

Pseudorandom Error-Correcting Codes

Miranda Christ*
Columbia University

Sam Gunn*†
UC Berkeley

Abstract

We construct *pseudorandom error-correcting codes* (or simply *pseudorandom codes*), which are error-correcting codes with the property that any polynomial number of codewords are pseudorandom to any computationally-bounded adversary. Efficient decoding of corrupted codewords is possible with the help of a decoding key.

We build pseudorandom codes that are robust to substitution and deletion errors, where pseudorandomness rests on standard cryptographic assumptions. Specifically, pseudorandomness is based on either $2^{O(\sqrt{n})}$ -hardness of LPN, or polynomial hardness of LPN and the planted XOR problem at low density.

As our primary application of pseudorandom codes, we present an undetectable watermarking scheme for outputs of language models that is robust to cropping and a constant rate of random substitutions and deletions. The watermark is undetectable in the sense that any number of samples of watermarked text are computationally indistinguishable from text output by the original model. This is the first undetectable watermarking scheme that can tolerate a constant rate of errors.

Our second application is to steganography, where a secret message is hidden in innocent-looking content. We present a constant-rate stateless steganography scheme with robustness to a constant rate of substitutions. Ours is the first stateless steganography scheme with provable steganographic security and *any* robustness to errors.

*Equal contribution. Email addresses: mchrist@cs.columbia.edu, gunn@berkeley.edu

†Supported by a Google PhD Fellowship.

Contents

1	Introduction	3
1.1	An approach to watermarking	3
1.2	Our contributions	5
1.3	Related work: short summary	7
1.4	Organization	8
1.5	Differences from a previous version	8
2	Technical overview	10
2.1	Pseudorandom code basics	10
2.2	Pseudorandom LDPC codes	11
2.3	Pseudorandom codes for the deletion channel	14
2.4	Constant-rate pseudorandom codes	16
2.5	Watermarks for language models	16
2.6	Watermarks with public attribution	18
2.7	Robust steganography	19
3	Preliminaries	21
3.1	Cryptography preliminaries	21
3.2	Coding theory preliminaries	22
4	Pseudorandom code basics	23
4.1	Definitions	23
4.2	Heuristic construction from permuted codes	24
5	Constructing pseudorandom codes from cryptographic assumptions	26
5.1	The zero-bit construction	26
5.2	Codeword detection (zero-bit decoding)	27
5.3	Pseudorandomness from the planted XOR assumption and LPN	29
5.4	Pseudorandomness from subexponential LPN	31
6	Boosting the rate and robustness of any pseudorandom code	38
6.1	Multi-bit pseudorandom codes	38
6.2	Constant-rate pseudorandom codes	40
6.3	Pseudorandom codes for the deletion channel	41
7	Application: watermarking for language models	46
7.1	Watermarking preliminaries	46
7.2	Robust watermarking from PRCs	49
7.3	Language model steganography	55
7.4	Watermarks with public attribution	55
8	Application: universal steganography	59
8.1	Steganography preliminaries	59
8.2	Robust stateless steganography	60
8.3	Stateless steganography from weaker modeling assumptions	61
A	Related work	66

1 Introduction

The proliferation of AI-generated content is one of the biggest issues facing the internet today. Recent breakthroughs in large language models have made it increasingly difficult to distinguish this influx of AI-generated content from human-generated content.

A promising solution for detecting AI-generated content is *watermarking*, where a hidden signal is embedded in the content. Several major companies, including OpenAI and Google, have pledged to embed watermarks in content output by their models [BH23]. Despite this explosion of interest in watermarking, there are very few techniques for building watermarking schemes that do not noticeably alter the generated content. Existing schemes incur trade-offs between the quality of generated content, the robustness of the watermark, and the computational complexity of detection.

In this work, we take a new cryptographic approach to this problem that allows us to avoid some of these trade-offs. Our approach is based on a new cryptographic primitive that we call a pseudorandom error-correcting code, or simply a pseudorandom code (PRC). A PRC is an error-correcting code that is parameterized by a decoding key. The pseudorandomness property states that, without this decoding key, any polynomial number of codewords are pseudorandom.

We find that the problem of building robust, quality-preserving watermarks reduces to the problem of building PRCs. Essentially, the watermarking strategy is to replace some of the randomness used by the generative algorithm with outputs from a PRC.

Building PRCs is challenging: Error-correcting codes are typically highly structured, while pseudorandomness implies a lack of discernible structure. Indeed, a priori it is not clear that such objects should exist. Nonetheless, we construct PRCs from standard (subexponential) cryptographic assumptions. Our constructions are related to low-density parity-check codes, and we base pseudorandomness on the Learning Parity with Noise assumption. We construct PRCs with strong robustness properties, including robustness to any constant rate of substitution and deletion errors.

Applying these PRCs to watermarking for language models, we obtain the first quality-preserving language model watermarking schemes that are robust to cropping and any constant rate of substitution and deletion errors. That is, the watermark detector will work as long as it is provided any sufficiently-long sequence of text, even if that text is subjected to any constant (less than 1/2) rate of random substitutions and any constant (less than 1) rate of random deletions.

1.1 An approach to watermarking

In this subsection we present a simple, general template for watermarking AI-generated content. The template described here can in principle be used to watermark arbitrary media, but we only present concrete instantiations in certain contexts (Sections 7 and 8).

In all generative AI settings there is a generative algorithm, `Generate`, that defines the behavior of the AI. A user provides a prompt, and `Generate` outputs some randomly-generated content. A watermarking scheme modifies `Generate` so that the generated content contains a hidden pattern, called a watermark. There are two essential requirements that any watermarking scheme should satisfy:

- *Quality*: embedding the watermark should not reduce the quality of generative algorithm; and
- *Robustness*: the watermark should be detectable in generated content, even if this content is corrupted.

Achieving both of these properties simultaneously is the central challenge of watermarking. Quality means that the watermark should not significantly alter the generated content, while robustness seems to require the watermark to change the content a great deal.

In this work, we propose a new strategy for watermarking: *replacing the randomness used by `Generate` with codewords from a pseudorandom error-correcting code*. A pseudorandom error-correcting code, or simply

pseudorandom code (PRC), is a new cryptographic primitive that we introduce in this work. A PRC is defined by algorithms `Encode`, `Decode` satisfying two properties:

- *Pseudorandomness*: Any efficient adversary, without knowledge of the decoding key, cannot distinguish between oracle access to `Encode` and an oracle that always outputs a fresh random string; and
- *Error correction* or *robustness*: For any message m , if $x \leftarrow \text{Encode}(m)$ and x' is a “corrupted” version of x where the amount of corruption is bounded, then $\text{Decode}(x') = m$.

For watermarking the message can be simply $m = 1$, indicating the presence of a watermark.¹

In order to make detection possible, we specify `Generate` in such a way that the detector can approximate the randomness used to produce any given content. To test for the presence of a watermark, the detector computes this approximation and then applies `Decode` to the result. If the content is watermarked and the approximation is close enough to the true randomness, then robustness of the PRC ensures that `Decode` returns 1. This indicates to the detector that the watermark is present. Stronger robustness of the PRC translates to stronger robustness of the watermark.

Pseudorandomness of the PRC guarantees that, without the decoding key, watermarked content is indistinguishable from content produced by the original generative algorithm — even if one is allowed to see many outputs. In particular, the quality of the generative algorithm is not deteriorated by the watermark. In [CGZ23], such a watermark is referred to as *undetectable*.

Therefore, the problem of building robust, quality-preserving watermarks reduces to building PRCs (and appropriately specifying `Generate`). Below, we describe the application of this template to watermarking for language models.

Watermarks for language models. A generative language model is a randomized algorithm that takes as input a prompt, and samples text constituting a response. This text consists of *tokens*. For simplicity, we assume here that tokens are binary and that the full response is of a fixed length n . Neither of these assumptions is important for our results, as discussed in Section 7.

Given any generative language model, it is not difficult to define an algorithm `Generate` that takes a prompt and a random seed $x \in \{0, 1\}^n$, and samples a response $t \in \{0, 1\}^n$ such that

- if x is uniformly random, then t is distributed identically to the original model, and
- each bit of t is correlated with the corresponding bit of x .

See Section 2.5 for an example of such an algorithm. Now it is easy to recover an approximation of the randomness x that was used to produce a given text: just use the text t itself.

One natural watermarking strategy is to use the *same* random seed $x \in \{0, 1\}^n$ for every call to `Generate`, storing x itself as the watermarking key. This is essentially the strategy used by [KTHL23], and it yields a highly robust watermark because the detector can compute the edit distance between the given text and x . The resulting quality guarantee is that a single response from the watermarked model is distributed identically to a single response from the original model.

However, this strategy results in redundancy of responses because the i^{th} token of every response is biased towards the i^{th} bit of x . This is problematic as it limits the variability of text that the language model generates. For instance, it should not be the case that certain words are preferred as the first word of every response. One can mitigate this issue by storing a family of seeds $x_1, \dots, x_\ell \in \{0, 1\}^n$ and randomly choosing one such seed for each response. Increasing ℓ improves the variability, but comes at the cost of a corresponding increase in both watermark key length *and* detector runtime. Now, the detector must

¹By instead encoding a longer message in the PRC, this technique extends to steganography — where messages are secretly communicated in innocent-looking content — as well.

compute the edit distance for each seed, resulting in a runtime of $O(n^2\ell)$. In particular, for any polynomial-time watermarking scheme using this approach, there exists a polynomial-time adversary that can distinguish watermarked content from un-watermarked content without needing the watermark detection key.

A PRC is exactly the object needed to avoid this tradeoff between response variability and efficiency. Our watermark detection key will be the decoding key for a PRC, and we will embed the watermark by sampling a fresh pseudorandom seed $x \leftarrow \text{Encode}(1)$ for each query to **Generate**. This results in no observable correlations between responses, regardless of the number of queries — i.e., the watermark is undetectable. Since the same hidden structure is present in every sample from **Encode**(1), our detector can simply apply **Decode** to check for this structure. Now, the detector’s runtime has no dependence on the response variability. Using PRCs that we construct in Sections 5 and 6, we also find that our watermarking scheme can be made robust to a constant fraction of random substitution and deletion errors.

A note on undetectability. Undetectability is a strong quality guarantee for watermarking. Outputs of the watermarked model must be *computationally indistinguishable* from outputs of the original model, even to a user who is allowed to make many queries. While this is imperative for steganography, its necessity in watermarking is less clear, as some noticeable changes to the model may be permissible as long as the outputs remain “high-quality.”

However, measuring the quality of watermarked content is often challenging or impossible — especially when the content is used in a wide range of applications. Computational indistinguishability is a strong quality guarantee that applies uniformly to every application: it implies that the watermark causes no observable loss in *any* efficiently-computable quality metric. Without such a guarantee, it is impossible to verify that a watermark is quality-preserving across all applications. We therefore focus on undetectability in this work.

1.2 Our contributions

Pseudorandom codes (Sections 4 to 6). Our first contribution is to identify PRCs as an interesting cryptographic primitive, with applications to robustly hiding messages in AI-generated content. Very roughly, the definition is as follows — for more details, see either the technical overview (Section 2.1) or the formal definitions (Section 4.1).

Definition (Definitions 1 and 2). A *pseudorandom code* (PRC) is an error-correcting code where codewords are pseudorandom to any computationally-bounded adversary who doesn’t hold the decoding key.

We consider both public- and secret-key variants of PRCs. For a public-key PRC, the encoding algorithm can be made public without sacrificing pseudorandomness. When the message space consists of only a single message, we call it a *zero-bit* PRC. Zero-bit PRCs can also be viewed as robust *backdoored* (or *trapdoor*) *pseudorandom generators* [VV83, DGG⁺15], and they are sufficient for our applications to watermarking.

We show how to build zero-bit public-key PRCs related to low-density parity-check (LDPC) codes, where pseudorandomness rests on standard (subexponential) cryptographic assumptions. All of the PRCs we construct are over a binary alphabet. Depending on the parameter choices, the pseudorandomness of these LDPC-related codes is based on either of two assumptions, which we state together as Assumption 1.

Assumption (Assumption 1). Either

- LPN is hard for any $2^{O(\sqrt{n})}$ -time adversary, or
- LPN and planted XOR are both hard for any polynomial-time adversary.

For descriptions of the LPN and planted XOR assumptions, see either Section 2.2 or Section 5. Under Assumption 1, we prove that there exist PRCs with robustness to channels that introduce a bounded number of substitution errors. We say that any channel that introduces at most a p fraction of substitution errors (and no other types of errors) is *p-bounded*.

Theorem (Theorems 1 and 2). *Let $p \in (0, 1/2)$ be any constant. Under Assumption 1, there exists a zero-bit public-key PRC that is robust to every p -bounded channel.*

For some applications it will be useful to have multi-bit PRCs. Any such construction should ideally have a high *rate*, which is the ratio of the number of message bits to the number of codeword bits. We prove that any zero-bit PRC can be combined with any error correcting code to give a multi-bit PRC.

Theorem (Theorem 3). *Suppose that there exists a zero-bit (public-key) PRC and a rate- R error-correcting code, that are both robust to every p -bounded channel. Then there exists a (public-key) PRC of rate $R - o(1)$ that is robust to every $(p - \varepsilon)$ -bounded channel, for every constant $\varepsilon > 0$.*

Applying this theorem with our zero-bit LDPC-based PRCs and the binary error-correcting codes of [ABN⁺92, NN93, Ta-17], we have the following corollary.

Corollary. *Let $p \in (0, 1/2)$ be any constant. Under Assumption 1, there exists a constant-rate PRC that is robust to every p -bounded channel.*

None of the PRCs mentioned so far can handle deletions. Deletions are a particularly important type of edit for text watermarks, because an adversary may try to remove the watermark by simply deleting some of the words.

Deletions are notoriously difficult to handle in error correction, and standard techniques involve significant structure — thus violating pseudorandomness. Nonetheless, we show that if a PRC has sufficiently strong robustness to substitutions, then it can be converted to a PRC with robustness to deletions (at the cost of a decreased rate).

Let BSC_p be the binary symmetric channel with error rate p , and BDC_q be the binary deletion channel with deletion rate q . That is, BSC_p randomly flips each bit with probability p and BDC_q randomly deletes each bit with probability q .

Theorem (Theorem 4). *For any constants $p \in (0, 1/2)$ and $q \in (0, 1)$, there exists $p' \in (0, 1/2)$ such that the following holds. If there exists a zero-bit (public-key) PRC with robustness to every p' -bounded channel, then there exists a zero-bit (public-key) PRC that is robust to the channel $BSC_p \circ BDC_q$.*

Together with our LDPC-based PRCs, we obtain the following result.

Corollary. *Let $p \in (0, 1/2)$ and $q \in (0, 1)$ be any constants. Under Assumption 1, there exists a zero-bit public-key PRC that is robust to the channel $BSC_p \circ BDC_q$.*

Watermarking and steganography for language models (Section 7). We apply our zero-bit PRCs to build the first undetectable watermarking scheme for language models with robustness to a constant rate of random substitutions and deletions. In this section, we assume that text output by the language model is represented as a bitstring. Arbitrary text can be mapped to a bitstring by either randomly assigning a single bit to each token, or by expanding the tokens into a binary representation.

Theorem (Theorem 5). *Let $p \in (0, 1/2)$ be any constant. Under Assumption 1, there exists an undetectable watermarking scheme \mathcal{W} such that the watermark appears in any sufficiently high-entropy text, even if the text is subjected to the channel BSC_p .*

Under an extra assumption about the generated text, which roughly corresponds to the text having few repeated words, we can strengthen this to handle deletions as well.

Theorem (Theorem 6). *Let $p \in (0, 1/2)$ and $q \in (0, 1)$ be any constants. Under Assumption 1, there exists an undetectable watermarking scheme \mathcal{W} such that the watermark appears in any sufficiently high-entropy and “variable” text, even if the text is subjected to the channel $BSC_p \circ BDC_q$.*

In all of our theorems the text can be cropped, as long as the remaining text is sufficiently high-entropy.

We also construct undetectable watermarking schemes with *unforgeable public attribution* and the same robustness as \mathcal{W} . Public attribution means that there is a public algorithm to identify which portion of a given text was output by the model. Unforgeability means that no efficient user can produce text that the attribution algorithm identifies as model-generated, but that was not output by the model. Interestingly, our schemes retain the standard robust secret-key detector in addition to this public attribution algorithm.

Theorem (Theorem 8). *Under Assumption 1, there exists a watermarking scheme \mathcal{W}_{att} that retains all properties of \mathcal{W} from Theorem 5 or Theorem 6, and additionally satisfies unforgeable public attribution.*

Finally, using PRCs with constant rate (Theorem 3), we also obtain the first *robust* language model steganography scheme.²

Theorem (Theorem 7). *Let $p \in (0, 1/2)$ be any constant. Under Assumption 1, there exists a language model steganography scheme with constant information rate, such that the message can be recovered from any sufficiently high-entropy text, even if the text is subjected to the channel BSC_p .*

Universal steganography (Section 8). In Section 8 we show that PRCs can be used to solve a long-standing open question in steganography: A simple application of PRCs yields the first robust, stateless universal steganography scheme. Universal steganography can be used for language model steganography, but it is more general [vAH04]. We take the rate of the steganography scheme to be the ratio of the number of stegotext symbols to the number of bits in the message being encoded.

Theorem (Theorem 9). *Suppose there is a hash function that is unbiased over the covertext channel. If PRC is any PRC, then there exists a stateless, public-key universal steganography scheme with the same rate and robustness as PRC.*

Finally, we show that this result can be extended to the setting where an unbiased hash function on the covertext channel is not known, with some loss in robustness.

Theorem (Theorem 10). *Suppose there is a hash function that has constant min-entropy over the covertext channel. Then under Assumption 1, for any $p \in (0, 1/2)$, there exists a constant-rate public-key stateless steganography scheme that is robust to a p rate of random substitutions.*

1.3 Related work: short summary

We briefly outline some of the related work here. See Appendix A for a more complete discussion.

- Code-based cryptography: Our work bears some similarity to the field of code-based cryptography. However, code-based cryptography is generally focused on building existing primitives from new assumptions — whereas PRCs are a new primitive that we base on existing assumptions.
- Trapdoor pseudorandom generators: Our zero-bit PRCs can be equivalently viewed as *robust* trapdoor (or equivalently, backdoored) pseudorandom generators [VV83, DGG⁺15]. That is, we require the additional property that the trapdoor (or secret key) can be used to detect even *corrupted* pseudorandom strings.
- Watermarking for language models: We build watermarking schemes for language models satisfying the strongest quality guarantee, undetectability. Undetectability was defined by [CGZ23], where undetectable watermarks for language models were also constructed. In that work and in [Aar22, KGW⁺23a], it is essential for watermark detection that many sufficiently-long contiguous substrings of the response remain *unchanged*. Therefore, these watermarks are easily removable by simple attacks (see the “robustness of our watermark” paragraph of Section 2.5). The watermarks of [KTHL23] are more robust — their robustness is more comparable to ours — but they sacrifice undetectability. Instead, their watermarks satisfy the weaker property of distortion-freeness, which is the single-query

²Since this scheme doesn’t rely on the decoder having access to the prompt, it can also be seen as an undetectable “multi-bit watermarking scheme.”

version of undetectability. [ZALW23] obtain even stronger robustness, at the cost of even further reduced quality.

- Impossibility of strong watermarks: [ZEF+23] explore the possibility of watermarking for language models in the presence of motivated adversaries. They argue that sampling a *random* response is easier when one is provided *any* response. Since a random response cannot be watermarked (or else there would be a high false-positive rate), they use this to argue that any watermarked language model necessarily provides some assistance in generating un-watermarked text.
- Steganography: Steganography is the study of concealing secret messages in innocent-looking content. Whereas encryption is about hiding the message, steganography is about hiding the *existence* of the message. Ever since steganography was formalized by [HLVA02], *robust* steganography schemes (that don't require a shared state) have remained elusive. We resolve this problem using PRCs.

1.4 Organization

In Section 2, we give a technical overview of the paper, which is self-contained and sufficient to understand the high-level ideas and results of each section. In Section 3, we state relevant preliminaries and notation. In Section 4, we formally define PRCs and provide a heuristic construction. In Section 5, we build PRCs from LDPC codes and prove their pseudorandomness from standard cryptographic assumptions. We then show in Section 6 how to improve the rate and robustness of any PRC, resulting in our constant-rate multi-bit PRCs and PRCs for deletion channels. In Section 7, we present our language model watermarking schemes from PRCs, including both our standard watermarks and our watermarks with unforgeable public attribution. Finally, in Section 8, we show how PRCs can be used to construct robust universal steganography schemes.

Appendix A gives a more comprehensive overview of related work.

1.5 Differences from a previous version

This version of the paper contains two significant technical updates relative to the previous version, as well as a few editorial updates. The two significant updates address subtle technical issues, but do not substantially change any of the ideas or messages of the paper.

The first significant update is to the choice of parameters for which we invoke the planted XOR assumption (Assumption 3). The previous version of the paper invoked the assumption with m , the number of samples, set to $\Omega(n)$. However, as pointed out to us by an anonymous CRYPTO 2024 reviewer, Theorem 4.26 of [ASS+23] (which is itself an updated version of a paper that we had cited in the old version of this paper) shows this assumption to be false. Fortunately we had only set $m = \Omega(n)$ because it made notation more convenient, and by instead setting $m = n^{1-\Omega(1)}$, we avoid the attack without sacrificing any of our results. We have updated all of the relevant theorem statements and proofs in this version. We emphasize that this issue did not affect our other main proof that our LDPC-based PRCs are pseudorandom, because that proof only relies on LPN.

The other significant update regards the completeness and robustness of our watermarking scheme in Section 7. To sample a long response, we repeatedly sample fresh PRC codewords to bias the text. However, in order to argue that the resulting text contains the watermark, we need to ensure that the channel applied to the text is p -bounded — which, according to our definition, means that the error channel is not allowed to depend on the PRC key. But if the response contains multiple PRC blocks, then the error channel on one block could depend on codewords embedded in the previous blocks, violating this condition. In the prior version we had missed this issue, making our completeness and robustness claims incorrect. In this updated version we apply a one-time pad to each PRC codeword before embedding it, to ensure that the error channel cannot depend on the PRC key — thus resolving the issue.

Acknowledgements. We thank Yael Kalai, Venkat Guruswami, Rainer Böhme, Or Zamir, Shyamal Patel, and Shivam Nadimpalli for helpful research conversations. We additionally thank Vinod Vaikuntanathan for pointing out the LDPC-based PRC variant that we use to present Lemma 9 in the Technical Overview. We thank Mihalis Yannakakis and Fermi Ma for helpful suggestions about the write-up. We thank Omar Arabiah for help in proving Lemma 9, as well as general assistance with coding theory. We thank an anonymous CRYPTO 2024 reviewer for helpful feedback, especially regarding our use of the planted XOR assumption (Assumption 3).

2 Technical overview

2.1 Pseudorandom code basics

Pseudorandom codes (PRCs) can be viewed as a combination of two related primitives:

- Pseudorandom encryption, where ciphertexts are indistinguishable from random under a chosen plaintext attack [RBB03]. Secret-key pseudorandom encryption is easy to build using a pseudorandom function F_{sk} — just encrypt m by sampling a random r and outputting $(r, m \oplus F_{\text{sk}}(r))$. Public-key pseudorandom encryption is also known from standard assumptions [vAH04]. However, none of these constructions have any nontrivial robustness.
- Robust encryption, where encryptions of messages are robust to errors. Robust encryption is easy to build by applying an error-correcting code to ciphertexts from any standard encryption scheme. Even if that encryption scheme is pseudorandom, the use of the error-correcting code will in general render the robust encryption scheme not pseudorandom.

A PRC is required to simultaneously satisfy *both* pseudorandomness and robustness — properties that are in direct tension with each other. Using the secret key, one should be able to discern the redundancy and structure that give ciphertexts their robustness. Without the secret key, ciphertexts must appear completely unstructured.

We define secret-key PRCs below. For public-key PRCs (Definition 2), we further require that the encoding algorithm can be made public without sacrificing pseudorandomness.

Definition (Definition 1). Let Σ be an alphabet and $\mathcal{E} : \Sigma^* \rightarrow \Sigma^*$ be a channel. A *secret-key PRC* with robustness to \mathcal{E} is described by algorithms $\text{Encode}_{\text{sk}} : \Sigma^k \rightarrow \Sigma^n$ and $\text{Decode}_{\text{sk}} : \Sigma^* \rightarrow \Sigma^k \cup \{\perp\}$, parameterized by a secret key sk , satisfying the following criteria for every security parameter λ :

- (Error correction, or robustness) For any message $m \in \Sigma^k$,

$$\Pr_{\text{sk}}[\text{Decode}_{\text{sk}}(\mathcal{E}(x)) = m : x \leftarrow \text{Encode}_{\text{sk}}(m)] \geq 1 - \text{negl}(\lambda).$$

- (Soundness) For any fixed $c \in \Sigma^*$,

$$\Pr_{\text{sk}}[\text{Decode}_{\text{sk}}(c) = \perp] \geq 1 - \text{negl}(\lambda).$$

- (Pseudorandomness) For any polynomial-time adversary \mathcal{A} ,

$$\left| \Pr_{\text{sk}}[\mathcal{A}^{\text{Encode}_{\text{sk}}}(1^\lambda) = 1] - \Pr_{\mathcal{U}}[\mathcal{A}^{\mathcal{U}}(1^\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $\mathcal{A}^{\mathcal{U}}$ means that the adversary has access to an oracle that, on any (even previously queried) input, responds with a freshly drawn uniform value in Σ^n .

If the scheme can only encode a singular message (i.e. $k = 0$), then we call it a *zero-bit* PRC. Soundness is a technical condition that we include only to ensure that zero-bit PRCs are non-trivial.

For a sufficiently weak channel \mathcal{E} , it is not hard to construct a secret-key PRC with robustness to \mathcal{E} where pseudorandomness rests on very mild assumptions. For instance, if $F_{\text{sk}} : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a pseudorandom function, we can build a zero-bit secret-key PRC with the following encoding algorithm:

$\text{Encode}_{\text{sk}}(1)$:

1. Randomly sample $x_1, \dots, x_\ell \leftarrow \{0, 1\}$.
2. For $i = \ell + 1, \dots, n$, let $x_i = F_{\text{sk}}(x_{i-\ell}, \dots, x_{i-1})$.

3. Output x_1, \dots, x_n .

The decoding algorithm $\text{Decode}_{\text{sk}}$ simply checks whether much more than a $1/2$ fraction of conditions $x_i = F_{\text{sk}}(x_{i-\ell}, \dots, x_{i-1})$ are satisfied. Pseudorandomness follows by taking ℓ to be the security parameter. It is immediate that this PRC is robust to any length-preserving channel that introduces at most a p fraction of errors, for $p \ll 1/\ell$. We call any such channel *p-bounded*.

This secret-key PRC construction can be seen as implicit in prior hash-based watermarking schemes [Aar22, KGW⁺23a, CGZ23], where essentially the same level of robustness to a $1/\ell$ error rate is obtained. Unfortunately, it has an inherent trade-off between pseudorandomness and robustness: After roughly $2^{\ell/2}$ samples from $\text{Encode}_{\text{sk}}(1)$, there will be repeated prefixes and therefore correlations between the samples. In particular, if we want this PRC to be robust to a constant rate of errors, we have to set $\ell = O(1)$, in which case even a constant number of queries are enough to observe correlations. We therefore turn to alternative constructions.

2.2 Pseudorandom LDPC codes

Fortunately, one of the most prominent assumptions in theoretical cryptography is a statement about codes: Learning Parity with Noise (LPN). The LPN assumption states that noisy samples from the codespace of a random linear code are pseudorandom, even to an adversary who knows a generator matrix for the code. In more detail, let n, g be integers and $G \leftarrow \mathbb{F}_2^{n \times g}$ be a random matrix. The LPN assumption (with noise rate η and secrets of size g) states that $(G, Gs \oplus e) \approx (G, u)$,³ where $s \leftarrow \mathbb{F}_2^g$, $e \leftarrow \text{Ber}(n, \eta)$, and $u \leftarrow \mathbb{F}_2^n$.

The LPN assumption suggests using noisy codewords from a random linear code as a PRC. That is, let $G \leftarrow \mathbb{F}_2^{n \times g}$ be the secret key and Encode_G be the following zero-bit secret-key PRC encoder:

- $\text{Encode}_G(1)$: Sample $s \leftarrow \mathbb{F}_2^g$ and $e \leftarrow \text{Ber}(n, \eta)$. Output $Gs \oplus e$.

The LPN assumption immediately implies that an arbitrary polynomial number of samples from $\text{Encode}_G(1)$ are pseudorandom. However, recall that the LPN assumption states that these samples are pseudorandom *even given G* — which means precisely that there does not exist an efficient zero-bit decoder $\text{Decode}_G(x)$!

While this random linear code construction does not work, it naturally suggests a strategy that does. If we find a sampling procedure that produces a random (or even pseudorandom) generator matrix G *together with a trapdoor for efficient decoding*, then we have a public-key PRC where the generator matrix is the public encoding key and the trapdoor is the secret decoding key. By the LPN assumption, $\text{Encode}_G(1)$ produces pseudorandom vectors even to an adversary who knows G , so the construction will satisfy pseudorandomness.

It turns out that low-weight parity checks can serve as such a trapdoor. That is, instead of sampling G uniformly at random, we first sample a “parity-check matrix” $P \in \mathbb{F}_2^{r \times n}$ with sparse rows (i.e., “low density”), and then sample $G \in \mathbb{F}_2^{n \times g}$ subject to $PG = 0$. For appropriate choice of n, g, t, r , we will show that the resulting marginal distribution on G is random or pseudorandom. The low-density parity-check matrix P will allow for efficient detection of near-codewords.

Codes defined by Low-Density Parity-Check matrices are called LDPC codes. For $n, g, t, r \in \mathbb{N}$, we define an (n, g, t, r) random LDPC code by the following distribution over parity-check and generator matrices:

LDPC $[n, g, t, r]$:

1. Sample a random matrix $P \in \mathbb{F}_2^{r \times n}$ subject to every row of P being t -sparse.
2. Sample a random matrix $G \in \mathbb{F}_2^{n \times g}$ subject to $PG = 0$.
3. Output (P, G) .

Zero-bit decoding works by counting the number of satisfied parity checks. For any fixed $x \in \mathbb{F}_2^n$, with high probability over $P \in \mathbb{F}_2^{r \times n}$ we expect that the number of unsatisfied parity checks, $\text{wt}(Px)$, is roughly $r/2$.

³Throughout this work we use \approx to refer to computational indistinguishability.

But if x is close to $\text{Im } G \subseteq \ker P$ in Hamming distance, then as long as the error and the sparsity t of the parity checks are not too high, we expect $\text{wt}(Px)$ to be significantly smaller than $r/2$.

Therefore our zero-bit pseudorandom LDPC code uses the following zero-bit decoding algorithm:

- $\text{Decode}_P(x)$: If $\text{wt}(Px) < (1/2 - r^{-1/4}) \cdot r$, output 1; otherwise output \perp .⁴

Encoding is exactly the same Encode_G algorithm as above — but now that G is sampled together with the trapdoor P , we have an efficient decoding algorithm. Observe that this is a zero-bit scheme, because the decoder only determines whether the input is related to the keys or not. Using belief propagation, it is possible to push this construction beyond a zero-bit PRC, although this results in lower robustness. We ultimately construct a constant-rate multi-bit PRC by other means, which we discuss in Section 6.2.

Let LDPC-PRC_0 be the zero-bit public-key PRC defined by $(\text{Encode}_G, \text{Decode}_P)$ for $(P, G) \leftarrow \text{LDPC}[n, g, t, r]$. In a moment we will outline our proofs that LDPC-PRC_0 is a public-key PRC with very strong robustness. First, let us see some restrictions on the sparsity parameter t that provide important context for these proofs.

If random noise of rate $1/2 - \varepsilon$ is applied to $x \in \text{Im } G$, then the probability of each parity check being satisfied for the noisy codeword is $1/2 - (2\varepsilon)^t/2$. So in order for $\text{Decode}_P(x)$ to output 1 with high probability, we need $(2\varepsilon)^t/2 > r^{-1/4}$, i.e., $t < 1 + \log r/4 \log(1/2\varepsilon) = O(\log r)$. We will always have $r = n^{\Omega(1)}$, so this restriction says that $t = O(\log n)$ for appropriate choice of constant.

On the other hand, if we set $t = O(1)$ then $\text{Encode}_G(1)$ cannot be pseudorandom. This is because it is possible to brute-force search over all $\binom{n}{t}$ possible parity checks of weight t , and one can test whether Encode_G is consistent with a given parity check $s \in \mathbb{F}_2^n$ by simply computing $s \cdot x$ for many samples $x \leftarrow \text{Encode}_G(1)$.

Therefore, we will choose $t = \Theta(\log n)$ in order to rely on the weakest possible cryptographic assumption for pseudorandomness, without sacrificing robustness to a constant noise rate.

Remark. *The LDPC codes considered in this work differ from the traditional Gallager’s LDPC ensemble in two important ways. First, our LDPC codes will have $t = \Theta(\log n)$ sparsity as opposed to constant sparsity. Unfortunately, the usual belief propagation decoder does not work for noise rates beyond $O(\log t/t)$; this is the reason why we only perform the simple zero-bit decoding. The second difference is that we use independent parity checks, which results in an irregular Tanner graph.*

Remark. *There is a well-known public-key encryption scheme, due to Alekhnovich [Ale03, Cryptosystem 1], based on a low-noise variant of LPN. This scheme is similar to ours, but the decoder cannot tolerate any constant rate of errors.*

Pseudorandom generator matrix (Lemma 8). For appropriate choices of parameters, it turns out that the generator matrix of $\text{LDPC}[n, g, t, r]$ is pseudorandom under the planted t -XOR assumption. The planted t -XOR problem (and its generalization, the planted t -SUM problem) is a natural and well-studied problem — see e.g. [ASS+23] for a more detailed discussion. Formally, the (n, m, t) *planted XOR problem* states that it is computationally hard to distinguish between

\mathcal{D}_0^m : a random m -dimensional linear subspace $V \subseteq \mathbb{F}_2^n$, and

\mathcal{D}_1^m : a random m -dimensional linear subspace $V_s \subseteq \mathbb{F}_2^n$ satisfying a random planted t -XOR relation s (i.e., s is a random t -sparse vector and $s \cdot v = 0$ for all $v \in V_s$).

Throughout this overview, we consider linear subspaces to be described by a random basis. Recalling the definition of $(P, G) \leftarrow \text{LDPC}[n, g, t, r]$, if $r = 1$ the (n, g, t) planted XOR assumption immediately implies that G is pseudorandom (by identifying V with $\text{Im } G$). But for the more interesting case that $r > 1$, we require a stronger version of the planted XOR assumption with many planted relations. That is, we need to assume that the following distribution is indistinguishable from \mathcal{D}_0^m :

⁴Throughout this work, wt will refer to Hamming weight.

\mathcal{D}_r^m : a random m -dimensional linear subspace $V_{s_1, \dots, s_r} \subseteq \mathbb{F}_2^n$ satisfying r random planted t -XOR relations s_1, \dots, s_r (i.e., s_1, \dots, s_r are random t -sparse vectors and $s_1 \cdot v = \dots = s_r \cdot v = 0$ for all $v \in V_{s_1, \dots, s_r}$).

