

A One-time Stegosystem and Applications to Efficient Covert Communication

Aggelos Kiayias^{*1}, Yona Raekow^{†2}, Alexander Russell^{‡1} and Narasimha Shashidhar^{§3}

¹Department of Computer Science and Engineering, University of Connecticut, Storrs, CT
`{akiayias, acr}@cse.uconn.edu`

²Fraunhofer Institute for Algorithms and Scientific Computing, St. Augustin, Germany
`yona.raekow@scai.fraunhofer.de`

³Department of Computer Science, Sam Houston State University, Huntsville, TX
`karpoor@shsu.edu`

Abstract

We present the first information-theoretic steganographic protocol with an asymptotically optimal ratio of key length to message length that operates on arbitrary covert distributions with constant min-entropy. Our results are also applicable to the computational setting: our stegosystem can be composed over a pseudorandom generator to send longer messages in a computationally secure fashion. In this respect our scheme offers a significant improvement in terms of the number of pseudorandom bits generated by the two parties in comparison to previous results known in the computational setting. Central to our approach for improving the overhead for general distributions is the use of combinatorial constructions that have been found to be useful in other contexts for derandomization: almost t -wise independent function families.

Keywords: Information hiding, steganography, data hiding, steganalysis, covert communication.

1 Introduction

Steganographic protocols enable one to “embed” covert messages into inconspicuous data over a public communication channel in such a way that no one, aside from the sender and the intended receiver, can even detect the presence of the secret message. The steganographic communication problem can be described using Simmons’ [13] formulation of the problem: In this scenario, prisoners Alice and Bob wish to communicate securely in the presence of an adversary, called the “Warden,” who monitors whether they exchange “conspicuous” messages. In particular, Alice and Bob may exchange messages that adhere to a certain channel distribution that represents “inconspicuous” communication. By controlling the messages that are transmitted over such a channel, Alice and Bob may exchange messages that cannot be detected by the Warden. There have been two approaches in formalizing this problem, one based on information theory [2, 15, 8] and one based on complexity theory [5]. Most steganographic constructions supported by provable security guarantees are instantiations of the following basic procedure (often referred to as “rejection-sampling”).

^{*}Supported by NSF CAREER grant CCR-0447808 and NSF grants CNS-0831304, CNS-0831306.

[†]This work was done while the author was in the Department of Computer Science and Engineering, University of Connecticut.

[‡]Supported by NSF CAREER grant CCR-0093065, and NSF grants CCR-0121277, CCR-0220264, CCR-0311368, and EIA-0218443.

[§]This work was done while the author was in the Department of Computer Science and Engineering, University of Connecticut.

The problem specifies a family of message distributions (the “channel distributions”) that provide a number of possible options for a so-called “covertext” to be transmitted. Additionally, the sender and the receiver possess some sort of private function indexed by the shared secret key (typically a keyed hash function, MAC, or other similar function) that maps channel messages to a single bit. In order to send a message bit m , the sender draws a covertext from the channel distribution, applies the function to the covertext and checks whether it happens to produce the “stegotext” m she originally wished to transmit. If this is the case, the covertext is transmitted. In case of failure, this procedure is repeated. While this is a fairly concrete procedure, there are a number of choices to be made with both practical and theoretical significance. From the security viewpoint, one is primarily interested in the choice of the function that is shared between the sender and the receiver. From a practical viewpoint, one is primarily interested in how the channel is implemented and whether it conforms to the various constraints that are imposed on it by the steganographic protocol specifications (e.g., are independent draws from the channel allowed? does the channel remember previous draws? etc.).

The shared key between Alice and Bob can be an expensive resource; to focus on this parameter, we define a notion of overhead equal to the ratio of the length of the secret-key to the length of the message. Prior work in statistically-secure steganography either gives high overhead or considers restricted covertext distributions. For instance, in the information-theoretic model, Cachin [2] demonstrated a steganographic protocol that works on restricted covertext distributions where the channel is a stationary distribution produced by a sequence of independent repetitions of the same experiment. Under this uniformity assumption, he uses sequences of covertexts to encode the message and obtains optimal overhead. In the complexity-theoretic setting, Hopper et al. [5, 6] provided a provably secure stegosystem that pairs rejection sampling with a pseudorandom function family to offer security for general (history-dependent) channel distributions with constant min-entropy. However, this protocol has a few drawbacks. First, casting their result in the information-theoretic setting, the length of the secret-key shared by Alice and Bob yields an overhead polynomial in the length of the message as this is the overhead required to share a suitable random function. In the complexity-theoretic setting, from an efficiency viewpoint, their construction required about 2 evaluations of a pseudorandom function per bit transmission. Constructing efficient pseudorandom functions is possible either generically [4] or, more efficiently, based on specific number-theoretic assumptions [10]. Nevertheless, pseudorandom function families are a conceptually complex and fairly expensive cryptographic primitive. For example, the evaluation of the Naor-Reingold pseudorandom function on an input x requires $O(|x|)$ modular exponentiations. Similarly, the generic construction [4] requires $O(k)$ PRG doublings of the input string where k is the length of the key.

