

Efficiently Decodable Codes for the Binary Deletion Channel*

Venkatesan Guruswami¹ and Ray Li²

- 1 Carnegie Mellon University, Pittsburgh, PA
venkatg@cs.cmu.edu
- 2 Carnegie Mellon University, Pittsburgh, PA
ryli@andrew.cmu.edu

Abstract

In the random deletion channel, each bit is deleted independently with probability p . For the random deletion channel, the *existence* of codes of rate $(1-p)/9$, and thus bounded away from 0 for any $p < 1$, has been known. We give an explicit construction with polynomial time encoding and deletion correction algorithms with rate $c_0(1-p)$ for an absolute constant $c_0 > 0$.

1998 ACM Subject Classification E.4 Coding and Information Theory

Keywords and phrases Coding theory, Combinatorics, Synchronization errors, Channel capacity

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2017.47

1 Introduction

We consider the problem of designing error-correcting codes for reliable and efficient communication on the binary deletion channel. The *binary deletion channel* (BDC) *deletes* each transmitted bit independently with probability p , for some $p \in (0, 1)$ which we call the *deletion probability*. Crucially, the location of the deleted bits are *not* known at the decoder, who receives a *subsequence* of the original transmitted sequence. The loss of synchronization in symbol locations makes the noise model of deletions challenging to cope with. As one indication of this, we still do not know the channel capacity of the binary deletion channel. Quoting from the first page of Mitzenmacher’s survey [17]: “Currently, we have no closed-form expression for the capacity, nor do we have an efficient algorithmic means to numerically compute this capacity.” This is in sharp contrast with the noise model of bit erasures, where each bit is independently replaced by a ‘?’ with probability p (the binary erasure channel (BEC)), or of bit errors, where each bit is flipped independently with probability p (the binary symmetric channel (BSC)). The capacity of the BEC and BSC equal $1-p$ and $1-h(p)$ respectively, and we know codes of polynomial complexity with rate approaching the capacity in each case.

The capacity of the binary deletion channel is clearly at most $1-p$, the capacity of the simpler binary erasure channel. Diggavi and Grossglauser [3] establish that the capacity of the deletion channel for $p \leq \frac{1}{2}$ is at least $1-h(p)$. Kalai, Mitzenmacher, and Sudan [11] proved this lower bound is tight as $p \rightarrow 0$, and Kanoria and Montanari [12] determined a series expansion that can be used to determine the capacity exactly. Turning to large p , Rahmati and Duman [18] prove that the capacity is at most $0.4143(1-p)$ for $p \geq 0.65$. Drinea and Mitzenmacher [4, 5] proved that the capacity of the BDC is at least $(1-p)/9$,

* Research supported in part by NSF grant CCF-1422045.



which is within a constant factor of the upper bound. In particular, the capacity is positive for every $p < 1$, which is perhaps surprising. The asymptotic behavior of the capacity of the BDC at both extremes of $p \rightarrow 0$ and $p \rightarrow 1$ is thus known.

This work is concerned with *constructive* results for coding for the binary deletion channel. That is, we seek codes that can be constructed, encoded, and decoded from deletions caused by the BDC, in polynomial time. Recently, there has been good progress on codes for *adversarial* deletions, including constructive results. Here the model is that the channel can delete an arbitrary subset of pn bits in the n -bit codeword. A code capable of correcting pn worst-case deletions can clearly also correct deletions caused by a BDC with deletion probability $(p - \epsilon)$ with high probability, so one can infer results for the BDC from some results for worst-case deletions. For small p , Guruswami and Wang [9] constructed binary codes of rate $1 - O(\sqrt{p})$ to efficiently correct a p fraction worst-case deletions. So this also gives codes of rate approaching 1 for the BDC when $p \rightarrow 0$. For larger p , Kash et al. [13] proved that randomly chosen codes of small enough rate $R > 0$ can correctly decode against pn adversarial deletions when $p \leq 0.17$. Even non-constructively, this remained the best achievability result in terms of correctable deletion fraction until the recent work of Bukh, Guruswami, and Håstad [2] who constructed codes of positive rate efficiently decodable against pn adversarial deletions for any $p < \sqrt{2} - 1$. For adversarial deletions, it is impossible to correct a deletion fraction of $1/2$, whereas the capacity of the BDC is positive for all $p < 1$. So solving the problem for the much harder worst-case deletions is not a viable approach to construct positive rate codes for the BDC for $p > 1/2$.

To the best of our knowledge, explicit efficiently decodable code constructions were not available for the binary deletion channel for arbitrary $p < 1$. We present such a construction in this work. Our rate is worse than the $(1 - p)/9$ achieved non-constructively, but has asymptotically the same dependence on p for $p \rightarrow 1$.

► **Theorem 1.** *Let $p \in (0, 1)$. There is an explicit a family of binary codes that (1) has rate $(1 - p)/110$, (2) is constructible in polynomial time, (3) encodable in time $O(N)$, and (3) decodable with high probability on the binary deletion channel with deletion probability p in time $O(N^2)$. (Here N is the block length of the code)*

1.1 Some other related work

One work that considers efficient recovery against random deletions is by Yazdi and Dolecek [20]. In their setting, two parties Alice and Bob are connected by a two-way communication channel. Alice has a string X , Bob has string Y obtained by passing X through a binary deletion channel with deletion probability $p \ll 1$, and Bob must recover X . They produce a polynomial-time synchronization scheme that transmits a total of $O(pn \log(1/p))$ bits and allows Bob to recover X with probability exponentially approaching 1.

For other models of random synchronization errors, Kirsch and Drinea [14] prove information capacity lower bounds for channels with i.i.d deletions and duplications. Fertoni et al. [6] prove capacity bounds for binary channels with i.i.d insertions, deletions, and substitutions.