We are not aware of any prior work on this assumption that $\mathcal{D}_0^m \approx \mathcal{D}_r^m$, so it is not immediately clear how reliable it is. Fortunately, it is *implied* by the $(n, m+r, t)$ planted XOR assumption.

We prove this with a hybrid argument. Suppose that an efficient adversary \mathcal{A} distinguishes between \mathcal{D}_0^m and \mathcal{D}_r^m with advantage $\varepsilon > 0$. By a telescoping argument, \mathcal{A} must distinguish between \mathcal{D}_i^m and \mathcal{D}_{i+1}^m with advantage ε/r , for some $i \in \{0, \dots, r-1\}$. For each i , the following efficient reduction Red_i satisfies $\text{Red}_i(\mathcal{D}_0^{m+r}) \equiv \mathcal{D}_i^m$ and $\text{Red}_i(\mathcal{D}_1^{m+r}) \equiv \mathcal{D}_{i+1}^m$, which implies that ε/r (and therefore ε) is negligible under the $(n, m+r, t)$ planted XOR assumption.

$\text{Red}_i(W)$:

1. Sample i random t -sparse vectors $s_1, \dots, s_i \in \mathbb{F}_2^n$ and let $S = \{v \in \mathbb{F}_2^n : v \cdot s_j = 0 \ \forall j \in [i]\}$.
2. Let $U = W \cap S$. Notice that $\dim U \geq \dim W - i$. Since $\dim W = m+r$ and $i < r$, this is at least m .
3. Output a random m -dimensional subspace of U .

It remains to see why $\text{Red}_i(\mathcal{D}_0^{m+r}) \equiv \mathcal{D}_i^m$ and $\text{Red}_i(\mathcal{D}_1^{m+r}) \equiv \mathcal{D}_{i+1}^m$. In fact both of these statements are true even for *fixed* planted relations.

- $\text{Red}_i(\mathcal{D}_0^{m+r}) \equiv \mathcal{D}_i^m$. Fix s_1, \dots, s_i sampled in Red_i and let $S = \{v \in \mathbb{F}_2^n : v \cdot s_j = 0 \ \forall j \in [i]\}$. Since W is a random subspace of \mathbb{F}_2^n , conditioned on any $d = \dim(W \cap S)$, $U = W \cap S$ is a random d -dimensional subspace of S . Therefore the output of $\text{Red}_i(\mathcal{D}_0^{m+r})$ is a random m -dimensional subspace of S .
- $\text{Red}_i(\mathcal{D}_1^{m+r}) \equiv \mathcal{D}_{i+1}^m$. Fix s_1, \dots, s_i sampled in Red_i and let $S = \{v \in \mathbb{F}_2^n : v \cdot s_j = 0 \ \forall j \in [i]\}$, as above. Suppose that $W \leftarrow \mathcal{D}_1^{m+r}$ is sampled with the planted relation s . Fix s and let $S' = \{v \in S : v \cdot s = 0\}$. Again, conditioned on any $d = \dim(W \cap S)$, $U = W \cap S = W \cap S'$ is a random d -dimensional subspace of S' . Therefore the output of $\text{Red}_i(\mathcal{D}_1^{m+r})$ is a random m -dimensional subspace of S' .

Narrow, statistically random generator matrix (Lemma 9). Since the planted XOR assumption is not a standard cryptographic assumption, we show in Lemma 9 that the generator matrix of $\text{LDPC}[n, c \log^2 n, \log n, 0.99n]$ is *statistically* random for some $c > 0$. This removes the need for the planted XOR assumption, but it comes at the cost of requiring a stronger version of the LPN assumption: When we invoke LPN to see that samples $(G, Gs \oplus e)$ are pseudorandom, the secrets s are now only of size $c \log^2 n$. Therefore, for this PRC we will rely on a subexponential version of the LPN assumption which states that LPN is $2^{O(\sqrt{n})}$ -hard.

For the purposes of this technical overview, we will show that the generator matrix of a closely related code is random. The proof for this version is significantly simpler, but the distribution is less natural and has worse error-correcting properties. The modified distribution on (P, G) is defined as follows:

1. Sample a uniformly random $G_0 \leftarrow \mathbb{F}_2^{0.01n \times g}$.
2. For $i \in [0.99n]$:
 - (a) Sample a random $(t-1)$ -sparse $s_i \in \mathbb{F}_2^{0.01n}$.
 - (b) Let G_i be the matrix G_{i-1} with the extra row $s_i^T G_0$ appended to the bottom,

$$G_i = \begin{bmatrix} G_{i-1} \\ s_i^T G_0 \end{bmatrix}.$$

- (c) Let $s'_i = [s_i^T, 0^{i-1}, 1, 0^{0.99n-i}]$.
3. Let P be the matrix whose rows are $s'_1, \dots, s'_{0.99n}$ and $G = G_{0.99n}$. Output (P, G) .

First observe that $PG = 0$, because $s'_i G = [s_i^T, 0^{i-1}, 1, 0^{0.99n-i}]G = s_i^T G_0 \oplus (s_i^T G_0) = 0$ for every $i \in [0.99n]$.

The leftover hash lemma immediately implies that G is statistically random. Recall that the leftover hash lemma states that if $A \leftarrow \mathbb{F}_2^{g \times \ell}$ is a uniformly random matrix and $s \in \mathbb{F}_2^\ell$ has min-entropy μ , then (A, As) is $2^{-(\mu-g)/2}$ -close to uniform in statistical distance. In our case, we use $A = G_0^T$ and $s = s_i$ to see that $s_i^T G_0$ is $2^{-(\log \binom{0.01n}{i} - g)/2}$ -close to uniform in statistical distance for each $i \in [n-g]$. If $t = \Omega(\log n)$, then there is a choice of $g = \Omega(\log^2 n)$ such that $2^{-(\log \binom{0.01n}{i} - g)/2} = 2^{-\Omega(\log^2 n)} = \text{negl}(n)$, completing the proof.

LDPC-PRC₀ is robust to any p -bounded channel (Lemma 4). Recall that we say that any length-preserving channel that introduces at most a p fraction of bit-flip errors is p -bounded. To prove robustness, we need to show two things:

1. any fixed $x \in \mathbb{F}_2^n$ decodes to \perp with high probability, and
2. for any p -bounded channel \mathcal{E} , samples from $\mathcal{E}(\text{Encode}_G(1))$ decode to 1 with high probability.

Unfortunately, (1) does not quite hold for the scheme presented above. The issue is that, while *most* fixed strings will decode to \perp , a small fraction of strings will decode to 1 regardless of P . For instance, 0 will always decode to 1 because $\text{wt}(P \cdot 0) = 0$ for any P . Therefore we modify our scheme slightly by using a one-time pad $z \in \mathbb{F}_2^n$, included in the public key. The modified $\text{Encode}_G(1)$ outputs $Gs \oplus e \oplus z$, and $\text{Decode}_P(x)$ computes $\text{wt}(Px \oplus Pz)$ instead of $\text{wt}(Px)$; this is the actual scheme we describe in Section 5.

Now, for any fixed $x \in \mathbb{F}_2^n$, $\text{wt}(Px \oplus Pz)$ is distributed identically to $\text{wt}(Pz)$ because z is uniform. In Section 5.4 we will see that P is full rank with high probability, so Pz is uniformly random. By a Chernoff bound, $\text{wt}(Pz) \geq (1/2 - r^{-1/4}) \cdot r$ with high probability and therefore $\text{Decode}_P(x)$ outputs \perp . This concludes the proof of (1), so we turn to (2).

Suppose that we sample $Gs \oplus e \oplus z \leftarrow \text{Encode}_G(1)$ and apply some p -bounded channel \mathcal{E} . The one-time pad effectively converts \mathcal{E} to a fixed error channel, independent of P, G, s, e : Suppose that $\mathcal{E}(x) = x \oplus e(x)$, where $e(x)$ is a random variable depending on x . Since \mathcal{E} is p -bounded, $\text{wt}(e(x)) \leq p \cdot n$. Letting $y = Gs \oplus e \oplus z$, we have

$$\begin{aligned} \mathcal{E}(Gs \oplus e \oplus z) \oplus z &= (Gs \oplus e \oplus z) \oplus e(Gs \oplus e \oplus z) \oplus z \\ &= Gs \oplus e \oplus e(y) \end{aligned}$$

where y is uniformly random in \mathbb{F}_2^n , independent of P, G, s, e because of z . Now it only remains to see that $\text{wt}(P(Gs \oplus e \oplus e(y))) = \text{wt}(P(e \oplus e(y))) < (1/2 - r^{-1/4}) \cdot r$ with high probability. Since e and $e(y)$ are independent errors, each of weight $(1/2 - \Omega(1)) \cdot n$, the combined error $e \oplus e(y)$ also has weight $(1/2 - \Omega(1)) \cdot n$. Therefore, if the row sparsity t of P is $c \log n$ for sufficiently small constant c , then we will have $\text{wt}(P(e \oplus e(y))) < (1/2 - r^{-1/4}) \cdot r$ with high probability, completing the proof of (2).

2.3 Pseudorandom codes for the deletion channel

So far, we have only considered PRCs for substitution channels. For our applications to watermarking and steganography, it will be useful to have PRCs for the noisy deletion channel as well. The noisy deletion channel randomly introduces both deletions and substitutions.

Unfortunately, existing error-correcting codes for the deletion channel introduce a large amount of structure into codewords that precludes pseudorandomness. For instance, the popular techniques of synchronization symbols or concatenation with constant-sized inner codes are immediately seen to be incompatible with pseudorandomness. Even further limiting the techniques available to us, we want our PRCs for the noisy deletion channel to have a binary alphabet in order to be useful for watermarking.

We therefore turn to alternative techniques. Surprisingly, we find that the repetition code — perhaps the simplest and most-structured error-correcting code — is a useful starting point.

For odd integer T , the rate- $(1/T)$ repetition code works by repeating each bit of the message T times. That is, for any message $\mathbf{m} = (\mathbf{m}_1 || \cdots || \mathbf{m}_k) \in \{0, 1\}^k$, the encoder RepEnc_T is defined by $\text{RepEnc}_T(\mathbf{m}) = (\mathbf{m}_1)^T || \cdots || (\mathbf{m}_k)^T$, where $(\mathbf{m}_i)^T$ denotes bit \mathbf{m}_i repeated T times. For example, the rate- $(1/3)$ repetition code encodes 010 as $\text{RepEnc}_T(010) = 000111000$.

Now suppose that the encoding $(\mathbf{m}_1)^T || \cdots || (\mathbf{m}_k)^T$ is subjected to the noisy deletion channel, resulting in a string z . A natural algorithm for decoding z is to partition z into k equal-length blocks z_1, \dots, z_k , and compute the majority of each block:

$\text{MajDec}_k(z)$:

1. Partition z into k equal-length blocks $z = (z_1 || \cdots || z_k)$.
2. Output $(\text{Maj}(z_1) || \cdots || \text{Maj}(z_k))$.

As long as the deletions are sufficiently balanced across the different blocks, the z_i will align well with the original blocks $(\mathbf{m}_i)^T$. Provided further that there are not too many substitutions in any block, we should have $\text{MajDec}_k(z) = \mathbf{m}$. The issue is that $\text{RepEnc}_T(\mathbf{m})$ is not pseudorandom even for random \mathbf{m} , because a random string is (extremely) unlikely to consist of T repeated bits.

On the other hand, a random string typically does have $\Theta(\sqrt{T})$ bias towards 0 or 1.⁵ So if we change or delete a small $O(\sqrt{T})$ number of bits of a random string, we expect the majority to stay the same. This observation brings us to the following encoder MajEnc_T , which encodes each bit in the majority of a random string. We refer to the code defined by $(\text{MajEnc}_T, \text{MajDec}_k)$ as the *majority code*.

$\text{MajEnc}_T(\mathbf{m})$:

1. For $i \in [k]$, let z_i be a random sample from $\{0, 1\}^T$ conditioned on $\text{Maj}(z_i) = \mathbf{m}_i$.
2. Output $(z_1 || \cdots || z_k)$.

Now if \mathbf{m} is random, then $z = \text{MajEnc}_T(\mathbf{m})$ is random as well. Furthermore, if we subject z to the noisy deletion channel to obtain z' , then the bits of $\mathbf{m}' = \text{MajDec}_k(z')$ are positively correlated with the bits of \mathbf{m} . This is because the deletions are at random locations, and are therefore (roughly) evenly-distributed across the different blocks z_i — meaning that MajDec_k will mostly use the correct locations to decode each bit. Since the bit-flip errors are random, they merely dilute the $\Theta(\sqrt{T})$ biases. If $T \gg k$ and the rates of deletions and bit-flip errors are constants below 1 and $1/2$ respectively, then we show in Lemma 15 that $\Pr[\mathbf{m}_i = \mathbf{m}'_i]$ is a constant greater than $1/2$. Therefore, the majority code has the effect of converting the constant-rate noisy deletion channel into some p -bounded channel.

Of course, the majority code is not itself a PRC. The first problem is that codewords for the majority code are only random if the message is random, whereas a PRC needs to allow encoding of any particular message. The second problem is that, even if the message is random, the majority code recovers a string that is only *correlated* with it.

But these are exactly the problems solved by PRCs for bounded-weight error channels! That is, if PRC is any PRC with robustness to every p -bounded channel (e.g. the PRCs from Section 2.2), then the combined code $\text{PRC}_{\text{del}} = (\text{MajEnc} \circ \text{PRC.Encode}, \text{PRC.Decode} \circ \text{MajDec})$ is a PRC with robustness to some constant-rate noisy deletion channel.⁶ Pseudorandomness follows from the pseudorandomness of PRC.Encode : Since $\text{PRC.Encode}(\mathbf{m})$ is pseudorandom for any message \mathbf{m} , $\text{MajEnc}(\text{PRC.Encode}(\mathbf{m}))$ is as well. Robustness follows from the fact that the majority code has the effect of converting the constant-rate noisy deletion channel into some p -bounded channel, which is handled by PRC.Decode .

⁵That is, a random string has $\Theta(\sqrt{T})$ more 0's than 1's, or 1's than 0's. This can be seen as a consequence of the fact that a one-dimensional simple random walk of length T will usually terminate $\Theta(\sqrt{T})$ away from the origin.

⁶As p approaches $1/2$, the combined PRC tolerates a noisy deletion channel with rates of deletions and bit-flip errors approaching 1 and $1/2$.

2.4 Constant-rate pseudorandom codes

So far we have only considered zero-bit PRCs, but for many applications it will be useful to have PRCs that can encode longer messages. There is a simple construction of a multi-bit PRC directly from any zero-bit PRC: Encode each bit of the message with either a zero-bit PRC codeword, or a uniformly random string. That is, if PRC is a zero-bit PRC, we encode a message $\mathbf{m} \in \{0, 1\}^k$ as $(x_1 || \dots || x_k)$ where for each $i \in [k]$, $x_i \leftarrow \{0, 1\}^n$ if $m_i = 0$ and $x_i \leftarrow \text{PRC.Encode}(1)$ if $m_i = 1$.

Unfortunately this scheme has a very low rate. If the zero-bit PRC has block length n , then the resulting multi-bit PRC has rate $1/n$. However, we show in Section 6.2 that one can use any such low-rate PRC to make any error-correcting code pseudorandom, with no asymptotic loss in rate.

The idea is to encode a seed for a one-time pad in the simple low-rate PRC just described, and then use the one-time pad to hide an error-correcting encoding of the message. More formally, let (Enc, Dec) be any (standard) error-correcting code and PRC be a low-rate PRC. We do not require (Enc, Dec) to have any cryptographic properties. We encode a message \mathbf{m} as⁷

$$\text{PRC.Encode}(r), \text{Enc}(\mathbf{m}) \oplus \text{PRG}(r),$$

where PRG is any pseudorandom generator and $r \leftarrow \{0, 1\}^k$ is a fresh uniformly random string.

By pseudorandomness of PRC, $\text{PRC.Encode}(r)$ is indistinguishable from a uniformly random string — even for a fixed choice of r . By security of PRG, the encoding is therefore indistinguishable from a totally random string.

Decoding works as long as $\text{PRC.Encode}(r)$ is not too corrupted for PRC to recover r , and $\text{Enc}(\mathbf{m}) \oplus \text{PRG}(r)$ is not too corrupted for Dec to recover \mathbf{m} .

2.5 Watermarks for language models

In this work, a generative language model is formally described by an algorithm `Model` that takes as input a prompt `PROMPT` and a sequence of tokens output thus far $\mathbf{t}_1, \dots, \mathbf{t}_{i-1}$, and produces a distribution over the next token. A full response is generated by iteratively sampling from these distributions, at each step providing `Model` with the partial response, and terminating once a special “done” token is sampled. For simplicity, we assume here that tokens are binary, which allows us to specify the distribution \mathbf{p}_i over the next token as a Bernoulli distribution $\text{Ber}(\hat{\mathbf{p}}_i)$ where $\hat{\mathbf{p}}_i := \mathbb{E}[\mathbf{p}_i] \in [0, 1]$. We also assume for the purposes of this technical overview that the model always generates a response of length n . In Section 7 we explain why neither of these assumptions is important for our results.

As defined in [CGZ23], a watermarking scheme for a language model consists of algorithms `Wat` and `Detect`, where `Wat` is the watermarked model and `Detect` is an algorithm used to detect the presence of the watermark. In this work we are interested in watermarks that are *undetactable*, *sound*, and *robust*, loosely defined as follows.

- *Undetectability*: Any polynomial number of responses from the watermarked model are computationally indistinguishable from those of the original model.
- *Soundness*: Text generated independently of the watermarked model is not falsely detected.
- *Robustness*: Sufficiently high-entropy text output by the model is detected as watermarked, even if it is altered.

We show that the watermarking strategy from Section 1, which replaces some of the model’s randomness with PRC codewords, yields a scheme that satisfies all of the above properties.

⁷In order to obtain stronger robustness guarantee, we actually randomly permute the symbols of this encoding. For the purposes of this technical overview we omit this detail.

Defining Generate for language models. Recall that the approach from Section 1 requires an algorithm `Generate` that takes as input a prompt and a random seed $x \in \{0, 1\}^n$, and samples a response $\mathbf{t} \in \{0, 1\}^n$ such that

- (1) if x is uniformly random, then \mathbf{t} is distributed identically to a response from `Model`, and
- (2) each bit of \mathbf{t} is correlated with the corresponding bit of x .

We define `Generate(PROMPT, x)` to sample the i^{th} bit \mathbf{t}_i of the response as follows. It first computes \hat{p}_i by querying `Model` with `PROMPT` and the response output thus far, then:

- If $\hat{p}_i \leq 1/2$, sample $\mathbf{t}_i \leftarrow \text{Ber}(2x_i\hat{p}_i)$.
- If $\hat{p}_i > 1/2$, sample $\mathbf{t}_i \leftarrow \text{Ber}(1 - 2(1 - x_i)(1 - \hat{p}_i))$.

For any $\hat{p}_i \in [0, 1]$, one can easily see that \mathbf{t}_i is distributed as $\text{Ber}(\hat{p}_i)$ since x_i is a uniformly random bit. This means that `Generate` satisfies Condition (1) above.

For Condition (2), the bias toward the seed is stronger the closer \hat{p}_i is to $1/2$. It is strongest when $\hat{p}_i = 1/2$ exactly, in which case \mathbf{t}_i is sampled from $\text{Ber}(x_i)$ and is therefore equal to x_i . At the other extreme, if $\hat{p}_i = 0$ or 1 , there is no bias. In general the response is a noisy version of the seed, where the amount of noise on the i^{th} token decays as the binary entropy of $\text{Ber}(\hat{p}_i)$ grows.

Replacing seeds with PRC codewords. We use PRC samples $x \leftarrow \text{PRC.Encode}(1)$, instead of random samples $x \leftarrow \{0, 1\}^n$, as the seeds in `Generate`. That is, if PRC is a zero-bit PRC, we let our watermarking scheme $\mathcal{W}[\text{PRC}]$ be defined by

`Wat(PROMPT)`: Sample $x \leftarrow \text{PRC.Encode}(1)$ and output a sample from `Generate(PROMPT, x)`.

`Detect(\mathbf{t})`: Compute `PRC.Decode(\mathbf{t})` and output the result.

By Condition (1) and the pseudorandomness property of PRC, the responses from `Wat` are computationally indistinguishable from those of the original model. By Condition (2) and the robustness property of PRC, the watermark will be detectable as long as the PRC is sufficiently powerful.

Remark. *Depending on the kind of robustness of the PRC, substituting the entire seed with a single PRC sample results in a watermark that may or may not be detectable from just a subsequence. This is easily fixed by using $x = (x_1 || \dots || x_m)$, where $x_i \leftarrow \text{PRC.Encode}(1)$ are independent PRC samples that are much shorter than the generated content. As long as the text contains at least one subsequence corresponding to a PRC sample x_i , the watermark will be detected.*

PRC error correction and watermark robustness. To understand how error correction of PRC translates to robustness of $\mathcal{W}[\text{PRC}]$, it is helpful to think of `Generate`’s sampling process as a noisy *embedding channel* applied to the seed. That is, for a seed $x \in \{0, 1\}^n$, let $\mathcal{E}_{\text{Emb}}(x) = \text{Generate}(\text{PROMPT}, x)$ be the “embedding channel” describing the noise in $x \mapsto \mathbf{t}$. For detection, it is sufficient for PRC to correct against the channel \mathcal{E}_{Emb} , since watermarked responses are exactly samples from $\mathcal{E}_{\text{Emb}}(x)$ for $x \leftarrow \text{PRC.Encode}(1)$.

Robustness of $\mathcal{W}[\text{PRC}]$ is determined by PRC’s ability to correct from additional errors on top of \mathcal{E}_{Emb} . Let \mathcal{E}_{adv} be a channel modeling the changes an adversary introduces to a watermarked response, so the overall error applied to \mathbf{t} follows $\mathcal{E}_{\text{adv}} \circ \mathcal{E}_{\text{Emb}}$. If PRC is robust to $\mathcal{E}_{\text{adv}} \circ \mathcal{E}_{\text{Emb}}$, the watermark is robust to this adversary’s modifications.

In Section 7.2, we show that as long as the text has non-zero entropy, \mathcal{E}_{Emb} introduces errors at a rate of less than $1/2$. Therefore, using any PRC with robustness to every p -bounded channel, we immediately obtain watermarks that are robust to a constant rate of random substitutions.

Robustness of our watermark. Hashing-based watermarking schemes — including all existing undetectable schemes — are removable by the simple “emoji attack” [Aar22, KGW⁺23a]. In this attack, an adversary asks the model to respond to its prompt and insert an emoji between every word of its response. The adversary then deletes the emojis from the response. This attack removes any watermark that relies on the detector seeing contiguous sequences of watermarked text.

It turns out that hashing-based schemes [Aar22, KGW⁺23a, CGZ23] can easily be made robust to this particular attack.⁸ However, if the adversary instead instructs the model to insert the emojis *randomly*, then we do not know how to make any hashing-based scheme robust. By constructing PRCs with robustness to random deletion channels, we give the first undetectable watermarking scheme that can resist this kind of attack.

In order to show this, we require a stronger assumption on the response: That \mathcal{E}_{Emb} behaves as the binary symmetric channel BSC_q for some $q \in (0, 1/2)$. The binary symmetric channel BSC_q is the channel that flips each bit of its input independently with probability q , so $\text{BSC}_q(x) = x \oplus \text{Ber}(q)$ for $x \in \{0, 1\}$. Essentially, this is equivalent to the assumption that the response has high entropy *and does not repeat words too often*.

Under this assumption, if $\mathcal{E}_{\text{adv}} = \text{BDC}_p$ is the binary deletion channel for some $p \in (0, 1)$, then we just need a PRC with robustness to $\mathcal{E}_{\text{adv}} \circ \mathcal{E}_{\text{Emb}} = \text{BDC}_p \circ \text{BSC}_q$. The binary deletion channel BDC_p is the channel that deletes each bit of its input independently with probability p . Indeed, we saw in Section 2.3 that there exist PRCs with robustness to $\text{BDC}_p \circ \text{BSC}_q$ for any $p \in (0, 1), q \in (0, 1/2)$.

2.6 Watermarks with public attribution

For the “standard” notions of watermarks considered so far, the goal is to determine whether a given text is a possibly-corrupted version of an output generated by a model. Standard watermarks are well-suited for applications such as detecting plagiarism, where one wishes to know if a model was used at all to produce a text, even if that text has been altered by the user.

A different use of watermarks is in *attributing* content to an LLM that generated it. For example, if harmful content generated by an LLM is found on social media, it would be useful to trace this content back to the model using the watermark. Ideally, anyone holding a *public* detection key should be able to trace the content. On the other hand, it should only be possible to embed the watermark by using a *secret* embedding key, in order to avoid falsely attributing text to any model.⁹ In other words, to an attacker who does not know the secret key, the watermark should be *unforgeable*.

In addition to unforgeability, watermarks with public attribution have subtly different detection properties than standard watermarks. Robustness of standard watermarks means that an LLM-generated text will be detected even if small modifications are made. If robust watermarks were used for attribution, an attacker could use a model to generate a benign watermark text, then change a few words to make it offensive. By robustness, the watermark would still be present in this now-offensive content. So whereas robustness is a useful feature for standard watermarking, in the context of attribution it is actually an issue.

We therefore define a watermark with public attribution to have a separate detection algorithm called `AttrText`, which is intentionally designed to *not* be robust. `AttrText`, given a text x and a public detection key, indicates whether the model output verbatim a significant part of that text, and outputs that portion of the text if so.

In order to preserve the benefits of robust watermarking for applications like detecting plagiarism, our publicly attributable watermarks also retain a `Detect` algorithm (in addition to the `AttrText` algorithm) with the robustness of our standard watermarking schemes. One can choose at detection time whether one wants to use `Detect` for standard detection, or `AttrText` for attribution.

⁸For instance, we could choose to only hash tokens whose index has the same parity as the token being sampled.

⁹Note that the roles of the public and secret keys are reversed here. For PRCs, a secret key is necessary for decoding, but anyone can encode with knowledge of a public key. Public-key PRCs are useful for public-key steganography.

Our watermarking scheme with public attribution, $\mathcal{W}_{\text{att}}[\text{PRC}]$, is a natural extension of our regular watermarking scheme $\mathcal{W}[\text{PRC}]$. Recall that $\mathcal{W}[\text{PRC}]$ embeds a codeword of a zero-bit PRC into the model’s response; the detector checks whether the given text is close to a codeword. Of course, if we use a PRC that encodes an arbitrary message (rather than only ‘1’ as in a zero-bit PRC), then $\mathcal{W}[\text{PRC}]$ will embed arbitrary messages in the text. $\mathcal{W}_{\text{att}}[\text{PRC}]$ does exactly this, where the message that it encodes is a *signature on the response output thus far*. `AttrText` decodes the given text to obtain this signature, and checks using the public detection key that it is a valid signature of a portion of the response. If so, this signed portion must have been generated by the model.

Concurrent work [FGJ⁺23] also constructs a watermark with public detection, although this scheme is designed to have mild robustness (comparable to that of [CGZ23]) and therefore is not appropriate for attribution as-is. Their scheme can easily be modified to satisfy our definition of unforgeable public attribution, but it would then lose all robustness guarantees for standard watermarking. Our scheme simultaneously functions as a highly robust standard watermark via `Detect`, while also satisfying unforgeable public attribution via `AttrText`.

2.7 Robust steganography

In steganography, the goal is to send a hidden message such that an observer cannot tell that a message is being sent at all. In the classic presentation, a prisoner wishes to secretly communicate with an outside party even though the warden is filtering their letters. If the warden detects any unusual language then the communication channel will be shut down, so the prisoner cannot simply encrypt the message: The warden should not only be unable to learn anything about the message, but should be unable to even detect that secret communication is occurring at all.

Steganography was formalized in [HLVA02]. In this presentation, there is some underlying *steganographic channel*,¹⁰ a distribution with which the sender wishes to conceal a message. The sender is given sample access to this steganographic channel and sends a *stegotext* to the receiver. *Steganographic secrecy* requires that the distribution of stegotexts is indistinguishable from the steganographic channel, except to the receiver who can recover the message with a secret key.

[HLVA02] proves the security of a steganography scheme of [AP98] that can be constructed using any encryption scheme (`Encode`, `Decode`) with pseudorandom ciphertexts. The key idea behind this scheme is to embed each bit x_i of a pseudorandom encryption of the message by drawing a sample d_i from the steganographic channel such that $f(d_i) = x_i$ for some hash function f :

`Steg.Encode`(sk, m) [AP98, HLVA02]:

1. Let $x = x_1 || \dots || x_n \leftarrow \text{Encode}(\text{sk}, m)$
2. For $i \in [n]$, sample a random d_i from the channel conditioned on $f(d_i) = x_i$
3. Output $d = d_1 || \dots || d_n$

The decoder `Steg.Decode` simply outputs `Decode`(sk, $f(d_1) || \dots || f(d_n)$) = `Decode`(sk, x) = m.

If x_i is uniform over $\{0, 1\}$, and f is perfectly unbiased for the channel, then d_i is sampled exactly from the channel distribution. Therefore, by pseudorandomness of the ciphertext, an observer cannot distinguish stegotexts from samples from the steganographic channel. The receiver, which knows the decoding key for the encryption scheme, can evaluate f on each block of the stegotext to obtain the ciphertext, then decrypt to recover the message.

However, this scheme is very brittle. If the stegotext is corrupted with any errors at all — even ones resulting from any small bias in the hash function f — the message cannot be recovered. A natural attempt

¹⁰In the steganography literature this is usually just called a “channel”; we call it a “steganographic channel” to differentiate it from the coding-theoretic channels we use in the context of robustness.

at achieving robustness is for the sender to apply an error-correcting code (Enc, Dec) to the ciphertext before embedding it. But this loses pseudorandomness and therefore steganographic secrecy! Consequently, the robust steganography schemes of prior work rely on stronger assumptions like the ability of the sender and receiver to share state. A shared state allows the sender to generate a fresh one-time pad r for each message it sends, making the task much easier: The sender embeds $x = \text{Enc}(\mathbf{m}) \oplus r$ by choosing d such that $f(d_i) = x_i$, and the receiver computes $\text{Dec}(\tilde{x} \oplus r)$, where (Enc, Dec) is any error-correcting code. However, if the sender and receiver become unsynchronized (as is likely in practice), the receiver can no longer decode the message.

We observe that PRCs are exactly the primitive needed for robust stateless steganography: Using a PRC as the pseudorandom encryption scheme in the above $(\text{Steg.Encode}, \text{Steg.Decode})$ construction *immediately* gives us a steganography scheme with the same robustness as the PRC. If we use a public-key PRC, the resulting steganography scheme is also public-key. Furthermore, the robustness of the PRC allows us to relax the assumption that f is perfectly unbiased on the steganographic channel.

Our main result of Section 8 is the first stateless steganography scheme with nontrivial robustness to errors. In particular, using our LDPC-based PRCs, we construct stateless steganography schemes that are robust to p -bounded channels for any constant p , or any constant-rate random deletion channel.

3 Preliminaries

Let $\mathbb{N} := \{1, 2, \dots\}$ denote the set of positive integers. We will write $[q] := \{1, \dots, q\}$. For a set X , we define $X^* := \{(x_1, \dots, x_k) \mid x_1, \dots, x_k \in X \wedge k \in \mathbb{Z}_{\geq 0}\}$ to be the set of all strings with alphabet X . For a binary string $s \in X^*$, we let s_i denote the i^{th} symbol of s and $\text{len } s$ denote the length of s . For a string $s \in X^*$ and positive integers $a \leq b \leq \text{len } s$, let $s[a : b]$ denote the substring (s_a, \dots, s_b) .

For a finite set X , we will use the notation $x \leftarrow X$ to denote a uniformly random sample x from X . If X is a set of n -dimensional column vectors, we will write X^m to refer to the set of $n \times m$ matrices whose columns take values in X . Unless otherwise specified, vectors are assumed to be column vectors. For matrices $M \in \mathbb{F}_2^{m \times n}$, let $\ker M$ and $\text{Im } M$ denote the kernel and image of M over \mathbb{F}_2 , respectively.

We use $\log(x)$ to denote the logarithm base 2 of x , and $\ln(x)$ to denote the natural logarithm of x .

Let $\text{Ber}(p)$ be the Bernoulli distribution on $\{0, 1\}$ with expectation p . Let $\text{Ber}(n, p)$ be the distribution on n -bit strings where each bit is an i.i.d sample from $\text{Ber}(p)$.

We let \circ denote the composition of functions, algorithms, or channels; that is, $f \circ g$ denotes (the function/algorithm/channel) obtained by applying g , then f .

Let λ denote the security parameter. A function f of λ is *negligible* if $f(\lambda) = O(\frac{1}{\text{poly}(\lambda)})$ for every polynomial $\text{poly}(\cdot)$. We write $f(\lambda) \leq \text{negl}(\lambda)$ to mean that f is negligible. We let \approx denote computational indistinguishability and \equiv denote statistical indistinguishability.

Lemma 1 (Azuma's inequality). *Let Z_0, \dots, Z_n be a martingale with respect to X_0, \dots, X_n . If the differences $|Z_i - Z_{i-1}|$ are all bounded by C with probability $1 - \varepsilon$, then for all $t > 0$,*

$$\Pr[|Z_n - Z_0| > t] \leq \exp\left(\frac{-t^2}{2nC^2}\right) + \varepsilon.$$

Lemma 2 (Chernoff bounds). *Let $X_1, \dots, X_n \in [0, 1]$ be independent random variables. Let $\mu = \mathbb{E}[\sum_{i=1}^n X_i]$. Then for any $\delta \in (0, 1)$:*

$$\Pr\left[\sum_{i=1}^n X_i \geq (1 + \delta)\mu\right] \leq \exp\left(-\frac{\mu\delta^2}{3}\right) \text{ and}$$

$$\Pr\left[\sum_{i=1}^n X_i \leq (1 - \delta)\mu\right] \leq \exp\left(-\frac{\mu\delta^2}{2}\right).$$

Let $\text{Hyp}(N, K, n)$ denote the hypergeometric distribution with a population of size N , with K success elements, and n draws. That is, a random variable $X \sim \text{Hyp}(N, K, n)$ is the number of success elements contained in n uniform draws from the population without replacement.

Lemma 3 (Hypergeometric tail bounds [Hoe94]). *Let $X \sim \text{Hyp}(N, K, n)$ and $p = K/N$. Then for any $0 < t < K/N$,*

$$\Pr[X \leq (p - t)n] \leq e^{-2t^2n}, \text{ and}$$

$$\Pr[X \geq (p + t)n] \leq e^{-2t^2n}.$$

3.1 Cryptography preliminaries

Pseudorandom function (PRF). Let $\mathcal{F} = \{F_{\text{sk}} : \{0, 1\}^{\ell_1(\lambda)} \rightarrow \{0, 1\}^{\ell_2(\lambda)} \mid \text{sk} \in \{0, 1\}^\lambda\}$ be a family of functions. \mathcal{F} is a PRF if F_{sk} is efficiently computable and for all polynomial-time distinguishers D ,

$$\left| \Pr_{\text{sk} \leftarrow \{0, 1\}^\lambda} \left[D^{F_{\text{sk}}(\cdot)}(1^\lambda) = 1 \right] - \Pr_f \left[D^{f(\cdot)}(1^\lambda) = 1 \right] \right| \leq \text{negl}(\lambda).$$

$\text{SigForge}_{\mathcal{A},\Pi}(\lambda)$ $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ $(\text{m}, \sigma) \leftarrow \mathcal{A}^{\text{Sign}_{\text{sk}}(\cdot)}(\text{pk})$ <p style="margin: 0;">Let \mathcal{Q} denote the set of queries made by \mathcal{A} to $\text{Sign}_{\text{sk}}(\cdot)$.</p> $\text{return Vrfy}_{\text{pk}}(\text{m}, \sigma) \wedge \text{m} \notin \mathcal{Q}$
--

Figure 1: The signature forgery experiment $\text{SigForge}_{\mathcal{A},\Pi}(\lambda)$

where f denotes a random function from $\{0, 1\}^{\ell_1(\lambda)}$ to $\{0, 1\}^{\ell_2(\lambda)}$.

