear function of x .

Note that when the insertion probability σ is positive, the received word $\mathbf{y} \leftarrow C(x)$ does not have an a priori bounded length. This is a minor annoyance can be handled in several different ways; we choose one way in the next section.

A.2 Mean Traces for the General Channel

We revisit some of our definitions and observations about mean traces from Section 3, in our new context of the general channel. We begin with (1), the definition of the mean trace. Since the length of a received word may now be arbitrarily large, the mean trace is now an infinite vector. We deal with this by truncating it at what we call the “effective trace length bound N ”.

Definition A.1. For the general channel C with insertion probability $0 \leq \sigma < 1$, we define the *effective trace length bound* $N = N(\sigma)$ to be $N = \left\lceil 10 \cdot \frac{n + \ln(1/(1-\sigma))}{1-\sigma} \right\rceil \leq \text{poly}(n, \frac{1}{1-\sigma})$.

Definition A.2. For the general channel C and a source string $x \in \{-1, 1\}^n$, we define the *idealized mean trace* to be the infinite sequence

$$\mu_C^{\text{ideal}}(x) = \mathbb{E}_{\mathbf{y} \leftarrow C(x)}[(\mathbf{y}, 0, 0, 0, \dots)] \in [-1, 1]^{\mathbb{N}}.$$

We define just the *mean trace* to be its truncation to length N :

$$\mu_C(x) = (\mu_C^{\text{ideal}}(x)_0, \mu_C^{\text{ideal}}(x)_1, \dots, \mu_C^{\text{ideal}}(x)_{N-1}) \in [-1, 1]^N.$$

² Here we use the convention that Geometric random variables take values $0, 1, 2, \dots$ (equal to the number of “failures”); i.e., $\Pr[G = t] = \sigma^t(1 - \sigma)$ for each $t \geq 0$.

Recalling (13), we see that the length n of a received word is stochastically dominated by $(G_1 + 1) + \dots + (G_n + 1)$, where the G_i 's are i.i.d. random variables distributed as $\text{Geometric}(1 - \sigma)$. We upper bound this using Janson's bound on the sum of independent Geometric random variables (Theorem 2.1 of [Jan14]), noting that his Geometric random variables count the number of "trials", which aligns precisely with our $(G_i + 1)$'s. His bound gives that $\Pr[n \geq N + j] \leq \exp(-(N + j)(1 - \sigma)/2)$ for any $j \geq 0$, and hence we have the following: for any $x \in [-1, 1]^n$,

$$\begin{aligned} \|\mu_C(x) - \mu_C^{\text{ideal}}(x)\|_1 &= \sum_{\ell=N}^{\infty} |\mu_C^{\text{ideal}}(x)|_{\ell} \\ &\leq \sum_{\ell=N}^{\infty} \Pr[n \geq \ell] \\ &= \sum_{j=0}^{\infty} \Pr[n \geq N + j] \\ &= \exp(-N(1 - \sigma)/2) \cdot \frac{1}{1 - \exp(-(1 - \sigma)/2)} \\ &< \frac{4 \exp(-N(1 - \sigma)/2)}{1 - \sigma} \\ &\leq 4 \exp(-n), \quad \text{by our choice of } N. \end{aligned} \quad (14)$$

The mean-based trace reconstruction model for the general channel. Definition 3.2 has a natural analogue for the general channel: an algorithm in the mean-based general-channel model specifies a cost parameter $T \in \mathbb{N}$ and is given an estimate $\hat{\mu}_C(x) \in [-1, 1]^N$ of the mean trace satisfying $\|\hat{\mu}_C(x) - \mu_C(x)\|_1 \leq 1/T$. It is clear that an algorithm in the mean-based general-channel trace reconstruction model with cost T_1 and postprocessing time T_2 may be converted into a normal trace reconstruction algorithm using $\text{poly}(N, T_1) = \text{poly}(n, \frac{1}{1-\sigma}, T_1)$ samples and $\text{poly}(n, \frac{1}{1-\sigma}, T_1) + T_2$ time. Note that since we will be studying algorithms with cost $T \ll 2^n$, by (14) there is no real difference between getting an estimate of $\mu_C(x)$ or of $\mu_C^{\text{ideal}}(x)$.