For deletion channels over non-binary alphabets, Rahmati and Duman [18] prove a capacity upper bound of $C_2(p) + (1 - p) \log(|\Sigma|/2)$, where $C_2(p)$ denotes the capacity of the binary deletion channel with deletion probability p , when the alphabet size $|\Sigma|$ is even. In particular, using the best known bound for $C_2(p)$ of $C_2(p) \leq 0.4143(1 - p)$, the upper bound is $(1 - p)(\log |\Sigma| - 0.5857)$.

In [8], the authors of this paper consider the model of *oblivious* deletions, which is in between the BDC and adversarial deletions in power. Here, the channel can delete any pn

bits of the codeword, but must do so without knowledge of the codeword. In this model, they prove the *existence* of codes of positive rate for correcting any fraction $p < 1$ of oblivious deletions.

1.2 Our construction approach

Our construction concatenates a high rate outer code over a large alphabet that is efficiently decodable against a small fraction of *adversarial* insertions and deletions, with a good inner binary code. For the outer code, we can use the recent construction of [10]. To construct the inner code, we first choose a binary code correcting a small fraction of adversarial deletions. By concentration bounds, duplicating bits of a codeword in a disciplined manner is effective against the random deletion channel, so we, for some constant B , duplicate every bit of the binary code $B/(1-p)$ times. We further ensure our initial binary code has only runs of length 1 and 2 to maximize the effectiveness of duplication. We add small buffers of 0s between inner codewords to facilitate decoding.

One might wonder whether it would be possible to use Drinea and Mitzenmacher's existential result [4, 5] of a $(1-p)/9$ capacity lower bound as a black box inner code to achieve a better rate together with efficient decodability. We discuss this approach in §3.2 and elaborate on what makes such a construction difficult to implement.

2 Preliminaries

General Notation. Throughout the paper, $\log x$ refers to the base-2 logarithm. We use interval notation $[a, b] = \{a, a+1, \dots, b\}$ to denote intervals of integers, and we use $[a] = [1, a] = \{1, 2, \dots, a\}$. Let $\text{Binomial}(n, p)$ denote the Binomial distribution.

Words. A *word* is a sequence of symbols from some *alphabet*. We denote explicit words using angle brackets, like $\langle 01011 \rangle$. We denote string concatenation of two words w and w' with ww' . We denote $w^k = ww \cdots w$ where there are k concatenated copies of w .

A *subsequence* of a word w is a word obtained by removing some (possibly none) of the symbols in w .

Let $\Delta_{i/d}(w_1, w_2)$ denote the *insertion/deletion distance* between w_1 and w_2 , i.e. the minimum number of insertions and deletions needed to transform w_1 into w_2 . By a lemma due to Levenshtein [15], this is equal to $|w_1| + |w_2| - 2\text{LCS}(w_1, w_2)$, where LCS denotes the length of the longest common subsequence.

Define a *run* of a word w to be a maximal single-symbol subword. That is, a subword w' in w consisting of a single symbol such that any longer subword containing w' has at least two different symbols. Note the runs of a word partition the word. For example, 110001 has 3 runs: one run of 0s and two runs of 1s.

We say that $c \in \{0, 1\}^m$ and $c' \in \{0, 1\}^m$ are *confusable under δm deletions* if it is possible to apply δm deletions to c and c' and obtain the same result. If δ is understood, we simply say c and c' are *confusable*.

Concentration Bounds. We use the following forms of Chernoff bound.

► **Lemma 2** (Chernoff). *Let A_1, \dots, A_n be i.i.d random variables taking values in $[0, 1]$. Let $A = \sum_{i=1}^n A_i$ and $\delta \in [0, 1]$. Then*

$$\Pr[A \leq (1 - \delta) \mathbb{E}[A]] \leq \exp(-\delta^2 \mathbb{E}[A]/2) \quad (1)$$

Furthermore,

$$\Pr[A \geq (1 + \delta) \mathbb{E}[A]] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\mathbb{E}[A]}. \quad (2)$$

We also have the following corollary, whose proof is in Appendix A.

► **Lemma 3.** *Let $0 < \alpha < \beta$. Let A_1, \dots, A_n be independent random variables taking values in $[0, \beta]$ such that, for all i , $\mathbb{E}[A_i] \leq \alpha$. For $\gamma \in [\alpha, 2\alpha]$, we have*

$$\Pr \left[\sum_{i=1}^n A_i \geq n\gamma \right] \leq \exp \left(-\frac{(\gamma - \alpha)^2 n}{3\alpha\beta} \right). \quad (3)$$

3 Efficiently decodable codes for random deletions with p approaching 1

3.1 Construction

We present a family of constant rate codes that decodes with high probability on a binary deletion channel with deletion fraction p (BDC_p). These codes have rate $c_0(1 - p)$ for an absolute positive constant c_0 , which is within a constant of the upper bound $(1 - p)$, which even holds for the erasure channel. By Drinea and Mitzenmacher [4] the maximum known rate of a non-efficiently correctable binary deletion channel code is $(1 - p)/9$.

The construction is based on the intuition that deterministic codes are better than random codes for the deletion channel. Indeed, for adversarial deletions, length n random codes correct at most $0.22n$ deletions [13], while explicitly constructed codes can correct close to $(\sqrt{2} - 1)n$ deletions [2].

We begin by borrowing a result from [9].

► **Lemma 4** (Corollary of Lemma 2.3 of [9]). *Let $0 < \delta < \frac{1}{2}$. For every binary string $c \in \{0, 1\}^m$, there are at most $\delta m \binom{m}{(1-\delta)m}^2$ strings $c' \in \{0, 1\}^m$ such that c and c' are confusable under δm deletions.*

The next lemma gives codes against a small fraction of adversarial deletions with an additional run-length constraint on the codewords.