Digital signature scheme. We use the definition of a digital signature from [KL07], with small modifications. A *digital signature scheme* is defined over a message space \mathcal{M} and consists of polynomial-time algorithms $(\text{Gen}, \text{Sign}, \text{Vrfy})$ such that:

Gen : takes as input a security parameter 1^λ and outputs a public-private key pair (pk, sk) .

Sign : takes as input a private key sk and a message $\text{m} \in \mathcal{M}$. It outputs a signature σ , which we write as $\sigma \leftarrow \text{Sign}_{\text{sk}}(\text{m})$.

Vrfy : takes as input a public key pk , a message m , and a signature σ . It outputs 1 if the signature is valid and 0 otherwise. We write $b := \text{Vrfy}_{\text{pk}}(\text{m}, \sigma)$.

It is required that except with negligible over (pk, sk) output by $\text{Gen}(1^\lambda)$, it holds that $\text{Vrfy}_{\text{pk}}(\text{m}, \text{Sign}_{\text{sk}}(\text{m})) = 1$ for every $\text{m} \in \mathcal{M}$.

Definition (Security of a digital signature scheme.). A digital signature scheme $\Pi = (\text{Gen}, \text{Sign}, \text{Vrfy})$ is *existentially unforgeable under an adaptive chosen-message attack*, or just *secure*, if for all polynomial-time adversaries \mathcal{A} ,

$$\Pr[\text{SigForge}_{\mathcal{A},\Pi}(\lambda) = 1] \leq \text{negl}(\lambda)$$

where the experiment $\text{SigForge}_{\mathcal{A},\Pi}(\lambda)$ is defined in Figure 1.

3.2 Coding theory preliminaries

A *channel* is a randomized map $\mathcal{E} : \Sigma^* \rightarrow \Sigma^*$. That is, for $x \in \Sigma^*$, $\mathcal{E}(x)$ is a random sample from Σ^* . We use channels to model errors introduced by the environment, or by an adversary attempting to e.g. remove a watermark. Two of the most important channels we consider are the *binary symmetric channel* BSC and the *binary deletion channel* BDC:

- $\text{BSC}_p : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is the binary symmetric channel with error rate $p \in (0, 1)$. That is, $\text{BSC}_p(x) = x \oplus e$ where $e \leftarrow \text{Ber}(\text{len } x, p)$.
- $\text{BDC}_q : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is the binary deletion channel with deletion rate $q \in (0, 1)$. That is, $\text{BDC}_p(x)$ randomly deletes each bit x_i independently with probability q .

For the purposes of this work, an error-correcting code with robustness to the channel \mathcal{E} is a pair of algorithms (Enc, Dec) where $\text{Enc}, \text{Dec} : \Sigma^* \rightarrow \Sigma^*$. Error-correction (or robustness) for the channel \mathcal{E} says that if $x \leftarrow \text{Enc}(\text{m})$, then $\text{Dec}(\mathcal{E}(x)) = \text{m}$. The *block length* of an error correcting code is the number of symbols in a codeword required to encode a message of a particular length. We therefore write the block length $n = n(k)$ as a function of the message length k . The *rate* of a code is the function $k \mapsto k/n(k)$, which may or may not depend on the message length k .

4 Pseudorandom code basics

4.1 Definitions

In this section we define pseudorandom codes (PRCs)¹¹ and related terminology. A PRC can be viewed as a family¹² of error-correcting codes indexed by encoding and decoding keys. For secret-key PRCs, the encoding and decoding keys are identical; for public-key PRCs, the encoding key is public and the decoding key is secret.

Formally, a PRC is specified by three algorithms. A key generation function samples the keys. An encoding function takes as input the encoding key and a message, and it outputs a codeword. A decoding function takes as input the decoding key and a perturbed codeword, and it outputs the message or \perp .

Our error correction guarantee is defined in terms of a channel. A PRC is robust to a channel \mathcal{E} if `Decode` can recover \mathbf{m} from $\mathcal{E}(\text{Encode}(\mathbf{m}))$ with overwhelming probability. In addition to requiring that we can recover messages from noisy codewords, we also require that `Decode` outputs \perp given a string that is unrelated to the code. That is, for any string c , `Decode`(c) outputs \perp with overwhelming probability over the choice of the decoding key. This property is important for applications such as watermarking, where we want the ability to distinguish codewords from uniformly random strings.

Pseudorandomness of the PRC ensures that an adversary without knowledge of the secret key cannot distinguish between an oracle for the `Encode` algorithm of the scheme, or an oracle that outputs independently drawn uniform strings. If a PRC is viewed as an encryption scheme, pseudorandomness is equivalent to indistinguishability from random bits against a chosen plaintext attack (IND-CPA security) [RBB03]. For public-key PRCs, pseudorandomness holds even against adversaries holding the encryption key.

Definition 1 (Secret-key PRC). Let Σ be a fixed alphabet. A *secret-key pseudorandom error-correcting code* (abbreviated as secret-key PRC) with robustness to a channel $\mathcal{E} : \Sigma^* \rightarrow \Sigma^*$ is a triple of polynomial-time randomized algorithms (`KeyGen`, `Encode`, `Decode`) satisfying

- (Syntax) There exist functions $\ell, n, k : \mathbb{N} \rightarrow \mathbb{N}$ such that for all $\lambda \in \mathbb{N}$, `KeyGen`(1^λ) $\in \{0, 1\}^{\ell(\lambda)}$, `Encode` : $\{1^\lambda\} \times \{0, 1\}^{\ell(\lambda)} \times \Sigma^{k(\lambda)} \rightarrow \Sigma^{n(\lambda)}$, and `Decode` : $\{1^\lambda\} \times \{0, 1\}^{\ell(\lambda)} \times \Sigma^* \rightarrow \Sigma^{k(\lambda)} \cup \{\perp\}$.
- (Error correction, or robustness) For any $\lambda \in \mathbb{N}$ and any message $\mathbf{m} \in \Sigma^{k(\lambda)}$,

$$\Pr_{\text{sk} \leftarrow \text{KeyGen}(1^\lambda)} [\text{Decode}(1^\lambda, \text{sk}, \mathcal{E}(x)) = \mathbf{m} : x \leftarrow \text{Encode}(1^\lambda, \text{sk}, \mathbf{m})] \geq 1 - \text{negl}(\lambda).$$

- (Soundness) For any fixed $c \in \Sigma^*$,

$$\Pr_{\text{sk} \leftarrow \text{KeyGen}(1^\lambda)} [\text{Decode}(1^\lambda, \text{sk}, c) = \perp] \geq 1 - \text{negl}(\lambda).$$

- (Pseudorandomness) For any polynomial-time adversary \mathcal{A} ,

$$\left| \Pr_{\text{sk} \leftarrow \text{KeyGen}(1^\lambda)} [\mathcal{A}^{\text{Encode}(1^\lambda, \text{sk}, \cdot)}(1^\lambda) = 1] - \Pr_{\mathcal{U}} [\mathcal{A}^{\mathcal{U}}(1^\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $\mathcal{A}^{\mathcal{U}}$ means that the adversary has access to an oracle that, on any (even previously queried) input, responds with a freshly drawn uniform value in $\Sigma^{n(\lambda)}$.

Definition 2 (Public-key PRC). Let Σ be a fixed alphabet. A *public-key pseudorandom error-correcting code* (abbreviated as public-key PRC) with robustness to a channel $\mathcal{E} : \Sigma^* \rightarrow \Sigma^*$ is a triple of polynomial-time randomized algorithms (`KeyGen`, `Encode`, `Decode`) satisfying

¹¹An unrelated notion of pseudorandom codes was defined by [KKRT16]. Their notion does not require efficient decoding or pseudorandomness of codewords; rather, they require only that codewords are far apart.

¹²We remark that if a PRC were a fixed code rather than a family of codes, pseudorandomness against non-uniform adversaries would be impossible since the adversary could have the decoding key hard-coded.

- (Syntax) There exist functions $\ell_{\text{Dec}}, \ell_{\text{Enc}}, n, k : \mathbb{N} \rightarrow \mathbb{N}$ such that for all $\lambda \in \mathbb{N}$, $\text{KeyGen}(1^\lambda) \in \{0, 1\}^{\ell_{\text{Dec}}(\lambda)} \times \{0, 1\}^{\ell_{\text{Enc}}(\lambda)}$, $\text{Encode} : \{1^\lambda\} \times \{0, 1\}^{\ell_{\text{Enc}}(\lambda)} \times \Sigma^{k(\lambda)} \rightarrow \Sigma^{n(\lambda)}$, and $\text{Decode} : \{1^\lambda\} \times \{0, 1\}^{\ell_{\text{Dec}}(\lambda)} \times \Sigma^* \rightarrow \Sigma^{k(\lambda)} \cup \{\perp\}$.
- (Error correction, or robustness) For any $\lambda \in \mathbb{N}$ and any message $m \in \Sigma^{k(\lambda)}$,

$$\Pr_{(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)} [\text{Decode}(1^\lambda, \text{sk}, \mathcal{E}(x)) = m : x \leftarrow \text{Encode}(1^\lambda, \text{pk}, m)] \geq 1 - \text{negl}(\lambda).$$

- (Soundness) For any fixed $c \in \Sigma^*$,

$$\Pr_{(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)} [\text{Decode}(1^\lambda, \text{sk}, c) = \perp] \geq 1 - \text{negl}(\lambda).$$

- (Pseudorandomness) For any polynomial-time adversary \mathcal{A} ,

$$\left| \Pr_{(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)} [\mathcal{A}^{\text{Encode}(1^\lambda, \text{pk}, \cdot)}(1^\lambda, \text{pk}) = 1] - \Pr_{(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)} [\mathcal{A}^{\mathcal{U}}(1^\lambda, \text{pk}) = 1] \right| \leq \text{negl}(\lambda),$$

where $\mathcal{A}^{\mathcal{U}}$ means that the adversary has access to an oracle that, on any (even previously queried) input, responds with a freshly drawn uniform value in $\Sigma^{n(\lambda)}$.

The *block length* of a (secret-key or public-key) PRC is $n(\lambda)$ and the *message length* is $k(\lambda)$. The *rate* is the function $\lambda \mapsto k(\lambda)/n(\lambda)$. We often drop the dependence on λ when it is clear from context.

For both secret-key and public-key PRCs, if there is only one possible message (i.e. $k(\lambda) = 0$), then we say that the scheme is a *zero-bit* PRC.

Remark. We present our constructions of PRCs for messages of fixed lengths. That is, for a given set of keys, the construction will only work for messages of a fixed length. However, we note that this can be easily remedied using a pseudorandom function (PRF) to select new keys for every message length. We do not include the PRF in our constructions in order to simplify presentation.

4.2 Heuristic construction from permuted codes

In this section we describe a natural *heuristic* transformation for building a candidate secret-key PRC from any error-correcting code. In Section 5 we will give provably secure constructions of PRCs from standard (subexponential) cryptographic assumptions.

Our heuristic construction can be applied to any binary error-correcting code, such as the polar code [Ari09]. The secret key will be a random permutation of the indices of the codewords. To encode a message, we encode it using the error-correcting code, then apply the secret permutation, and finally add a small amount of random Bernoulli noise.

There are two drawbacks of this generic permuted code construction relative to our pseudorandom LDPC codes:

- In general the pseudorandomness of a permuted code is based on non-standard, ad-hoc conjectures. For certain codes, such as Gallager's LDPC ensemble [Gal62], the permuted code construction is *not* pseudorandom.
- Permuted codes are inherently secret-key PRCs, whereas our pseudorandom LDPC codes are public-key.

For any error-correcting code (Enc, Dec) and parameter $\eta > 0$, we formally define the corresponding permuted code $\text{Permuted-PRC}_\eta[\text{Enc}, \text{Dec}]$ as follows. Let $n = n(\cdot)$ be the block length for (Enc, Dec) , as a function of the

message length. (Recall that the block length is the number of codeword symbols needed to encode a given message.) $\text{Permuted-PRC}_\eta[\text{Enc}, \text{Dec}]$ will encode messages of length k into codewords of length $n = n(k + \lambda)$, where λ is a security parameter.

Construction 1 ($\text{Permuted-PRC}_\eta[\text{Enc}, \text{Dec}]$).

- Let λ be a security parameter, let k be the length of messages we wish to encode, and let $n = n(k + \lambda)$.
- $\text{KeyGen}(1^\lambda)$: Sample a random permutation $\pi : [n] \rightarrow [n]$ and output $\text{sk} = \pi$.
- $\text{Encode}(\text{sk}, \mathbf{m})$ for $\mathbf{m} \in \Sigma^k$:
 1. Sample a random string $r \leftarrow \Sigma^\lambda$ and a noise vector $e \leftarrow \text{Ber}(n, \eta)$.
 2. Compute $c = \text{Enc}(r || \mathbf{m}) \oplus e$.
 3. Output $(c_{\pi(1)} || \cdots || c_{\pi(n)})$.
- $\text{Decode}(\text{sk}, x)$ for $x \in \Sigma^n$:
 1. Compute $c = (x_{\pi^{-1}(1)} || \cdots || x_{\pi^{-1}(n)})$.
 2. Compute $\mathbf{m}' = \text{Dec}(c)$.
 3. Output the last k symbols of \mathbf{m}' .

A permuted code has the same robustness to substitutions, and nearly the same rate, as (Enc, Dec) .

A natural choice of error-correcting codes to use in this permuted code construction is polar codes [Ari09]. Since polar codes have linear rate and tolerate a constant rate of adversarial errors, permuted polar codes are a candidate linear-rate secret-key PRC with robustness to a constant rate of adversarial errors. We do not know whether such codes satisfy pseudorandomness.

5 Constructing pseudorandom codes from cryptographic assumptions

In this section, we introduce a public-key PRC based on LDPC codes. We present the zero-bit version of our scheme, LDPC-PRC₀, in Section 5.1; we will see in Section 6 that this immediately implies a many-bit scheme with essentially the same robustness. In Section 5.2 we show that LDPC-PRC₀ is robust to every error channel of bounded weight, as long as the parity checks have sufficiently low weight. Finally, we prove pseudorandomness of LDPC-PRC₀ in two different parameter regimes under different cryptographic assumptions: In Section 5.3, we prove pseudorandomness under LPN and a certain planted XOR assumption; in Section 5.4, we prove pseudorandomness under a subexponential-query variant of LPN.

5.1 The zero-bit construction

Let

$$\mathcal{S}_{t,n} = \{s \in \mathbb{F}_2^n : \text{wt}(s) = t\}$$

be the set of all t -sparse vectors in \mathbb{F}_2^n , and

$$\mathcal{S}_{t,r,n} = \{P \in \mathbb{F}_2^{r \times n} : \text{wt}(P_{i,:}) = t \forall i \in [r]\}$$

be the set of all t -row-sparse matrices in $\mathbb{F}_2^{r \times n}$.

Our zero-bit pseudorandom LDPC codes are parameterized by a public generator matrix $G \in \mathbb{F}_2^{n \times g}$ and a secret parity-check matrix $P \in \mathbb{F}_2^{r \times n}$. The sampling process for these matrices is described in Definition 3.

Definition 3 (Random LDPC code, LDPC $[n, g, t, r]$). For $n, g, t, r \in \mathbb{N}$, define the distribution LDPC $[n, g, t, r]$ over $\mathbb{F}_2^{r \times n} \times \mathbb{F}_2^{n \times g}$ as follows:

LDPC $[n, g, t, r]$:

1. Sample $P \leftarrow \mathcal{S}_{t,r,n}$, i.e. $P \in \mathbb{F}_2^{r \times n}$ is chosen to have i.i.d random t -sparse rows.
2. Sample $G \leftarrow (\ker P)^g$, i.e. $G \in \mathbb{F}_2^{n \times g}$ is a random matrix subject to $PG = 0$.
3. Output (P, G) .

An (n, g, t, r) random LDPC code is a pair of matrices $(P, G) \leftarrow \text{LDPC}[n, g, t, r]$.

The focus of this section will be on the following *zero-bit* PRC. Recall that a zero-bit PRC is one whose message space is just $\{1\}$. We will see in Section 6.2 that a constant-rate PRC can be generically constructed from any zero-bit PRC.

Construction 2 (Zero-bit public-key pseudorandom LDPC code, LDPC-PRC₀ $[n, g, t, r, \eta, \zeta]$). Let $n, g, t, r : \mathbb{N} \rightarrow \mathbb{N}$ and $\eta, \zeta : \mathbb{N} \rightarrow [0, 1/2)$ be efficiently-computable functions of the security parameter. We define LDPC-PRC₀ $[n, g, t, r, \eta, \zeta]$ by the following algorithms, where we leave the dependence of n, g, t, r, η, ζ on λ implicit:

- **KeyGen**(1^λ): Sample $(P, G) \leftarrow \text{LDPC}[n, g, t, r]$ and $z \leftarrow \mathbb{F}_2^n$. Output $(\text{sk} = (P, z), \text{pk} = (G, z))$.
- **Encode**($1^\lambda, (G, z)$): Sample $u \leftarrow \mathbb{F}_2^g$, $e \leftarrow \text{Ber}(n, \eta)$. Output $Gu \oplus z \oplus e$.
- **Decode**($1^\lambda, (P, z), x$): If $\text{wt}(Px \oplus Pz) < (\frac{1}{2} - \zeta) \cdot r$, output 1; otherwise output \perp .

For the remainder of this section, we will identify the security parameter λ with the dimension of the code n . We will therefore write g, t, r, η, ζ as functions of n , with the understanding that $n(\lambda) = \lambda$.

5.2 Codeword detection (zero-bit decoding)

We say that a length-preserving binary channel $\mathcal{E} : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is p -bounded if there exists a negligible function $\text{negl}(\cdot)$ such that for all $n \in \mathbb{N}$, $\Pr_{x \leftarrow \{0, 1\}^n} [\text{wt}(\mathcal{E}(x) \oplus x) > pn] \leq \text{negl}(n)$.

Lemma 4. *For any $p, \eta \in [0, 1/2)$ and $\varepsilon \in (0, 1)$, there exists $\delta > 0$ such that the following holds. For any $t \leq \delta \log n$, $g > 0$, and $n^\varepsilon \leq r \leq 0.99n$, LDPC-PRC₀ $[n, g, t, r, \eta, r^{-1/4}]$ is robust to every p -bounded channel.*

Proof. Lemma 6 shows that any fixed string $c \in \mathbb{F}_2^n$ decodes to \perp with probability $1 - \text{negl}(n)$. It remains to show that codewords subject to any p -bounded channel decode to 1.

Let \mathcal{E} be any p -bounded channel. We need to show that $\text{wt}(Px \oplus Pz) < (\frac{1}{2} - r^{-1/4})r$ with probability $1 - \text{negl}(n)$ over $(P, G) \leftarrow \text{LDPC}[n, g, t, r], u \leftarrow \mathbb{F}_2^g, z \leftarrow \mathbb{F}_2^n, e \leftarrow \text{Ber}(n, \eta), x \leftarrow \mathcal{E}(Gu \oplus z \oplus e)$.

We will show that there exists a constant $\alpha \in (0, 1/2)$ such that $\text{wt}(Gu \oplus z \oplus x) \leq (\frac{1}{2} - \alpha)n$ with probability $1 - \text{negl}(n)$. Then we can apply Lemma 5 with $y = Gu \oplus z \oplus x$ to see that $\text{wt}(Py) = \text{wt}(Px \oplus Pz) < (\frac{1}{2} - r^{-1/4})r$ with probability $1 - \text{negl}(n)$ for appropriate choices of t, r .

Let $c = Gu \oplus z \oplus e$. Then

$$\begin{aligned} y &= Gu \oplus z \oplus x \\ &= Gu \oplus z \oplus \mathcal{E}(Gu \oplus z \oplus e) \\ &= e \oplus c \oplus \mathcal{E}(c), \end{aligned}$$

where c is uniformly distributed because of z . Let $e' = c \oplus \mathcal{E}(c)$, so that $y = e \oplus e'$. Then since $e \leftarrow \text{Ber}(n, \eta)$ is independent from e' , we have

$$\begin{aligned} \mathbb{E}[\text{wt}(y)] &= \mathbb{E}[\text{wt}(e \oplus e')] \\ &= (1 - \eta) \mathbb{E}[\text{wt}(e')] + \eta \mathbb{E}[n - \text{wt}(e')] \\ &\leq (1 - \eta)pn + \eta(1 - p)n + \text{negl}(n) \end{aligned}$$

where the inequality holds because \mathcal{E} is p -bounded, so $\text{wt}(e') \leq pn$ with probability $1 - \text{negl}(n)$. Conditioned on $\text{wt}(e') \leq pn$, a Chernoff bound over the random choice of e implies that

$$\text{wt}(e \oplus e') \leq \left[\frac{(1 - \eta)p + \eta(1 - p)}{2} + \frac{1}{4} \right] \cdot n$$

with probability $1 - \text{negl}(n)$. Therefore we apply Lemma 5 with $\alpha = \left[\frac{(1 - \eta)p + \eta(1 - p)}{2} + \frac{1}{4} \right]$. \square

Lemma 5. *Let $\alpha \in (0, 1/2)$ be some constant and suppose that $r = n^{\Omega(1)}$ and $t \leq \frac{1}{5} \log_{1/2\alpha} r$. If $y \in \mathbb{F}_2^n$ satisfies $\text{wt}(y) \leq (\frac{1}{2} - \alpha)n$, then*

$$\Pr_{P \leftarrow \mathcal{S}_{t, r, n}} \left[\text{wt}(Py) < \left(\frac{1}{2} - r^{-1/4} \right) r \right] \geq 1 - \text{negl}(n).$$

Proof. Let P_i denote the i^{th} row of P . We will show that

$$\Pr_{P_i \leftarrow \mathcal{S}_{t, 1, n}} [P_i y = 0] \geq \frac{1}{2} + \frac{2}{r^{1/4}}, \quad (1)$$

and the independence of the rows of P will imply the lemma by a Chernoff bound.

Let $j_1, \dots, j_t \in [n]$, $Y_1, \dots, Y_t \in \{-1, 1\}$ be random variables, defined as follows, for $\ell = 1, \dots, t$:

1. Sample $j_\ell \leftarrow [n] \setminus \{j_1, \dots, j_{\ell-1}\}$.

2. Let $Y_\ell = (-1)^{x_{j_\ell}}$.

Then the bit $P_i y$ is distributed as $(1 - Y_1 \cdots Y_t)/2$, so

$$\begin{aligned} \Pr_{P_i \leftarrow \mathcal{S}_{t,1,n}} [P_i y = 0] &= \Pr[Y_1 \cdots Y_t = 1] \\ &= \frac{1}{2} + \frac{1}{2} \mathbb{E}[Y_1 \cdots Y_t]. \end{aligned}$$

The remainder of the proof is devoted to showing that

$$\mathbb{E}[Y_1 \cdots Y_t] \geq (2\alpha)^t - 2t^2/n. \quad (2)$$

For large enough n , this will imply Equation (1) by the assumption that $t \leq \frac{1}{5} \log_{1/2\alpha} r$.

Let $\alpha' = 1/2 - \text{wt}(y)/n$, and note that $\alpha' \geq \alpha$. To prove Equation (2) we first show that for all $m \in [t]$,

$$|\mathbb{E}[Y_1 \cdots Y_m] - 2\alpha' \mathbb{E}[Y_1 \cdots Y_{m-1}]| \leq 2t/n. \quad (3)$$

By the tower property and linearity of expectation,

$$|\mathbb{E}[Y_1 \cdots Y_m] - 2\alpha' \mathbb{E}[Y_1 \cdots Y_{m-1}]| = |\mathbb{E}[Y_1 \cdots Y_{m-1} \cdot (\mathbb{E}[Y_m | Y_{<m}] - 2\alpha')]|. \quad (4)$$

For all possible assignments of $Y_{<m}$,

$$\begin{aligned} \mathbb{E}[Y_m | Y_{<m}] &= \Pr_{j_m} [x_{j_m} = 0 | Y_{<m}] - (1 - \Pr_{j_m} [x_{j_m} = 0 | Y_{<m}]) \\ &= 2 \Pr_{j_m} [x_{j_m} = 0 | Y_{<m}] - 1 \end{aligned}$$

Recall that j_m is chosen to be a uniformly random index from $[n] \setminus \{j_{<m}\}$. Since $m \leq t$, $[n] \setminus \{j_{<m}\}$ contains at least $(\frac{1}{2} + \alpha')n - t$ and at most $(\frac{1}{2} + \alpha')n$ indices for which the corresponding bit of y is 0. Therefore, $\frac{1}{2} + \alpha' - t/n \leq \Pr_{j_m} [x_{j_m} = 0 | Y_{<m}] \leq \frac{1}{2} + \alpha' + t/n$.

We now have that $|\mathbb{E}[Y_m | Y_{<m}] - 2\alpha'| \leq 2t/n$, and $|Y_1 \cdots Y_{m-1}| = 1$. Plugging these two facts into Equation (4) gives us Equation (3).

We complete the proof by showing by induction that $\mathbb{E}[Y_1 \cdots Y_m] \geq (2\alpha')^m - 2mt/n$. Observe first that $\mathbb{E}[Y_1] = (\frac{1}{2} + \alpha') - (\frac{1}{2} - \alpha') > 2\alpha' - 2t/n$. Assume that the inequality holds for some $m < t$; we'll show that it holds for $m + 1$. We have $\mathbb{E}[Y_1 \cdots Y_{m+1}] \geq 2\alpha' \mathbb{E}[Y_1 \cdots Y_m] - 2t/n$. Therefore, since $2\alpha' \leq 1$,

$$\begin{aligned} \mathbb{E}[Y_1 \cdots Y_{m+1}] &\geq 2\alpha' \left((2\alpha')^m - \frac{2mt}{n} \right) - \frac{2t}{n} \\ &= (2\alpha')^{m+1} - \frac{2(m+1)t}{n} \end{aligned}$$

as desired. □

Lemma 6. *If $n^{\Omega(1)} \leq r \leq 0.99n$, then for any fixed $c \in \mathbb{F}_2^n$,*

$$\Pr_{\substack{P \leftarrow \mathcal{S}_{t,r,n} \\ z \leftarrow \mathbb{F}_2^n}} \left[\text{wt}(Pc \oplus Pz) \geq \left(\frac{1}{2} - r^{-1/4} \right) r \right] \geq 1 - \text{negl}(n).$$

Proof. With probability $1 - \text{negl}(n)$, $P \leftarrow \mathcal{S}_{t,r,n}$ is full rank by Lemma 12 (invoking the lemma with $g = 0$). For any such P , over the random choice of $z \leftarrow \mathbb{F}_2^n$, the vector $P(c \oplus z) \in \mathbb{F}_2^r$ is uniformly random. The lemma follows from a Chernoff bound. □

5.3 Pseudorandomness from the planted XOR assumption and LPN

In this subsection we prove that the generator matrix (public key) of LDPC-PRC₀ is pseudorandom under the planted XOR assumption. This is Lemma 8. The number of columns of the generator matrix and the number of rows of the parity check matrix will depend on the particular planted XOR assumption one is willing to make. Then, since the generator matrix is pseudorandom, pseudorandomness of LDPC-PRC₀ (Theorem 1) follows directly from the standard LPN assumption.

Cryptographic assumptions. The existence of our LDPC-based PRCs relies on either of two assumptions. We state the two assumptions together as Assumption 1.

Assumption 1. At least one of the following two statements is true:

- There exists a constant $\eta \in (0, 1/2)$ such that, for any function $g(n) = \Omega(\log^2 n)$, the LPN $_{g,\eta}$ assumption (Assumption 2) holds.
- There exist constants $\eta \in (0, 1/2)$ and $\varepsilon \in (0, 1)$ such that, for any function $t(n) = \Theta(\log n)$, both the LPN $_{n^\varepsilon,\eta}$ assumption (Assumption 2) and the XOR $_{2n^\varepsilon,t}$ assumption (Assumption 3) hold.

We now define the LPN $_{g,\eta}$ and XOR $_{m,t}$ assumptions. Let us first recall the LPN assumption. We only state the decisional, constant-noise version, since all of our results pertain to that variant. For $\eta \in (0, 1/2)$ and $g : \mathbb{N} \rightarrow \mathbb{N}$, the LPN $_{g,\eta}$ problem is to distinguish between $(A, As \oplus e)$ and (A, u) , where $A \leftarrow \mathbb{F}_2^{n \times g(n)}$, $s \leftarrow \mathbb{F}_2^{g(n)}$, $e \leftarrow \text{Ber}(n, \eta)$, and $u \leftarrow \mathbb{F}_2^n$.

Assumption 2 (LPN assumption). For $\eta \in (0, 1/2)$ and $g : \mathbb{N} \rightarrow \mathbb{N}$, the LPN $_{g,\eta}$ assumption states that for every $n \in \mathbb{N}$ and every polynomial-time adversary \mathcal{A} ,

$$\left| \Pr_{\substack{A \leftarrow \mathbb{F}_2^{n \times g(n)} \\ s \leftarrow \mathbb{F}_2^{g(n)} \\ e \leftarrow \text{Ber}(n, \eta)}} [\mathcal{A}(A, As \oplus e) = 1] - \Pr_{\substack{A \leftarrow \mathbb{F}_2^{n \times g(n)} \\ u \leftarrow \mathbb{F}_2^n}} [\mathcal{A}(A, u) = 1] \right| \leq \text{negl}(n).$$

By a standard hybrid argument, the LPN $_{g,\eta}$ assumption implies that any polynomial number of samples of the form $(A, As \oplus e)$ are indistinguishable from uniformly random samples.

The planted XOR assumption is lesser-known than LPN, but there is still precedent [ASS⁺23]. This assumption states that a random linear subspace over \mathbb{F}_2 is indistinguishable from one containing a planted sparse vector. Formally, it says that the distributions $\mathcal{D}_0(n, m)$ and $\mathcal{D}_1(n, m, t)$ are computationally indistinguishable, where the “null” distribution \mathcal{D}_0 is the uniform distribution over $\mathbb{F}_2^{n \times m}$, and the “planted” distribution \mathcal{D}_1 is defined as follows.

$\mathcal{D}_1(n, m, t)$:

1. Sample $s \leftarrow \mathcal{S}_{t,n}$.
2. Sample a random matrix $G \in \mathbb{F}_2^{n \times m}$ subject to $s^T G = 0$.
3. Output G .

Assumption 3 (Planted XOR assumption). For $m, t : \mathbb{N} \rightarrow \mathbb{N}$, the XOR $_{m,t}$ assumption states that for every $n \in \mathbb{N}$ and every polynomial-time adversary \mathcal{A} ,

$$\left| \Pr_{G \leftarrow \mathcal{D}_0(n, m(n))} [\mathcal{A}(G) = 1] - \Pr_{G \leftarrow \mathcal{D}_1(n, m(n), t(n))} [\mathcal{A}(G) = 1] \right| \leq \text{negl}(n).$$

Remark. A previous version of this paper made use of the planted XOR assumption with $m = \Omega(n)$. However, an anonymous CRYPTO 2024 reviewer pointed out to us that this assumption is false by Theorem 4.26 of [ASS⁺23]. All of our results are therefore updated in this version of the paper to use $m = n^{\Omega(1)}$; none of our results significantly change with the new choice of parameters. See Theorem 1 for the updated theorem statement. Note also that Theorem 2 is not affected, as it does not make use of the planted XOR assumption.

Note that when $m = n - O(\log n)$, the XOR _{m,t} assumption is false because there will be only $n^{O(1)}$ vectors $v \in \mathbb{F}_2^n$ such that $v^T G = 0$, and one can therefore brute force search for the planted relation s . On the other hand, Claim 7 shows that when (for instance) $t \sim \log n$ and $m \sim \log^2 n$ the XOR _{m,t} assumption holds against even *unbounded* adversaries. This claim follows from Theorem 4.2 of [ASS⁺23].

Claim 7. *If $t = O(\log n)$ and $m \leq (1 - \Omega(1))t \log n$, then*

$$SD(\mathcal{D}_0(n, m), \mathcal{D}_1(n, m, t)) = n^{-\Omega(t)}.$$

For larger values of m , $\mathcal{D}_0(n, m)$ and $\mathcal{D}_1(n, m, t)$ are no longer statistically close, but the planted XOR assumption says that they remain computationally indistinguishable.

Lemma 8. *Let $m, t, r : \mathbb{N} \rightarrow \mathbb{N}$ be such that $m(n) + r(n) \leq n - \omega(\log n)$. The XOR _{$m+r,t$} assumption (Assumption 3) implies that the marginal distribution on G for $(P, G) \leftarrow \text{LDPC}[n, m, t, r]$ is pseudorandom.*

Proof. The proof closely mirrors that in the technical overview (Section 2.1), with the main difference being that here we deal with generator matrices instead of the linear subspaces themselves. For $i \in \{0, \dots, r-1\}$ and $m' \in [m+r]$, let $\mathcal{D}_i(n, m', t)$ be defined as follows.

$\mathcal{D}_i(n, m', t)$:

1. Sample $s_1, \dots, s_i \leftarrow \mathcal{S}_{t,n}$.
2. Sample a random matrix $G \in \mathbb{F}_2^{n \times m'}$ subject to $s_j^T G = 0$ for all $j \in [i]$.
3. Output G .

Observe that $\mathcal{D}_0(n, m', t) = \mathcal{D}_0(n, m')$, and $\mathcal{D}_1(n, m', t)$ is consistent with the definition given earlier.

For each $i \in \{0, \dots, r-1\}$ and $m' \in [m+r]$, since $m' \leq m+r \leq n - \omega(\log n)$, the matrix $G \leftarrow \mathcal{D}_i(n, m', t)$ has full rank with probability $1 - \text{negl}(n)$. Therefore, $\mathcal{D}_i(n, m', t)$ is $\text{negl}(n)$ -close in statistical distance to the following distribution:

$\hat{\mathcal{D}}_i(n, m', t)$:

1. Sample $s_1, \dots, s_i \leftarrow \mathcal{S}_{t,n}$.
2. Sample a random *full-rank* matrix $G \in \mathbb{F}_2^{n \times m'}$ subject to $s_j^T G = 0$ for all $j \in [i]$.
3. Output G .

Since $\hat{\mathcal{D}}_0(n, m+r, t)$ (resp. $\hat{\mathcal{D}}_1(n, m+r, t)$) is $\text{negl}(n)$ -close to $\mathcal{D}_0(n, m+r)$ (resp. $\mathcal{D}_1(n, m+r, t)$) in statistical distance, the $(n, m+r, t)$ planted XOR assumption implies that $\hat{\mathcal{D}}_0(n, m+r, t)$, and $\hat{\mathcal{D}}_1(n, m+r, t)$ are computationally indistinguishable.