Our protocol remedies these shortcomings. We show how it is possible to attain constant overhead for general channel distributions with constant min-entropy. The only assumptions employed in our analysis are merely that the channel alphabet is polynomial in the length of the message m and the security required is $2^{-|m|}$. Furthermore, our protocol in the computational setting is much more efficient: in particular, while the Hopper et al. stegosystem requires 2 evaluations *per bit* of a pseudorandom function, amounting to a linear (in the key-size) number of applications of the underlying PRG (in the standard construction for pseudorandom functions of [4]), in our stegosystem we require a constant number of PRG applications *per bit*. So the number of cryptographic operations per bit transmitted drops from linear to constant.

Central to our approach for improving the efficiency and overhead for general distributions is the use of combinatorial constructs such as almost t -wise independent function families given by Alon et al. [1]. Our protocol is based on the rejection sampling technique outlined above in combination with an explicit almost t -wise independent family of functions. We note that such combinatorial constructions have been extremely useful for derandomization methods and here, to the best of our knowledge, are employed for the first time in the design of steganographic protocols. The present paper is an extended version based on preliminary work that appeared in [7]; the present version includes a full security analysis that works for any constant min-entropy (as opposed to min-entropy of 1 bit that was assumed in this previous work).

2 Definitions and tools

The security of a steganography protocol is measured by the adversary’s ability to distinguish between “normal” and “covert” messages over a communication channel. To characterize normal communication we need to define and formalize the communication channel. We follow the standard terminology used in the literature [5, 2, 14, 6]: Let $\Sigma = \{\sigma_1, \dots, \sigma_s\}$ denote an alphabet and treat the *channel* as a family of random variables $\mathcal{C} = \{C_h\}_{h \in \Sigma^*}$; each C_h is supported on Σ . These channel distributions model a history-dependent notion of channel data that captures the notion of real-life communication. Such a channel induces a natural distribution on Σ^n for any n : σ_1 is drawn from C_ϵ , and each subsequent σ_i is drawn from $C_{\sigma_1 \dots \sigma_{i-1}}$. (Here we let ϵ denote the empty string.) Recall that the *min-entropy* of a random variable X , taking values in a set V , is the quantity

$$H_\infty(X) \triangleq \min_{v \in V} (-\log \Pr[X = v]) .$$

We say that a channel \mathcal{C} has min-entropy δ if for all $h \in \Sigma^*$, $H_\infty(C_h) \geq \delta$.

2.1 One-time stegosystems; the steganographic models

Steganography has been studied in two natural (but implicit up to now) communication models differing in Alice’s ability to sample from the channel.

Current History Model: The first model we study was that adopted by Hopper et al. [5]. In this model, Alice—and consequently the steganographic encoding protocol—has access to a channel oracle that provides samples from the channel *for the current history*. Alice is given no means of sampling from C_h for other histories. We call this the *current history* model. In this case, one can imagine that the channel is determined by a complex environment: while Alice is permitted to sample from the channel determined by the current environment, she cannot simulate potential future environments. Naturally, the communication history is updated when a symbol is transmitted on the wire from Alice to Bob. Formally, if $h_1, h_2, \dots, h_\ell \in \Sigma$ have been *transmitted* along the channel thus far, Alice may sample solely from $C_{h_1 \circ \dots \circ h_\ell}$ and send an element of her choice.