► **Lemma 5.** *Let $\delta > 0$. There exists a length m binary code of rate $\mathcal{R} = 0.6942 - 2h(\delta) - O(\log(\delta m)/m)$ correcting a δ fraction of adversarial insertions and deletions such that each codeword contains only runs of size 1 and 2. Furthermore this code is constructible in time $\tilde{O}(2^{(0.6942 + \mathcal{R})m})$.*

Proof. It is easy to show that the number of codewords with only runs of 1 and 2 is F_m , the m th Fibonacci number, and it is well known that $F_m = \varphi^m + o(1) \approx 2^{0.6942m}$ where φ is the golden ratio. Now we construct the code by choosing it greedily. Each codeword is confusable with at most $\delta m \binom{m}{(1-\delta)m}^2$ other codewords, so the number of codewords we can choose is at least

$$\frac{2^{0.6942m}}{\delta m \binom{m}{(1-\delta)m}^2} = 2^{m(0.6942 - 2h(\delta) - O(\log(\delta m)/m))}. \quad (4)$$

We can find all words of length m whose run lengths are only 1 and 2 by recursion in time $O(F_m) = O(2^{0.6942m})$. Running the greedy algorithm, we need to, for at most $F_m \cdot 2^{\mathcal{R}m}$ pairs

of such words, determine whether the pair is confusable (we only need to check confusability of a candidate word with words already added to the code). Checking confusability of two words under adversarial deletions reduces to checking whether the longest common subsequence is at least $(1 - \delta)m$, which can be done in time $O(m^2)$. This gives an overall runtime of $O(m^2 \cdot F_m \cdot 2^{\mathcal{R}m}) = \tilde{O}(2^{(0.6942 + \mathcal{R})m})$. ◀

► **Corollary 6.** *There exists a constant m_0^* such that for all $m \geq m_0^*$, there exists a length m binary code of rate $\mathcal{R}_{in} = 0.555$ correcting a $\delta_{in} = 0.0083$ fraction of adversarial insertions and deletions such that each codeword contains runs of size 1 and 2 only and each codeword starts and ends with a 1. Furthermore this code is constructible in time $O(2^{1.25m})$.*

Our construction utilizes the following result as a black box for efficiently coding against an arbitrary fraction of insertions and deletions with rate approaching capacity.

► **Theorem 7** (Theorem 1.1 of [10]). *For any $0 \leq \delta < 1$ and $\epsilon > 0$, there exists a code C over alphabet Σ , with $|\Sigma| = \text{poly}(1/\epsilon)$, with block length n , rate $1 - \delta - \epsilon$, and is efficiently decodable from δn insertions and deletions. The code can be constructed in time $\text{poly}(n)$, encoded in time $O(n)$, and decoded in time $O(n^2)$.*

We apply Theorem 7 for small δ , so we also could use the high rate binary code construction of [7] as an outer code.

We now turn to our code construction for Theorem 1.

The code. Let $B = 60, B^* = 1.43\bar{B} = 86, \eta = \frac{1}{1000}, \delta_{out} = \frac{1}{1000}$. Let

$$m_0 = \max(\alpha \log(1/\delta_{out})/\eta, m_0^*),$$

where α is a sufficiently large constant and where m_0^* is given by Corollary 6. Let $\epsilon_{out} > 0$ be small enough such that the alphabet Σ , given by Theorem 7 with $\epsilon = \epsilon_{out}$ and $\delta = \delta_{out}$, satisfies $|\Sigma| \geq m_0$, and let C_{out} be the corresponding code.

Let $C_{in} : |\Sigma| \rightarrow \{0, 1\}^m$ be the code given by Corollary 6, and let $\mathcal{R}_{in} = 0.555, \delta_{in} = 0.0083$, and $m = \frac{1}{\mathcal{R}_{in}} \log |\Sigma| = O(\log(1/\epsilon))$ be the rate, tolerable deletion fraction, and block length of the code, respectively (\mathcal{R}_{in} and δ_{in} are given by Corollary 6). Each codeword of C_{in} has runs of length 1 and 2 only, and each codeword starts and ends with a 1. This code is constructed greedily.

Our code is a modified concatenated code. We encode our message as follows.

- *Outer Code.* First, encode the message into the outer code, C_{out} , to obtain a word $c^{(out)} = \sigma_1 \dots \sigma_n$.
- *Concatenation with Inner Code.* Encode each outer codeword symbol $\sigma_i \in \Sigma$ by the inner code C_{in} .
- *Buffer.* Insert a buffer of ηm 0s between adjacent inner codewords. Let the resulting word be $c^{(cat)}$. Let $c_i^{(in)} = C_{in}(\sigma_i)$ denote the encoded inner codewords of $c^{(cat)}$.
- *Duplication.* After concatenating the codes and inserting the buffers, replace each character (including characters in the buffers) with $\lceil B/(1-p) \rceil$ copies of itself to obtain a word of length $N := Bnm/(1-p)$. Let the resulting word be c , and the corresponding inner codewords be $\{c_i^{(dup)}\}$.

Rate. The rate of the outer code is $1 - \delta_{out} - \epsilon_{out}$, the rate of the inner code is \mathcal{R}_{in} , the buffer and duplications multiply the rate by $\frac{1}{1+\eta}$ and $(1-p)/B$ respectively. This gives a total rate that is slightly greater than $(1-p)/110$.