Now suppose that an efficient adversary \mathcal{A} distinguishes between $\hat{\mathcal{D}}_0(n, m, t)$ and $\hat{\mathcal{D}}_r(n, m, t)$ with advantage $\varepsilon > 0$. By a telescoping argument, \mathcal{A} must distinguish between $\hat{\mathcal{D}}_i(n, m, t)$ and $\hat{\mathcal{D}}_{i+1}(n, m, t)$ with advantage ε/r , for some $i \in \{0, \dots, r-1\}$. For each $i \in \{0, \dots, r-1\}$, the following efficient reduction Red_i satisfies $\text{Red}_i(\hat{\mathcal{D}}_0(n, m+r, t)) \equiv \hat{\mathcal{D}}_i(n, m, t)$ and $\text{Red}_i(\hat{\mathcal{D}}_1(n, m+r, t)) \equiv \hat{\mathcal{D}}_{i+1}(n, m, t)$. Therefore, the $(n, m+r, t)$ planted XOR assumption implies that $\varepsilon/r = \text{negl}(n)$, which will complete the proof.

$\text{Red}_i(W)$:

1. Sample i random t -sparse vectors $s_1, \dots, s_i \in \mathbb{F}_2^n$ and let $S = \{v \in \mathbb{F}_2^n : v \cdot s_j = 0 \ \forall j \in [i]\}$.
2. Let $U = \text{Im}(W) \cap S$. Since $i < r$ and $\dim \text{Im } W = m + r$, we have $\dim U > m$.
3. Sample a random full-rank matrix $G \in \mathbb{F}_2^{n \times m}$ such that $\text{Im}(G) \subseteq U$.
4. Output G .

It remains to see why $\text{Red}_i(\hat{\mathcal{D}}_0(n, m + r, t)) \equiv \hat{\mathcal{D}}_i(n, m, t)$ and $\text{Red}_i(\hat{\mathcal{D}}_1(n, m + r, t)) \equiv \hat{\mathcal{D}}_{i+1}(n, m, t)$. In fact both of these statements are true even for *fixed* planted relations.

Proof that $\text{Red}_i(\hat{\mathcal{D}}_0(n, m + r, t)) \equiv \hat{\mathcal{D}}_i(n, m, t)$. Suppose that $W \leftarrow \hat{\mathcal{D}}_0(n, m + r, t)$. Fix s_1, \dots, s_i sampled in $\text{Red}_i(W)$ and let $S = \{v \in \mathbb{F}_2^n : v \cdot s_j = 0 \ \forall j \in [i]\}$. For any $d \in \{m+1, \dots, m+r\}$, consider the distribution of the subspace $U = \text{Im}(W) \cap S$ conditioned on the event that $\dim U = d$. Before the conditioning $\text{Im } W$ was a random subspace of \mathbb{F}_2^n , so after the conditioning U is a random d -dimensional subspace of S . The output of $\text{Red}_i(W)$ is a random full-rank matrix $G \in \mathbb{F}_2^{n \times m}$ such that $\text{Im}(G) \subseteq U \subseteq S$. But we have just seen that U is a uniformly random $d > m$ dimensional subspace of S , so it follows that G is a random matrix subject to $\text{Im } G \subseteq S$ — i.e., $s_j^T G = 0$ for all $j \in [i]$.

Proof that $\text{Red}_i(\hat{\mathcal{D}}_1(n, m + r, t)) \equiv \hat{\mathcal{D}}_{i+1}(n, m, t)$. Suppose that $W \leftarrow \hat{\mathcal{D}}_1(n, m + r, t)$ is sampled with the planted relation s . Fix s and s_1, \dots, s_i sampled in Red_i . Let $S = \{v \in \mathbb{F}_2^n : v \cdot s_j = 0 \ \forall j \in [i]\}$ and $S' = \{v \in \mathbb{F}_2^n : v \cdot s = 0\}$. Again, for each $d \in \{m+1, \dots, m+r\}$, consider the distribution of $U = \text{Im}(W) \cap S$ conditioned on the event that $\dim U = d$. Before the conditioning $\text{Im } W$ was a random subspace of S' , so after the conditioning U is a random d -dimensional subspace of $S \cap S'$. The output of $\text{Red}_i(W)$ is a random full-rank matrix $G \in \mathbb{F}_2^{n \times m}$ such that $\text{Im}(G) \subseteq U \subseteq S \cap S'$. But we have just seen that U is a uniformly random $d > m$ dimensional subspace of $S \cap S'$, so it follows that G is a random matrix subject to $\text{Im } G \subseteq S \cap S'$ — i.e., $s^T G = 0$ and $s_j^T G = 0$ for all $j \in [i]$. \square

Claim 7 and Lemma 8 together imply that the generator matrix from LDPC $[n, \log^2 n, 4 \log n, \log^2 n]$ is statistically uniform. In Lemma 9 we improve this to show that the generator matrix from LDPC $[n, g, 4 \log n, r]$ is statistically uniform for any $r \leq 0.99n$ (and some $g = \Omega(\log^2 n)$); this forms the basis of our PRC construction based only on LPN (Theorem 2).

Theorem 1. *For any constants $p, \eta \in (0, 1/2)$ and $\varepsilon \in (0, 1)$, there exists a function $t = \Theta(\log n)$ such that LDPC-PRC $_0[n, n^\varepsilon, t, n^\varepsilon, \eta, n^{-\varepsilon/4}]$ is a zero-bit public-key PRC (Definition 2) that is robust to every p -bounded channel, where pseudorandomness rests on the LPN $_{n^\varepsilon, \eta}$ assumption (Assumption 2) and the XOR $_{2n^\varepsilon, t}$ assumption (Assumption 3).*

Proof. By Lemma 4, there exists $t = \Theta(\log n)$ such that LDPC-PRC $_0[n, n^\varepsilon, t, n^\varepsilon, \eta, n^{-\varepsilon/4}]$ is robust to every p -bounded channel.

By Lemma 8, the XOR $_{2n^\varepsilon, t}$ assumption implies that for $(P, G) \leftarrow \text{LDPC}[n, n^\varepsilon, t, n^\varepsilon]$, the marginal distribution on G is pseudorandom. By the LPN $_{n^\varepsilon, \eta}$ assumption, it follows that LDPC-PRC $_0[n, n^\varepsilon, t, n^\varepsilon, \eta, n^{-\varepsilon/4}]$ is a public-key PRC. \square

5.4 Pseudorandomness from subexponential LPN

In Section 5.3, we showed that the generator matrix $G \in \mathbb{F}_2^{n \times g}$ of LDPC-PRC $_0$ was pseudorandom under the planted XOR assumption. In Lemma 9 of this section, we will show that G is *statistically* random if we set g to be sufficiently small. Therefore, the planted XOR assumption is no longer necessary; however, the number of columns will be only $g = O(\log^2 n)$, so we must rely on a stronger, sub-exponential variant of LPN than in Section 5.3.

Remark. Before we see Lemma 9, note that we use a different proof strategy here than the technical overview. The reason we use this more involved proof here is that the proof outlined in the technical overview results in a different ensemble of parity-check matrices P with a triangular restriction and non-independent rows. By proving Lemma 9, we are able to use the same ensemble as in Section 5.3 — that is, the rows of P are still independent and uniform t -sparse vectors.

Lemma 9. If $r \leq 0.99n$ and $\omega(\sqrt{\log n}) \leq t \leq O(\log n)$, then there is a $g = \Omega(t^2)$ such that the marginal distribution on G for $(P, G) \leftarrow \text{LDPC}[n, g, t, r]$ is $\text{negl}(n)$ -close to uniform in statistical distance.

Proof. Recall the definitions of $\mathcal{S}_{t,n}$ and $\mathcal{S}_{t,r,n}$ from Section 5.1. Let \mathcal{D} be the marginal distribution on $G \in \mathbb{F}_2^{n \times g}$ for $(P, G) \leftarrow \text{LDPC}[n, g, t, r]$.

We will show that \mathcal{D} is $\text{negl}(n)$ -close to the uniform distribution in statistical distance. For any $G^* \in \mathbb{F}_2^{n \times g}$,

$$\begin{aligned} \Pr_{G \leftarrow \mathcal{D}} [G = G^*] &= \sum_{\substack{P^* \in \mathcal{S}_{t,r,n}: \\ P^* G^* = 0}} \Pr_{P \leftarrow \mathcal{S}_{t,r,n}} [P = P^*] \cdot \Pr_{G \leftarrow (\ker P^*)^g} [G = G^*] \\ &= \frac{1}{\binom{n}{t}^r} \sum_{\substack{P^* \in \mathcal{S}_{t,r,n}: \\ P^* G^* = 0}} \frac{1}{|\ker P^*|^g} \\ &\leq \frac{1}{\binom{n}{t}^r} \sum_{\substack{P^* \in \mathcal{S}_{t,r,n}: \\ P^* G^* = 0}} \frac{1}{2^{(n-r)m}} \\ &= \frac{\eta(G^*)^r}{\binom{n}{t}^r \cdot 2^{(n-r)m}} \end{aligned}$$

where $\eta(G^*)$ is the number of collections of t rows of G^* that sum to 0. Since valid rows of P^* correspond to collections of t rows of G^* that sum to 0, $|\{P^* \in \mathbb{F}_2^{r \times n} : P^* G^* = 0\}| = \eta(G^*)^r$, which gives us the last equality above.

Let $\mathcal{F} \subset \mathbb{F}_2^{r \times n}$ be the set of t -row-sparse matrices of full rank,

$$\mathcal{F} := \{P \in \mathcal{S}_{t,r,n} \mid \text{rank}(P) = r\} = \{P \in \mathcal{S}_{t,r,n} \mid \dim \ker(P) = n - r\}.$$

Then we also have

$$\begin{aligned} \Pr_{G \leftarrow \mathcal{D}} [G = G^*] &= \frac{1}{\binom{n}{t}^r} \sum_{\substack{P^* \in \mathcal{S}_{t,r,n}: \\ P^* G^* = 0}} \frac{1}{|\ker P^*|^g} \\ &\geq \frac{1}{\binom{n}{t}^r} \sum_{\substack{P^* \in \mathcal{F}: \\ P^* G^* = 0}} \frac{1}{2^{(n-r)g}} \\ &= \frac{\eta(G^*)^r}{\binom{n}{t}^r \cdot 2^{(n-r)g}} \cdot \Pr_{P^* \leftarrow \mathbb{F}_2^{r \times n}} [P^* \text{ has full rank} \mid P^* G^* = 0]. \end{aligned}$$

Let $\mathcal{P}(G^*) = \{P^* \in \mathcal{S}_{t,r,n} \mid P^* G^* = 0\}$. The above two inequalities give us a point-wise approximation to the density of \mathcal{D} ,

$$\Pr_{P^* \leftarrow \mathcal{P}(G^*)} [P^* \text{ has full rank}] \cdot \frac{\eta(G^*)^r}{\binom{n}{t}^r \cdot 2^{(n-r)g}} \leq \Pr_{G \leftarrow \mathcal{D}} [G = G^*] \leq \frac{\eta(G^*)^r}{\binom{n}{t}^r \cdot 2^{(n-r)g}}. \quad (5)$$

We complete the proof with the following three facts:

- Lemma 10 is a general statistical fact, which implies that if

$$\Pr_{G^* \leftarrow \mathbb{F}_2^{n \times g}} \left[\left| \Pr_{G \leftarrow \mathcal{D}} [G = G^*] - \frac{1}{2^{ng}} \right| \leq \frac{\text{negl}(n)}{2^{ng}} \right] \geq 1 - \text{negl}(n),$$

then \mathcal{D} is $\text{negl}(n)$ -close to uniform in statistical distance. Crucially, Lemma 10 and Equation (5) reduce the problem to reasoning about *the uniform distribution* over G^* , rather than \mathcal{D} .

- Lemma 11 implies that

$$\Pr_{G^* \leftarrow \mathbb{F}_2^{n \times g}} \left[\left| \frac{\eta(G^*)^r}{\binom{n}{t}^r \cdot 2^{(n-r)g}} - \frac{1}{2^{ng}} \right| \leq \frac{\text{negl}(n)}{2^{ng}} \right] \geq 1 - \text{negl}(n).$$

The proof is a simple invocation of Chebyshev's inequality.

- Lemma 12 implies that

$$\Pr_{G^* \leftarrow \mathbb{F}_2^{n \times g}} \left[\Pr_{P^* \leftarrow \mathcal{P}(G^*)} [P^* \text{ has full rank}] \geq 1 - \text{negl}(n) \right] \geq 1 - \text{negl}(n).$$

Using Lemmas 11 and 12 with Equation (5), the condition of Lemma 10 is satisfied. This completes the proof of the theorem. \square

Lemma 10. *Let p be a probability distribution on a finite set \mathcal{X} . If*

$$\Pr_{x \leftarrow \mathcal{X}} \left[\left| p(x) - \frac{1}{|\mathcal{X}|} \right| > \frac{\alpha}{|\mathcal{X}|} \right] \leq \beta,$$

then the statistical distance between p and the uniform distribution is at most $\alpha + \beta$.

Proof. We just compute

$$\begin{aligned} & \sum_{\substack{x \in \mathcal{X}: \\ p(x) < \frac{1}{|\mathcal{X}|}}} \left(\frac{1}{|\mathcal{X}|} - p(x) \right) \\ &= \sum_{\substack{x \in \mathcal{X}: \\ \frac{1-\alpha}{|\mathcal{X}|} \leq p(x) < \frac{1}{|\mathcal{X}|}}} \left(\frac{1}{|\mathcal{X}|} - p(x) \right) + \sum_{p(x) < \frac{1-\alpha}{|\mathcal{X}|}} \left(\frac{1}{|\mathcal{X}|} - p(x) \right) \\ &\leq \alpha + \sum_{p(x) < \frac{1-\alpha}{|\mathcal{X}|}} \frac{1}{|\mathcal{X}|} \\ &\leq \alpha + \beta. \end{aligned}$$

\square

Lemma 11. *For any $\varepsilon > 0$,*

$$\Pr_{G^* \leftarrow \mathbb{F}_2^{n \times g}} \left[\left| \eta(G^*) - \binom{n}{t} \cdot 2^{-g} \right| > \varepsilon \cdot \binom{n}{t} \cdot 2^{-g} \right] \leq \frac{2^g}{\binom{n}{t} \cdot \varepsilon^2}.$$

Proof. By linearity of expectation,

$$\mathbb{E}_{G^* \leftarrow \mathbb{F}_2^{n \times g}} \eta(G^*) = \sum_{w \in \mathcal{S}_{t,n}} \Pr_{G^* \leftarrow \mathbb{F}_2^{n \times g}} [w^T G^* = 0] = \binom{n}{t} \cdot 2^{-g}.$$

Furthermore, $\{\mathbb{1}[w^T G^* = 0]\}_{w \in \mathcal{S}_{t,n}}$ are pairwise independent random variables over $G^* \leftarrow \mathbb{F}_2^{n \times g}$, so the lemma follows from Chebyshev's inequality. \square

Lemma 12. *If $r \leq 0.99n$ and $\omega(\sqrt{\log n}) \leq t \leq O(\log n)$, then there is a $\tilde{g} = \Omega(t^2)$ such that for all $g \leq \tilde{g}$,*

$$\Pr_{G^* \leftarrow \mathbb{F}_2^{n \times g}} \left[\Pr_{P^* \leftarrow \mathcal{P}(G^*)} [P^* \text{ has full rank}] \geq 1 - \text{negl}(n) \right] \geq 1 - \text{negl}(n).$$

Proof. We say that a subset of rows $S \subseteq [r]$ of P^* forms a “simple dependency” if those rows sum to 0, and no subset of them sums to 0. For a $1 - \text{negl}(n)$ fraction of $G^* \leftarrow \mathbb{F}_2^{n \times g}$, we will show that over $P^* \leftarrow \mathcal{P}(G^*)$ the expected number of simple dependencies in P^* is $\text{negl}(n)$.

Letting W_i denote the i th row of the matrix W , we define

$$\text{SD}_\ell = \left\{ W \in \mathcal{S}_{t,\ell,n} \mid \bigoplus_{i \in [\ell]} W_i = 0 \text{ and } \forall T \subsetneq [\ell], \bigoplus_{i \in T} W_i \neq 0 \right\}$$

and

$$\text{SD}_\ell(G^*) = \{W \in \text{SD}_\ell \mid W G^* = 0\}.$$

For any G^* , over $P^* \leftarrow \mathcal{P}(G^*)$ the expected number of simple dependencies in P^* is

$$\begin{aligned} \mathbb{E}_{P^* \leftarrow \mathcal{P}(G^*)} \sum_{\ell=1}^r \sum_{S \in \binom{[r]}{\ell}} \mathbb{1}[P_S^* \in \text{SD}_\ell] &= \sum_{\ell=1}^r \sum_{S \in \binom{[r]}{\ell}} \Pr_{P^* \leftarrow \mathcal{P}(G^*)} [P_S^* \in \text{SD}_\ell] \\ &= \sum_{\ell=1}^r \binom{r}{\ell} \frac{|\text{SD}_\ell(G^*)|}{\eta(G^*)^\ell}. \end{aligned} \quad (6)$$

Now since any $\ell - 1$ of the rows of any simple dependency are linearly independent, we have

$$\Pr_{G^* \leftarrow \mathbb{F}_2^{n \times g}} [W G^* = 0] = \frac{1}{2^{(\ell-1)g}}$$

for any $W \in \text{SD}_\ell$. Therefore,

$$\begin{aligned} \mathbb{E}_{G^* \leftarrow \mathbb{F}_2^{n \times g}} |\text{SD}_\ell(G^*)| &= \sum_{W \in \text{SD}_\ell} \Pr_{G^* \leftarrow \mathbb{F}_2^{n \times g}} [W G^* = 0] \\ &= \frac{|\text{SD}_\ell|}{2^{(\ell-1)g}}. \end{aligned}$$

By Markov's inequality it follows that for any $q > 0$, with probability $1 - 1/q$ over $G^* \leftarrow \mathbb{F}_2^{n \times g}$,

$$|\text{SD}_\ell(G^*)| \leq \frac{|\text{SD}_\ell|}{2^{(\ell-1)g}} \cdot q.$$

Together with Lemma 11, we have that for any $q, \varepsilon > 0$ the expected number of simple dependencies computed in Equation (6) is at most

$$\begin{aligned} \sum_{\ell=1}^r \binom{r}{\ell} \frac{|\text{SD}_\ell(G^*)|}{\eta(G^*)^\ell} &\leq q \sum_{\ell=1}^r \binom{r}{\ell} \frac{|\text{SD}_\ell|}{\eta(G^*)^\ell \cdot 2^{(\ell-1)g}} \\ &\leq \frac{q}{(1-\varepsilon)^n} \sum_{\ell=1}^r \binom{r}{\ell} \frac{2^{g\ell} |\text{SD}_\ell|}{\binom{n}{t}^\ell 2^{(\ell-1)g}} \\ &= \frac{q \cdot 2^g}{(1-\varepsilon)^n} \sum_{\ell=1}^r \binom{r}{\ell} \frac{|\text{SD}_\ell|}{\binom{n}{t}^\ell} \end{aligned}$$

with probability $1 - \frac{1}{q} - \frac{2^g}{\binom{n}{t} \cdot \varepsilon^2}$ over G^* . Since

$$\frac{|\text{SD}_\ell|}{\binom{n}{t}^\ell} \leq \Pr_{w_1, \dots, w_\ell \leftarrow \mathcal{S}_{t,n}} [w_1 \oplus \dots \oplus w_\ell = 0],$$

the remainder of the proof is devoted to showing that

$$\sum_{\ell=1}^r \binom{r}{\ell} \Pr_{w_1, \dots, w_\ell \leftarrow \mathcal{S}_{t,n}} [w_1 \oplus \dots \oplus w_\ell = 0] \leq 2^{-ct^2}. \quad (7)$$

for some constant $c > 0$. Setting $\tilde{g} = ct^2/4$, $q = 2^{ct^2/4}$, and $\varepsilon = 2^{-t \log(n/t)/8}$ will then complete the proof of Lemma 12.

Let A be the transition matrix for the random walk on \mathbb{F}_2^n where, at each step, we sample a random $w \leftarrow \mathcal{S}_{t,n}$ and move $x \mapsto x \oplus w$. Observe that $\Pr_{w_1, \dots, w_\ell \leftarrow \mathcal{S}_{t,n}} [w_1 \oplus \dots \oplus w_\ell = 0]$ is equal to the probability that ℓ steps of this walk form a (not necessarily simple) cycle, i.e.,

$$\Pr_{w_1, \dots, w_\ell \leftarrow \mathcal{S}_{t,n}} [w_1 \oplus \dots \oplus w_\ell = 0] = \frac{1}{2^n} \text{Tr}(A^\ell).$$

Let H be the transition matrix for the random walk on the hypercube graph on \mathbb{F}_2^n . In Claim 13, we bound the probability that ℓ steps of A form a cycle in terms of the probability that ℓt steps of H form a cycle.

Claim 13. *If $\ell = O(n)$ and $t = o(n^{1/3})$, then $\text{Tr}(A^\ell) \leq O\left(e^{t^2 \ell/n} \cdot \text{Tr}(H^{\ell t})\right)$.*

Proof. Let \mathcal{E}_H denote the event that ℓt steps of H form a cycle, and let \mathcal{E}_A denote the event that ℓ steps of A form a cycle. Observe that

$$\frac{1}{2^n} \text{Tr}(A^\ell) = \Pr[\mathcal{E}_A] = \Pr[\mathcal{E}_H \mid \text{every } t\text{-block is distinct}]$$

where we consider the walk in H as consisting of ℓ consecutive blocks of t steps each, and we say a block is “distinct” if a different index is changed in each step.

Furthermore,

$$\frac{1}{2^n} \text{Tr}(H^{\ell t}) = \Pr[\mathcal{E}_H] \geq \Pr[\mathcal{E}_H \mid \text{every } t\text{-block is distinct}] \cdot \Pr[\text{every } t\text{-block is distinct}]$$

So $\Pr[\mathcal{E}_A] \leq \frac{\Pr[\mathcal{E}_H]}{\Pr[\text{every } t\text{-block is distinct}]}$. Thus, it suffices to see that

$$\begin{aligned} \Pr[\text{every } t\text{-block is distinct}] &= \Pr[\text{a given } t\text{-block is distinct}]^\ell \\ &\geq (1 - t/n)^{\ell t} \\ &\geq \left(e^{-t} \left(1 - \frac{t^2}{n}\right)\right)^{\ell t/n} \\ &\geq e^{-t^2 \ell/n} \cdot (1 - t^3 \ell/n^2) \\ &\geq \Omega(e^{-t^2 \ell/n}). \end{aligned}$$

□

Applying Claim 13, we have reduced the problem of proving Equation (7) to showing that

$$\sum_{\ell=1}^r \binom{r}{\ell} \cdot \frac{e^{t^2 \ell/n}}{2^n} \text{Tr}(H^{\ell t}) \leq 2^{-\Omega(t^2)}. \quad (8)$$

Fortunately, the eigenvalues of H are simple to compute:

$$\frac{1}{2^n} \text{Tr}(H^{\ell t}) = \mathbb{E}_{x \leftarrow \mathbb{F}_2^n} \left[\left(1 - \frac{2 \text{wt}(x)}{n} \right)^{\ell t} \right] \quad (9)$$

where $\text{wt}(x)$ denotes the Hamming weight of $x \in \mathbb{F}_2^n$. We analyze this expression separately depending on how large ℓ is.

Small ℓ ($\ell \leq (e - \Omega(1)) \cdot n/t$). We can rewrite Equation (9) as a moment of a simple random walk:

$$\mathbb{E}_{x \leftarrow \mathbb{F}_2^n} \left[\left(1 - \frac{2 \text{wt}(x)}{n} \right)^{\ell t} \right] = n^{-\ell t} \mathbb{E}_{X_1, \dots, X_n \leftarrow \{1, -1\}} \left[\left(\sum_{i=1}^n X_i \right)^{\ell t} \right].$$

Let $X = \sum_{i=1}^n X_i$ where $X_1, \dots, X_n \leftarrow \{1, -1\}$. For small ℓ , the Gaussian approximation to these moments is good enough. For even p ,

$$\begin{aligned} \mathbb{E}[X^p] &= \sum_{i_1, \dots, i_p \in [n]} \mathbb{E}[X_{i_1} \cdots X_{i_p}] \\ &= |\{(i_1, \dots, i_p) \in [n]^p \mid \text{each } i_j \text{ appears an even number of times}\}| \\ &\leq n^{p/2} \cdot (p-1)!! \end{aligned}$$

For odd p , $\mathbb{E}[X^p] = 0$. Therefore, for $\ell = O(n/t)$ we have

$$\begin{aligned} \binom{r}{\ell} \cdot \frac{e^{t^2 \ell/n}}{2^n} \text{Tr}(H^{\ell t}) &\leq e^{O(t)} \binom{n}{\ell} \cdot \frac{(\ell t - 1)!!}{n^{\ell t/2}} \\ &\leq e^{O(t)} \left(\frac{en}{\ell} \right)^\ell \cdot \left(\frac{\ell t}{en} \right)^{\ell t/2} \cdot O(\sqrt{\ell t}) \\ &= e^{O(t)} \left(\frac{\ell}{en} \right)^{\ell(t/2-1)} \cdot t^{\ell t/2} \cdot O(\sqrt{\ell t}) \end{aligned}$$

If $\ell \leq t$, then this yields

$$\binom{r}{\ell} \cdot \frac{e^{t^2 \ell/n}}{2^n} \text{Tr}(H^{\ell t}) \leq e^{O(t)} \left(\frac{t^{t-1}}{(en)^{t/2-1}} \right)^\ell = n^{-\Omega(t)}$$

and if $t < \ell \leq (e - \Omega(1)) \cdot n/t$,

$$\binom{r}{\ell} \cdot \frac{e^{t^2 \ell/n}}{2^n} \text{Tr}(H^{\ell t}) \leq e^{O(t)} \left(\frac{e - \Omega(1)}{et} \right)^{\ell(t/2-1)} \cdot t^{\ell t/2} \cdot O(\sqrt{\ell t}) = 2^{-\Omega(t^2)}.$$

In either case we have a bound of $2^{-\Omega(t^2)}$.

Large ℓ ($\ell \geq (e - o(1)) \cdot n/t$). Using the binomial theorem,

$$\begin{aligned} \frac{1}{2^n} \text{Tr}(H^{\ell t}) &= \mathbb{E}_{x \leftarrow \mathbb{F}_2^n} \left[\left(1 - \frac{2 \text{wt}(x)}{n} \right)^{\ell t} \right] \\ &= \frac{1}{2^n} \sum_{s=0}^n \binom{n}{s} \cdot \left(1 - \frac{2s}{n} \right)^{\ell t} \\ &\leq \frac{1}{2^n} \sum_{s=0}^n \binom{n}{s} \cdot e^{-2s\ell t/n} \\ &= \left(\frac{1 + e^{-2\ell t/n}}{2} \right)^n. \end{aligned}$$

For $\ell \geq (e - o(1)) \cdot n/t$, this is at most $\left(\frac{1 + e^{o(1) - 2e}}{2} \right)^n$. Since $\left(\frac{1 + e^{-2e}}{2} \right) < 2^{-0.99}$, we have for $r \leq 0.99n$ that

$$\sum_{\ell=1}^r \binom{r}{\ell} \cdot \frac{e^{t^2 \ell/n}}{2^n} \text{Tr}(H^{\ell t}) \leq 2^r \cdot e^{t^2} \cdot \left(\frac{1 + e^{o(1) - 2e}}{2} \right)^n = 2^{-\Omega(n)}. \quad \square$$

Theorem 2. For any $p, \eta \in (0, 1/2)$, there exists $g = \Omega(\log^2 n)$, $t = \Theta(\log n)$ such that $\text{LDPC-PRC}_0[n, g, t, 0.99n, \eta, (0.99n)^{-1/4}]$ is a zero-bit public-key PRC (Definition 2) that is robust to every p -bounded channel, where pseudorandomness rests on the $\text{LPN}_{g, \eta}$ assumption (Assumption 2).

Proof. By Lemma 4, there exists $t = \Theta(\log n)$ such that for any $g > 0$, $\text{LDPC-PRC}_0[n, g, t, 0.99n, \eta, (0.99n)^{-1/4}]$ is robust to every p -bounded channel.

By Lemma 9, there exists a function $g = \Omega(\log^2 n)$ such that the generator matrix of this code — that is, G for $(P, G) \leftarrow \text{LDPC}[n, g, t, 0.99n]$ — is $\text{negl}(n)$ -close to uniformly random.

Since G has dimensions $n \times g$, the $\text{LPN}_{g, \eta}$ assumption implies that $\text{LDPC-PRC}_0[n, g, t, 0.99n, \eta, (0.99n)^{-1/4}]$ is a public-key PRC. \square

6 Boosting the rate and robustness of any pseudorandom code

6.1 Multi-bit pseudorandom codes

We show how to construct a multi-bit PRC from any zero-bit PRC. The high-level idea is to encode a given message bit-by-bit, where we use codewords from the zero-bit PRC to represent bits of the message that are 1, and uniformly random strings to represent bits of the message that are 0. We say this encoding consists of many *blocks*, where each block is either a codeword from the zero-bit PRC or a random string. Since our zero-bit PRC allows a decoder with the secret key to distinguish uniform strings from noisy codewords, we can use this decoder to recover each bit of the message from each block.

However, as described so far, there are two issues with this scheme. The first is that this scheme encodes the all-0 string as a uniformly random string, but the error correction property of a PRC requires that the decoder can distinguish encodings (of any message) from random strings. Thus, we modify the above scheme to append a codeword from the zero-bit PRC to the end of every encoding.

The other issue is that this scheme may lose the zero-bit PRC's robustness. In particular, consider a zero-bit PRC that is robust to all p -bounded channels; we would like for our multi-bit PRC to retain this robustness. However, consider the channel that flips each bit in only the first block of the encoding, independently with probability $\frac{1}{2}$. The decoder will now be unable to recover the first bit of the message, and this channel is p -bounded since it changes only a sub-constant fraction of the bits. The issue here is that our p -bounded channel was not bounded at all on the first block; we'd like for the channel's effect on every block of the encoding to be p -bounded. We solve this issue by randomly permuting the encoding, which ensures that the errors introduced by the channel cannot be too concentrated in any block. The decoder now inverts the permutation before decoding block-by-block as before.

Although the constructions are essentially the same, we separately present multi-bit secret-key and public-key PRCs as they have slightly different syntax.

Construction 3 (Multi-bit secret-key PRC). Let $\text{PRC}_0 = (\text{KeyGen}_0, \text{Encode}_0, \text{Decode}_0)$ be a zero-bit secret-key PRC with block length n . We define an ℓ -bit secret-key PRC, $\text{PRC}_\ell = (\text{KeyGen}_\ell, \text{Encode}_\ell, \text{Decode}_\ell)$, as follows:

- $\text{KeyGen}_\ell(1^\lambda)$: Sample $\text{sk}_0 \leftarrow \text{KeyGen}_0(1^\lambda)$. Sample a random permutation $\pi : [n \cdot (\ell + 1)] \rightarrow [n \cdot (\ell + 1)]$. Output $\text{sk} = (\text{sk}_0, \pi)$.
- $\text{Encode}_\ell(\text{sk}, \mathbf{m})$: Given as input a message $\mathbf{m} \in \{0, 1\}^\ell$, for each $i \in [\ell + 1]$, let

$$c_i = \begin{cases} c \leftarrow \{0, 1\}^n & \text{if } \mathbf{m}_i = 0 \text{ and } i \neq \ell + 1 \\ c \leftarrow \text{Encode}_0(\text{sk}_0) & \text{if } \mathbf{m}_i = 1 \text{ and } i \neq \ell + 1 \\ c \leftarrow \text{Encode}_0(\text{sk}_0) & \text{if } i = \ell + 1 \end{cases}$$

Let $y = c_1 || \dots || c_{\ell+1} \in \{0, 1\}^{n(\ell+1)}$. Output $x = y_{\pi(1)} || \dots || y_{\pi(n(\ell+1))}$.

- $\text{Decode}_\ell(\text{sk}, x_1 || \dots || x_{n(\ell+1)})$: For each $i \in [\ell + 1]$, let $\hat{y}_i = x_{\pi^{-1}(1+(i-1)n)} || x_{\pi^{-1}(2+(i-1)n)} || \dots || x_{\pi^{-1}(n+(i-1)n)}$.

$$\hat{x}_i = \begin{cases} 1 & \text{if } \text{Decode}_0(\text{sk}_0, \hat{y}_i) = 1 \\ 0 & \text{otherwise} \end{cases}$$

If $\hat{x}_{\ell+1} \neq 1$, output \perp . Otherwise, output $\hat{\mathbf{m}} = \hat{x}_{\pi^{-1}(1)} || \dots || \hat{x}_{\pi^{-1}(\ell)}$.

Construction 4 (Multi-bit public-key PRC). Let $\text{PRC}_0 = (\text{KeyGen}_0, \text{Encode}_0, \text{Decode}_0)$ be a zero-bit public-key PRC with block length n . We define a ℓ -bit public-key PRC, $\text{PRC}_\ell = (\text{KeyGen}_\ell, \text{Encode}_\ell, \text{Decode}_\ell)$, as follows:

- $\text{KeyGen}_\ell(1^\lambda)$: Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{KeyGen}_0(1^\lambda)$. Sample a random permutation $\pi : [n \cdot (\ell + 1)] \rightarrow [n \cdot (\ell + 1)]$. Output $\text{sk} = (\text{sk}_0, \pi)$ and $\text{pk} = (\text{pk}_0, \pi)$.

- $\text{Encode}_\ell(\text{pk}, m)$: Given as input a message $m \in \{0, 1\}^\ell$, for each $i \in [\ell + 1]$, let

$$c_i = \begin{cases} c \leftarrow \{0, 1\}^n & \text{if } m_i = 0 \text{ and } i \neq \ell + 1 \\ c \leftarrow \text{Encode}_0(\text{pk}_0) & \text{if } m_i = 1 \text{ and } i \neq \ell + 1 \\ c \leftarrow \text{Encode}_0(\text{pk}_0) & \text{if } i = \ell + 1 \end{cases}$$

Let $y = c_1 || \dots || c_{\ell+1} \in \{0, 1\}^{n(\ell+1)}$. Output $x = y_{\pi(1)} || \dots || y_{\pi(n(\ell+1))}$.

- $\text{Decode}_\ell(\text{sk}, x_1 || \dots || x_{n(\ell+1)})$: For each $i \in [\ell+1]$, let $\hat{y}_i = x_{\pi^{-1}(1+(i-1)n)} || x_{\pi^{-1}(2+(i-1)n)} || \dots || x_{\pi^{-1}(n+(i-1)n)}$.

$$\hat{x}_i = \begin{cases} 1 & \text{if } \text{Decode}_0(\text{sk}_0, \hat{y}_i) = 1 \\ 0 & \text{otherwise} \end{cases}$$

If $\hat{x}_{\ell+1} \neq 1$, output \perp . Otherwise, output $\hat{m} = \hat{x}_{\pi^{-1}(1)} || \dots || \hat{x}_{\pi^{-1}(\ell)}$.