Look-Ahead Model: The second model we study—the *look-ahead* model—was adopted by von Ahn and Hopper [14]. This model is a relaxation of the “current history” model: Alice is now provided with a means for sampling “deep into the channel.” In particular, Alice and, consequently, the steganographic encoding protocol, has access to a channel oracle that can sample from the channel for *any* history. Formally, during the embedding process, Alice may sample from $C_{h_1 \circ \dots \circ h_\ell}$ for any future history $h = h_1 \circ \dots \circ h_\ell$ she wishes (though Alice is constrained to be efficient and so can make no more than polynomially many queries of polynomial length). This more generous model allows Alice to transform a channel C with min-entropy δ into a channel $C^{(\tau)}$ with min-entropy $\tau\delta$. Specifically, the channel $C^{(\tau)}$ is defined over the alphabet Σ^τ , whose elements we write as vectors $\mathbf{h} = (h_1, \dots, h_\tau)$. The distribution $C_{\mathbf{h}^1, \dots, \mathbf{h}^n}^{(\tau)}$ is determined by the channel C with history $\bar{h} = h_1^1 \circ \dots \circ h_\tau^1 \circ h_1^2 \circ \dots \circ h_\tau^n$. Below we give the definition of a *one-time stegosystem* that works in either of the above models and is a steganographic system that enables the one-time steganographic transmission of a message provided that the two parties share a suitable key.

Definition 1. A *one-time stegosystem* consists of three probabilistic polynomial time algorithms

$$S = (SK, SE, SD) ,$$

where:

- SK is the *key generation algorithm*; we write $SK(1^k) = \kappa$. It produces a key κ of length k .
- SE is the *embedding procedure* and has access to the channel; $SE(\kappa, m; \mathcal{O}) = s \in \Sigma^*$. The embedding procedure takes into account the history h of communication that has taken place between Alice and

Bob thus far and begins its operation corresponding to this history. It takes as input the key κ of length k , a message m of length $\ell = \ell(k)$ and a (probabilistic) oracle \mathcal{O} that allows SE to draw independent samples repeatedly from C_h in the current history model. In the look-ahead model, the oracle \mathcal{O} accepts as input a (polynomial-length) history $h' \in \Sigma^*$ and allows SE to draw independent samples repeatedly from $C_{h \circ h'}$. The output is the stegotext $s \in \Sigma^*$. Observe that in a one-time stegosystem, once a security parameter k is chosen, the length of the message ℓ is a fixed function of k . As described above, the access that SE has to the channel is dictated by the model of communication.

- SD is the *extraction procedure*; $SD(\kappa, c) = m$ or fail. It takes as input the key κ of length k , and some $c \in \Sigma^*$. The output is a message m or the token fail.

We next define a notion of correctness for a one-time stegosystem.

Definition 2 (Correctness). A one-time stegosystem (SK, SE, SD) is said to be (ϵ, δ) -correct provided that for all channels \mathcal{C} of min-entropy δ , it holds that $\forall h \in \Sigma^*$

$$\forall m \in \{0, 1\}^{\ell(k)} \Pr[SD(\kappa, SE(\kappa, m; \mathcal{O})) \neq m \mid \kappa \leftarrow SK(1^k)] \leq \epsilon .$$

In general, we treat both $\epsilon = \epsilon(k)$ and $\delta = \delta(k)$ as functions of k , the security parameter and the oracle \mathcal{O} as a function of the history h .

In the following paragraphs, we talk about the security for a one-time stegosystem. One-time stegosystem security is based on the indistinguishability between a transmission that contains a steganographically embedded message and a transmission that contains no embedded messages. The adversarial game discussed next is meant to model the behavior of a warden in the Simmons' formulation of the problem discussed earlier.

An adversary \mathcal{A} against a one-time stegosystem $S = (SK, SE, SD)$ is a pair of algorithms $\mathcal{A} = (SA_1, SA_2)$, that plays the following game, denoted $G^{\mathcal{A}}(1^k)$:

1. A key κ is generated by $SK(1^k)$.
2. Algorithm SA_1 receives as input the security parameter k and outputs a triple $(m, \text{aux}, h_c) \in M_\ell \times \{0, 1\}^* \times \Sigma^*$, where m is the challenge plaintext, h_c is the history of the channel that the adversary wishes to use for the steganographic embedding to start, and aux is some auxiliary information that will be passed to SA_2 . Note that SA_1 is provided access to \mathcal{C} via an oracle $\mathcal{O}(h)$, which takes the history h as input. $\mathcal{O}(\cdot)$, on input h , returns to SA_1 an element c selected according to C_h . This way, the warden can learn about the channel distribution for any history.
3. A bit b is chosen uniformly at random.
 - If $b = 0$ let $c^* \leftarrow SE(\kappa, m; \mathcal{O})$ where \mathcal{O} samples covertexts from C_{h_c} .
 - If $b = 1$ let $c^* = c_1 \circ \dots \circ c_\lambda$ where $\lambda = |SE(\kappa, m; \mathcal{O})|$ and $c_i \stackrel{r}{\leftarrow} C_{h_c \circ c_1 \circ \dots \circ c_{i-1}}$.
4. The input for SA_2 is $1^k, h_c, c^*$ and aux . SA_2 outputs a bit b' . If $b' = b$ then we say that (SA_1, SA_2) *succeeded* and write $G^{\mathcal{A}}(1^k) = \text{success}$.