Notation. Let s denote the received word after the codeword c is passed through the deletion channel. Note that (i) every bit of c can be identified with a bit in $c^{(cat)}$, and (ii) each bit in the received word s can be identified with a bit in c . Thus, we can define relations $f^{(dup)} : c^{(cat)} \rightarrow c$, and $f^{(del)} : c \rightarrow s$ (that is, relations on the indices of the strings). These are not functions because some bits may be mapped to multiple (for $f^{(dup)}$) or zero (for $f^{(del)}$) bits. Specifically, $f^{(del)}$ and $f^{(dup)}$ are the inverses of total functions. In this way, composing these relations (i.e. composing their inverse functions) if necessary, we can speak about the *image* and *pre-image* of bits or subwords of one of $c^{(cat)}$, c , and s under these relations. For example, during the Duplication step of encoding, a bit $\langle b_j \rangle$ of $c^{(cat)}$ is replaced with $B/(1-p)$ copies of itself, so the corresponding string $\langle b_j \rangle^{B/(1-p)}$ in c forms the *image* of $\langle b_j \rangle$ under $f^{(dup)}$, and conversely the *pre-image* of the duplicated string $\langle b_j \rangle^{B/(1-p)}$ is that bit $\langle b_j \rangle$.


Decoding algorithm

- *Decoding Buffer.* First identify all runs of 0s in the received word with length at least $B\eta m/2$. These are our *decoding buffers* that divide the word into *decoding windows*, which we identify with subwords of s .
- *Deduplication.* Divide each decoding window into runs. For each run, if it has strictly more than B^* copies of a bit, replace it with as two copies of that bit, otherwise replace it with one copy. For example, $\langle 0 \rangle^{2B}$ gets replaced with $\langle 00 \rangle$ while $\langle 0 \rangle^B$ gets replaced with $\langle 0 \rangle$. For each decoding window, concatenate these runs of length 1 and 2 in their original order in the decoding window to produce a *deduplicated* decoding window.
- *Inner Decoding.* For each deduplicated decoding window, decode an outer symbol $\sigma \in \Sigma_{out}$ from each decoding window by running the brute force deletion correction algorithm for C_{in} . That is, for each deduplicated decoding window $s_*^{(in)}$, find by brute force a codeword $c_*^{(in)}$ in C_{in} that such that $\Delta_{i/d}(c_*^{(in)}, s_*^{(in)}) \leq \delta_{in}m$. If $c_*^{(in)}$ is not unique or does not exist, do not decode an outer symbol σ from this decoding window. Concatenate the decoded symbols σ in the order in which their corresponding decoding windows appear in the received word s to obtain a word $s^{(out)}$.
- *Outer Decoding.* Decode the message \mathbf{m} from $s^{(out)}$ using the decoding algorithm of C_{out} in Theorem 7.

For purposes of analysis, label as $s_i^{(dup)}$ the decoding window whose pre-image under $f^{(del)}$ contains indices in $c_i^{(dup)}$. If this decoding window is not unique (that is, the image of $c_i^{(dup)}$ contains bits in multiple decoding windows), then assign $s_i^{(dup)}$ arbitrarily. Note this labeling may mean some decoding windows are unlabeled, and also that some decoding windows may have multiple labels. In our analysis, we show both occurrences are rare. For a decoding window $s_i^{(dup)}$, denote the result of $s_i^{(dup)}$ after Deduplication to be $s_i^{(in)}$.

The following diagram depicts the encoding and decoding steps.

The pair $(\{c_i^{(in)}\}_i, c^{(cat)})$ indicates that, at that step of encoding, we have produced the word $c^{(cat)}$, and the sequence $\{c_i^{(in)}\}_i$ are the “inner codewords” of $c^{(cat)}$ (that is, the words in between what would be identified by the decoder as decoding buffers). The pair $(\{c_i^{(dup)}\}_i, c)$ is used similarly.



Runtime. The outer code is constructible in $\text{poly}(n)$ time and the inner code is constructible in time $O(2^{1.25m}) = \text{poly}(1/\epsilon)$, which is a constant, so the total construction time is $\text{poly}(N)$.

Encoding in the outer code is linear time, each of the n inner encodings is constant time, and adding the buffers and applying duplications each can be done in linear time. The overall encoding time is thus $O(N)$.

The Buffer step of the decoding takes linear time. The Deduplication step of each inner codeword takes constant time, so the entire step takes linear time. For each inner codeword, Inner Decoding takes time $O(m^2 2^m) = \text{poly}(1/\epsilon)$ by brute force search over the 2^m possible codewords: checking each of the 2^m codewords is a longest common subsequence computation and thus takes time $O(m^2)$, giving a total decoding time of $O(m^2 2^m)$ for each inner codeword. We need to run this inner decoding $O(n)$ times, so the entire Inner Decoding step takes linear time. The Outer Decoding step takes $O(n^2)$ time by Theorem 7. Thus the total decoding time is $O(N^2)$.

Correctness. Note that, if an inner codeword is decoded incorrectly, then one of the following holds.

1. *Spurious Buffer.* A spurious decoding buffer is identified in the corrupted codeword during the Buffer step.
2. *Deleted Buffer.* A decoding buffer neighboring the codeword is deleted.
3. *Inner Decoding Failure.* Running the Deduplication and Inner Decoding steps on $s_i^{(dup)}$ computes the inner codeword incorrectly.

We show that, with high probability, the number of occurrences of each of these events is small.

The last case is the most nontrivial, so we deal with it first, assuming the codeword contains no spurious decoding buffers and the neighboring decoding buffers are not deleted. In particular, we consider an i such that our decoding window $s_i^{(dup)}$ whose pre-image under $f^{(del)}$ only contains bits in $c_i^{(dup)}$ (because no deleted buffer) and no bits in the image of $c_i^{(dup)}$ appear in any other decoding window (because no spurious buffer).

Recall that the inner code C_{in} can correct against $\delta_{in} = 0.0083$ fraction of adversarial insertions and deletions. Suppose an inner codeword $c_i^{(in)} = r_1 \dots r_k \in C_{in}$ has k runs r_j each of length 1 or 2, so that $m/2 \leq k \leq m$.

► **Definition 8.** A subword of α identical bits in the received word s is

- *type-0* if $\alpha = 0$,
- *type-1* if $\alpha \in [1, B^*]$,
- *type-2* if $\alpha \in [B^* + 1, \infty)$.