Claim 14. *If $\text{PRC}_0 = (\text{KeyGen}_0, \text{Encode}_0, \text{Decode}_0)$ is a zero-bit secret-key (public-key) PRC robust to p -bounded channels, PRC_ℓ is a ℓ -bit secret-key (public-key) PRC robust to $(p - \varepsilon)$ -bounded channels for any constant $\varepsilon \in (0, p)$.*

Proof. Error correction. We first show that with overwhelming probability, a $(p - \varepsilon)$ -bounded channel applied to a codeword of Encode_ℓ is p -bounded on each of its blocks. Since PRC_0 is robust against p -bounded channels, this implies that each block will be correctly decoded by Decode_0 . First, by definition of p -boundedness, the entire codeword from PRC_ℓ will have at most $(p - \varepsilon)n(\ell + 1)$ errors with overwhelming probability. Let $\alpha \leq p - \varepsilon$ be the actual fraction of errors. Now, fix any n -length block of the codeword after the permutation has been inverted. Observe that over the randomness of the permutation, the number of errors in this block is a random variable $X \sim \text{Hyp}(n(\ell + 1), \alpha n(\ell + 1), n)$. By Lemma 3,

$$\Pr[X \geq (\alpha + \varepsilon)n] \leq e^{-2\varepsilon^2 n},$$

which is negligible in n . Since $\alpha \leq p - \varepsilon$, the probability that there are at least pn errors in our block is at most the probability that there are $(\alpha + \varepsilon)n$ errors, which we've just shown is negligible. By a union bound, the probability that *any* block has more than pn errors is negligible.

We now show the second property of error correction, that an unrelated string c decodes to \perp with overwhelming probability. Consider the last block $\hat{y}_{\ell+1}$ of c after the permutation has been inverted. By error correction of PRC_0 , the probability over sk_0 that $\text{Decode}_0(\text{sk}_0, \hat{y}_{\ell+1}) = 1$ is negligible. Therefore, with overwhelming probability, $\text{Decode}_\ell(\text{sk}, c) = \perp$.

Pseudorandomness. We prove pseudorandomness of public-key PRC_ℓ . We observe that the same proof holds for secret-key PRC_ℓ , by omitting the public keys as input to the adversaries \mathcal{A} and \mathcal{B} , and instead giving them oracle access to $\text{Encode}_0(\text{sk}_0, \cdot)$ and $\text{Encode}_\ell(\text{sk}, \cdot)$ respectively.

Suppose that an adversary \mathcal{A} given pk can distinguish between oracle access to $\text{Encode}_\ell(\text{pk}, \cdot)$ and the uniform distribution. Then \mathcal{B} given pk_0 can distinguish between oracle access to Encode_0 and the uniform distribution as follows. \mathcal{B} samples a random permutation $\pi : [n(\ell + 1)] \rightarrow [n(\ell + 1)]$ and gives $\text{pk} = (\text{pk}_0, \pi)$ to \mathcal{A} as input. Whenever \mathcal{A} queries m to its oracle, \mathcal{B} computes a response x by following Encode_ℓ , but querying its own oracle when Encode_ℓ requires a call to Encode_0 , and using π as the permutation. Observe that if \mathcal{B} 's oracle is the uniform distribution, the resulting x is permutation of a uniform string, so x is uniform. If \mathcal{B} 's oracle is Encode_0 , the resulting x is drawn from $\text{Encode}_\ell(\text{sk}, m)$. Therefore, \mathcal{B} 's advantage is exactly \mathcal{A} 's advantage. \square

Remark. *Unfortunately, Constructions 3 and 4 yield codes of rate roughly $1/n$: If the underlying zero-bit PRC has block length n , then in order to encode a ℓ -bit message we need codewords of length $n \cdot (\ell + 1)$. Ideally, we would like to have constant rate — i.e., to encode ℓ -bit messages into $O(\ell)$ -bit codewords. While*

typical LDPC codes have constant rate, and there is a simple modification of our pseudorandom LDPC codes that have constant rate, at densities of $\Omega(\log n)$ it is not known how to decode from a constant rate of errors. Instead, we build constant-rate PRCs generically from any multi-bit PRC in Section 6.2.

6.2 Constant-rate pseudorandom codes

We now show how to build constant-rate PRCs from any multi-bit PRC (as in Section 6.1) and any constant-rate error-correcting code. We state the public-key version of the following construction only; as usual, the secret-key version is similar.

Construction 5 (Constant-rate public-key PRC). Let λ be a security parameter and PRC_λ be a λ -bit public-key PRC with block length n' . Let (Enc, Dec) be any error-correcting code with block length $n > \lambda$ and messages of length k . Let $\text{PRG} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^n$ be any pseudorandom generator. We define $\text{PRC}[\text{PRC}_\lambda, (\text{Enc}, \text{Dec})]$ which is a k -bit public-key PRC as follows:

- $\text{PRC.KeyGen}(1^\lambda)$: Sample $(\text{sk}', \text{pk}') \leftarrow \text{PRC}_\lambda.\text{KeyGen}(1^\lambda)$. Sample a random permutation $\pi : [n' + n] \rightarrow [n' + n]$. Let $\text{sk} = (\text{sk}', \pi)$ and $\text{pk} = (\text{pk}', \pi)$; output (sk, pk) .
- $\text{PRC.Encode}(\text{pk}, \text{m})$: Given as input a message $\text{m} \in \{0, 1\}^k$, let $r \leftarrow \{0, 1\}^\lambda$, and let

$$x \leftarrow \text{PRC}_\lambda.\text{Encode}(\text{pk}, r) \parallel \text{PRG}(r) \oplus \text{Enc}(\text{m}).$$

Let $x_1, \dots, x_{n'+n}$ denote the bits of x . The output is $x_{\pi(1)} \parallel x_{\pi(2)} \parallel \dots \parallel x_{\pi(n'+n)}$.

- $\text{PRC.Decode}(\text{sk}, c)$: Letting $c_1, \dots, c_{n'+n}$ denote the bits of c , the decoder first computes

$$y = c_{\pi^{-1}(1)} \parallel \dots \parallel c_{\pi^{-1}(n'+n)}.$$

The decoder then parses y as $y_1 \parallel y_2$, where $|y_1| = n'$ and $|y_2| = n$. It computes $r \leftarrow \text{PRC}_\lambda.\text{Decode}(\text{sk}, y_1)$ and outputs $\text{m} = \text{Dec}(\text{PRG}(r) \oplus y_2)$.

The block length of the resulting PRC is $n + n'$, so the rate is $k/(n + n')$. Since n' is only a function of the security parameter λ and does not depend on k , for large k the rate approaches the rate k/n of the underlying error-correcting code (Enc, Dec) .

Theorem 3. Let PRC_λ be a λ -bit public-key PRC with block length n' , and let (Enc, Dec) be any error-correcting code with block length $n > \lambda$ and messages of length k .

Then $\text{PRC} = \text{PRC}[\text{PRC}_\lambda, (\text{Enc}, \text{Dec})]$ of Construction 5 is a public-key PRC where

- the rate of PRC is $k/(n + n')$; and
- for any constants $p, \varepsilon \in (0, 1/2)$, if PRC_λ and (Enc, Dec) are robust to channels that introduce at most a p fraction of errors at random locations, then PRC is robust to all $(p - \varepsilon)$ -bounded channels.

Proof. Pseudorandomness. By pseudorandomness of PRC_λ , no polynomial-time adversary can distinguish between PRC.Encode and a hybrid where $x \leftarrow \text{PRC}_\lambda.\text{Encode}(\text{pk}, r)$ is replaced by a uniformly random $x \leftarrow \{0, 1\}^{n'}$. By pseudorandomness of the PRG, this is computationally indistinguishable from a hybrid where $\text{PRG}(r)$ is replaced by a uniformly random $s \leftarrow \{0, 1\}^n$, making $(x \parallel s \oplus \text{Enc}(\text{m}))$ uniform. Therefore, in this hybrid, which is indistinguishable from oracle access to PRC.Encode , each query to the oracle outputs an independently drawn uniform string in $\{0, 1\}^{n+n'}$.

Robustness. We first show that for any p -bounded channel \mathcal{E} , if the decoder receives $c \leftarrow \mathcal{E}(\text{PRC.Encode}(\text{pk}, \text{m}))$, then the resulting y_1 and y_2 are equal to $\mathcal{E}_1(\text{Encode}(r))$ and $\mathcal{E}_2(\text{Enc}(\text{m}) \oplus \text{PRG}(r))$ respectively, where \mathcal{E}_1 and \mathcal{E}_2 are both p -bounded. Observe that the number of errors introduced by \mathcal{E}_1 is distributed as a hypergeometric random variable with a population of size $n + n'$, $p(n + n')$ success elements (representing the errors), and n' draws. By Lemma 3, the probability that \mathcal{E}_1 introduces at least $(p - \varepsilon)n'$ errors is at most $e^{-2\varepsilon^2 n'}$,

which is negligible. Similarly, the number of errors introduced by \mathcal{E}_2 is distributed as a hypergeometric random variable with a population of size $n + n'$, $p(n + n')$ success elements (representing the errors), and n draws. By Lemma 3, the probability that \mathcal{E}_2 introduces at least $(p - \varepsilon)n'$ errors is at most $e^{-2\varepsilon^2 n}$, which is negligible as $n > \lambda$. By a union bound, with overwhelming probability both \mathcal{E}_1 and \mathcal{E}_2 introduce errors with at most a rate of $(p - \varepsilon)$, and therefore they are $(p - \varepsilon)$ -bounded.

Furthermore, because of the random permutation π , the locations of the errors on y_2 are random. Therefore by robustness of PRC_λ , we have that $\text{Decode}(\text{sk}, y_1) = r$ with overwhelming probability; and by error correction of (Enc, Dec) , $\text{Dec}(\text{PRG}(r) \oplus y_2) = \text{Dec}(\text{BSC}_p(\text{Enc}(m))) = m$ with overwhelming probability as well. \square

If we instantiate Construction 5 with

- PRC_λ from Construction 4, using the LDPC-based PRCs of Construction 2 as PRC_0 ; and
- the error-correcting codes (Enc, Dec) of [ABN⁺92, NN93, Ta-17],

then we obtain constant-rate PRCs with robustness to every p -bounded channel for $p \in (0, 1/2)$.

6.3 Pseudorandom codes for the deletion channel

In this section, our primary goal is to construct a PRC, PRC_{del} , that is robust against a constant-rate deletion channel. That is, the message can be recovered from an encoding under PRC_{del} when each of its bits is deleted independently with some constant probability p . The actual robustness guarantee we obtain is even stronger: we can recover the message even when this constant-rate deletion channel is composed with a constant-rate binary symmetric channel. Our construction makes black-box use of any PRC that is robust against p -bounded channels.

Our high-level idea is to take an encoding of this original PRC and make it redundant, while preserving pseudorandomness. Let $x = x_1, \dots, x_n$ be a PRC encoding of some message. We could make x robust to the deletion channel by duplicating each bit x_i some number of times m . However, the result would no longer be pseudorandom. Instead of duplicating x_i , we sample a random $y \in \{0, 1\}^m$, conditioned on the majority of the bits of y being x_i . Observe that if x_i is $\text{Ber}(\frac{1}{2})$, y is uniform over $\{0, 1\}^m$. Since x_i is a bit of our PRC, it is pseudorandom, and y is indistinguishable from uniform. The resulting codeword is y_1, \dots, y_n , where each $y_i \in \{0, 1\}^m$. We call this the “majority code.”

Observe that given $y = y_1, \dots, y_n$, one can recover x by computing the majority of each length- m block y_i . If each bit of y is deleted independently with some probability p to yield y' , we can recover an approximate version of x as follows. We partition y' into equal-length blocks y'_1, \dots, y'_n , and we compute the majority of each block. Intuitively, we expect most of the new blocks y'_i to be subsets of the corresponding original blocks y_i . Since the deletions are random, we expect them to preserve the majority for most blocks. Therefore, the message x' that we recover should be approximately x , with some bounded number of bits flipped. Recalling that x itself was a PRC codeword, and our PRC is robust to bounded-weight channels, we can recover the original message from x' .

It may be tempting to apply this majority code to other encoding schemes such as pseudorandom encryption, and the result would indeed be pseudorandom. However, it would not be robust against a constant-rate deletion channel. Since decoding from the deletion channel introduces bounded-weight errors, the message that the majority code is applied to must be error-correcting against bounded-weight channels. Therefore, our PRC is crucial to this approach.

We first introduce some notation. Let $\text{Maj}(y)$ denote the majority function for bitstrings y , where ties are resolved by choosing a random bit. For a given positive odd integer $m \in \mathbb{N}$, let $\text{MajEnc}_m : \{0, 1\} \rightarrow \{0, 1\}^k$ be a randomized function that takes an input bit $b \in \{0, 1\}$ and outputs a random string $y \in \{0, 1\}^m$ conditioned on $\text{Maj}(y) = b$. For $x \in \{0, 1\}^n$, we define $\text{MajEnc}_m(x) := (\text{MajEnc}_m(x_1), \dots, \text{MajEnc}_m(x_n)) \in \{0, 1\}^{nm}$.

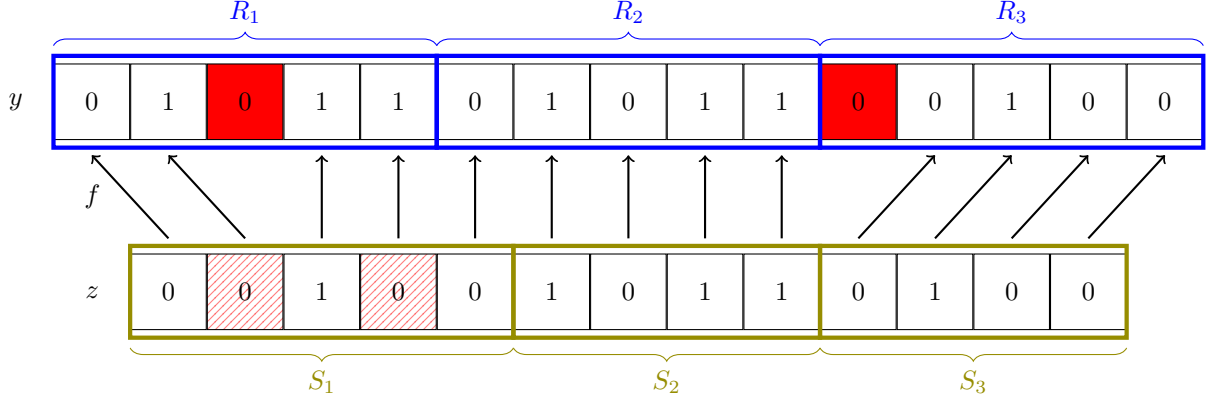


Figure 2: Objects from the proof of Lemma 15: $y \leftarrow \text{MajEnc}_m(x)$, $z \leftarrow \text{BSC}_q \circ \text{BDC}_p(y)$, the function f , and the partitions \mathcal{R}, \mathcal{S} for $n = 3, m = 5$. In this illustration $x = (1, 1, 0)$ and $\text{MajDec}_3(z) = (0, 1, 0)$. There are deletions at locations $D = \{3, 11\}$ (indicated by solid red), and there are errors from BSC_q at locations 2 and 4 (indicated by hatched red). The arrows indicate the mapping f , i.e. an arrow points from an index in z to the index in y from which it originated.

The decoding algorithm $\text{MajDec}_n : \{0, 1\}^* \rightarrow \{0, 1\}^n$ evenly partitions the received string into n blocks and outputs a string of the blocks' majorities.

$\text{MajEnc}_m(x)$:

1. For each $i \in [\text{len } x]$, let $y^i = \text{MajEnc}_m(x_i)$.
2. Output $(y^1 || \dots || y^{\text{len } x}) \in \{0, 1\}^{\text{len } x \cdot m}$.

$\text{MajDec}_n(z)$:

1. Partition z into n blocks $z = (z^1 || \dots || z^n)$ such that $|\text{len}(z^i) - \text{len}(z)/n| \leq 1$ for all $i \in [n]$.
2. Output $(\text{Maj}(z^1), \dots, \text{Maj}(z^n))$.

In Lemma 15, we show that majority encoding a message allows us to recover an approximation of the message with bounded-weight error after a deletion channel is applied. This holds even if this deletion channel is composed with a binary symmetric channel.

Lemma 15. *Assume $m = \Omega(n \log^6 n)$. For any constant deletion rate $p \in (0, 1)$ and error rate $q \in (0, 1/2)$, there exists a constant $\varepsilon > 0$ such that*

$$\Pr_{x \leftarrow \{0, 1\}^n} [\text{wt}(x \oplus \text{MajDec}_n \circ \text{BSC}_q \circ \text{BDC}_p \circ \text{MajEnc}_m(x)) \leq (1/2 - \varepsilon) \cdot n] \geq 1 - \text{negl}(n).$$

Proof. Let $y \leftarrow \text{MajEnc}_m(x)$ and let $z \leftarrow \text{BSC}_q \circ \text{BDC}_p(y)$. Let $D \subseteq [nm]$ be the set of indices deleted by BDC_p ; D is a random variable where each $i \in [nm]$ is included in D independently with probability p . Let $f : [nm - |D|] \rightarrow [nm]$ be the function that maps indices of z to those of y , i.e., $z_j = \text{BSC}_q(y_{f(j)})$.

Let $\mathcal{R} = R_1, \dots, R_n \subseteq [nm]$ be the partition of $[nm]$ into n blocks of length m . Let $\mathcal{S} = S_1, \dots, S_n \subseteq [nm - |D|]$ be the partition of $[nm - |D|]$ into almost-equal-sized blocks from the definition of $\text{MajDec}_n(z)$. \mathcal{R} and \mathcal{S} define partitions of y and z , respectively. All of

Let $a, b \in [nm]$ be such that $a \leq b$. Applying Hoeffding's inequality to $|[a, b] \setminus D| = \sum_{j=a}^b \mathbb{1}[j \notin D]$ yields

$$\Pr \left[\left| (b - a + 1)(1 - p) - |[a, b] \setminus D| \right| > \sqrt{nm \log n} \right] \leq e^{-\Omega(\log^2 n)}.$$

By a union bound, with probability $1 - \text{negl}(n)$ we have that for every $a, b \in [nm]$

$$\left| (b - a + 1)(1 - p) - |[a, b] \setminus D| \right| \leq \sqrt{nm} \log n. \quad (10)$$

That is, the number of deletions in every contiguous interval $[a, b]$ of y concentrates and is roughly $p(b - a) \pm \sqrt{nm} \log n$.

We now argue in Claim 16 that because of this concentration, each S_i in the decoder's partition corresponds to a subset of the original message y that is largely contained in R_i . That is, the symmetric set difference between $f(S_i)$ and $(R_i \setminus D)$ is small.

Claim 16. *If Equation (10) holds for all $a, b \in [nm]$, then for all $i \in [n]$ we have $|f(S_i) \Delta (R_i \setminus D)| = O(\sqrt{nm} \log n)$ (where Δ denotes the symmetric set difference).*

Proof. Observe that $|f(S_i) \Delta (R_i \setminus D)| \leq |(i - 1)m - \min f(S_i)| + |im - \max f(S_i)|$. We will show that $|(i - 1)m - \min f(S_i)| = O(\sqrt{nm} \log n)$. A nearly identical proof shows that $|im - \max f(S_i)| = O(\sqrt{nm} \log n)$ as well, so this will complete the proof.

Setting $[a, b] = [nm]$ to be all of y in Equation (10), we have that $|(1 - p)nm - \text{len } z| \leq \sqrt{nm} \log n$. Therefore, $|(1 - p)m - |S_i|| \leq \left\lceil \frac{\sqrt{nm} \log n}{n} \right\rceil$ for every $i \in [n]$. It follows that for every $i \in [n]$,

$$|(1 - p)(i - 1)m - \min S_i| = O(\sqrt{nm} \log n). \quad (11)$$

For any $j \in [nm - |D|]$, setting $[a, b] = [f(j)]$ in Equation (10) we have that $|j - (1 - p)f(j)| = O(\sqrt{nm} \log n)$. Applying this fact to $j = \min S_i$ for any $i \in [n]$,

$$|\min S_i - (1 - p) \min f(S_i)| = O(\sqrt{nm} \log n).$$

Together with Equation (11) and the fact that $p \in (0, 1)$ is a constant, we have

$$|(i - 1)m - \min f(S_i)| = O(\sqrt{nm} \log n)$$

as desired. □

So far, we've argued that each block z^i in the decoder's partition consists mostly of bits from block y^i of the encoding. We now argue that with high probability, the bits of y^i that were excluded and the bits from other blocks that were erroneously included do not affect the majority of z^i .

To do so, we define Δ_i to be the difference of the erroneously excluded bits and the erroneously included bits, converted to values in $\{-1, 1\}$ so we can later reason about this as a random walk. Let

$$\Delta_i = \sum_{j \in (R_i \setminus D) \setminus f(S_i)} (-1)^{y_j} - \sum_{j \in f(S_i) \setminus (R_i \setminus D)} (-1)^{y_j}.$$

Conditioned on Equation (10) holding for all $a, b \in [nm]$, Claim 16 implies that Δ_i is a sum of $O(\sqrt{nm} \log n)$ random $\{1, -1\}$ variables. Furthermore, these random variables are independent because y is uniform over $\{0, 1\}^{nm}$. By Hoeffding's inequality,

$$\Pr[|\Delta_i| \geq (nm)^{1/4} \log^{3/2} n] \leq 2 \exp\left(\frac{-\sqrt{nm} \log^3 n}{O(\sqrt{nm} \log n)}\right) \leq \text{negl}(n). \quad (12)$$

By a union bound, the probability that $|\Delta_i| \leq (nm)^{1/4} \log^{3/2} n$ for all $i \in [n]$ is at least $1 - \text{negl}(n)$.

We will complete the proof by applying Claim 17 with $N = |S_i|$, $N' = |R_i \cap D|$, and $\Delta = \Delta_i$. The variables X_j are the values of y_j for $j \in f(S_i) \cup (R_i \cap D)$ and Z_j are the errors introduced by BSC_q . Since the sets $f(S_i) \cup (R_i \cap D)$ are disjoint for distinct i , the events that $\text{Maj}(z_i) = x_i$ are independent once we condition on the values of y_j for $j \notin \bigcup_{i \in [n]} f(S_i) \cup (R_i \cap D)$; therefore the result will follow from a Chernoff bound.

Claim 17. Let $N, N' \in \mathbb{N}$ and $\Delta \in \mathbb{R}$ be such that $N' = O(N)$ and $\Delta = O(\sqrt{N})$. Let $X_1, \dots, X_N, Y_1, \dots, Y_{N'}$ be independent, uniform $\{-1, 1\}$ random variables and Z_1, \dots, Z_N be independent $\{-1, 1\}$ random variables with expectation $\delta = \Omega(1)$. Then for sufficiently large N ,

$$\Pr \left[\operatorname{sgn} \left(\sum_{j \in [N]} X_j \cdot Z_j \right) = \operatorname{sgn} \left(\sum_{j \in [N]} X_j + \sum_{j \in [N']} Y_j + \Delta \right) \right] = \frac{1}{2} + \Omega(1),$$

where we define $\operatorname{sgn}(0)$ to be a random value in $\{-1, 1\}$.

Proof. We can equivalently view Z_j as a random variable which is 1 with probability δ , and a uniformly random $\{-1, 1\}$ value with probability $1 - \delta$. Let F be the set of indices $j \in [N]$ where Z_j is deterministically 1. By a Chernoff bound, $|\delta N - |F|| \leq \delta N/2$ with probability $1 - \operatorname{negl}(n)$; fix such an F . We can write

$$\begin{aligned} \sum_{j \in [N]} X_j \cdot Z_j &= \sum_{j \in F} X_j + \sum_{j \in [N] \setminus F} X_j \cdot Z_j \\ &=: \sum_{j \in F} X_j + \sum_{j \in [N_a]} X_j^a \end{aligned}$$

where we have defined $N_a := N - |F|$ and $\{-1, 1\}$ variables $\{X_j^a\}_{j \in [N_a]} = \{X_j \cdot Z_j\}_{j \in [N] \setminus F}$. We also write

$$\begin{aligned} \sum_{j \in [N]} X_j + \sum_{j \in [N']} Y_j + \Delta &= \sum_{j \in F} X_j + \left(\sum_{j \in [N] \setminus F} X_j + \sum_{j \in [N']} Y_j \right) + \Delta \\ &=: \sum_{j \in F} X_j + \sum_{j \in [N_b]} X_j^b + \Delta \end{aligned}$$

where we have defined $N_b := N - |F| + N'$ and $\{-1, 1\}$ variables $\{X_j^b\}_{j \in [N_b]} = \{X_j\}_{j \in [N] \setminus F} \cup \{Y_j\}_{j \in [N']}$. Observe that $\{X_j\}_{j \in F} \cup \{X_j^a\}_{j \in [N_a]} \cup \{X_j^b\}_{j \in [N_b]}$ are all independent, uniformly random $\{-1, 1\}$ variables. We wish to show that

$$\Pr \left[\operatorname{sgn} \left(\sum_{j \in F} X_j + \sum_{j \in [N_a]} X_j^a \right) = \operatorname{sgn} \left(\sum_{j \in F} X_j + \sum_{j \in [N_b]} X_j^b + \Delta \right) \right] = \frac{1}{2} + \Omega(1).$$

Let $X = \sum_{j \in F} X_j$, $A = \sum_{j \in [N_a]} X_j^a$, and $B = \sum_{j \in [N_b]} X_j^b$. The above equation we wish to show becomes

$$\Pr[\operatorname{sgn}(X + A) = \operatorname{sgn}(X + B + \Delta)] = \frac{1}{2} + \Omega(1).$$

If $|X| > \max\{|A|, |B + \Delta|\}$, then $\operatorname{sgn}(X + A) = \operatorname{sgn}(X + B + \Delta) = \operatorname{sgn}(X)$. On the other hand,

$$\Pr \left[\operatorname{sgn}(X + A) = \operatorname{sgn}(X + B + \Delta) \mid |X| \leq \max\{|A|, |B + \Delta|\} \right] = \frac{1}{2}.$$

Therefore, it suffices to show that $\Pr[|X| > \max\{|A|, |B + \Delta|\}] = \Omega(1)$. We do this by bounding each of the three factors below separately:

$$\Pr \left[|X| > \max\{|A|, |B + \Delta|\} \right] \geq \Pr \left[|X| > \sqrt{|F|} + |\Delta| \right] \cdot \Pr \left[|A| \leq \sqrt{|F|} + |\Delta| \right] \cdot \Pr \left[|B + \Delta| \leq \sqrt{|F|} + |\Delta| \right].$$

- Since $\sqrt{|F|} + |\Delta| = O(\sqrt{|F|})$, the central limit theorem implies that $\Pr \left[|X| > \sqrt{|F|} + |\Delta| \right] = \Omega(1)$.
- Since $N_a = O(|F|)$, Hoeffding's inequality implies that $\Pr \left[|A| \leq \sqrt{|F|} + |\Delta| \right] \geq \Pr \left[|A| \leq \sqrt{|F|} \right] = \Omega(1)$.

- By the triangle inequality, $\Pr \left[|B + \Delta| \leq \sqrt{|F|} + |\Delta| \right] \geq \Pr \left[|B| \leq \sqrt{|F|} \right]$. Since $N_b = O(|F|)$, Hoeffding's inequality again implies that $\Pr \left[|B| \leq \sqrt{|F|} \right] = \Omega(1)$. \square

It only remains to show that the conditions of Claim 17 are satisfied with probability $1 - \text{negl}(n)$. That is, for $N = |S_i|$, we need to show that $N' = |R_i \cap D| = O(N)$ and $\Delta = \Delta_i = O(\sqrt{N})$ with probability $1 - \text{negl}(n)$. Conditioned on Equation (10) holding for all $a, b \in [nm]$, we have that $|pm - |R_i \cap D|| \leq \sqrt{nm} \log n$ and $|(1-p)m - |S_i|| = O(\sqrt{nm} \log n)$ for all $i \in [n]$. Therefore $|R_i \cap D| = \Theta(m)$ and $|S_i| = \Theta(m)$. By Equation (12), $|\Delta_i| = O((nm)^{1/4} \log^{3/2} n)$ for all $i \in [n]$ with probability $1 - \text{negl}(n)$. Therefore, $|\Delta_i| = O(\sqrt{|S_i|})$ with probability $1 - \text{negl}(n)$ if $m = \Omega(n \log^6 n)$. \square

We now show that we can compose this majority code with a PRC for bounded-weight channels to yield a PRC for the deletion channel.

Construction 6 (PRC for deletions). Let $\text{PRC} = (\text{KeyGen}, \text{Encode}, \text{Decode})$ be a PRC with block length n . $\text{PRC}_{\text{del}}[m, \text{PRC}] = (\text{KeyGen}_{\text{del}}, \text{Encode}_{\text{del}}, \text{Decode}_{\text{del}})$ is defined as follows:

- $\text{KeyGen}_{\text{del}}(1^\lambda)$: Output $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$.
- $\text{Encode}_{\text{del}}(\text{sk}, \mathbf{m})$: Output $\text{MajEnc}_m(\text{Encode}(\text{sk}, \mathbf{m}))$.
- $\text{Decode}_{\text{del}}(\text{sk}, z)$: Output $\text{Decode}(\text{sk}, \text{MajDec}_n(z))$.

Theorem 4. For any constants $p \in (0, 1)$ and $q \in (0, 1/2)$, there exists $\varepsilon > 0$ such that the following holds. If PRC is a PRC for $(\frac{1}{2} - \varepsilon)$ -bounded channels with block length n , there exists an $m \in \mathbb{N}$, where $m = \Omega(n \log^6 n)$, such that $\text{PRC}_{\text{del}}[m, \text{PRC}]$ is a PRC for the channel $\text{BSC}_q \circ \text{BDC}_p$.

Proof. Error correction. Let \mathbf{m} be any message, and let $x = \text{Encode}(\text{sk}, \mathbf{m})$ for $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$. By pseudorandomness of PRC, x is indistinguishable from a random string in $\{0, 1\}^n$. By Lemma 15, there exists $\varepsilon \in (0, 1/2)$ such that

$$\Pr_{x \leftarrow \{0, 1\}^n} [\text{wt}(x \oplus \text{MajDec}_n \circ \text{BSC}_q \circ \text{BDC}_p \circ \text{MajEnc}_m(x)) \leq (1/2 - \varepsilon) \cdot n] \geq 1 - \text{negl}(n).$$

Let $\mathcal{E}' = \text{MajDec}_n \circ \text{BSC}_q \circ \text{BDC}_p \circ \text{MajEnc}_m$ be a channel, and observe that the above implies that \mathcal{E}' is $(1/2 - \varepsilon)$ -bounded.

Since PRC is error-correcting against $(1/2 - \varepsilon)$ -bounded channels,

$$\Pr_{\text{sk} \leftarrow \text{KeyGen}(1^\lambda)} [\text{Decode}(\text{sk}, \mathcal{E}'(x)) = s : x \leftarrow \text{Encode}(\text{sk}, \mathbf{m})] \geq 1 - \text{negl}(\lambda).$$

Now, we rewrite this guarantee to be in terms of PRC_{del} . Observe that by definition, for all \mathbf{m} , $\text{Decode}_{\text{del}}(\text{sk}, \text{BSC}_q \circ \text{BDC}_p(\text{Encode}_{\text{del}}(\text{sk}, \mathbf{m})))$ is distributed identically to $\text{Decode}(\text{sk}, \mathcal{E}'(\text{Encode}(\text{sk}, \mathbf{m})))$. Therefore,

$$\Pr_{\text{sk} \leftarrow \text{KeyGen}_{\text{del}}(1^\lambda)} [\text{Decode}_{\text{del}}(\text{sk}, \text{BSC}_q \circ \text{BDC}_p(x)) = \mathbf{m} : x \leftarrow \text{Encode}_{\text{del}}(\text{sk}, s)] \geq 1 - \text{negl}(\lambda).$$

We finally show that unrelated strings c decode to \perp under PRC_{del} . Let $c' = \text{MajDec}_n(c)$. By the analogous property for PRC,

$$\Pr_{\text{sk} \leftarrow \text{KeyGen}(1^\lambda)} [\text{Decode}(\text{sk}, c') = \perp] \geq 1 - \text{negl}(\lambda)$$

as desired.

Pseudorandomness. By pseudorandomness of PRC, its codewords are indistinguishable from random strings from $\{0, 1\}^n$. Thus, let x be drawn uniformly from $\{0, 1\}^n$. The distribution $\text{MajEnc}_m(x)$ is uniform over $\{0, 1\}^{nm}$ because for odd $m \in \mathbb{N}$, $|\{y \in \{0, 1\}^m : \text{Maj}(y) = 0\}| = |\{y \in \{0, 1\}^m : \text{Maj}(y) = 1\}| = 2^{m-1}$. \square

7 Application: watermarking for language models

In this section we show that PRCs can be used to build quality-preserving (undetectable) and robust watermarks. For an overview of our approach, see either the introduction or Section 2.5.

7.1 Watermarking preliminaries

We follow [CGZ23] in our definition of a *language model*. We will often refer to language models simply as *models*.

Definition 4. A language model Model over token set \mathcal{T} is a deterministic algorithm that takes as input a prompt PROMPT and tokens previously output by the model $\mathbf{t} = (t_1, \dots, t_{i-1})$, and outputs a probability distribution $\mathbf{p}_i = \text{Model}(\text{PROMPT}, \mathbf{t})$ over \mathcal{T} .

We write $\mathbf{p}_i(t)$ to denote the probability that the model outputs a token $t \in \mathcal{T}$ as specified by the distribution \mathbf{p}_i . When \mathbf{p}_i is supported on $\{0, 1\}$, we write $\hat{\mathbf{p}}_i := \mathbb{E}[\mathbf{p}_i] \in [0, 1]$.

A language model Model is used to generate text as a response to a prompt by iteratively sampling from the returned distribution until a special terminating token $\text{done} \in \mathcal{T}$ is drawn. We assume an upper bound L^* on the length of the token sequence comprising any response.

Definition 5. A language model’s *response* to PROMPT is a random variable $\overline{\text{Model}}(\text{PROMPT}) \in \mathcal{T}^*$ that is defined algorithmically as follows. We begin with an empty list of tokens $\mathbf{t} = ()$. As long as the last token in \mathbf{t} is not done , we draw a token t_i from the distribution $\text{Model}(\text{PROMPT}, \mathbf{t})$ and append it to \mathbf{t} . Finally, we set $\overline{\text{Model}}(\text{PROMPT}) = \mathbf{t}$.

For a probability distribution D over elements of a finite set X , we define the Shannon *entropy* of D as

$$H(D) = \mathbb{E}_{x \sim D} [-\log D(x)],$$

where $D(x)$ is the probability of x in the distribution D . The *empirical entropy* (also known as Shannon information or surprisal) of x in D is simply $-\log D(x)$. The expected empirical entropy of $x \sim D$ is exactly $H(D)$. Intuitively, the empirical entropy of x (with respect to D) is the number of random bits that were required to draw x out of the distribution D .