The *advantage* of the adversary \mathcal{A} over a stegosystem S is defined as:

$$\mathbf{Adv}_S^{\mathcal{A}}(k) = \left| \Pr [G^{\mathcal{A}}(1^k) = \text{success}] - \frac{1}{2} \right| .$$

The probability includes the coin tosses of \mathcal{A} and SE , as well as the coin tosses of $G^{\mathcal{A}}(1^k)$. The (information-theoretic) insecurity of the stegosystem is defined as

$$\mathbf{InSec}_S(k) = \max_{\mathcal{A}} \{ \mathbf{Adv}_S^{\mathcal{A}}(k) \} ,$$

this maximum taken over all (time unbounded) adversaries \mathcal{A} .

Definition 3 (Security). We say that a stegosystem is (ϵ, δ) -secure if for all channels with min-entropy δ we have $\text{InSec}_S(k) \leq \epsilon$.

As above, in general we treat both $\epsilon = \epsilon(k)$ and $\delta = \delta(k)$ as functions of k , the security parameter.

Overhead. The *overhead* of a one-time stegosystem expresses the relation of the key length k and message length ℓ variables; specifically, we adopt the ratio $\beta = k/\ell$ as a measure for overhead.

This paper is an extended version of a previous abstract that appeared in [7]. The work presented by Kiayias et al. [7] considers the scenario where the communication channel has min-entropy at least 1 in the current history model. In this paper, we present steganography protocols both in the current history and the look-ahead model. We also present explicit constructions of error-correcting codes using Forney [3] concatenation scheme. Furthermore, our protocols operate on any communication channel with min-entropy $\delta > 0$.

2.2 Error-correcting codes

Our steganographic construction requires an efficient family of codes that can recover from errors introduced by certain binary symmetric channels. In particular, we require a version of the Shannon coding theorem [12, 11] that yields explicit control on the various parameters of the code as the rate approaches the capacity of the channel. We present this theorem in this section.

For an element $x \in \{0, 1\}^n$, we let $B_p(x)$ be the random variable equal to $x \oplus e$, where $e \in \{0, 1\}^n$ is a random error vector defined by independently assigning each $e_i = 1$ with probability p . (Here $x \oplus e$ denotes the vector with the i^{th} coordinate equal to $x_i \oplus e_i$.) The classical coding theorem [12] asserts that for every pair of real numbers $0 < R < C \leq 1$ and $n \in \mathbb{N}$, there is a binary code $A \subset \{0, 1\}^n$, with $\log |A|/n \geq R$ and $\theta \in \mathbb{R}$, so that for each $a \in A$, maximum-likelihood decoding recovers a from $B_p(a)$ with probability $1 - e^{-\theta \cdot n}$, where p is determined from C as

$$H(p) = p \log p^{-1} + (1 - p) \log(1 - p)^{-1} = 1 - C.$$

The quantity C is called the *capacity* of the binary symmetric channel and determines the random variable B_p ; the quantity $R = \log |A|/n$ is the *rate* of the code A . In this language, the coding theorem asserts that at transmission rates lower than the capacity of the channel, there exist codes that correct random errors with exponentially decaying failure probability (in n , the length of the code). We formalize our requirements below.

Definition 4. An error-correcting code of rate r is a pair of functions $\mathbf{E} = (\text{Enc}, \text{Dec})$, where $\text{Enc} : \{0, 1\}^{r \cdot n} \rightarrow \{0, 1\}^n$ is the *encoding algorithm* and $\text{Dec} : \{0, 1\}^n \rightarrow \{0, 1\}^{r \cdot n}$ the corresponding *decoding algorithm*. Specifically, we say that \mathbf{E} is a (r, p, ϵ) -code if for all $m \in \{0, 1\}^{r \cdot n}$,

$$\Pr[\text{Dec}(\text{Enc}(m) \oplus e) = m] \geq 1 - \epsilon,$$

where $e = (e_1, \dots, e_n)$ and each e_i is independently distributed in $\{0, 1\}$ so that $\Pr[e_i = 1] \leq p$. We say that \mathbf{E} is *efficient* if both Enc and Dec are computable in polynomial time in n .