By abuse of notation, we say that a length 1 or 2 run r_j of the inner codeword $c_i^{(in)}$ has *type- t_j* if the image of r_j in s under $f^{(del)} \circ f^{(dup)}$ forms a type- t_j subword.

Let t_1, \dots, t_k be the types of the runs r_1, \dots, r_k , respectively. The image of a run r_j under $f^{(del)} \circ f^{(dup)}$ has length distributed as $\text{Binomial}(B|r_j|/(1-p), 1-p)$. Let $\delta = 0.4\bar{3}$ be such that $B^* = (1 + \delta)B$. By the Chernoff bounds in Lemma 2, the probability that a run r_j of length 1 is type-2 is

$$\Pr_{Z \sim \text{Binomial}(B/(1-p), 1-p)} [Z > B^*] < (e^\delta / (1 + \delta)^{1+\delta})^B < 0.0071. \quad (5)$$

Similarly, the probability that a run r_j of length-2 is type-1 is at most

$$\Pr_{Z \sim \text{Binomial}(2B/(1-p), 1-p)} [Z \leq B^*] < e^{-((1-\delta)/2)^2 B} < 0.0081. \quad (6)$$

The probability any run is type-0 is at most $\Pr_{Z \sim \text{Binomial}(B/(1-p), 1-p)} [Z = 0] < e^{-B} < 10^{-10}$.

We now have established that, for runs r_j in $c_i^{(in)}$, the probability that the number of bits in the image of r_j in s under $f^{(del)} \circ f^{(dup)}$ is “incorrect” (between 1 and B^* for length 2 runs, and greater than B^* for length 1 runs), is at most 0.0081, which is less than δ_{in} . If the only kinds of errors in the Local Decoding step were runs of c of length 1 becoming runs of length 2 and runs of length 2 become runs of length 1, then we have that, by concentration bounds, with probability $1 - 2^{-\Omega(m)}$, the number of insertions deletions needed to transform $s_i^{(in)}$ back into $c_i^{(in)}$ is at most $\delta_{in}m$, in which case $s_i^{(in)}$ gets decoded to the correct outer symbol using C_{in} .

However, we must also account for the fact that some runs r_j of $c_i^{(in)}$ may become deleted completely after duplication and passing through the deletion channel. That is, the image of r_j in s under $f^{(del)} \circ f^{(dup)}$ is empty, or, in other words, r_j is type-0. In this case the two neighboring runs r_{j-1} and r_{j+1} appear merged together in the Deduplication step of decoding. For example, if a run of 1s was deleted completely after duplication and deletion, its neighboring runs of 0s would be interpreted by the decoder as a single run. Fortunately, as we saw, the probability that a run is type-0 is extremely small ($< 10^{-10}$), and we show each type-0 run only increases $\Delta_{i/d}(c_i^{(in)}, s_i^{(in)})$ by a constant. We show this constant is at most 6.

To be precise, let Y_j be a random variable that is 0 if $|r_j| = t_j$, 1 if $\{|r_j|, t_j\} = \{1, 2\}$, and 6 if $t_j = 0$. We claim $\sum_{j=1}^k Y_j$ is an upper bound on $\Delta_{i/d}(c_i^{(in)}, s_i^{(in)})$. To see this, first note that if $t_j \neq 0$ for all i , then the number of runs of $c_i^{(in)}$ and $s_i^{(in)}$ are equal, so we can transform $c_i^{(in)}$ into $s_i^{(in)}$ by adding a bit to each length-1 type-2 run of $c_i^{(in)}$ and deleting a bit from each length-2 type-1 run of $s_i^{(in)}$.

Now, if some number, ℓ , of the t_j are 0, then at most 2ℓ of the runs in $c_i^{(in)}$ become merged with some other run (or a neighboring decoding buffer) after duplication and deletion. Each set of consecutive runs $r_j, r_{j+2}, \dots, r_{j+2j'}$ that are merged after duplication and deletion gets replaced with 1 or 2 copies of the corresponding bit. For example, if $r_1 = \langle 11 \rangle, r_2 = \langle 0 \rangle, r_3 = \langle 11 \rangle$, and if after duplication and deletion, $2B$ bits remain in the image of each of r_1 and r_3 , and r_2 is type-0, then the image of $r_1 r_2 r_3$ under $f^{(del)} \circ f^{(dup)}$ is $\langle 1 \rangle^{4B}$, which gets decoded as $\langle 11 \rangle$ in the Deduplication step because $\langle 1 \rangle^{4B}$ is type-2. To account for the type-0 runs in transforming $c_i^{(in)}$ into $s_i^{(in)}$, we (i) delete at most two bits from each of the ℓ type-0 runs in $c_i^{(in)}$ and (ii) delete at most two bits for each of at most 2ℓ merged runs in $c_i^{(in)}$. The total number of additional insertions and deletions required to account for type-0 runs of c is thus at most 6ℓ , so we need at most 6 insertions and deletions to account for each type-0 run.

Our analysis covers the case when some bits in the image of $c_i^{(in)}$ under $f^{(del)} \circ f^{(dup)}$ are interpreted as part of a decoding buffer. Recall that inner codewords start and end with a 1, so that $r_1 \in \{\langle 1 \rangle, \langle 11 \rangle\}$ for every inner codeword. If, for example, $t_1 = 0$, that is, the image under $f^{(del)} \circ f^{(dup)}$ of the first run of 1s, r_1 , is the empty string, then the bits of r_2 are interpreted as part of the decoding buffer. In this case too, our analysis tells us that the type-0 run r_1 increases $\Delta_{i/d}(c_i^{(in)}, s_i^{(in)})$ by at most 6.

We conclude $\sum_{j=1}^k Y_j$ is an upper bound for $\Delta_{i/d}(c_i^{(in)}, s_i^{(in)})$.