Definition 6 (Empirical entropy). For a language model Model , a prompt PROMPT , and a possible response $\mathbf{t} \in \mathcal{T}^*$, we define the *empirical entropy* of Model responding with \mathbf{t} to PROMPT as

$$H_e(\text{Model}, \text{PROMPT}, \mathbf{t}) := -\log \Pr \left[\overline{\text{Model}}(\text{PROMPT}) = \mathbf{t} \right].$$

We next generalize the definition of empirical entropy from whole outputs to *substrings* of a model’s output. This will quantify how much entropy was involved in the generation of a particular contiguous substring of the output.

Definition 7. For a language model Model , a prompt PROMPT , a possible response $\mathbf{t} \in \mathcal{T}^*$, and indices $i, j \in [\text{len } \mathbf{t}]$ with $i \leq j$ we define the *empirical entropy on substring* $[i, j]$ of $\overline{\text{Model}}$ responding with \mathbf{t} to PROMPT as

$$H_e^{[i,j]}(\text{Model}, \text{PROMPT}, \mathbf{t}) := -\log \Pr \left[\overline{\text{Model}}(\text{PROMPT})[i:j] = \mathbf{t}[i:j] \right. \\ \left. \mid \overline{\text{Model}}(\text{PROMPT})[1:(i-1)] = \mathbf{t}[1:(i-1)] \right].$$

For convenience, in settings where the model and prompt are clear we simply write $H_e^{[i,j]}(\mathbf{t})$ to mean $H_e^{[i,j]}(\text{Model}, \text{PROMPT}, \mathbf{t})$. We sometimes write $H_e^i := H_e^{[i,i]}$ to denote the empirical entropy of a single token i .

We formally define a watermarking scheme as follows.

Definition 8 (Watermarking scheme). A *watermarking scheme* for a model Model over \mathcal{T} is a tuple of polynomial-time algorithms $\mathcal{W} = (\text{Setup}, \text{Wat}, \text{Detect})$ where:

- $\text{Setup}(1^\lambda) \rightarrow \text{sk}$ outputs a secret key, with respect to a security parameter λ .
- $\text{Wat}_{\text{sk}}(\text{PROMPT})$ is a randomized algorithm that takes as input a prompt PROMPT and generates a response in \mathcal{T}^* .
- $\text{Detect}_{\text{sk}}(\mathbf{t}) \rightarrow \{\text{true}, \text{false}\}$ is an algorithm that takes as input a sequence $\mathbf{t} \in \mathcal{T}^*$ outputs true or false.

Ideally, $\text{Detect}_{\text{sk}}(\mathbf{t})$ should output **true** if \mathbf{t} is generated by $\text{Wat}_{\text{sk}}(\text{PROMPT})$, and should output **false** if \mathbf{t} is generated independently of sk . The former property is called *completeness* and the latter *soundness*.

Definition 9 (Soundness). A watermarking scheme \mathcal{W} is *sound* if for every security parameter λ and token sequence $\mathbf{t} \in \mathcal{T}^*$ of length $\text{poly}(\lambda)$,

$$\Pr_{\text{sk} \leftarrow \text{Setup}(1^\lambda)} [\text{Detect}_{\text{sk}}(\mathbf{t}) = \text{true}] \leq \text{negl}(\lambda).$$

Definition 10 (Completeness). A watermarking scheme \mathcal{W} is *b(L)-complete* if for every security parameter λ and prompt PROMPT of length $\text{poly}(\lambda)$,

$$\Pr_{\substack{\text{sk} \leftarrow \text{Setup}(1^\lambda) \\ \mathbf{t} \leftarrow \text{Wat}_{\text{sk}}(\text{PROMPT})}} [\text{Detect}_{\text{sk}}(\mathbf{t}) = \text{false} \text{ and } H_e(\text{Model}, \text{PROMPT}, \mathbf{t}) \geq b(\text{len } \mathbf{t})] \leq \text{negl}(\lambda).$$

Definition 11 (Substring completeness). A watermarking scheme \mathcal{W} is *b(L)-substring-complete* if for every prompt PROMPT and security parameter λ ,

$$\Pr_{\substack{\text{sk} \leftarrow \text{Setup}(1^\lambda) \\ \mathbf{t} \leftarrow \text{Wat}_{\text{sk}}(\text{PROMPT})}} \left[\exists i, L \in [\text{len } \mathbf{t}] \text{ such that } \text{Detect}_{\text{sk}}(\mathbf{t}[i : i + L]) = \text{false} \right. \\ \left. \text{and } H_e^{[i:i+L]}(\text{Model}, \text{PROMPT}, \mathbf{t}) \geq b(L) \right] \leq \text{negl}(\lambda).$$

If the watermarked model is indistinguishable from the un-watermarked model (and therefore perfectly quality-preserving), we say that the watermarking scheme is *undetectable*.

Definition 12 (Undetectability). A watermarking scheme $\mathcal{W} = (\text{Setup}, \text{Wat}, \text{Detect})$ is *undetectable* if for every security parameter λ and all polynomial-time distinguishers D ,

$$\left| \Pr[D^{\text{Model}, \overline{\text{Model}}}(1^\lambda) \rightarrow 1] - \Pr_{\text{sk} \leftarrow \text{Setup}(1^\lambda)} [D^{\text{Model}, \text{Wat}_{\text{sk}}}(1^\lambda) \rightarrow 1] \right| \leq \text{negl}(\lambda),$$

where the notation $D^{\mathcal{O}_1, \mathcal{O}_2}$ means that D is allowed to adaptively query both \mathcal{O}_1 and \mathcal{O}_2 with arbitrary prompts.

We also define robustness for a watermarking scheme. For robustness, the detector should be able to identify the watermark even if the text has undergone corruption by some channel \mathcal{E} . Substring robustness says that the watermark should be robust to both cropping as well as \mathcal{E} .

Definition 13 (Substring robustness). A watermarking scheme $\mathcal{W} = (\text{Setup}, \text{Wat}, \text{Detect})$ is *b(L)-substring-robust* against a channel \mathcal{E} if for every security parameter λ and prompt PROMPT of length $\text{poly}(\lambda)$,

$$\Pr_{\substack{\text{sk} \leftarrow \text{Setup}(1^\lambda) \\ \mathbf{t} \leftarrow \text{Wat}_{\text{sk}}(\text{PROMPT})}} \left[\exists i, L \in [\text{len } \mathbf{t}] \text{ such that } \text{Detect}_{\text{sk}}(\mathcal{E}(\mathbf{t}[i : i + L])) = \text{false} \right. \\ \left. \text{and } H_e^{[i:i+L]}(\text{Model}, \text{PROMPT}, \mathbf{t}) \geq b(L) \right] \leq \text{negl}(\lambda).$$

Reducing to a binary alphabet. Recall that a language model operates over an arbitrary token alphabet \mathcal{T} . In constructing our PRC-based watermarks, it will be convenient to take \mathcal{T} to be $\{0, 1\}$, since our pseudorandom LDPC codes have binary codewords. Prior work [CGZ23] suggests a black-box transformation from any model with an arbitrary-token alphabet to a model with a binary-token alphabet, by reinterpreting the probability vectors output by the model. Here, we describe that transformation in more detail.

Let Model be any model with token alphabet \mathcal{T} . Model , when queried with a prompt and token sequence, outputs a probability vector $\mathbf{p} \in [0, 1]^{|\mathcal{T}|}$. We construct a model Model' that has token alphabet $\mathcal{T}' = \{0, 1\}$; that is, Model' outputs probability vectors $\mathbf{p}' \in [0, 1]^2$. This construction will make only black-box use of Model and will ensure that there is some decoding function f converting binary outputs of $\overline{\text{Model}'}$ into sequences of tokens from \mathcal{T} , such that the distributions of $\overline{\text{Model}}$ and $f(\overline{\text{Model}'})$ are identical.

Let $\text{Enc} : \mathcal{T} \rightarrow \{0, 1\}^*$ be any prefix-free encoding function, and let $\text{Dec} : \{0, 1\}^* \rightarrow \mathcal{T}^* \times \{0, 1\}^*$ be the corresponding decoding function such that for any $t \in \mathcal{T}$, $\text{Dec}(\text{Enc}(t)) = (t, \perp)$. Furthermore, if Dec is given as input an encoded token $\text{Enc}(t)$ concatenated with some binary string $s \in \{0, 1\}^*$ such that no prefix of s is a valid codeword, Dec outputs the token and the string; that is, $\text{Dec}(\text{Enc}(t)||s) = (t, s)$. We will construct Model' such that the distribution of $\text{Dec}(\overline{\text{Model}'})$ is identical to that of $\overline{\text{Model}}$. That is, $\overline{\text{Model}'}$ will output a binary encoding of a token sequence in \mathcal{T} . So Model' will compute its distribution \mathbf{p}' over the next binary token by querying Model for its distribution \mathbf{p} over \mathcal{T} , and computing (according to \mathbf{p}) the distribution over the next bit of the binary encoding of the next token.

Suppose that Model' is given as input a prompt PROMPT and a binary token sequence $(t'_1, \dots, t'_\ell) \in \{0, 1\}^\ell$. Model' must now output a probability distribution over the next (binary) token, which must match Model under the appropriate transformation. Let $((t_1, \dots, t_i), s) \leftarrow \text{Dec}(t'_1, \dots, t'_\ell)$, and let $\mathbf{p} = \text{Model}(\text{PROMPT}, (t_1, \dots, t_i))$, where $\mathbf{p} \in [0, 1]^{|\mathcal{T}|}$. Model' computes \mathbf{p}' as

$$\begin{aligned} \mathbf{p}'(0) &= \sum_{\substack{t \in \mathcal{T} \\ \text{Enc}(t)[1:\text{len } s+1] = s||0}} \mathbf{p}(t) \\ \mathbf{p}'(1) &= 1 - \mathbf{p}'(0) \end{aligned}$$

where $\text{Enc}(t)[1 : \text{len } s + 1]$ denotes the first $(\text{len } s + 1)$ bits of $\text{Enc}(t)$. Observe that $\mathbf{p}'(0)$ is exactly the probability that the next bit of $\text{Enc}(t)$ is 0, where $t \leftarrow \mathbf{p}$ from Model , conditioned on the event that the previous bits of $\text{Enc}(t)$ match s . Therefore, the probability that t is drawn as the next token from Model is exactly the same as the probability that $\text{Enc}(t)$ is drawn as the next sequence of binary tokens from Model' .

This transformation exactly preserves the empirical entropy of the response, since the distribution of $\text{Dec}(\overline{\text{Model}'})$ is identical to that of $\overline{\text{Model}}$. However, it does reduce the rate of entropy per token, since responses from $\overline{\text{Model}'}$ are longer. This reduction will depend on the length of codewords, and using an efficient encoding such as a Huffman encoding can help mitigate this rate reduction.

For the remainder of this paper, we assume that the token alphabet of the underlying model is binary. In practice, we can embed the watermark in a model Model with an arbitrary token alphabet \mathcal{T} by using this transformation to compute Model' , generating a watermarked binary response from Model' , and transforming this response back to tokens in \mathcal{T} using the decoding function D . This transformation preserves the undetectability of our watermark. Recall that by undetectability, the distribution of watermarked (binary) responses from Model' is indistinguishable from the original (binary) distribution of Model' . Let Wat denote the distribution of watermarked responses from Model' . Since our transformation guarantees that $\text{Dec}(\overline{\text{Model}'}) \equiv \overline{\text{Model}}$, we have that $\text{Dec}(\text{Wat}) \approx \overline{\text{Model}}$, meaning that our watermarked responses after this transformation are indistinguishable from those of the original model.

Since Model' outputs distributions \mathbf{p} over $\{0, 1\}$, we often specify these distributions by their expectations $\hat{\mathbf{p}} := \mathbb{E}[\mathbf{p}] \in [0, 1]$.

Remark. In our robustness theorems, we assume that the bit-errors are random. Using the above transformation, this may not be realistic because bit-errors within the representation of a given token may be correlated. We note that an alternative transformation that avoids this issue simply uses the first bit of the above representation.

7.2 Robust watermarking from PRCs

Watermarking schemes for language models typically embed the watermark by sampling each token with some bias. In this section, we show that using a PRC to determine this bias yields a watermarking scheme that inherits robustness from the PRC. In more detail, our watermark generator first samples a PRC codeword $x \in \{0, 1\}^n$. It then samples each token $t_i \in \{0, 1\}$ to be biased toward x_i , yielding a response that is a noisy codeword. By pseudorandomness of PRCs, this biased sampling does not noticeably change the distribution of t_i in expectation over x_i . Yet the detector, which knows the PRC secret key, can check if the given text is a noisy codeword to detect the watermark. Furthermore, one can use our scheme to embed an arbitrary message m in the response, by choosing x to be $\text{Encode}(\text{pk}, m)$ under a multi-bit PRC. This yields a language model steganography scheme, since steganographic secrecy is implied by undetectability.

Our watermarking scheme can tolerate additional changes to the response depending on the robustness of the PRC used. In particular, if the PRC is robust to deletions, our scheme evades the *emoji attack* which succeeds against all existing undetectable watermarking schemes. Here, the attacker asks the model to answer the prompt and randomly insert an emoji in between words, then deletes the emojis. This attack succeeds because the detectors in existing schemes must be given *contiguous* text from the watermarked model. However, modeling this attack as a random deletion channel, our codes from Section 6.3 give rise to watermarking schemes that are robust to this attack.

Construction 7 (Watermark from PRCs). Let PRC be a PRC of block length n with security parameter λ . $\mathcal{W}[\text{PRC}] = (\text{Setup}, \text{Wat}, \text{Detect})$ is the watermarking scheme whose algorithms are specified in Figure 3.

Remark. In Algorithm 3, we specify the detector to do a brute-force search over $O(Ln)$ index pairs to find the start and end locations of the perturbed codeword, since deletions may change its length. If the perturbed codeword length is known, e.g., when one wants robustness only to substitutions, one can use a more efficient detector that searches only over $O(L)$ indices to find the start location of the codeword.

Recall that we write \hat{p} for the expectation $\mathbb{E}[p]$ of a distribution p over $\{0, 1\}$.

Soundness and undetectability We first show that $\mathcal{W}[\text{PRC}]$ is sound and undetectable, which follows quickly from pseudorandomness and error correction of PRCs.

Lemma 18 (Soundness). $\mathcal{W}[\text{PRC}]$ is sound.

Proof. Let $t \in \mathcal{T}^*$ be any token sequence of length $(\text{len } t) \leq \text{poly}(\lambda)$. $\text{Detect}_{\text{sk}}(t)$ iterates over all $i \in [\text{len } t]$, $j \in [i, \min\{i + n - 1, \text{len } t\}]$, checks if $\text{PRC.Decode}(\text{PRC.sk}, (t_i, \dots, t_j) \oplus a_\ell) = 1$ for any i, j, ℓ , and outputs true if so. Recall that the error correction of PRC ensures that for any $c \in \Sigma^* = \mathcal{T}^*$,

$$\Pr_{\text{PRC.sk} \leftarrow \text{PRC.KeyGen}(1^\lambda)} [\text{Decode}(\text{PRC.sk}, c) \neq \perp] \leq \text{negl}(\lambda).$$

Setting $c_i = (t_i, \dots, t_j) \oplus a_\ell$ for each i, j, ℓ , the probability that Decode returns 1 for any c_i is negligible by a union bound. Consequently, the probability over choice of sk that $\text{Detect}_{\text{sk}}$ outputs true is negligible. \square

Lemma 19 (Undetectability). $\mathcal{W}[\text{PRC}]$ is undetectable.

Proof. We prove that if the output of Encode is replaced with uniform randomness, this watermark is undetectable. It then follows from pseudorandomness of the PRC that the actual watermark is undetectable.

Algorithm 1: Watermark setup procedure $\text{Setup}(1^\lambda)$

Input: A security parameter λ

Result: A watermark key sk

```

1  $\text{PRC.sk} \leftarrow \text{PRC.KeyGen}(1^\lambda)$ ;
2  $a_1, \dots, a_{\lceil \frac{L^*}{n} \rceil} \leftarrow \{0, 1\}^n$ ;
3 return  $\text{sk} = (\text{PRC.sk}, (a_1, \dots, a_{\lceil \frac{L^*}{n} \rceil}))$ ;

```

Algorithm 2: Watermarked text generator Wat_{sk}

Input: A prompt (PROMPT) and a key sk

Result: Watermarked text $\mathbf{t}_1, \dots, \mathbf{t}_L$

```

1  $(\text{PRC.sk}, a) \leftarrow \text{sk}$ ;
2  $x \leftarrow \text{PRC.Encode}(\text{PRC.sk}) \oplus a_1$ ;
3  $i := 1$ ;
4  $j := 1$ ;
5 while  $\text{done} \notin \{\mathbf{t}_1, \dots, \mathbf{t}_{i-1}\}$  do
6    $\mathbf{p}_i := \text{Model}(\text{PROMPT}, \mathbf{t}_1, \dots, \mathbf{t}_{i-1})$ ;
7    $\mathbf{t}_i \leftarrow \text{Ber}(\hat{\mathbf{p}}_i - (-1)^{x_j} \cdot \min\{\hat{\mathbf{p}}_i, 1 - \hat{\mathbf{p}}_i\})$ ;
8    $i \leftarrow i + 1$ ;
9    $j \leftarrow j + 1$ ;
10  if  $j > n$  then
11    // Re-sample from the PRC
12     $x \leftarrow \text{PRC.Encode}(\text{PRC.sk}) \oplus a_{\lceil \frac{i}{n} \rceil}$ ;
13     $j \leftarrow 1$ ;
13 return  $\mathbf{t}_1, \dots, \mathbf{t}_L$ ;

```

Algorithm 3: Watermark detector $\text{Detect}_{\text{sk}}$

Input: Text $\mathbf{t}_1, \dots, \mathbf{t}_L$ and watermark key sk

Result: true or false

```

1  $(\text{PRC.sk}, a) \leftarrow \text{sk}$ ;
2 for  $i \in [L], j \in [i, \min\{i + n - 1, L\}], \ell \in [\lceil \frac{L^*}{n} \rceil]$  do
3   if  $\text{PRC.Decode}(\text{PRC.sk}, (\mathbf{t}_i, \dots, \mathbf{t}_j) \oplus a_\ell) \neq \perp$  then
4     return true;
5 return false;

```

Figure 3: The algorithms (Setup, Wat, Detect) of $\mathcal{W}[\text{PRC}]$.

We begin by showing that if $x_j \leftarrow \text{Ber}(\frac{1}{2})$, then $\Pr[t_i = 1] = \hat{\rho}_i$, where $t_i \leftarrow \text{Ber}(\hat{\rho}_i - (-1)^{x_j} \cdot \min\{\hat{\rho}_i, 1 - \hat{\rho}_i\})$. Suppose first that $\min\{\hat{\rho}_i, 1 - \hat{\rho}_i\} = \hat{\rho}_i$. Then

$$\hat{\rho}_i - (-1)^{x_j} \cdot \min\{\hat{\rho}_i, 1 - \hat{\rho}_i\} = \begin{cases} 0 & \text{if } x_j = 0 \\ 2\hat{\rho}_i & \text{if } x_j = 1 \end{cases}$$

so $\Pr_{x_j}[t_i = 1] = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 2\hat{\rho}_i = \hat{\rho}_i$. Now, suppose that $\min\{\hat{\rho}_i, 1 - \hat{\rho}_i\} = 1 - \hat{\rho}_i$. Then

$$\hat{\rho}_i - (-1)^{x_j} \cdot \min\{\hat{\rho}_i, 1 - \hat{\rho}_i\} = \begin{cases} 2\hat{\rho}_i - 1 & \text{if } x_j = 0 \\ 1 & \text{if } x_j = 1 \end{cases}$$

so $\Pr_{x_j}[t_i = 1] = \frac{1}{2} \cdot (2\hat{\rho}_i - 1) + \frac{1}{2} \cdot 1 = \hat{\rho}_i$. Therefore, if the x_j 's are independent $\text{Ber}(\frac{1}{2})$, the distribution of every t_i under the watermark is identical to its distribution under the original model.

Suppose for the sake of contradiction that \mathcal{W} is not undetectable; that is, there is a polynomial-time distinguisher D such that for some constant $c > 0$,

$$\left| \Pr[D^{\text{Model}, \overline{\text{Model}}}(1^\lambda) \rightarrow 1] - \Pr_{\text{sk} \leftarrow \text{Setup}(1^\lambda)}[D^{\text{Model}, \text{Wat}_{\text{sk}}}(1^\lambda) \rightarrow 1] \right| \geq \frac{1}{\lambda^c}.$$

We'll construct an adversary \mathcal{A} that uses D to break pseudorandomness of PRC. Recall that in the PRC pseudorandomness experiment, \mathcal{A} has access to \mathcal{O} , which is either an oracle for $\text{Encode}(\text{PRC.sk}, \cdot)$ or the uniform distribution \mathcal{U} , and its goal is to distinguish between these two cases. \mathcal{A} interacts with D and simulates D 's oracle \mathcal{O}_2 , which is either $\overline{\text{Model}}$ or Wat_{sk} . When D queries a prompt PROMPT to \mathcal{O}_2 , \mathcal{A} returns a response according to Algorithm 2, *with one key modification*: It draws each x by querying its oracle \mathcal{O} . If \mathcal{O} is $\text{Encode}(\text{PRC.sk})$, the responses returned by \mathcal{A} are exactly those of Wat_{sk} . By our earlier observation that if the x_j 's are uniform, the t_i 's distributions are unchanged, if \mathcal{O} is \mathcal{U} the responses returned by \mathcal{A} are exactly those of $\overline{\text{Model}}$. Therefore, if we set \mathcal{A} to output 1 if and only if D outputs 1, we have that

$$\left| \Pr_{(\text{PRC.sk}) \leftarrow \text{KeyGen}(1^\lambda)}[\mathcal{A}^{\text{Encode}(\text{PRC.sk}, \cdot)}(1^\lambda) = 1] - \Pr_{\mathcal{U}}[\mathcal{A}^{\mathcal{U}}(1^\lambda) = 1] \right| \geq \frac{1}{\lambda^c},$$

contradicting pseudorandomness of the PRC. \square

Robustness of our scheme In the watermarked text generator (Algorithm 2), we embed a PRC output r into the model's response to a prompt PROMPT by sampling t_i from $\text{Ber}(\hat{\rho}_i - (-1)^{x_j} \cdot \min\{\hat{\rho}_i, 1 - \hat{\rho}_i\})$. Observe that if $\hat{\rho}_i = \frac{1}{2}$, this is equal to $\text{Ber}(x_j)$, so t_i will *always* equal x_j . If $\hat{\rho}_i \neq \frac{1}{2}$, then t_i is sampled with bias toward x_j , but it sometimes will not equal x_j . One can think of this embedding process as taking as input x , applying some random error, and returning the noisy response x as output.

We model this process as an *embedding channel* $\mathcal{E}_{\text{Emb}} : \{0, 1\}^n \rightarrow \{0, 1\}^n$, where for $x \leftarrow \{0, 1\}^n$, $\mathcal{E}_{\text{Emb}}(x) \sim \text{Model}(\text{PROMPT})[i : i + n - 1]$ (i.e., the n bits of the watermarked model's response to PROMPT where x is embedded). For ease of notation, we write $\mathbf{t} \leftarrow \mathcal{E}_{\text{Emb}}(x)$, leaving PROMPT and the index i implicit. We will show that \mathcal{E}_{Emb} introduces a bounded number of errors when the text has non-zero entropy.

The relevant notion of entropy is the empirical entropy (Definition 6). We will also make use of the *truncated empirical entropy* \bar{H}_e , which is defined as follows. For a model output \mathbf{t} (given some prompt that we leave implicit), $\bar{H}_e(\mathbf{t}) = \sum_{i \in \text{len } \mathbf{t}} \min\{1, H_e^i(\mathbf{t})\}$. As with empirical entropy, we define truncated empirical entropy for partial responses: $\bar{H}_e^{[i, j]}(\mathbf{t}) = \sum_{\ell=i}^j \min\{1, H_e^\ell(\mathbf{t})\}$, and $\bar{H}_e^i(\mathbf{t}) = \min\{1, H_e^i(\mathbf{t})\}$.

Lemma 20. *For any $c > 0$ and any indices $a, b \in \mathbb{N}$ where $a \leq b$,*

$$\Pr_{\mathbf{t} \leftarrow \text{Model}} \left[H_e^{[a, b]}(\mathbf{t}) > c \cdot (b - a) \text{ and } \bar{H}_e^{[a, b]}(\mathbf{t}) \leq \frac{c^2}{2}(b - a) \right] \leq \text{negl}(b - a).$$

Proof. We will show that with probability $1 - \text{negl}(b - a)$ over $\mathbf{t} \leftarrow \text{Model}$,

$$\sum_{i=a}^b \left(\frac{H_e^i(\mathbf{t})}{\bar{H}_e^i(\mathbf{t})} \right)^2 \leq 2(b - a).$$

By the Cauchy-Schwarz inequality, we have

$$\sum_{i=a}^b H_e^i(\mathbf{t}) \leq \sqrt{\sum_{i=a}^b \left(\frac{H_e^i(\mathbf{t})}{\bar{H}_e^i(\mathbf{t})} \right)^2} \cdot \sqrt{\sum_{i=a}^b \bar{H}_e^i(\mathbf{t})^2}.$$

Combining these two inequalities, it will follow that with probability $1 - \text{negl}(b - a)$ if $\sum_{i=a}^b H_e^i(\mathbf{t}) \geq c \cdot (b - a)$ then

$$\begin{aligned} \sum_{i=a}^b \left(\bar{H}_e^i(\mathbf{t}) \right)^2 &\geq \left(\sum_{i=a}^b H_e^i(\mathbf{t}) \right)^2 \left(\sum_{i=a}^b \left(\frac{H_e^i(\mathbf{t})}{\bar{H}_e^i(\mathbf{t})} \right)^2 \right)^{-1} \\ &\geq \frac{c^2 \cdot (b - a)^2}{2(b - a)} \\ &= \frac{c^2}{2} (b - a). \end{aligned}$$

And since $\bar{H}_e^i(\mathbf{t}) \leq 1$ by definition, $\bar{H}_e^i(\mathbf{t}) \geq \left(\bar{H}_e^i(\mathbf{t}) \right)^2$ for every i , and $\sum_{i=a}^b \bar{H}_e^i(\mathbf{t}) \geq \frac{c^2}{2} (b - a)$ as desired.

It remains to show that with overwhelming probability, $\sum_{i=a}^b \left(\frac{H_e^i(\mathbf{t})}{\bar{H}_e^i(\mathbf{t})} \right)^2 \leq 2(b - a)$. For $i \in [a, b]$, let $X_i := (H_e^i(\mathbf{t})/\bar{H}_e^i(\mathbf{t}))^2$ and $\mathbf{q}_i := \min\{\mathbf{p}_i(0), \mathbf{p}_i(1)\}$. If $\mathbf{q}_i \geq 1/e$, $X_i = 1$; if $\mathbf{q}_i < 1/e$ then

$$X_i = \begin{cases} 1 & \text{w.p. } 1 - \mathbf{q}_i \\ \ln^2 \mathbf{q}_i & \text{w.p. } \mathbf{q}_i \end{cases}$$

and $\mathbb{E}[X_i \mid \mathbf{t}_a, \dots, \mathbf{t}_{i-1}] = 1 - \mathbf{q}_i + \mathbf{q}_i \ln^2 \mathbf{q}_i \leq 3/2$. For any fixed $\mathbf{t}_a, \dots, \mathbf{t}_{i-1}$,

$$\Pr_{\mathbf{t}_i \leftarrow \mathbf{p}_i} [X_i \geq \ln^4(b - a)] = \mathbf{q}_i \cdot \mathbb{1}[\mathbf{q}_i \leq e^{-\ln^2(b-a)}] \leq e^{-\ln^2(b-a)}.$$

By a union bound, $\Pr[\exists i \in [a, b] \text{ s.t. } X_i \geq \ln^4(b - a)] = \text{negl}(b - a)$.

Let $Z_{a-1} = 0$ and for $i \in [a, b]$ let $Z_i = Z_{i-1} + X_i - (1 - \mathbf{q}_i + \mathbf{q}_i \ln^2 \mathbf{q}_i)$. Observe that Z_a, \dots, Z_b is a martingale with respect to $\mathbf{t}_a, \dots, \mathbf{t}_b$. Since the differences $|Z_i - Z_{i-1}|$ are bounded by $\ln^4(b - a)$ with probability $1 - \text{negl}(b - a)$, Azuma's inequality (Lemma 1) implies that

$$\Pr[Z_b > (b - a)/2] \leq \exp\left(\frac{-(b - a)}{8 \ln^8(b - a)}\right) + \text{negl}(b - a) \leq \text{negl}(b - a).$$

Since $Z_b = \sum_{i=a}^b [X_i - (1 - \mathbf{q}_i + \mathbf{q}_i \ln^2 \mathbf{q}_i)] \geq \sum_{i=a}^b X_i - \frac{3}{2}(b - a)$, it follows that

$$\Pr\left[\sum_{i=a}^b X_i > 2(b - a)\right] \leq \text{negl}(b - a). \quad \square$$

We show in the following lemma that the embedding channel introduces bounded-weight errors into any contiguous substring of a response with sufficient empirical entropy.

Lemma 21 (Entropy and \mathcal{E}_{Emb}). *Let $\kappa > 0$ be any constant, and let $a, b \in \mathbb{N}$ be indices such that $b = a + \lceil \kappa n \rceil - 1$. For the embedding channel \mathcal{E}_{Emb} of our watermarked model,*

$$\Pr_{x \leftarrow \{0,1\}^{\lceil \kappa n \rceil}} \left[\begin{array}{l} \text{wt}(\mathbf{t} \oplus x) > \left(\frac{1}{2} - \frac{c^2}{16}\right) \cdot \text{len } \mathbf{t} \text{ and } H_e^{[a,b]}(\mathbf{t}') > c \cdot \text{len } \mathbf{t} \\ \mathbf{t}' \leftarrow \mathcal{E}_{\text{Emb}}(x) \\ \mathbf{t} \leftarrow \mathbf{t}'[a : b] \end{array} \right] \leq \text{negl}(n).$$

Proof. For ease of notation, in this proof we use $\mathbf{p}_i(\hat{\mathbf{t}}_i)$ to denote the probability that the model outputs a bit $\hat{\mathbf{t}}_i \in \{0,1\}$ as the next $(a+i)^{\text{th}}$ token given that it has output tokens $\mathbf{t}'[: a+i-1]$ so far. We also let $H_e(\mathbf{t})$ denote $H_e^{[a,b]}(\mathbf{t}')$ and $\bar{H}_e(\mathbf{t})$ denote $\bar{H}_e^{[a,b]}(\mathbf{t}')$.

For each $i \in \text{len } \mathbf{t}$, recall that the watermark samples $\mathbf{t}_i \leftarrow \text{Ber}(\hat{\rho}_i - (-1)^{x_i} \cdot \min\{\hat{\rho}_i, 1 - \hat{\rho}_i\})$, and for $\hat{\mathbf{t}}_i \in \{0,1\}$, $\Pr_{x_i \leftarrow \{0,1\}}[\mathbf{t}_i = \hat{\mathbf{t}}_i] = \mathbf{p}_i(\hat{\mathbf{t}}_i)$.

For fixed $\hat{\mathbf{t}}_i \in \{0,1\}$, we have

$$\begin{aligned} \frac{\Pr[\mathbf{t}_i = \hat{\mathbf{t}}_i \mid x_i = \hat{\mathbf{t}}_i]}{2} &= \Pr[x_i = \hat{\mathbf{t}}_i \text{ and } \mathbf{t}_i = \hat{\mathbf{t}}_i] \\ &= \Pr[\mathbf{t}_i = \hat{\mathbf{t}}_i] \cdot \Pr[x_i = \hat{\mathbf{t}}_i \mid \mathbf{t}_i = \hat{\mathbf{t}}_i] \\ &= \mathbf{p}_i(\hat{\mathbf{t}}_i) \cdot \Pr[x_i = \hat{\mathbf{t}}_i \mid \mathbf{t}_i = \hat{\mathbf{t}}_i]. \end{aligned}$$

Because \mathbf{t}_i is distributed as $\text{Ber}(\hat{\rho}_i - (-1)^{x_i} \cdot \min\{\hat{\rho}_i, 1 - \hat{\rho}_i\}) = \text{Ber}(\hat{\rho}_i - (-1)^{x_i} \cdot \min\{\hat{\rho}_i, 1 - \hat{\rho}_i\})$, we also have

$$\Pr[\mathbf{t}_i = \hat{\mathbf{t}}_i \mid x_i = \hat{\mathbf{t}}_i] = \mathbf{p}_i(\hat{\mathbf{t}}_i) + \min\{\mathbf{p}_i(\hat{\mathbf{t}}_i), 1 - \mathbf{p}_i(\hat{\mathbf{t}}_i)\}.$$

Therefore,

$$\begin{aligned} \Pr[x_i = \hat{\mathbf{t}}_i \mid \mathbf{t}_i = \hat{\mathbf{t}}_i] &= \frac{1}{2} + \frac{1}{2} \min \left\{ 1, \frac{1}{\mathbf{p}_i(\hat{\mathbf{t}}_i)} - 1 \right\} \\ &\geq \frac{1}{2} + \frac{1}{2} \min \{1, -\ln \mathbf{p}_i(\hat{\mathbf{t}}_i)\} \end{aligned}$$

where for the inequality we have used the fact that $\ln z \geq 1 - 1/z$ for $z > 0$.

Let $Y_1, \dots, Y_{\text{len } \mathbf{t}}$ be Bernoulli random variables where $Y_i = 1$ if and only if $x_i = \hat{\mathbf{t}}_i$. The expected number of correct bits given that $\mathbf{t} = \hat{\mathbf{t}}$ is

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^{\text{len } \mathbf{t}} Y_i \mid \mathbf{t} = \hat{\mathbf{t}} \right] &= \sum_{i=1}^{\text{len } \mathbf{t}} \Pr[Y_i = 1 \mid \mathbf{t}_i = \hat{\mathbf{t}}_i] \\ &\geq \frac{\text{len } \mathbf{t}}{2} + \frac{1}{2} \sum_{i=1}^{\text{len } \mathbf{t}} \min\{1, -\ln \mathbf{p}_i(\hat{\mathbf{t}}_i)\} \\ &\geq \frac{\text{len } \mathbf{t}}{2} + \frac{\ln 2}{2} \sum_{i=1}^{\text{len } \mathbf{t}} \min\{1, -\log \mathbf{p}_i(\hat{\mathbf{t}}_i)\} \\ &\geq \frac{\text{len } \mathbf{t}}{2} + \frac{1}{4} \bar{H}_e(\mathbf{t}). \end{aligned}$$

By Lemma 20, if $H_e(\mathbf{t}) > c \cdot \text{len } \mathbf{t}$, this is at least $\left(\frac{1}{2} + \frac{c^2}{8}\right) \text{len } \mathbf{t}$. Since the events $x_i = \hat{\mathbf{t}}_i$ are independent for fixed $\hat{\mathbf{t}}_i$, we can apply a Chernoff bound to see that $\text{wt}(x \oplus \mathbf{t}) \leq \left(\frac{1}{2} - \frac{c^2}{16}\right) \cdot \text{len } \mathbf{t}$ with probability $1 - \text{negl}(n)$. \square

In the following proofs of substring-completeness and substring-robustness, it will be useful to consider the model's response as consisting of *blocks*. A block is a contiguous substring of the response where a PRC

codeword was embedded. More formally, recall that the generator Wat_{sk} first chooses a PRC codeword of length n that corresponds to the first n tokens of the response, $(\mathbf{t}_1, \dots, \mathbf{t}_n)$, and applies a one-time pad a_1 . It then chooses a new PRC codeword for the next n tokens of the response and applies a fresh one-time pad, continuing to do so until it terminates, having output the final response $\mathbf{t}_1, \dots, \mathbf{t}_{cn+i}$ for some $c, i \in \mathbb{Z}_{\geq 0}$. This response consists of $(c-1)$ full blocks $(\mathbf{t}_1, \dots, \mathbf{t}_n), \dots, (\mathbf{t}_{(c-1)n+1}, \dots, \mathbf{t}_{cn})$. It also consists of a possibly incomplete block, $(\mathbf{t}_{cn+1}, \dots, \mathbf{t}_{cn+i})$. The substring $\mathbf{t}_{n-1}, \dots, \mathbf{t}_{2n+2}$ of this response consists of two incomplete blocks: $(\mathbf{t}_{n-1}, \mathbf{t}_n)$ and $(\mathbf{t}_{2n+1}, \mathbf{t}_{2n+2})$, and one full block: $(\mathbf{t}_{n+1}, \dots, \mathbf{t}_{2n})$.