We record the theorem below with the proof in Appendix A.

Theorem 1 (Based on Shannon [12, 11], Forney [3]). *For any $p \in [1/4, 1/2]$ and any $R < 1 - H(p)$, there is an efficient (R, p, ϵ) -code for which there is $\theta \in [0, 1]$, $n_0 \in \mathbb{N}$ such that $\epsilon \leq 2^{-\theta n / \log n}$ for any $n \geq n_0$. Furthermore, it holds that $\theta^{-1} = \Theta(Z^{-1})$ and $\log n_0 = \Theta(Z^{-2} \log Z^{-1})$ as $Z \rightarrow 0$ where $Z = 1 - H(p) - R$.*

2.3 Function families and almost t -wise independence

We will employ the notion of (almost) t -wise independent function families (cf., [1], [9]).

Definition 5. A family \mathcal{F} of Boolean functions on $\{0, 1\}^v$ is said to be ϵ -away from t -wise independent or (v, t, ϵ) -independent if for any t distinct domain elements q_1, q_2, \dots, q_t we have

$$\sum_{\alpha \in \{0, 1\}^t} \left| \Pr_f [f(q_1)f(q_2) \cdots f(q_t) = \alpha] - \frac{1}{2^t} \right| \leq \epsilon, \quad (1)$$

where f is chosen uniformly from \mathcal{F} .

The above is equivalent to the following formulation quantified over all computationally unbounded adversaries \mathcal{A} :

$$\left| \Pr_{f \leftarrow \mathcal{F}} [\mathcal{A}^{f^{[t]}}(1^v) = 1] - \Pr_{f \leftarrow \mathcal{R}} [\mathcal{A}^{f^{[t]}}(1^v) = 1] \right| \leq \epsilon, \quad (2)$$

where \mathcal{R} is the collection of *all* functions from $\{0, 1\}^v$ to $\{0, 1\}$ and $\mathcal{A}^{f^{[t]}}$ is an unbounded adversary that is allowed to determine up to t queries to the function f before it outputs its bit. The equivalence is formally stated (without proof) as follows.

Lemma 2. \mathcal{F}_κ is ϵ' -away from t -wise independence according to equation (1) if and only if \mathcal{F}_κ is ϵ' -away from t -wise independence according to equation (2) above.

We employ the construction of almost t -wise independent sample spaces given by Naor and Naor [9], and Alon et al. [1]. The following theorem is a restatement of theorem 3 from Alon et al. [1].

Theorem 3 ([9], [1]). *There exist families of Boolean functions $\mathcal{F}_{t, \epsilon}^v$ on $\{0, 1\}^v$ that are ϵ -away from t -wise independent, are indexed by keys of length $(2 + o(1))(\log v + t/2 + \log(\epsilon^{-1}))$, and are computable in polynomial time.*

2.4 Rejection sampling

A common method used in steganography employing a channel distribution is that of *rejection sampling*, described below (cf., [2, 5]).

Rejection sampling in the current history model In the current history model, assuming that one wishes to transmit a single bit m and employs a random function $f : \{0, 1\}^d \times \Sigma \rightarrow \{0, 1\}$ that is secret from the adversary, one performs the following “rejection sampling” process:

$\text{rejsam}_h^f(m)$ $c \leftarrow C_h$ if $f(c) \neq m$ then $c \leftarrow C_h$ Output: c

Here, Σ denotes the output alphabet of the channel, h denotes the history of the channel at the start of the process, and C_h denotes the distribution on Σ given by the channel with history h . The receiver (also privy to the function f) applies the function to the received message $c \in \Sigma$ and recovers m with probability greater than $1/2$. The sender and the receiver may employ a joint state (e.g., a counter), that need not be secret from the adversary. Note that the above process performs only two draws from the channel with the *same* history (more draws could, in principle, be performed, but we justify our choice of two draws in Lemma 8 of Section 3.1.2). These draws are assumed to be independent. One basic property of rejection sampling that we will prove and is helpful for our construction is the following:

Lemma 4. *If f is drawn uniformly at random from the collection of all functions $\mathcal{R} = \{f : \Sigma \rightarrow \{0, 1\}\}$ and \mathcal{C} has min-entropy δ , then*

$$\Pr_{f \leftarrow \mathcal{R}} [f(\text{rejsam}_h^f(m)) \neq m] \leq p,$$

where $p = (1 + 2^{-\delta})/4$.