Note that if r_j has length 1, then by (5) we have

$$\mathbb{E}[Y_j] = 1 \cdot \Pr[r_j \text{ is type-2}] + 6 \cdot \Pr[r_j \text{ is type-0}] < 1 \cdot 0.0071 + 6 \cdot 10^{-9} < 0.0082. \quad (7)$$

Similarly, if r_j has length 2, then by (6) we have

$$\mathbb{E}[Y_j] = 1 \cdot \Pr[r_j \text{ is type-1}] + 6 \cdot \Pr[r_j \text{ is type-0}] < 1 \cdot 0.0081 + 6 \cdot 10^{-9} < 0.0082. \quad (8)$$

Thus $\mathbb{E}[Y_j] < 0.0082$ for all i . We know the word $s_i^{(in)}$ is decoded incorrectly (i.e. is not decoded as σ_i) in the Inner Decoding step only if $\Delta_{i/d}(c_i^{(in)}, s_i^{(in)}) > \delta_{in}m$. The Y_j are independent, so Lemma 3 gives

$$\begin{aligned} \Pr[s_i^{(in)} \text{ decoded incorrectly}] &\leq \Pr[Y_1 + Y_2 + \cdots + Y_k \geq \delta_{in}m] \\ &\leq \Pr[Y_1 + Y_2 + \cdots + Y_k \geq \delta_{in}k] \\ &\leq \exp\left(-\frac{(\delta_{in} - 0.0082)^2 k}{3 \cdot 6 \cdot \delta_{in}}\right) \\ &\leq \exp(-\Omega(m)) \end{aligned} \tag{9}$$

where the last inequality is given by $k \geq m/2$. Since our $m \geq \Omega(\log(1/\delta_{out}))$ is sufficiently large, we have the probability $s_i^{(in)}$ is decoded incorrectly is at most $\delta_{out}/10$. If we let $Y_j^{(i)}$ denote the Y_j corresponding to inner codeword $c_i^{(in)}$, the events E_i given by $\sum_j Y_j^{(i)} \geq \delta_{in}m$ are independent. By concentration bounds on the events E_i , we conclude the probability that there are at least $\delta_{out}n/9$ incorrectly decoded inner codewords that are not already affected by spurious buffers and neighboring deleted buffers is $2^{-\Omega(n)}$.

Our aim is to show that the number of spurious buffers, deleted buffers, and inner decoding failures is small with high probability. So far, we have shown that, with high probability, assuming a codeword is not already affected by spurious buffers and neighboring deleted buffers, the number of inner decoding failures is small. We now turn to showing the number of spurious buffers is likely to be small.

A spurious buffer appears inside an inner codeword if many consecutive runs of 1s are type-0. A spurious buffer requires at least one of the following: (i) a codeword contains a sequence of at least $\eta m/5$ consecutive type-0 runs of 1s, (ii) a codeword contains a sequence of $\ell \leq \eta m/5$ consecutive type-0 runs of 1s, such that, for the $\ell + 1$ consecutive runs of 0s neighboring these type-0 runs of 1s, their image under $f^{(del)} \circ f^{(dup)}$ has at least $0.5\eta m$ 0s. We show both happen with low probability within a codeword.

A set of ℓ consecutive type-0 runs of 1s occurs with probability at most $10^{-10\ell}$. Thus the probability an inner codeword has a sequence of $\eta m/5$ consecutive type-0 runs of 1s is at most $m^2 \cdot 10^{-10\eta m/5} = \exp(-\Omega(\eta m))$. Now assume that in an inner codeword, each set of consecutive type-0 runs of 1s has size at most $\eta m/5$. Each set of ℓ consecutive type-0 runs of 1s merges $\ell + 1$ consecutive runs of 0s in c , so that they appear as a single longer run in s . The sum of the lengths of these $\ell + 1$ runs is some number ℓ^* that is at most $2\ell + 2$. The number of bits in the image of these runs of $c_i^{(in)}$ under $f^{(del)} \circ f^{(dup)}$ is distributed as $\text{Binomial}(\ell^*B/(1-p), 1-p)$. This has expectation $\ell^*B \leq 0.41B\eta m$, so by concentration bounds, the probability this run of s has length at least $0.5B\eta m$, i.e. is interpreted as a decoding buffer, is at most $\exp(-\Omega(\eta m))$. Hence, conditioned on each set of consecutive type-0 runs of 1s having size at most $\eta m/5$, the probability of having no spurious buffers in a codeword is at least $1 - \exp(-\Omega(\eta m))$. Thus the overall probability there are no spurious buffers a given inner codeword is at least $(1 - \exp(-\Omega(\eta m)))(1 - \exp(-\Omega(\eta m))) = 1 - \exp(-\Omega(\eta m))$. Since each inner codeword contains at most m candidate spurious buffers (one for each type-0 run of 1s), the expected number of spurious buffers in an inner codeword is thus at most $m \cdot \exp(-\Omega(\eta m))$. By our choice of $m \geq \Omega(\log(1/\delta_{out})/\eta)$, this is at most $\delta_{out}/10$. The occurrence of conditions (i) and (ii) above are independent between buffers. The total number of spurious buffers thus is bounded by the sum of n independent random variables each with expectation at most $\delta_{out}/10$. By concentration bounds, the probability that there are at least $\delta_{out}n/9$ spurious buffers is $2^{-\Omega(n)}$.

A deleted buffer occurs only when the image of the ηm 0s in a buffer under $f^{(del)} \circ f^{(dup)}$ is at most $B\eta m/2$. The number of such bits is distributed as $\text{Binomial}(B\eta m/(1-p), 1-p)$. Thus, each buffer is deleted with probability $\exp(-B\eta m) < \delta_{out}/10$ by our choice of $m \geq \Omega(\log(1/\delta_{out})/\eta)$. The events of a buffer receiving too many deletions are independent across buffers. By concentration bounds, the probability that there are at least $\delta_{out}n/9$ deleted buffers is thus $2^{-\Omega(n)}$.