Lemma 22 (Substring completeness). *Let $\varepsilon > 0$ be any constant. If PRC is a zero-bit PRC with block length n and robustness to every $(1/2 - \varepsilon)$ -bounded channel, then $\mathcal{W}[\text{PRC}]$ is $(4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n)$ -substring-complete.*

Proof. Let $\mathbf{t} \leftarrow \text{Wat}_{\text{sk}}(\text{PROMPT})$ be a watermarked response for some arbitrary prompt PROMPT, and let \mathbf{t}' be a length- L contiguous substring of \mathbf{t} , where $H_\varepsilon(\mathbf{t}') \geq 4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n$.

The start and end of \mathbf{t}' may be part of incomplete blocks, with some number of full blocks in between. Recall that the truncated empirical entropy of any single-bit distribution is at most 1, so each of these possibly-incomplete blocks contains at most n truncated empirical entropy. Consider padding these possibly-incomplete blocks, appending deterministic bits so they are each of length n but have the same truncated empirical entropy. It now follows from Lemma 20 that with overwhelming probability in the length of these blocks (which is now n), they each have at most $n\sqrt{2}$ empirical entropy. The full blocks therefore contain at least $4\sqrt{\varepsilon} \cdot L$ empirical entropy.

By an averaging argument, \mathbf{t}' must contain some full block $\tilde{\mathbf{t}} \subseteq \mathbf{t}'$ with at least $4\sqrt{\varepsilon} \cdot n$ empirical entropy. By Lemma 21, the embedding channel applied to this block is $(1/2 - \varepsilon)$ -bounded. Furthermore, because of the one-time pads, these channels' errors do not depend on other codewords. Since the PRC used is robust against such bounded channels, the detector will see that $\text{PRC.Decode}(\text{PRC.sk}, \tilde{\mathbf{t}} \oplus \alpha_\ell) \neq \perp$ for some ℓ and output true. \square

Lemma 23 (Substring-robustness against substitutions). *Let $\varepsilon, \delta > 0$ be any constants. If PRC is a zero-bit PRC with block length n and robustness to every $(1/2 - \varepsilon \cdot \delta)$ -bounded channel, then $\mathcal{W}[\text{PRC}]$ is $(4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n)$ -substring-robust against $\text{BSC}_{1/2-\delta}$.*

Proof. As in the proof of Lemma 22, the decoder receives the output of the composition of the error channel and embedding channel on the sufficiently high-entropy block. The embedding channel is again $(1/2 - \varepsilon)$ -bounded on this block by Lemma 21. Therefore, the composition of the error channel $\text{BSC}_{1/2-\delta}$ with the embedding channel has expected error rate

$$(1/2 - \delta) \cdot (1/2 + \varepsilon) + (1/2 - \varepsilon) \cdot (1/2 + \delta) = \frac{1}{2} - 2\varepsilon \cdot \delta$$

and is in particular $(1/2 - \varepsilon \cdot \delta)$ -bounded. Since PRC is robust to such channels, it follows that $\mathcal{W}[\text{PRC}]$ is $(4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n)$ -substring-robust against $\text{BSC}_{1/2-\delta}$. \square

We showed above that when a substring \mathbf{t}' of a response has at least $4\sqrt{\varepsilon} \cdot \mathbf{t}'$ empirical entropy, the embedding channel \mathcal{E}_{Emb} applied to some block in that substring is $(\frac{1}{2} - \varepsilon)$ -bounded. For the following theorem, we need to assume something further about \mathcal{E}_{Emb} : restricted to that block in the substring, \mathcal{E}_{Emb} is equal to $\text{BSC}_{\alpha(\varepsilon)}$ for some $\alpha(\varepsilon) \in (0, \frac{1}{2} - \varepsilon]$.

Assumption 4. There is a function $\alpha : [0, 1/2) \rightarrow [0, 1/2)$ such that for any constant $\varepsilon \in [0, 1/2)$, if a response substring of length L has at least $(4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n)$ empirical entropy, the embedding channel can be modeled as $\text{BSC}_{\alpha(\varepsilon)}$.

As noted in the remark at the end of Section 7.1, this assumption is not realistic if we use the standard transformation from tokens to bits because the bit-errors within the representation of a given token may be correlated. However, an alternative transformation that avoids this issue simply uses one bit for each token.

In this case, Assumption 4 roughly states that the entropy is spread throughout the substring, and that the substring does not repeat too many tokens.

Lemma 24 (Substring-robustness against deletions). *Suppose that Assumption 4 holds with function α . Let $\varepsilon > 0$, $q \in (0, 1/2)$, and $p \in (0, 1)$ be any constants. If PRC_{del} is a zero-bit PRC with block length n and robustness to $\text{BDC}_p \circ \text{BSC}_q \circ \text{BSC}_{\alpha(\varepsilon)} = \text{BDC}_p \circ \text{BSC}_{q+\alpha(\varepsilon)-q\cdot\alpha(\varepsilon)}$, then $\mathcal{W}[\text{PRC}_{\text{del}}]$ is $(4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n)$ -substring-robust against $\text{BDC}_p \circ \text{BSC}_q$.*

Proof. Consider the block of the response over which \mathcal{E}_{Emb} is equivalent to $\text{BSC}_{\alpha(\varepsilon)}$. The decoder receives the output of the composition of the error channel $\text{BDC}_p \circ \text{BSC}_q$ and embedding channel $\text{BSC}_{\alpha(\varepsilon)}$ applied to the codeword used in this block. The composition of these channels is $\text{BDC}_p \circ \text{BSC}_q \circ \text{BSC}_{\alpha(\varepsilon)}$, which PRC_{del} is robust to.

Therefore, the detector of $\mathcal{W}[\text{PRC}_{\text{del}}]$ will output true when it runs Decode on this block of the response. \square

By applying the PRCs of Theorems 1 and 2 to Lemmas 23 and 24, we obtain the following results.

Theorem 5. *Let $\delta > 0$ be any constant and n be a security parameter. Under Assumption 1, there exists a language model watermarking scheme that is $O(L + n)$ -substring-robust against $\text{BSC}_{1/2-\delta}$.*

Theorem 6. *Suppose that Assumption 4 holds with function α , let $q \in (0, 1/2)$ and $p \in (0, 1)$ be any constants, and let n be a security parameter. Under Assumption 1, there exists a language model watermarking scheme that is $O(L + n)$ -substring-robust against $\text{BDC}_p \circ \text{BSC}_q$.*

We remark that although we state our theorems for error rates in $(0, 1/2)$ for convenience, one can easily extend these schemes to error rates for arbitrary constants in $(0, 1) \setminus \{1/2\}$.

7.3 Language model steganography

We have proven completeness, robustness, soundness, and undetectability of \mathcal{W} , where the random strings x used in Wat_{sk} were computed as $\text{PRC.Encode}(\text{PRC.sk})$. These proofs relied only on our ability to distinguish PRC codewords from unrelated strings, so it sufficed to use a zero-bit PRC. Of course, one could obtain a multi-bit watermarking scheme by replacing $\text{PRC.Encode}(\text{PRC.sk})$ with $\text{PRC}'.\text{Encode}(\text{PRC}'.\text{sk}, m)$ for any message m , using PRC' which is multi-bit and has the same robustness as PRC. Lemmas 23 and 24 both apply identically in this case. This yields a robust language model steganography scheme, where steganographic secrecy is implied by undetectability of $\mathcal{W}[\text{PRC}']$. This steganography scheme has the same robustness as the watermarking scheme.

Applying this observation and using the constant-rate PRCs of Section 6.2, we obtain the following result.

Theorem 7. *Let $\delta > 0$ be any constant and n be a security parameter. Under Assumption 1, there exists a language model steganography scheme that is $O(L + n)$ -substring-robust against $\text{BSC}_{1/2-\delta}$.*

In Section 7.4, we will see how this observation about encoding arbitrary messages can be used to build watermarks with public attribution. Since the proofs of security and robustness of the language model steganography scheme are completely identical to those for the watermarking scheme $\mathcal{W}[\text{PRC}]$, we do not formally prove them separately. However, in Section 8 we show that PRCs can be used for *universal* steganography (which is more general than language model steganography).

7.4 Watermarks with public attribution

See Section 2.6 for a description of publicly attributable watermarks.

We define public attribution in terms of a function called AttrText . Recall that we intentionally design this function to *not* be robust. AttrText , given a text t and a public detection key, outputs a pair (t', b) , where $b \in \{\text{true}, \text{false}\}$ indicates whether t contains verbatim a significant prefix of text output by the model, and

if so, \mathbf{t}' is that prefix. Although `AttrText` is intentionally not robust, our scheme's `Detect` function retains all properties (undetectability, robustness, soundness) of our standard watermarks.

Algorithms 3 to 6, given in Figure 4, comprise a watermarking scheme with unforgeable public attribution from any secret-key PRC and any digital signature scheme.

The watermarked text generator of \mathcal{W}_{att} (Algorithm 5) is exactly the same as that of \mathcal{W} (Algorithm 2), but x is now generated as the output of `PRC.Encode` on specific messages. In particular, the first x is sampled as `PRC.Encode(sk)`, and every subsequent x is the encoding of a signature on the response output thus far.

Definition 14 (Unforgeable public attribution). A watermarking scheme \mathcal{W}_{att} for a model `Model` has $b(L)$ *unforgeable public attribution* if there is a function `AttrTextpk(t)` satisfying:

- (Syntax): `AttrTextpk(t) → (t', b)` takes in a token sequence \mathbf{t} and outputs a (token sequence, boolean) pair. $\mathbf{b} = \text{true}$ indicates that a prefix of \mathbf{t} was output verbatim by the model; the corresponding \mathbf{t}' that it outputs is this prefix. $\mathbf{b} = \text{false}$ indicates that no sufficiently long prefix of \mathbf{t} was output verbatim by the model; in this case $\mathbf{t}' = \perp$.
- ($b(L)$ public attribution): For every security parameter λ and prompt `PROMPT` of length $\text{poly}(\lambda)$,

$$\Pr_{\substack{\text{pk, sk} \leftarrow \text{Setup}(1^\lambda) \\ \mathbf{t} \leftarrow \text{Wat}_{\text{sk}}(\text{PROMPT})}} \left[\exists i, L \in [\text{len } \mathbf{t}] \text{ such that } \text{len } \mathbf{t}' < i - 1 \text{ and} \right. \\ \left. H_e^{[i:i+L]}(\text{Model}, \text{PROMPT}, \mathbf{t}) \geq b(L) \mid (\mathbf{t}', \mathbf{b}) \leftarrow \text{AttrText}_{\text{pk}}(\mathbf{t}) \right] \leq \text{negl}(\lambda).$$

- (Unforgeability): For all polynomial-time adversaries \mathcal{A} ,

$$\Pr[\text{AttrForge}_{\mathcal{A}, \mathcal{W}_{\text{att}}}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where the `AttrForge` experiment is given in Figure 5.

Construction 8 (Watermark with public attribution). Let `Sig` be a digital signature scheme with signatures of length n , and let `PRC` be a PRC of block length n . $\mathcal{W}_{\text{att}}[\text{PRC}]$ is the watermarking scheme whose algorithms are specified in Figure 4. Its detector is the same as that of $\mathcal{W}[\text{PRC}]$.

Lemma 25 (Unforgeability). $\mathcal{W}_{\text{att}}[\text{PRC}]$ is unforgeable if the underlying signature scheme is unforgeable.

Proof. Suppose there exists an adversary \mathcal{A} that wins `AttrForgeA, Watt(λ)` with non-negligible probability. We construct \mathcal{B} that uses \mathcal{A} to break unforgeability of the underlying signature scheme `Sig`. \mathcal{B} acts as the challenger in `AttrForgeA, Watt(λ)`. It receives the signature public key `Sig.pk` used in \mathcal{W}_{att} from its own challenger for `SigForgeB, Sig`. \mathcal{B} generates the other parameters for \mathcal{W}_{att} (that is, the PRC parameters) itself. Let `pk` denote the resulting public key of the watermarking scheme; \mathcal{B} passes \mathcal{A} `pk` as input. When responding to \mathcal{A} 's queries to `Watsk`, \mathcal{B} generates the necessary signature using its signing oracle in `SigForgeB, Sig`. Notice that every query that \mathcal{B} makes to its signing oracle is a contiguous substring of some response $\mathbf{t} \in \mathcal{Q}$, where \mathcal{Q} is the set of responses returned to \mathcal{A} by `Watsk`. At the end of `AttrForgeA, Watt(λ)`, \mathcal{A} outputs \mathbf{t}^* . \mathcal{B} computes $(\mathbf{m}, \mathbf{b}) \leftarrow \text{AttrText}(\text{pk}, \mathbf{t}^*)$; if $\mathbf{b} = \text{true}$, \mathcal{B} outputs the corresponding \mathbf{m} and verifying signature σ .

Recall that if \mathcal{A} wins, it outputs \mathbf{t}^* such that $(\mathbf{m}, \mathbf{b}) \leftarrow \text{AttrText}(\text{pk}, \mathbf{t}^*)$ where $\mathbf{b} = \text{true}$ and \mathbf{m} is not a contiguous substring of any response in \mathcal{Q} . Since every query made by \mathcal{B} to its signing oracle was a substring of some response in \mathcal{Q} , \mathcal{B} did not query this \mathbf{m} . However, \mathcal{B} is able to output a verifying signature for \mathbf{m} , so \mathcal{B} wins `SigForgeB, Sig`. \square

Lemma 26 (Public attribution). Let $\varepsilon > 0$ be any constant. If `PRC` has block length n and is robust to $(\frac{1}{2} - \varepsilon)$ -bounded channels, then $\mathcal{W}_{\text{att}}[\text{PRC}]$ satisfies $(4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n)$ public attribution.

Algorithm 4: Publicly attributable watermark setup procedure $\text{Setup}(1^\lambda)$

Input: A security parameter λ .

Result: A watermark public key pk and secret key sk

```

1 PRC.sk  $\leftarrow$  PRC.KeyGen( $1^\lambda$ );
2 Sig.pk, Sig.sk  $\leftarrow$  Sig.KeyGen( $1^\lambda$ );
3  $a_1, \dots, a_{\lceil \frac{L^*}{n} \rceil} \leftarrow \{0, 1\}^n$ ;
4  $\text{pk} \leftarrow \left( \text{PRC.sk}, \text{Sig.pk}, (a_1, \dots, a_{\lceil \frac{L^*}{n} \rceil}) \right)$ ;
5  $\text{sk} \leftarrow \text{Sig.sk}$ ;
6 return  $\text{pk}, \text{sk}$ ;
```

Algorithm 5: Publicly attributable watermarked text generator Wat_{sk}

Input: A prompt (PROMPT), a watermark public key $\text{pk} = \left(\text{PRC.sk}, \text{Sig.pk}, (a_1, \dots, a_{\lceil \frac{L^*}{n} \rceil}) \right)$, and a watermark secret key $\text{sk} = \text{Sig.sk}$.

Result: Watermarked text $\mathbf{t}_1, \dots, \mathbf{t}_L$

```

1  $(\text{PRC.sk}, \text{Sig.pk}, a) \leftarrow \text{pk}$ ;
2  $x \leftarrow \text{PRC.Encode}(\text{PRC.sk}) \oplus a_1$ ;
3  $i := 1, j := 1$ ;
4 while  $\text{done} \notin \{\mathbf{t}_1, \dots, \mathbf{t}_{i-1}\}$  do
    // Embed the current message if first block, or signature otherwise
5    $\mathbf{p}_i := \text{Model}(\text{PROMPT}, \mathbf{t}_1, \dots, \mathbf{t}_{i-1})$ ;
6    $\mathbf{t}_i \leftarrow \text{Ber}(\hat{\mathbf{p}}_i - (-1)^{x_j} \cdot \min\{\hat{\mathbf{p}}_i, 1 - \hat{\mathbf{p}}_i\})$ ;
7    $i \leftarrow i + 1, j \leftarrow j + 1$ ;
    // If done embedding the current codeword, generate a new signature
8   if  $j > n$  then
9      $\mathbf{m} \leftarrow \mathbf{t}_1, \dots, \mathbf{t}_i$ ;
10     $\sigma \leftarrow \text{Sign}_{\text{Sig.sk}}(\mathbf{m})$ ;
11     $x \leftarrow \text{PRC.Encode}(\text{PRC.sk}, \sigma) \oplus a_{\lceil \frac{i}{n} \rceil}$ ;
12     $j \leftarrow 1$ ;
13 return  $\mathbf{t}_1, \dots, \mathbf{t}_L$ ;
```

Algorithm 6: Publicly attributable watermarked text attributor $\text{AttrText}_{\text{pk}}$

Input: Text $\mathbf{t}_1, \dots, \mathbf{t}_L$ and a watermark public key $\text{pk} = \left(\text{PRC.sk}, \text{Sig.pk}, (a_1, \dots, a_{\lceil \frac{L^*}{n} \rceil}) \right)$.

Result: true or false

```

1  $(\text{PRC.sk}, \text{Sig.pk}, a) \leftarrow \text{pk}$ ;
2 for  $i \in [L - n + 1], \ell \in \left[ \lceil \frac{L^*}{n} \rceil \right]$  do
3    $\mathbf{m} \leftarrow \mathbf{t}_1, \dots, \mathbf{t}_{i-1}$ ;
4    $\sigma \leftarrow \text{PRC.Decode}(\text{PRC.sk}, (\mathbf{t}_i, \dots, \mathbf{t}_{i+n-1}) \oplus a_\ell)$ ;
5   if  $\text{Vrfy}_{\text{Sig.pk}}(\mathbf{m}, \sigma)$  then
6     return  $(\mathbf{m}, \text{true})$ ;
7 return  $(\perp, \text{false})$ ;
```

Figure 4: Algorithms (Setup , Wat , AttrText) of $\mathcal{W}_{\text{att}}[\text{PRC}]$. The detector Detect is the same as that of $\mathcal{W}[\text{PRC}]$, given in Algorithm 3. Here, we let n be a codeword length sufficient to encode signatures from Sig using PRC .

$\text{AttrForge}_{\mathcal{A}, \mathcal{W}_{\text{att}}}(\lambda)$ $\text{pk}, \text{sk} \leftarrow \text{KeyGen}(1^\lambda)$ $\mathbf{t}^* \leftarrow \mathcal{A}^{\text{Model}, \text{Wat}_{\text{sk}}}(1^\lambda, \text{pk})$ <p style="margin: 0;">Let \mathcal{Q} denote the set of responses returned to \mathcal{A} by Wat_{sk}</p> $(\mathbf{m}, \mathbf{b}) \leftarrow \text{AttrText}(\text{pk}, \mathbf{t}^*)$ $\text{return } (\mathbf{b} = \text{true}) \wedge (\forall \mathbf{t} \in \mathcal{Q}, \mathbf{m} \text{ is not a contiguous substring of } \mathbf{t})$
--

Figure 5: The attribution forgery experiment $\text{AttrForge}_{\mathcal{A}, \mathcal{W}_{\text{att}}}(\lambda)$

Proof. Let $\mathbf{t} \leftarrow \text{Wat}_{\text{sk}}(\text{PROMPT})$ be a response, and let i be such that $H_e^{[i:i+L]}(\mathbf{t}) \geq 4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n$.

By the exact same proof as in Lemma 22, the robustness of PRC implies that there is some block \mathbf{t}' with starting index at least i , which is correctly decoded by PRC.Decode in Algorithm 6. That is, PRC.Decode yields a signature computed on the substring of the response up to the start of block \mathbf{t}' . Since \mathbf{t}' starts at index at least i , this signed portion has length at least $i - 1$. \square

Theorem 8 (A watermark with unforgeable public attribution). *Let $\varepsilon > 0$ be any constant. If the underlying signature scheme is unforgeable, and PRC is a PRC with block length n and robustness to every $(\frac{1}{2} - \varepsilon)$ -bounded channel, $\mathcal{W}_{\text{att}}[\text{PRC}]$ is unforgeable and $(4\sqrt{\varepsilon} \cdot L + 2\sqrt{2} \cdot n)$ -publicly-attributable.*

Furthermore, $\mathcal{W}_{\text{att}}[\text{PRC}]$ retains the same soundness, undetectability, substring-completeness, and substring-robustness properties as $\mathcal{W}[\text{PRC}]$. That is, Lemmas 18, 19 and 22 to 24 all apply to $\mathcal{W}_{\text{att}}[\text{PRC}]$.

Proof. Unforgeable public attribution follows from Lemmas 25 and 26.

As noted earlier, the proofs of Lemmas 18, 19 and 22 to 24 all hold for any choice of message input to PRC.Encode . The only difference between $\mathcal{W}[\text{PRC}]$ and $\mathcal{W}_{\text{att}}[\text{PRC}]$ is the choice of message input to PRC.Encode . Therefore, those proofs hold for $\mathcal{W}_{\text{att}}[\text{PRC}]$. \square

Related public watermarking work. Concurrent work [FGJ⁺23] constructs a watermark that similarly can be detected with a public key but requires a secret key for generation. Their scheme, like \mathcal{W}_{att} , embeds a digital signature in the response and includes the signature verification key in the public key. However, their scheme has two key differences from \mathcal{W}_{att} :

- They assume that each token sequence of a fixed length ℓ has at least α min-entropy.
- Their scheme has a single detector, analogous to AttrText , which is robust only to cropping. In addition to AttrText , our scheme also has Detect , which is robust to deletions and/or the binary symmetric channel, depending on the underlying PRC.

[FGJ⁺23] also discuss using error-correcting codes to improve robustness. They suggest applying an error-correcting code to the message and signature before embedding them. However, this would make the scheme no longer undetectable or distortion-free.

8 Application: universal steganography

Our main result in this section is the construction of a robust stateless steganography scheme using PRCs. We first prove that robustness follows immediately from applying a PRC to a non-robust steganography scheme appearing in several prior works [AP98, HLVA02, vAH04]. We then show that the assumptions necessary for this scheme can be weakened, which sacrifices some robustness but still yields the most robust scheme in its regime.

8.1 Steganography preliminaries

Setting (from [HLVA02]). A *steganographic channel* (or *covert text distribution*) \mathcal{C} is a distribution on sequences of symbols from some alphabet Σ . These symbols $d \in \Sigma$ are often called *documents*, and sequences of symbols are often called *covert texts*. It is often convenient to consider the conditional distribution over the next symbol, given some subsequence already output by \mathcal{C} . More precisely, given a *history* $h \in \Sigma^*$, we use \mathcal{C}_h to denote the conditional distribution of the next symbol from \mathcal{C} given that thus far \mathcal{C} has output h . For the steganographic encoder, we assume existence of an efficient oracle $M(\cdot)$ that can sample \mathcal{C}_h for any $h \in \Sigma^*$ of the encoder's choice. Our steganographic encoder makes use of a rejection sampling function denoted $RS^{M,f} : \{0, 1\} \times \mathbb{N} \rightarrow \Sigma$, which on input (x, κ) , samples $c \leftarrow M$ until $f(c) = x$, taking at most κ samples. If the sampler reaches the κ^{th} sample c_κ , it outputs c_κ regardless of whether $f(c_\kappa) = x$. We say a function $f : \Sigma \rightarrow R$ is an *unbiased function on a steganographic channel* \mathcal{C} if for all $r \in R$ and all histories $h \in \Sigma^*$, $\Pr_{d \leftarrow \mathcal{C}_h}[f(d) = r] = \frac{1}{|R|}$.

We recall the definition of a symmetric steganography scheme as presented in [KJGR21], which is equivalent to that of [HLVA02].

Definition 15 (Symmetric steganography scheme [KJGR21]). A symmetric steganography scheme $\Sigma_{\mathcal{C}}$ is a triple of possibly probabilistic algorithms $\Pi_{\mathcal{C}} = (\text{KeyGen}_{\mathcal{C}}, \text{Encode}_{\mathcal{C}}, \text{Decode}_{\mathcal{C}})$ parameterized by a covert text channel distribution \mathcal{C} .

- $\text{KeyGen}_{\mathcal{C}}(1^\lambda)$ generates the key sk .
- $\text{Encode}_{\mathcal{C}}(\text{sk}, m, h)$ is a (possibly probabilistic) algorithm that takes the key sk and a plaintext message m , called a *hiddentext*. Additionally, the algorithm can optionally take in a message history $h = \{h_0, h_1, \dots, h_{|h|-1}\} \in \Sigma^*$ of covert text messages. It returns a symbol sequence $c_i \in \Sigma^*$, called a *stegotext*.
- $\text{Decode}_{\mathcal{C}}(\text{sk}, c, h)$ is a (possibly probabilistic) algorithm that takes as input the key sk , a stegotext c , and optionally a history $h \in \Sigma^*$. It returns a plaintext message m on success or the empty string ε on failure.

The subscript \mathcal{C} indicates that these algorithms have access to \mathcal{C} via the oracle $M(\cdot)$; we often omit this subscript for convenience.

A scheme in which the decoding algorithm requires the history is *stateful*; a scheme where decoding is possible without knowledge of the history is *stateless*. This shared history in stateful schemes is very powerful, allowing the sender and receiver to maintain a shared counter of messages sent thus far, which they can use to generate a shared one-time pad for each message.

In a *public-key steganography* scheme, KeyGen outputs a public-secret key pair. Analogously to public-key encryption, the encoder takes as input the public key (and not the secret key), and the decoder still takes as input the secret key. We present formal definitions for secret-key steganography here, and refer the reader to [vAH04] for formal definitions of public-key steganography.

A steganography scheme must satisfy *correctness* and *security*.

Definition 16 (Steganographic correctness). A steganography scheme $\Pi_{\mathcal{C}} = (\text{KeyGen}_{\mathcal{C}}, \text{Encode}_{\mathcal{C}}, \text{Decode}_{\mathcal{C}})$ is *correct* if for any history $h \in \Sigma^*$ and any message m ,

$$\Pr_{\text{sk} \leftarrow \text{KeyGen}(1^\lambda)} [\text{Decode}_{\mathcal{C}}(\text{sk}, \text{Encode}_{\mathcal{C}}(\text{sk}, m, h), h) = m] \geq 1 - \text{negl}(\lambda).$$

We note that some other works require this probability to be 1.

Security requires that an adversary cannot distinguish between oracle access to $\text{Encode}_{\mathcal{C}}(\text{sk}, \cdot, \cdot)$ or a sampling oracle $O_{\mathcal{C}}(\cdot, \cdot)$ for the channel distribution.

Definition 17 (Steganographic security against chosen hiddentext attacks). $\Pi_{\mathcal{C}}$ is secure against *chosen hiddentext attacks* if for all polynomial-time adversaries \mathcal{A} , for all $\text{sk} \leftarrow \text{KeyGen}_{\mathcal{C}}(1^\lambda)$,

$$\left| \Pr \left[\mathcal{A}^{\text{Encode}_{\mathcal{C}}(\text{sk}, \cdot, \cdot), M(\cdot)} = 1 \right] - \Pr \left[\mathcal{A}^{O_{\mathcal{C}}(\cdot, \cdot), M(\cdot)} = 1 \right] \right| \leq \text{negl}(\lambda).$$

This definition can be modified for public-key steganography by giving the adversary the public key as input.

8.2 Robust stateless steganography

We first present **Steg**, a steganography scheme that appeared originally in [AP98]. **Steg** is parameterized by an encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ and a function f .

Construction 9 ($\text{Steg}[(\text{KeyGen}, \text{Enc}, \text{Dec}), f, \kappa]$). Let $\text{Steg.KeyGen} = \text{KeyGen}$. Let Steg.Encode and Steg.Decode be defined as in Figure 6.

[HLVA02] and [vAH04] prove security of the secret- and public-key versions of **Steg**, respectively.

Claim 27 ([HLVA02, vAH04]). *Let $(\text{KeyGen}, \text{Enc}, \text{Dec})$ be a (public-key) encryption scheme with ciphertexts indistinguishable from random bits under a chosen plaintext attack, and let $f : \Sigma \rightarrow \{0, 1\}$ be a function which is unbiased on \mathcal{C} . Then $\text{Steg}[(\text{KeyGen}, \text{Enc}, \text{Dec}), f, \lambda]$ is a secure (public-key) steganography scheme.*

Our first application of PRCs to steganography will be in showing that **Steg** is robust to errors when a PRC is used as the encryption scheme. To our knowledge, this gives the first robust stateless steganography scheme.

Here, we define robustness of a steganography scheme against a channel \mathcal{E} analogously to robustness of a PRC. In particular, \mathcal{E} may be a deletion channel, rather than just making substitutions. Robustness requires that given the output of the error channel applied to a stegotext, the hiddentext can be recovered with overwhelming probability.

Definition 18 (Robustness to an error channel). Let \mathcal{C} be a steganographic channel with symbol alphabet Σ , and let $\mathcal{E} : \Sigma^* \rightarrow \Sigma^*$ be an error channel. A steganography scheme $(\text{KeyGen}, \text{Encode}, \text{Decode})$ for \mathcal{C} is robust to \mathcal{E} if for all messages m and all histories h ,

$$\Pr_{\text{sk}, \text{pk} \leftarrow \text{KeyGen}(1^\lambda)} [\text{Decode}(\text{sk}, \mathcal{E}(x), h) = m : x \leftarrow \text{Encode}(\text{pk}, m, h)] = 1 - \text{negl}(\lambda).$$

This definition can be adapted for secret-key schemes by replacing pk with sk .

We say that the *rate* of a steganography scheme is the ratio of the number of message bits to the number of symbols of stegotext that are needed to decode the message.

Theorem 9 (Robust stateless steganography). *Let $\text{PRC} = (\text{KeyGen}, \text{Encode}, \text{Decode})$ be a (public-key) PRC that is robust to some binary channel \mathcal{E} . Let $f : \Sigma \rightarrow \{0, 1\}$ be a public function which is unbiased on \mathcal{C} . Then $\text{Steg}[\text{PRC}, f, \lambda]$ is a (public-key) steganography scheme, with the same rate as PRC , that is robust to any \mathcal{E}' such that $f \circ \mathcal{E}' = \mathcal{E} \circ f$.*

Algorithm 7: Steg.Encode

Input: Key pk , message m , history h .

```

1 Let  $x \leftarrow \text{Enc}(pk, m)$ ;
2 for  $i = 1, \dots, n$  do
3    $c_i \leftarrow RS^{M(h),f}(x_i, \kappa)$ ;
4   Set  $h = h||c_i$ ;
5 return  $c_1||c_2||\dots||c_n$ 

```

Algorithm 8: Steg.Decode

Input: Key sk and stegotext c' .

```

1 Parse  $c'$  as  $c'_1||c'_2||\dots||c'_n$ ;
2 for  $i = 1, \dots, n$  do
3   Set  $x'_i = f(c'_i)$ ;
4 Set  $x' = x'_1||x'_2||\dots||x'_n$ ;
5 return  $\text{Dec}(sk, x')$ 

```

Figure 6: Encoding and decoding algorithms of $\text{Steg}[(\text{KeyGen}, \text{Enc}, \text{Dec}), f, \kappa]$.

Remark. We use $\mathcal{E} \circ f$ to denote the channel where f is applied individually to each symbol of the input, and then \mathcal{E} is applied to the resulting bitstring. Similarly, $f \circ \mathcal{E}'$ is the channel where \mathcal{E}' is applied, and then f is applied individually to each symbol of the resulting Σ -string. The condition that $f \circ \mathcal{E}' = \mathcal{E} \circ f$ means that \mathcal{E}' acts as \mathcal{E} in the binary representation. For instance, if \mathcal{E}' is a deletion channel over Σ , then \mathcal{E} is the corresponding deletion channel over $\{0, 1\}$; if \mathcal{E}' is the channel that independently replaces each symbol with a random symbol with probability p , and f is a random function, then \mathcal{E} is approximately $BSC_{p/2}$.

Proof. Steganographic security follows immediately from Claim 27 and the fact that a PRC is a pseudorandom encryption scheme.

For robustness, let m be an arbitrary message and h an arbitrary history, and let $c = c_1||\dots||c_n \leftarrow \text{Steg.Encode}(pk, m, h)$. Recall that each c_i is chosen as $RS^{M(h),f}(x_i, \lambda)$, where $x = x_1||\dots||x_n$ is the output of PRC.Encode . Since f is unbiased, each sample x' from RS satisfies $f(x') = x_i$ independently with probability $1/2$, and therefore with overwhelming probability within λ draws the sampler will output c_i such that $f(c_i) = x_i$.

Given $c' = \mathcal{E}'(c)$, decoder first computes $x' = f(c'_1)||\dots||f(c'_n)$. Observe that $x' = f \circ \mathcal{E}'(c)$, which by assumption is equal to $\mathcal{E} \circ f(c) = \mathcal{E}(x)$, where x is the output of PRC.Encode . By robustness of the PRC, $\text{PRC.Decode}(sk, x', h) = m$. \square

Applying Theorem 9 with our PRCs from Section 6.3, we obtain a stateless public-key steganography scheme that is robust to random deletions and substitutions.

8.3 Stateless steganography from weaker modeling assumptions

The assumption of the existence of an unbiased function f over \mathcal{C} is quite strong. One way to relax this assumption is to assume instead that $f(\mathcal{C})$ meets some min-entropy requirement. However, now f may be unbalanced; for example, it could be the case that $\Pr_{c_i \leftarrow \mathcal{C}_h}[f(c_i) = 1] = 2/3$ for all h . Now, the encoder in Figure 6 would no longer be secure: Consider sampling each symbol c_i in the stegotext to *always* satisfy $f(c_i) = x_i$. Then, since each x_i is uniform in $\{0, 1\}$, half of the symbols c_i in the stegotext would satisfy $f(c_i) = 1$. Therefore, one could distinguish between stegotexts and samples from \mathcal{C} using f .

A natural way (used by [HLVA02]) to build a secure stateful steganography scheme under only this min-entropy assumption, is to let c be an error-corrected version of the message and use in the encoder a rejection sampler that takes at most two samples, sometimes outputting $f(c_i) \neq x_i$. This rejection sampler preserves the channel distribution as long as c_i is random, which is not the case for generic error-correcting codes. Therefore, [HLVA02] relies on shared state to let the sender and receiver generate a fresh one-time pad per stegotext, letting c be an error-corrected message with this one-time pad applied. This shared state significantly simplifies the problem and is a strong assumption of its own.