The complexity of mean-based trace reconstruction for the general channel. Regarding the complexity of mean-based trace reconstruction, for the general channel we define $\epsilon_C(n)$ and $\epsilon_C^{\text{frac}}(n)$ in the obvious way, replacing each occurrence of the length- n vector $\mu_{\text{Del}_\delta}(\cdot)$ in Definition 3.3 with the length- N vector $\mu_C(\cdot)$. As in Section 3.2, to show that trace reconstruction can be performed under the general channel in time $\text{poly}(N, M) = \text{poly}(n, \frac{1}{1-\sigma}, M)$ it suffices to show that $\epsilon_C^{\text{frac}}(n) \geq 1/M$.³

Reduction to complex analysis for the general channel. For $x \in \{-1, 1\}^n$ the *general-channel polynomial* is defined entirely analogously to Definition 3.4:

$$P_{C,x}(z) = \sum_{j < N} \mu_C(x)_j \cdot z^j;$$

³Again, to carry out the linear-programming algorithm, we can either assume that the channel parameters δ, σ, γ are known to the algorithm, or else they should be estimated; we omit the details here.

note that this is a polynomial of degree less than N . This definition extends to $x \in [-1, 1]^n$ using the linearity of μ_C . Similarly, we may define the *idealized general-channel "polynomial"* by

$$P_{C,x}^{\text{ideal}}(z) = \sum_{j \in \mathbb{N}} \mu_C^{\text{ideal}}(x)_j \cdot z^j;$$

this will actually be a rational function of z .

Entirely analogous to Proposition 3.5, we get that for every $b \in [-1, 1]^n$,

$$\max_{z \in \partial D_1(0)} |P_{C,b}(z)| \leq \|\mu_C(b)\|_1 \leq \sqrt{N} \max_{z \in \partial D_1(0)} |P_{C,b}(z)|.$$

Similar to Section 3.3, a factor of $\sqrt{N} = \text{poly}(n, \frac{1}{1-\sigma})$ is negligible compared to the bounds we will prove, so it suffices to analyze $\max_{z \in \partial D_1(0)} |P_{C,b}(z)|$ rather than $\|\mu_C(b)\|_1$ in the definitions of $\epsilon_C(n)$ and $\epsilon_C^{\text{frac}}(n)$. Moreover, since by (14) we have that $|P_{C,b}^{\text{ideal}}(z) - P_{C,b}(z)| \leq 2^{-n}$ for all $b \in [-1, 1]^n$ and all $z \in \partial D_1(0)$, it suffices to analyze $\max_{z \in \partial D_1(0)} |P_{C,b}^{\text{ideal}}(z)|$; we do this in the next subsection.

A.3 Channel Polynomial for General Channels

We now compute the ideal channel polynomial for the general channel defined in Section A.1, using the same technique as in Section 4 and recalling the discussion around the alternative channel description (13). As usual, let $\rho = 1 - \delta$. Let J_i be the random variable whose value is $\perp$ if x_i is either deleted (probability δ) or is replaced by a random bit (probability $(1 - \delta) \cdot \gamma$), or else is the position j such that coordinate x_i of the source string ends up in coordinate j in the received string $\mathbf{y}$. As before we let J_i denote the random variable J_i conditioned on not being $\perp$. Since $\Pr[J_i \neq \perp] = (1 - \delta) \cdot (1 - \gamma)$, a derivation identical to that of (3) yields

$$P_{C,x}^{\text{ideal}}(z) = (1 - \delta)(1 - \gamma) \sum_{i < n} x_i \cdot \mathbb{E}[z^{\tilde{J}_i}]. \quad (15)$$

To compute $\mathbb{E}[z^{\tilde{J}_i}]$, it is straightforward to see that each coordinate $x_{i'}$ with $i' < i$ independently generates a random number of received positions distributed as $G + B$, where $G \sim \text{Geometric}(1 - \sigma)$ and independently $B \sim \text{Bernoulli}(\rho)$. Further, conditioned on x_i not being deleted, x_i generates a number of received positions distributed as $G + 1$, where the final "+1" is for x_i (or $-x_i$) itself. Thus $\tilde{J}_i$ is distributed as

$$G_0 + \dots + G_i + B_0 + \dots + B_{i-1},$$

where the G_k 's are independent copies of G and the B_k 's are independent copies of B . We therefore obtain

$$\mathbb{E}[z^{\tilde{J}_i}] = \mathbb{E}[z^G]^{i+1} \cdot \mathbb{E}[z^B]^i = (\mathbb{E}[z^G] \cdot \mathbb{E}[z^B])^i \cdot \mathbb{E}[z^G].$$