Each inner decoding failure, spurious buffer, and deleted buffer increases the distance $\Delta_{i/d}(c_i^{(out)}, s_i^{(out)})$ by at most 3: each inner decoding failure causes up to 1 insertion and 1 deletion; each spurious buffer causes up to 1 deletion and 2 insertions; and each deleted buffer causes up to 2 deletions and 1 insertion. Our message is decoded incorrect if $\Delta_{i/d}(c_i^{(out)}, s_i^{(out)}) > \delta_{out}n$. Thus, there is a decoding error in the outer code only if at least one of (i) the number of incorrectly decoded inner codewords without spurious buffers or neighboring deleted buffers, (ii) the number of spurious buffers, or (iii) the number of deleted buffers is at least $\delta_{out}n/9$. However, by the above arguments, each is greater than $\delta_{out}n/9$ with probability $2^{-\Omega(n)}$, so there is a decoding error with probability $2^{-\Omega(n)}$. This concludes the proof of Theorem 1.

3.2 Possible Alternative Constructions

As mentioned in the introduction, Drinea and Mitzenmacher [4, 5] proved that the capacity of the BDC_p is at least $(1-p)/9$. However, their proof is nonconstructive and they do not provide an efficient decoding algorithm.

One might think it is possible to use Drinea and Mitzenmacher's construction as a black box. We could follow the approach in this paper, concatenating an outer code given by [10] with the rate $(1-p)/9$ random-deletion-correcting code as a black box inner code. The complexity of the Drinea and Mitzenmacher's so-called *jigsaw decoding* is not apparent from [5]. However, the inner code has constant length, so construction, encoding, and decoding would be constant time. Thus, the efficiency of the inner code would not affect the asymptotic runtime.

The main issue with this approach is that, while the inner code can tolerate random deletions with probability p , inner codeword bits are *not* deleted in the concatenated construction according to a BDC_p ; the 0 bits closer to the buffers between the inner codewords are deleted with higher probability because they might be "merged" with a buffer. For example, if an inner codeword is $\langle 101111 \rangle$, then because the codeword is surrounded by buffers of 0s, deleting the leftmost 1 effectively deletes two bits because the 0 is interpreted as part of the buffer. While this may not be a significant issue because the distributions of deletions in this deletion process and BDC_p are quite similar, much more care would be needed to prove correctness.

Our construction does not run into this issue, because our transmitted codewords tend to have *many* 1s on the ends of the inner codewords. In particular, each inner codeword of C_{in} has 1s on the ends, so after the Duplication step each inner codeword has $B/(1-p)$ or $2B/(1-p)$ 1s on the ends. The 1s on the boundary of the inner codeword will all be deleted with probability $\approx \exp(-B)$, which is small. Thus, in our construction, it is far more unlikely that bits are merged with the neighboring decoding buffer, than if we were to use a general inner code construction. Furthermore, we believe our construction based on bit duplication of a worst-case deletion correcting code is conceptually simpler than appealing to an existential code.

As a remark, we presented a construction with rate $(1-p)/110$, but using a randomized encoding we can improve the constant from $1/110$ to $1/60$. We can modify our construction

so that, during the Duplication step of decoding, instead of replacing each bit of $c^{(cat)}$ with a fix number $B/(1-p)$ copies of itself, we instead replaced each bit independently with $\text{Poisson}(B/(1-p))$ copies of itself. Then the image of a run r_j under duplication and deletion is distributed as $\text{Poisson}(B)$, which is independent of p . Because we don't have a dependence on p , we can tighten our bounding in (5) and (6). To obtain $(1-p)/60$, we can take $B = 28.12$ and set $B^* = 40$, where B^* is the threshold after which runs are decoded as two bits instead of one bit in the Deduplication step. The disadvantage of this approach is that we require our encoding to be randomized, whereas the construction presented above uses deterministic encoding.

4 Future work and open questions

A lemma due to Levenshtein [15] states that a code C can decode against pn adversarial deletions if and only if it can decode against pn adversarial insertions and deletions. While this does not automatically preserve the efficiency of the decoding algorithms, all the recent efficient constructions of codes for worst-case deletions also extend to efficient constructions with similar parameters for recovering from insertions and deletions [1, 7].

In the random error model, decoding deletions, insertions, and insertions and deletions are not the same. Indeed, it is not even clear how to define random insertions. One could define insertions and deletions via the Poisson repeat channel where each bit is replaced with a Poisson many copies of itself (see [4, 17]). However, random insertions do not seem to share the similarity to random deletions that adversarial deletions share with adversarial insertions; we can decode against arbitrarily large Poisson duplication rates, whereas for codes of block length n we can decode against a maximum of n adversarial insertions or deletions [5]. Alternatively one can consider a model of random insertions and deletions where, for every bit, the bit is deleted with a fixed probability p_1 , a bit is inserted after it with a fixed probability p_2 , or it is transmitted unmodified with probability $1 - p_1 - p_2$ [19]. One could also investigate settings involving memoryless insertions, deletions, and substitutions [16].

There remain a number of open questions even concerning codes for deletions only. Here are a few highlighted by this work.

1. Can we close the gap between $\sqrt{2} - 1$ and $\frac{1}{2}$ on the maximum correctable fraction of adversarial deletions?
2. Can we construct efficiently decodable codes for the binary deletion channel with better rate, perhaps reaching or beating the best known existential capacity lower bound of $(1-p)/9$?
3. Can we construct efficient codes for the binary deletion channel with rate $1 - O(h(p))$ for $p \rightarrow 0$?