Relaxing the assumption of an unbiased function poses more challenges for *stateless* steganography, where this error-correction approach fails. Although [HvAL08] does construct a stateless steganography scheme

under only a min-entropy assumption, this scheme has poor robustness. The idea behind this scheme is that the encoder first samples a symbol sequence d that is long enough to have λ min-entropy, and includes this sequence in the stegotext. Since both the encoder and decoder know d , d can now act as their shared state, and the remainder of the stegotext is formed using a stateful scheme such as that of [HLVA02]. The min-entropy assumption ensures that the state will never be used twice.

Theorem 10. *Let \mathcal{C} be a steganographic channel with alphabet Σ , and let $f : \Sigma \rightarrow \{0, 1\}$ be a function. Suppose that there exists some $\alpha > 0$ such that for all histories $h \in \Sigma^*$, $\min_{b \in \{0, 1\}} \Pr_{c \leftarrow \mathcal{C}_h} [f(c) = b] \geq \alpha$. Then for any $p \in (0, 1/2)$, there exists $q \in (0, 1/2)$ such that if PRC is any (public-key) PRC for every q -bounded channel, then $\text{Steg}[\text{PRC}, f, 2]$ is a (public-key) stateless steganography scheme for \mathcal{C} , with the same rate as PRC, that is robust to any channel \mathcal{E} such that $f \circ \mathcal{E} = \text{BSC}_p \circ f$.*

Proof. We begin by showing that $\text{Steg}[\text{PRC}, f, 2]$ is steganographically secure. As in Figure 6, suppose that we wish to encode a message m , let $x \leftarrow \text{PRC.Encode}(\text{pk}, m)$, and let $c_i \leftarrow \text{RS}^{M(h), f}(x_i, 2)$ be the i^{th} symbol output by the steganography scheme. Let c_i^1 denote the first sample from $M(h)$ and let c_i^2 denote the second sample (which is irrelevant in the event that $f(c_i^1) = x_i$).

By pseudorandomness of $x \leftarrow \text{PRC.Encode}(\text{pk}, m)$, it suffices to show that for any $i \in [n]$, $h \in \Sigma^{i-1}$, and $c^* \in \Sigma$,

$$\Pr_{\substack{x_i \leftarrow \{0, 1\} \\ c_i \leftarrow \text{RS}^{M(h), f}(x_i, 2)}} [c_i = c^*] = \Pr_{c \leftarrow M(h)} [c = c^*].$$

Indeed,

$$\begin{aligned} \Pr_{\substack{x_i \leftarrow \{0, 1\} \\ c_i \leftarrow \text{RS}^{M(h), f}(x_i, 2)}} [c_i = c^*] &= \Pr_{\substack{x_i \leftarrow \{0, 1\} \\ c_i^1 \leftarrow M(h)}} [c_i^1 = c^* \wedge f(c_i^1) = x_i] + \Pr_{\substack{x_i \leftarrow \{0, 1\} \\ c_i^1, c_i^2 \leftarrow M(h)}} [c_i^2 = c^* \wedge f(c_i^1) \neq x_i] \\ &= \frac{1}{2} \Pr_{c_i^1 \leftarrow M(h)} [c_i^1 = c^*] + \frac{1}{2} \Pr_{c_i^2 \leftarrow M(h)} [c_i^2 = c^*] \\ &= \Pr_{c \leftarrow M(h)} [c = c^*]. \end{aligned}$$

Now we turn to showing that $\text{Steg}[\text{PRC}, f, 2]$ is robust to any channel \mathcal{E} such that $f \circ \mathcal{E} = \text{BSC}_p \circ f$, provided that PRC is robust to every q -bounded channel for appropriate choice of q .

Again, pseudorandomness of $x \leftarrow \text{PRC.Encode}(\text{pk}, m)$ allows us to assume x is random. Observe that the bits $f(c_i)$ are correlated with the bits x_i :

$$\begin{aligned} \Pr_{\substack{x_i \leftarrow \{0, 1\} \\ c_i \leftarrow \text{RS}^{M(h), f}(x_i, 2)}} [f(c_i) = x_i] &= \Pr_{\substack{x_i \leftarrow \{0, 1\} \\ c_i^1 \leftarrow M(h)}} [f(c_i^1) = x_i] + \Pr_{\substack{x_i \leftarrow \{0, 1\} \\ c_i^1, c_i^2 \leftarrow M(h)}} [f(c_i^2) = x_i \wedge f(c_i^1) \neq x_i] \\ &= \frac{1}{2} + \frac{1}{2} \Pr_{\substack{x_i \leftarrow \{0, 1\} \\ c_i^1, c_i^2 \leftarrow M(h)}} [f(c_i^2) = x_i \mid f(c_i^1) \neq x_i] \\ &\geq \frac{1}{2} + \frac{1}{2} \alpha. \end{aligned}$$

For $i \in [n]$ let \hat{c} be the content generated by Steg (before the error channel \mathcal{E} is applied), defined by $\hat{c} = \hat{c}_1 \parallel \dots \parallel \hat{c}_n$ for $\hat{c}_i = f(c_i)$. By Azuma's inequality (Lemma 1), $\text{wt}(\hat{c} \oplus x) \leq (1/2 - \alpha/4) \cdot n$ with probability $1 - \text{negl}(n)$. The decoder for the steganography scheme will compute $(f \circ \mathcal{E})(c)$, which is distributed identically to $(\text{BSC}_p \circ f)(c) = \text{BSC}_p(\hat{c})$ since $f \circ \mathcal{E} = \text{BSC}_p \circ f$. By a Chernoff bound and straightforward computation,

$$\text{wt}(\text{BSC}_p(\hat{c}) \oplus x) \leq \left(\frac{1}{2} - \frac{\alpha}{8}(1 - 2p) \right) \cdot n$$

with probability $1 - \text{negl}(n)$. Therefore, if PRC is robust to any q -bounded channel for $q = 1/2 - (\alpha/8)(1 - 2p)$, then $\text{Steg}[\text{PRC}, f, 2]$ is robust to \mathcal{E} . \square

References

- [Aar22] Scott Aaronson. My AI Safety Lecture for UT Effective Altruism. <https://scottaaronson.blog/?p=6823>, November 2022. Accessed May 2023.
- [ABN⁺92] Noga Alon, Jehoshua Bruck, Joseph Naor, Moni Naor, and Ron M. Roth. Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *IEEE Trans. Inf. Theory*, 38(2):509–516, 1992.
- [ACI⁺20] Thomas Agrikola, Geoffroy Couteau, Yuval Ishai, Stanisław Jarecki, and Amit Sahai. On pseudorandom encodings. In *Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part III 18*, pages 639–669. Springer, 2020.
- [ADI⁺17] Benny Applebaum, Ivan Damgård, Yuval Ishai, Michael Nielsen, and Lior Zichron. Secure arithmetic computation with constant computational overhead. In *Annual International Cryptology Conference*, pages 223–254. Springer, 2017.
- [Ale03] Michael Alekhovich. More on average case vs approximation complexity. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11–14 October 2003, Cambridge, MA, USA, Proceedings*, pages 298–307. IEEE Computer Society, 2003.
- [AP98] Ross J Anderson and Fabien AP Petitcolas. On the limits of steganography. *IEEE Journal on selected areas in communications*, 16(4):474–481, 1998.
- [ARC⁺01] Mikhail J Atallah, Victor Raskin, Michael Crogan, Christian Hempelmann, Florian Kerschbaum, Dina Mohamed, and Sanket Naik. Natural language watermarking: Design, analysis, and a proof-of-concept implementation. In *Information Hiding: 4th International Workshop, IH 2001 Pittsburgh, PA, USA, April 25–27, 2001 Proceedings 4*, pages 185–200. Springer, 2001.
- [ARH⁺02] Mikhail J Atallah, Victor Raskin, Christian F Hempelmann, Mercan Karahan, Radu Sion, Umut Topkara, and Katrina E Triezenberg. Natural language watermarking and tamperproofing. In *International workshop on information hiding*, pages 196–212. Springer, 2002.
- [Ari09] Erdal Arikan. Channel polarization: a method for constructing capacity-achieving codes for symmetric binary-input memoryless channels. *IEEE Trans. Inf. Theory*, 55(7):3051–3073, 2009.
- [ASS⁺23] Shweta Agrawal, Sagnik Saha, Nikolaj Ignatieff Schwartzbach, Akhil Vanukuri, and Prashant Nalini Vasudevan. k -sum in the sparse regime. Cryptology ePrint Archive, Paper 2023/488, 2023. <https://eprint.iacr.org/2023/488>.
- [BC05] Michael Backes and Christian Cachin. Public-key steganography with active attacks. In *Theory of Cryptography Conference*, pages 210–226. Springer, 2005.
- [BCG⁺19] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent ot extension and more. In *Advances in Cryptology—CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part III 39*, pages 489–518. Springer, 2019.
- [BH23] Diane Bartz and Krystal Hu. OpenAI, Google, others pledge to watermark AI content for safety, White House says, July 2023. <https://www.reuters.com/technology/openai-google-others-pledge-watermark-ai-content-safety-white-house-2023-07-21>.
- [Cac98] Christian Cachin. An information-theoretic model for steganography. In *International Workshop on Information Hiding*, pages 306–318. Springer, 1998.
- [CGZ23] Miranda Christ, Sam Gunn, and Or Zamir. Undetectable watermarks for language models. *IACR Cryptol. ePrint Arch.*, page 763, 2023.

- [Cra98] Scott Craver. On public-key steganography in the presence of an active warden. In *International Workshop on Information Hiding*, pages 355–368. Springer, 1998.
- [DGG⁺15] Yevgeniy Dodis, Chaya Ganesh, Alexander Golovnev, Ari Juels, and Thomas Ristenpart. A formal treatment of backdoored pseudorandom generators. In *Advances in Cryptology—EUROCRYPT 2015: 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26–30, 2015, Proceedings, Part I 34*, pages 101–126. Springer, 2015.
- [DIRR09] Nenad Dedić, Gene Itkis, Leonid Reyzin, and Scott Russell. Upper and lower bounds on black-box steganography. *Journal of Cryptology*, 22:365–394, 2009.
- [dWSK⁺22] Christian Schroeder de Witt, Samuel Sokota, J Zico Kolter, Jakob Foerster, and Martin Strohmeier. Perfectly secure steganography using minimum entropy coupling. *arXiv preprint arXiv:2210.14889*, 2022.
- [FGJ⁺23] Jaiden Fairoze, Sanjam Garg, Somesh Jha, Saeed Mahloujifar, Mohammad Mahmoody, and Mingyuan Wang. Publicly detectable watermarking for language models. *Cryptology ePrint Archive*, 2023.
- [Fri09] Jessica Fridrich. *Steganography in digital media: principles, algorithms, and applications*. Cambridge University Press, 2009.
- [Gal62] Robert G. Gallager. Low-density parity-check codes. *IRE Trans. Inf. Theory*, 8(1):21–28, 1962.
- [GKVZ22] Shafi Goldwasser, Michael P Kim, Vinod Vaikuntanathan, and Or Zamir. Planting undetectable backdoors in machine learning models. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 931–942. IEEE, 2022.
- [HLVA02] Nicholas J Hopper, John Langford, and Luis Von Ahn. Provably secure steganography. In *Advances in Cryptology—CRYPTO 2002: 22nd Annual International Cryptology Conference Santa Barbara, California, USA, August 18–22, 2002 Proceedings 22*, pages 77–92. Springer, 2002.
- [Hoe94] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding*, pages 409–426, 1994.
- [Hop05] Nicholas Hopper. On steganographic chosen coverttext security. In *International Colloquium on Automata, Languages, and Programming*, pages 311–323. Springer, 2005.
- [HvAL08] Nicholas Hopper, Luis von Ahn, and John Langford. Provably secure steganography. *IEEE Transactions on Computers*, 58(5):662–676, 2008.
- [KAAL23] Jan Hendrik Kirchner, Lama Ahmad, Scott Aaronson, and Jan Leike. New ai classifier for indicating AI-written text. <https://openai.com/blog/new-ai-classifier-for-indicating-ai-written-text>, January 2023. Accessed May 2023.
- [KGW⁺23a] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A watermark for large language models. *arXiv preprint arXiv:2301.10226*, 2023.
- [KGW⁺23b] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Manli Shu, Khalid Saifullah, Kezhi Kong, Kasun Fernando, Aniruddha Saha, Micah Goldblum, and Tom Goldstein. On the reliability of watermarks for large language models. *arXiv preprint arXiv:2306.04634*, 2023.
- [KJGR21] Gabriel Kaptchuk, Tushar M. Jois, Matthew Green, and Aviel D. Rubin. Meteor: Cryptographically secure steganography for realistic distributions. In Yongdae Kim, Jong Kim, Giovanni Vigna, and Elaine Shi, editors, *CCS ’21: 2021 ACM SIGSAC Conference on Computer and*

- Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021*, pages 1529–1548. ACM, 2021.
- [KKRT16] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious prf with applications to private set intersection. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 818–829, 2016.
- [KL07] Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography: principles and protocols*. Chapman and hall/CRC, 2007.
- [KTHL23] Rohith Kuditipudi, John Thickstun, Tatsunori Hashimoto, and Percy Liang. Robust distortion-free watermarks for language models. *arXiv preprint arXiv:2307.15593*, 2023.
- [MLK⁺23] Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D. Manning, and Chelsea Finn. DetectGPT: Zero-shot machine-generated text detection using probability curvature. *CoRR*, abs/2301.11305, 2023.
- [MTSB13] Rafael Misoczki, Jean-Pierre Tillich, Nicolas Sendrier, and Paulo SLM Barreto. MDPC-McEliece: New McEliece variants from moderate density parity-check codes. In *2013 IEEE international symposium on information theory*, pages 2069–2073. IEEE, 2013.
- [NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM J. Comput.*, 22(4):838–856, 1993.
- [PSF⁺23] Julien Piet, Chawin Sitawarin, Vivian Fang, Norman Mu, and David Wagner. Mark my words: Analyzing and evaluating language model watermarks. *arXiv preprint arXiv:2312.00273*, 2023.
- [RBB03] Phillip Rogaway, Mihir Bellare, and John Black. Ocb: A block-cipher mode of operation for efficient authenticated encryption. *ACM Transactions on Information and System Security (TISSEC)*, 6(3):365–403, 2003.
- [Sim84] Gustavus J Simmons. The prisoners’ problem and the subliminal channel. In *Advances in Cryptology: Proceedings of Crypto 83*, pages 51–67. Springer, 1984.
- [Ta-17] Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 238–251. ACM, 2017.
- [TCH23] Ruixiang Tang, Yu-Neng Chuang, and Xia Hu. The science of detecting LLM-generated texts. *arXiv preprint arXiv:2303.07205*, 2023.
- [Tia23] Edward Tian. GPTZero update V1. <https://gptzero.substack.com/p/gptzero-update-v1>, January 2023. Accessed May 2023.
- [TTDI05] Mercan Topkara, Cunevt M Taskiran, and Edward J Delp III. Natural language watermarking. In *Security, Steganography, and Watermarking of Multimedia Contents VII*, volume 5681, pages 441–452. SPIE, 2005.
- [vAH04] Luis von Ahn and Nicholas J. Hopper. Public-key steganography. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, volume 3027 of *Lecture Notes in Computer Science*, pages 323–341. Springer, 2004.
- [VV83] Umesh V. Vazirani and Vijay V. Vazirani. Trapdoor pseudo-random number generators, with applications to protocol design. In *24th Annual Symposium on Foundations of Computer Science (sfcs 1983)*, pages 23–30, 1983.

- [ZALW23] Xuandong Zhao, Prabhanjan Ananth, Lei Li, and Yu-Xiang Wang. Provable robust watermarking for AI-generated text. *arXiv preprint arXiv:2306.17439*, 2023.
- [Zam24] Or Zamir. Excuse me, sir? your language model is leaking (information). *arXiv preprint arXiv:2401.10360*, 2024.
- [ZDR19] Zachary M. Ziegler, Yuntian Deng, and Alexander M. Rush. Neural linguistic steganography. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 1210–1215. Association for Computational Linguistics, 2019.
- [ZEF⁺23] Hanlin Zhang, Benjamin L Edelman, Danilo Francati, Daniele Venturi, Giuseppe Ateniese, and Boaz Barak. Watermarks in the sand: Impossibility of strong watermarking for generative models. *arXiv preprint arXiv:2311.04378*, 2023.
- [ZHR⁺19] Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. Defending against neural fake news. *Advances in neural information processing systems*, 32, 2019.

A Related work

Comparison to code-based cryptography. Our approach to constructing pseudorandom codes is closely related to code-based cryptography, and in particular the McEliece cryptosystem. The McEliece cryptosystem is a public-key encryption scheme in which the public key is a generator matrix for a linear code and the secret key is a trapdoor for efficient decoding. In this sense, our constructions of pseudorandom codes can be viewed as instantiations of the McEliece cryptosystem. However, there are two key differences between our setting and that of code-based cryptography that render code-based cryptography results inapplicable for our purposes.

The first difference is in the source of the errors, or noise. PRCs are required to correct uncontrolled errors introduced by an adversarial channel; in McEliece, all of the errors are introduced by an honest user for the purpose of securely sending the message over a perfect channel. Therefore, the channel itself is specified as part of the cryptosystem and the code only needs to be robust to *this channel*. In pseudorandom codes, it is crucial that we can correct from an arbitrary constant rate $p < 1/2$ of additional errors introduced by a noisy channel on a binary alphabet. Existing code-based cryptosystems (such as binary Goppa codes, Reed-Solomon codes, Reed-Muller codes, or Algebraic Geometry codes) either do not enjoy such strong robustness or rely on a larger alphabet.

The second difference is that we require noisy codewords to be *pseudorandom*, rather than just hiding. If every (noisy) codeword satisfies some efficiently-computable property P , but a uniformly random string does not satisfy P , then the code cannot be a PRC. However, such a code might still define a secure encryption scheme, as long as P does not distinguish between codewords for different messages.

Our primary construction of PRCs is based on low-density parity-check (LDPC) codes, with somewhat-higher $\Omega(\log n)$ density than is typically considered. The work of [MTSB13] consider “moderate-density parity-check” codes, with density $\Omega(\sqrt{n \log n})$. While their codes suffice for code-based cryptography, they do not enjoy as strong robustness as ours because of the higher density. In particular, they are not robust to any constant rate of substitutions.

PRC-like cryptographic tools. Error-correcting codes with some pseudorandomness properties have proven useful in cryptographic applications such as pseudorandom correlation generators [BCG⁺19] and oblivious linear evaluation (OLE) [ADI⁺17]. However, in all of these applications, decoding requires new side information for each message, such as the locations of the errors. The fundamental difficulty in constructing

our pseudorandom codes is that the only extra information provided to the decoder is a single, unchanging secret key; without this key the messages must be pseudorandom.

Pseudorandom encodings, defined and studied in [ACI+20], are similar in name to PRCs but are very different objects. A pseudorandom encoding for a distribution X is a pair of unkeyed algorithms (`Encode`, `Decode`) such that `Encode`(x) is pseudorandom (i.e., appears uniform) for a random $x \leftarrow X$, and `Decode`(`Encode`(x)) = x with overwhelming probability. For pseudorandom encodings there is no robustness requirement. Furthermore, pseudorandomness for PRCs requires that `Encode`(x) is pseudorandom for any x of the adversary’s choice.

Language watermarking schemes. Classical watermarking considers the problem of embedding a signal into a *fixed* object, such as a given text, in such a way that the watermark is hard to remove, and the quality of the original object is preserved. In that setting, studied specifically for natural language in [TTDI05, ARC+01, ARH+02], planting the watermark must alter the text, and the goal is to minimize some distance measure between the original text and the watermarked text. In contrast, since generative models are randomized, the quality goal of watermarks for generative models is to preserve certain properties of the *distribution* of the model’s outputs. Due to this difference, new techniques have been developed for watermarking AI-generated content, particularly text output by language models, which is the focus of our watermarking work in this paper.

A language model takes as input a *prompt* and randomly samples a *response* using an iterative process. Given a prompt, it computes a probability distribution p_1 for the first *token* and samples a token t_1 from this distribution. At each subsequent step, it takes as input the token sequence t_1, \dots, t_{i-1} output thus far, computes the distribution p_i , and samples the i^{th} token in the response t_i from p_i . It continues doing so until it samples a special “done” token, at which point it terminates and outputs the response. The randomness in this process is important for the usefulness of the model—a model should produce a wide variety of useful responses given the same prompt—and it is crucial for watermarking.

Recently, there have been several language model watermarking schemes proposed (e.g., [Aar22, KGW+23a, CGZ23, ZALW23, KTHL23, FGJ+23]), which embed the watermark by changing the way that each t_i is sampled from p_i . At a very high level, these watermark embedders give preference to certain tokens in the sampling process. The corresponding detectors check for the presence of more preferred tokens than would be expected in independently generated text.

The simplest of these schemes, [ZALW23], fixes a random partition of tokens into equal-sized green and red lists, and it embeds the watermark by increasing the probabilities of tokens on the green list when it samples its response. The detector computes the fraction of green tokens, which should be appreciably greater than $\frac{1}{2}$ for watermarked text and roughly $\frac{1}{2}$ for all other text. This red-green partition is used for every sampled token across all responses, which results in a significant change to the model’s distribution. For example, if “computer” is on the red list, the model now prefers not to talk about computers.

[Aar22, KGW+23a] mitigate this distributional shift by generating the red-green partitions dynamically. Rather than using the same red-green partition for every token, [Aar22, KGW+23a] select a partition for each token using a PRF evaluated at the previous k tokens. That is, these schemes use a PRF F_{sk} with secret key sk , and add weight to tokens t_i such that $F_{sk}(t_{i-k}, \dots, t_{i-1}, t_i) = 1$. This effectively generates a new red-green partition for each token, assuming that no prefix t_{i-k}, \dots, t_{i-1} ever appears twice. The detector computes the fraction of tokens on these green lists, which it can compute by evaluating the PRF, provided that the seeds are intact in the given text. The length k of the token sequence used as these seeds can be tuned to trade off between robustness and frequency of seed reuse. Longer seeds result in less frequent seed reuse but make the watermark easier to remove, since the adversary can destroy every seed and prevent the detector from computing the green lists by changing only $1/k$ tokens in a watermarked response.

Although [KGW+23a] and [Aar22] both generate randomness using this seeded PRF, they use this randomness to sample from p_i differently, achieving different quality guarantees as a result. The watermarking

algorithm in [Aar22] samples tokens such that, for each token of the response, the expected watermarked distribution (over the randomness of sk) is identical to the original model’s distribution. However, while individual tokens’ distributions are preserved, the watermark introduces correlations between tokens’ distributions as seeds can be reused (even within the same response). For example, if the same prompt is queried multiple times, the watermark will result in the same bias for the first token of the response each time. [KGW+23a] achieves a weaker guarantee and changes the distribution of the model’s outputs.

Distortion-freeness, a quality notion defined by [KTHL23], is the property that the expected distribution of a single response from the watermarked model is identical to the distribution of a single response from the original model. [KTHL23] constructs a distortion-free watermarking scheme with an interesting level of robustness. To generate each response, this watermarking scheme samples a string x^* from a fixed collection of seeds x_1, \dots, x_ℓ . The i^{th} token in the response is chosen to be correlated with the i^{th} bit of x^* . The detector checks whether the given text is watermarked by computing its edit distance from each of the strings x_1, \dots, x_ℓ , returning true if any of these distances is below some threshold.

This detection algorithm is quite slow in practice: It runs in time $O(n^2\ell)$, where n is the length of the random strings. While using a smaller ℓ improves the detector’s efficiency, it deteriorates the model’s quality by reducing the response variability. Our watermark avoids this tradeoff between variability of the model and detector efficiency, while still offering robustness.

The strongest quality guarantee is *undetectability*, defined by [CGZ23], which is the quality notion we focus on in this work. Undetectability requires that the watermarked model and original model are computationally indistinguishable to an adversary without knowledge of the detection key, even if the adversary can make an unbounded number of adaptive queries. In contrast, distortion-freeness of [KTHL23] says nothing about multiple responses from the model. In particular, a watermark can be distortion-free but render the response to a given prompt entirely deterministic — even if the original model had a great deal of variability on the same prompt.

Like [KGW+23a, Aar22], the watermark in [CGZ23] uses a PRF seeded with previously output tokens to generate the randomness used to bias the token sampler. A crucial observation used in [CGZ23] to avoid the seed reuse issue of [KGW+23a, Aar22], is that if the tokens constituting the seed contain enough entropy, seed reuse becomes unlikely. Because the watermarking algorithm has access to the token distributions, it can compute the amount of entropy in a token sequence and use it as a seed only once this entropy has exceeded a certain threshold. This allows [CGZ23] to achieve undetectability and a robustness guarantee they call *substring-completeness*, that any sufficiently high-entropy substring of a response will be detected as watermarked. This is essentially the same robustness guarantee as [KGW+23a, Aar22], but it is weaker than that of [KTHL23, ZALW23], since changing one token in every seed removes the watermark.

[FGJ+23] has similar quality to [CGZ23], and embeds a digital signature in the response to make it detectable by anyone with a public detection key. A separate secret key is required to embed the watermark. Similarly to [CGZ23], their watermark generates the initial portion of the response from the original model. It then uses this portion as a seed for a PRF and encodes each bit of a digital signature by sampling a block of tokens such that the PRF evaluated on that block is equal to that bit. Like [CGZ23], changing any token in the seed removes the watermark.

We compare some of these watermarking schemes in Table 1. We refer the reader to [PSF+23] for a more detailed empirical comparison of some of these schemes, and to [KGW+23b] for an empirical study of the robustness of [KGW+23a].

On removing watermarks. [CGZ23] describe an attack that removes any undetectable watermark and preserves response quality, but this attack is very impractical because it involves querying the model once for each token in the response. [ZEF+23] describe an attack that can remove any watermark in a way that preserves response quality, assuming that the attacker has access to a quality oracle and that random walks over the graph of responses mix sufficiently quickly. These assumptions are quite strong, and it is not

Paper	Single-query undetectable?	Undetectable?	Robust?
[Aar22]	✓*	✗	✗
[KGW ⁺ 23a]	✗	✗	✗
[CGZ23]	✓	✓	✗
[ZALW23]	✗	✗	✓
[KTHL23]	✓	✗	✓
[FGJ ⁺ 23]	✓	✓	✗
This work	✓	✓	✓

Table 1: Comparison between various watermarking schemes for language models. “Single-query undetectable” means that a single query to the watermarked model is computationally indistinguishable from one to the original model, and “undetectable” is the multi-query analogue as defined in [CGZ23]. In this table, “robust” means that the watermark is resilient to a constant rate of substitutions.

* [Aar22] is single-query undetectable provided that every short contiguous sequence of response tokens has high enough entropy.

clear whether there are fast, reliable quality oracles that don’t already yield a full language model. Still, in light of these impossibility results, one cannot hope to have a watermarking scheme that is robust against all possible attacks — indeed, an adversary with sufficient knowledge of the language could always write a high-quality text without even consulting the watermarked model. Therefore, we instead define robustness against particular classes of adversaries, including those that delete and substitute tokens.

Other techniques for detecting machine-generated text. Another approach for detecting AI-generated text is to train a machine learning classifier to distinguish between AI-generated and natural text [ZHR⁺19, MLK⁺23, Tia23, KAAL23]. However, these classifiers lack transparency, can be easily evaded, have high false-positive rates, and have unpredictable biases. OpenAI retracted its classifier-based detector due to these issues.

Instead of relying on existing idiosyncrasies of AI-generated text, watermarks embed patterns themselves, making detection more reliable and transparent. See, e.g., [TCH23, PSF⁺23] for an overview of the various approaches to detecting AI-generated text.

Undetectable backdoors for models. [GKVZ22] shows how to embed a hidden “backdoor” during training of a model. Using a secret key, one can “activate” the backdoor by slightly perturbing inputs to alter their classification under the backdoored model. These backdoors are undetectable in the sense that, without the secret key, one cannot distinguish between an honestly trained or a backdoored model. While not closely related to our work in technical content, [GKVZ22] is similar in spirit as an application of cryptographic definitions and techniques to machine learning.

Universal steganography. Steganography was introduced by Simmons [Sim84], who presented it as problem where two prisoners wish to communicate in the presence of a warden, hiding not only the content of their communication but also the fact that this communication is occurring. [HLVA02] first formalized steganography in the computational setting. Here, there is some distribution of *coverttexts* in which the sender wishes to conceal its *message*. In Section 8, we consider *universal* steganography, where the coverttext distribution is an arbitrary distribution to which the sender has sample access. The sender uses this sample access and a secret key to construct a *stegotext* that encodes the message, which the receiver decodes using the secret key.

Loosely speaking, steganographic security requires that an outside observer cannot distinguish stegotexts from coverttexts, and furthermore cannot learn anything about the message. There are several security

Scheme	Stateless?	No unbiased function?	Public-key?	Robust?
[HLVA02] Construction 2*	✓	✗	✗	✗
[HLVA02] Constructions 3,4	✗	✓	✗	✓
[HvAL08] NoState Construction	✓	✗	✗	✗
[vAH04] Construction 2	✓	✗	✓	✗
This work	✓	✓	✓	✓

Table 2: Comparison of universal steganography schemes. “Stateless” means that the sender and receiver do not need to share a synchronized state. “No unbiased function” means that the scheme does not require a function that is unbiased over the coartext distribution; instead, schemes with a check mark in this column rely on an assumption about the entropy of the text. As in Table 1, “robust” means robustness to a constant rate of substitutions.

* Although this scheme was also proposed by [AP98, Cac98], we refer to the construction from [HLVA02], whose setting and terminology we follow.

variants in the literature, including information-theoretic security [Cac98] and security against active attacks [BC05, Hop05]. We focus on (computational) *security against chosen message attacks* (CMA) as defined in [HLVA02], which is analogous to CPA security of encryption. There are several constructions of CMA-secure secret-key steganography schemes under certain assumptions, including [HLVA02, HvAL08, DIRR09]. We present a comparison of these schemes’ properties in Table 2. All of these schemes either make the very strong assumption that there is a known hash function that is *unbiased* on the coartext distribution, or else rely on a synchronized shared state between the sender and receiver.

An additional desirable property of a stegosystem, and one that is challenging to achieve, is *robustness*: Even if the stegotext is corrupted by an adversary, the receiver should be able to recover the message. To the best of our knowledge, prior to this work there was no provably secure *stateless* secret-key stegosystem with nontrivial robustness.

Our robustness notion is different from that of [HLVA02], called *substitution-robustness*. In substitution-robustness, an adversary may make substitutions of some symbols of the stegotext before it is given to the receiver, which should still recover the message. The set of substitutions that their adversary is allowed to make is parameterized by a relation R . That is, the adversary can change any symbol s to a symbol s' provided that $(s, s') \in R$. This definition breaks down when the alphabet is binary, since if R is nontrivial, containing without loss of generality $(0, 1)$, the adversary can change all stegotexts to the all-one string. Furthermore, it does not capture an adversary that has a bound on the number of symbols it may change, but that can change each symbol to any other symbol. Our definition is thus incomparable to theirs, and our channel may introduce deletions rather than just substitutions.

While [HLVA02] constructs a robust steganography scheme under this relation-based definition, they assume that the sender and receiver can share some state; in their scheme, this state is a synchronized counter N of the number of messages sent so far. See Section 2.7 for a discussion on synchronized states in steganography.

As with encryption, steganography has a public-key analogue, which was defined formally by [vAH04]. That is, encoding is possible with a public key, and decoding requires a secret key. While [vAH04] was the first to define and formally prove security of public-key steganography, public-key schemes existed in prior work [AP98, Cra98]. The main steganography construction in this work is public-key as well.

While we focus on steganography with provable security in the models of [HLVA02, vAH04], we note that there is a large body of work that constructs steganography schemes with heuristic guarantees. We refer the interested reader to [Fri09].

Language model steganography. The formal setting of steganography from [HLVA02] assumes that the sender has sample access to the coartext distribution. It is unclear how to realize this assumption

in practice: If the sender wishes to conceal its message in casual conversation, it must be able to sample a random casual conversation. One solution is to use a language model as this sampler for the covertext distribution, and there are several *language model steganography schemes* tailored to this setting where the sender interacts with a language model to craft its stegotexts [KJGR21, dWSK⁺22, ZDR19, Zam24]. This is a relaxation of universal steganography, since these language model steganography schemes leverage the sender’s explicit access to the covertext distribution via the model. Therefore, our universal steganography scheme in Section 8 is incomparable to these schemes.

Language model steganography and undetectable watermarks for language models are closely related, as both are concerned with secretly embedding a signal in the output of a language model. Indeed, we note in Section 7.2 that our watermarking scheme yields a language model steganography scheme, since undetectability implies steganographic secrecy. Our resulting steganography scheme has the strongest robustness of any existing scheme, and in particular is the only language model steganography scheme with robustness to a constant rate of random substitutions.

[Zam24] presents a language model steganography scheme derived from the watermarking scheme of [CGZ23]. This scheme is also stateless and relies on a minimal entropy assumption about the text, but it is not public-key or robust.

In Meteor [KJGR21], the sender and receiver share a generative language model, and they maintain a shared history of the prompt and all tokens output thus far by the model. To encode a message, the sender queries the model to obtain the distribution p_i over the next token. It then samples the next token in a way that encodes some information about the message. Because the sender and receiver share the model description and the token history *including the prompt*, the receiver can compute p_i . Meteor crucially uses the ability of the receiver to compute p_i , in order to decode the message. Any change to the text history at all (even removing the prompt) can destroy the receiver’s ability to compute p_i , and therefore the receiver’s ability to decode the message.

[dWSK⁺22] constructs a steganography scheme using minimum entropy couplings, in the information theoretic setting of [Cac98]. In their scheme, the decoder also must know the prompt used, as it requires access to an explicit description of the covertext distribution which is determined by the prompt. Furthermore, the adversary is not allowed to tamper with the stegotexts.

Backdoored/trapdoor PRGs. A *trapdoor* or *backdoored* PRG [VV83, DGG⁺15] is a pseudorandom generator that outputs a sequence of bits that are pseudorandom to an observer, but where a party holding a secret key can distinguish this sequence from random. In the context of backdoored PRGs, the secret key is viewed as a potential vulnerability of the PRG. A PRC is in particular a trapdoor/backdoored PRG, since codewords appear uniformly random to an outside observer but can be detected with a secret key.

However, a PRC comes with the additional property of robustness. This is especially interesting in the context of backdoored PRGs, where an *immunizer* may be applied to the PRG output in an attempt to thwart the adversary holding the backdoor. For example, one might apply a hash to the PRG output; [DGG⁺15] show that this immunizer is effective for certain PRGs. In that work they consider three classes of immunizers: public immunizers where the adversary can construct the trapdoor PRG based on knowledge of the seed of the function to be applied, semi-private immunizers where the adversary does not know the seed when constructing the PRG but does know it at attack time, and private immunizers where the adversary does not know the seed at all. Our PRCs show a strong impossibility result for immunizers against an adversary distinguishing PRG outputs from uniform randomness. For any immunizers in the class of channels that the PRC is robust to, the adversary can distinguish a single immunized output from uniform randomness with overwhelming probability. This applies even to private immunizers, since the randomness of the channel is not known to the code detection algorithm.