Let $F_G(z)$ denote $\mathbb{E}[z^G]$ and let $F_B(z)$ denote $\mathbb{E}[z^B]$. It is easy to calculate that $F_G(z) = \frac{1-\sigma}{1-\sigma z}$, and we saw earlier that $F_B(z) = (1 - \rho) + \rho z = \delta + \rho z$. For brevity, let us write

$$w = F_G(z)F_B(z) = \frac{(1 - \sigma) \cdot (\delta + \rho z)}{1 - \sigma z},$$

which is a Möbius transformation of z . Thus w ranges over a complex circle as z ranges over $\partial D_1(0)$. More specifically, as z ranges over $\partial D_1(0)$ we have that w ranges over $\partial D_r(1 - r)$, where

$$r = \frac{\rho + \delta\sigma}{1 + \sigma}.$$

Plugging this back into (15) using $E[z^{\tilde{J}_i}] = F_G(z) \cdot w^i$, we obtain

$$\begin{aligned} p_{C,x}^{\text{ideal}}(z) &= (1 - \delta) \cdot (1 - \gamma) \cdot F_G(z) \cdot \sum_{i < n} x_i \cdot w^i \\ &= (1 - \gamma) \cdot (1 - \delta) \cdot \frac{1 - \sigma}{1 - \sigma z} \cdot \sum_{i < n} x_i \cdot w^i. \end{aligned}$$

We use the bound $\left| \frac{1 - \sigma}{1 - \sigma z} \right| \geq \frac{1 - \sigma}{2}$ for $z \in \partial D_1(0)$. Now by the analysis of $\kappa_{\text{bounded}}^{\text{frac}}(r, d)$ given in Section 4 we get the following algorithmic result for general-channel trace reconstruction, which is our most general positive result:

Theorem 1.4, restated. *Let C be the general channel described in Section A.1 with deletion probability $\delta = 1 - \rho$, insertion probability σ , and bit-flip probability $\gamma/2$. Define*

$$r := \frac{\rho + \delta\sigma}{1 + \sigma}.$$

Then there is an algorithm for C -channel trace reconstruction using samples and running time bounded by

$$\text{poly}\left(\frac{1}{1-\delta}, \frac{1}{1-\sigma}, \frac{1}{1-\gamma}\right) \cdot \begin{cases} \exp(O(n/r)^{1/3}) & \text{if } C/n^{1/2} \leq r \leq 1/2, \\ \exp(O((1-r)n)^{1/3}) & \text{if } O(\log^3 n)/n \leq 1-r \leq 1/2. \end{cases}$$

Let us make some observations about this result. First, our Theorem 1.1 for the deletion channel is the special case of Theorem 1.4 obtained by setting $\sigma = \gamma = 0$. Next, for fixed δ ,

$$\begin{aligned} \text{if } \delta \leq 1/2, \quad & r \text{ ranges from } 1 - \delta \text{ down to } 1/2 \text{ as} \\ & \sigma \text{ ranges from } 0 \text{ up to } 1; \\ \text{if } \delta \geq 1/2, \quad & r \text{ ranges from } 1 - \delta \text{ up to } 1/2 \text{ as} \\ & \sigma \text{ ranges from } 0 \text{ up to } 1. \end{aligned}$$

The second statement is rather peculiar: it implies that when the deletion rate is high, the ability to perform trace reconstruction actually *improves*, the more insertions there are. Indeed, when we have deletions only, our ability to do trace reconstruction in time $\exp(O(n^{1/3}))$ is limited to retention probability $\rho \geq \Omega(1)$. But as soon as the insertion rate σ satisfies $\sigma \geq \Omega(1)$, we can do trace reconstruction in time $\exp(O(n^{1/3}))$ as long as the retention rate $\rho = 1 - \delta$ satisfies $\rho \geq \exp(-O(n^{1/3}))$.

ACKNOWLEDGMENTS

The authors would like to thanks the Simons Foundation for sponsoring the symposium on analysis of Boolean functions where the authors began work on this project. A. D. is grateful to Aravindan Vijayaraghavan for useful discussions about this problem.