References

- 1 Joshua Brakensiek, Venkatesan Guruswami, and Samuel Zbarsky. Efficient low-redundancy codes for correcting multiple deletions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1884–1892, 2016.
- 2 Boris Bukh, Venkatesan Guruswami, and Johan Håstad. An improved bound on the fraction of correctable deletions. *IEEE Trans. Information Theory*, 63(1):93–103, 2017.
- 3 Suhas Diggavi and Matthias Grossglauser. On transmission over deletion channels. In *Proceedings of the 39th Annual Allerton Conference on Communication, Control, and Computing*, pages 573–582, 2001.

- 4 Eleni Drinea and Michael Mitzenmacher. On lower bounds for the capacity of deletion channels. *IEEE Transactions on Information Theory*, 52(10):4648–4657, 2006.
- 5 Eleni Drinea and Michael Mitzenmacher. Improved lower bounds for the capacity of i.i.d. deletion and duplication channels. *IEEE Trans. Information Theory*, 53(8):2693–2714, 2007.
- 6 Dario Fertonani, Tolga M. Duman, and M. Fatih Erden. Bounds on the capacity of channels with insertions, deletions and substitutions. *IEEE Trans. Communications*, 59(1):2–6, 2011. doi:10.1109/TCOMM.2010.110310.090039.
- 7 Venkatesan Guruswami and Ray Li. Efficiently decodable insertion/deletion codes for high-noise and high-rate regimes. In *IEEE International Symposium on Information Theory, ISIT 2016, Barcelona, Spain, July 10-15, 2016*, pages 620–624, 2016. doi:10.1109/ISIT.2016.7541373.
- 8 Venkatesan Guruswami and Ray Li. Coding against deletions in oblivious and online models, 2017. Manuscript; arXiv abs/1612.06335. URL: <http://arxiv.org/abs/1612.06335>.
- 9 Venkatesan Guruswami and Carol Wang. Deletion codes in the high-noise and high-rate regimes. *IEEE Trans. Information Theory*, 63(4):1961–1970, 2017. doi:10.1109/TIT.2017.2659765.
- 10 Bernhard Haeupler and Amirbehshad Shahrashbi. Synchronization strings i: Codes for insertions and deletions approaching the singleton bound. *To appear in STOC'17*. <http://arxiv.org/abs/1704.00807>.
- 11 Adam Kalai, Michael Mitzenmacher, and Madhu Sudan. Tight asymptotic bounds for the deletion channel with small deletion probabilities. In *IEEE International Symposium on Information Theory, ISIT 2010, June 13-18, 2010, Austin, Texas, USA, Proceedings*, pages 997–1001, 2010. doi:10.1109/ISIT.2010.5513746.
- 12 Yashodhan Kanoria and Andrea Montanari. Optimal coding for the binary deletion channel with small deletion probability. *IEEE Trans. Information Theory*, 59(10):6192–6219, 2013. doi:10.1109/TIT.2013.2262020.
- 13 Ian Kash, Michael Mitzenmacher, Justin Thaler, and John Ullman. On the zero-error capacity threshold for deletion channels. In *Information Theory and Applications Workshop (ITA)*, pages 1–5, January 2011.
- 14 Adam Kirsch and Eleni Drinea. Directly lower bounding the information capacity for channels with i.i.d. deletions and duplications. *IEEE Trans. Information Theory*, 56(1):86–102, 2010. doi:10.1109/TIT.2009.2034883.
- 15 Vladimir I. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Dokl. Akad. Nauk*, 163(4):845–848, 1965 (Russian). English translation in *Soviet Physics Doklady*, 10(8):707–710, 1966.
- 16 Hugues Mercier, Vahid Tarokh, and Fabrice Labeau. Bounds on the capacity of discrete memoryless channels corrupted by synchronization and substitution errors. *IEEE Trans. Information Theory*, 58(7):4306–4330, 2012. doi:10.1109/TIT.2012.2191682.
- 17 Michael Mitzenmacher. A survey of results for deletion channels and related synchronization channels. *Probability Surveys*, 6:1–33, 2009.
- 18 Mojtaba Rahmati and Tolga M. Duman. Upper bounds on the capacity of deletion channels using channel fragmentation. *IEEE Trans. Information Theory*, 61(1):146–156, 2015. doi:10.1109/TIT.2014.2368553.
- 19 Ramji Venkataramanan, Sekhar Tatikonda, and Kannan Ramchandran. Achievable rates for channels with deletions and insertions. *IEEE Trans. Information Theory*, 59(11):6990–7013, 2013. doi:10.1109/TIT.2013.2278181.

- 20 S. M. Sadegh Tabatabaei Yazdi and Lara Dolecek. A deterministic polynomial-time protocol for synchronizing from deletions. *IEEE Trans. Information Theory*, 60(1):397–409, 2014. doi:10.1109/TIT.2013.2279674.

A Proof of Lemma 3

Proof. For each i , we can find a random variable B_i such that $B_i \geq A_i$ always, B_i takes values in $[0, \beta]$, and $\mathbb{E}[B_i] = \alpha$. Applying Lemma 2 gives

$$\begin{aligned}
 \Pr \left[\sum_{i=1}^n A_i \geq n\gamma \right] &\leq \Pr \left[\sum_{i=1}^n B_i \geq n\gamma \right] \\
 &\leq \Pr \left[\sum_{i=1}^n \frac{B_i}{\beta} \geq \left(1 + \left(\frac{\gamma - \alpha}{\alpha} \right) \right) \frac{n\alpha}{\beta} \right] \\
 &\leq \exp \left(- \frac{\left(\frac{\gamma - \alpha}{\alpha} \right)^2 \cdot \frac{n\alpha}{\beta}}{3} \right) \\
 &= \exp \left(- \frac{(\gamma - \alpha)^2 n}{3\alpha\beta} \right). \quad \blacktriangleleft
 \end{aligned}$$