REFERENCES

- [ADHR12] Alexandr Andoni, Constantinos Daskalakis, Avinatan Hassidim, and Sebastien Roch. Global alignment of molecular sequences via ancestral state reconstruction. *Stochastic Processes and their Applications*, 122(12):3852–3874, 2012.
- [BE97] Peter Borwein and Tamás Erdélyi. Littlewood-type polynomials on subarcs of the unit circle. *Indiana University Mathematics Journal*, 46(4):1323–1346, 1997.
- [BEK99] Peter Borwein, Tamás Erdélyi, and Géza Kós. Littlewood-type problems on $[0, 1]$. *Proceedings of the London Mathematical Society*, 3(79):22–46, 1999.
- [BKKM04] Tuğkan Batu, Sampath Kannan, Sanjeev Khanna, and Andrew McGregor. Reconstructing strings from random traces. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 910–918, 2004.
- [CK97] Christian Choffrut and Juhani Karhumäki. Combinatorics of words. In *Handbook of Formal Languages, Volume I*, pages 329–438. Springer, 1997.
- [DOS16] Anindya De, Ryan O'Donnell, and Rocco A. Servedio. Optimal mean-based algorithms for trace reconstruction. Technical Report <https://arxiv.org/pdf/1612.03148.pdf>, 9 Dec 2016.
- [DS03] Miroslav Dudík and Leonard Schulman. Reconstruction from subsequences. *Journal of Combinatorial Theory, Series A*, 103(2):337–348, 2003.
- [HMPW08] Thomas Holenstein, Michael Mitzenmacher, Rina Panigrahy, and Udi Wieder. Trace reconstruction with constant deletion probability and related results. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 389–398, 2008.
- [Jan14] Svante Janson. Tail bounds for sums of geometric and exponential variables, 2014. <http://www2.math.uu.se/~svante/papers/sjN14.pdf>.
- [Kal73] V. V. Kalashnik. Reconstruction of a word from its fragments. *Computational Mathematics and Computer Science (Vychislitel'naya matematika i vychislitel'naya tekhnika)*, Kharkov, 4:56–57, 1973.
- [KM05] Sampath Kannan and Andrew McGregor. More on reconstructing strings from random traces: Insertions and deletions. In *IEEE International Symposium on Information Theory*, pages 297–301, 2005.
- [KR97] Ilia Krasikov and Yehuda Roditty. On a reconstruction problem for sequences. *Journal of Combinatorial Theory, Series A*, 77(2):344–348, 1997.
- [Lev01a] Vladimir Levenshtein. Efficient reconstruction of sequences. *IEEE Transactions on Information Theory*, 47(1):2–22, 2001.
- [Lev01b] Vladimir Levenshtein. Efficient reconstruction of sequences from their subsequences or supersequences. *Journal of Combinatorial Theory Series A*, 93(2):310–332, 2001.
- [Mit09] Michael Mitzenmacher. A survey of results for deletion channels and related synchronization channels. *Probability Surveys*, 6:1–33, 2009.
- [MMS⁺91] Bennet Manvel, Aaron Meyerowitz, Allen Schwenk, Ken Smith, and Paul Stockmeyer. Reconstruction of sequences. *Discrete Mathematics*, 94(3):209–219, 1991.
- [Mos] Elchanan Mossel. Personal communication, October 2016.
- [Mos13] Elchanan Mossel. MSRI open problem session, 2013. https://www.msri.org/c/document_library/get_file?uuid=4a885484-bcdd-4238-a3da-21c05713034c&groupId=14404.
- [MPV14] Andrew McGregor, Eric Price, and Sofya Vorotnikova. Trace reconstruction revisited. In *Proceedings of the 22nd Annual European Symposium on Algorithms*, pages 689–700, 2014.
- [Sco97] Alexander Scott. Reconstructing sequences. *Discrete Mathematics*, 175(1):231–238, 1997.
- [Smy08] Chris Smyth. The Mahler measure of algebraic numbers: A survey. In *Number Theory and Polynomials*, pages 322–349. London Mathematical Society Lecture Note Series 352, 2008.
- [VS08] Krishnamurthy Viswanathan and Ram Swaminathan. Improved string reconstruction over insertion-deletion channels. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 399–408, 2008.