Proof. Define the event E to be

$$E = [f(\mathbf{c}_1) = m] \vee [f(\mathbf{c}_1) \neq m \wedge f(\mathbf{c}_2) = m];$$

thus E is the event that rejection sampling is successful for m .

Here $\mathbf{c}_1, \mathbf{c}_2$ are two independent random variables distributed according to the channel distribution C_h and h is determined by the history of channel usage. Recalling that $\Sigma = \{\sigma_1, \dots, \sigma_s\}$ is the support of the channel distribution C_h , let $p_i = \Pr[C_h = \sigma_i]$ denote the probability that σ_i occurs. As f is chosen uniformly at random,

$$\Pr[f(\mathbf{c}_1) = m] = \frac{1}{2}.$$

Then $\Pr[E] = 1/2 + \Pr[A]$, where A is the event that $f(\mathbf{c}_1) \neq m \wedge f(\mathbf{c}_2) = m$. To bound $\Pr[A]$, let D denote the event that $\mathbf{c}_1 \neq \mathbf{c}_2$. Observe that conditioned on D , A occurs with probability exactly $1/4$; on the other hand, A cannot occur simultaneously with \bar{D} . Thus

$$\Pr[E] = \frac{1}{2} + \Pr[A | D] \cdot \Pr[D] + \Pr[A | \bar{D}] \cdot \Pr[\bar{D}] = \frac{1}{2} + \frac{1}{4} \Pr[D].$$

To bound $\Pr[D]$, note that

$$\Pr[\bar{D}] = \sum_i p_i^2 \leq \max_i p_i \sum_i p_i = \max_i p_i$$

and hence that $\Pr[D] \geq 1 - \max_i p_i$. Considering that $H_\infty(C) \geq \delta$, we have $\max_i p_i \leq 1/2^\delta$ and hence the success probability is

$$\Pr[E] \geq \frac{1}{2} + \frac{1}{4} \cdot (1 - \max_i p_i) \geq \frac{1}{2} + \frac{1}{4} \left(1 - \frac{1}{2^\delta}\right) = 1 - p,$$

from which the statement of the lemma follows. \square

The above lemma is a generalization of a similar result that appeared in [5] and dealt with the special case that the min-entropy is 1 bit. The application of the rejection sampling procedure as described above in our stegosystem implies that a message bit transmitted over the communication channel can be flipped with a certain probability p . This can be viewed as an overlaid binary symmetric channel with a cross-over probability p . The error-correcting code is introduced to recover from these cross-over errors.

Rejection sampling in the Look-Ahead model In the look-ahead model, the rejection sampling procedure above can be coupled with the channel transformation described in Section 2.1. In particular, transforming a channel C with min-entropy δ into a channel $C^{(\tau)}$ with min-entropy $\delta\tau$, one can carry out the rejection sampling process above with samples drawn from Σ^τ . The binary symmetric channel cross-over probability is then $p = (1 + 2^{-\delta\tau})/4$.

3 The construction

In this section we outline our construction of a one-time stegosystem as an interaction between Alice (the sender) and Bob (the receiver). First, we focus on the construction in the current history model and defer the discussion for the look-ahead model to Section 3.2.

3.1 Our stegosystem for the current history model

In the current history model, Alice and Bob wish to communicate over a channel with distribution \mathcal{C} over an alphabet Σ . We assume that \mathcal{C} has min-entropy δ , so that $\forall h \in \Sigma^*$, $H_\infty(C_h) \geq \delta$. As in the statement of Lemma 4, let $p = (1 + 2^{-\delta})/4$. The construction we describe below uses two parameters: $r \in (0, 1 - H(p))$ and $\epsilon_{\mathcal{F}} \in (0, 1)$. Alice and Bob agree on the following:

An error-correcting code. Let $E = (\text{Enc}, \text{Dec})$ be an efficient $(r, p, \epsilon_{\text{enc}})$ -code of length n from Theorem 1. The theorem asserts that p and r determine a bound for the decoding error probability ϵ_{enc} assuming that n is suitably large.

A pseudorandom function family. Let \mathcal{F} be the function family that is $(\log n + \log |\Sigma|, 2n, \epsilon_{\mathcal{F}})$ -independent indexed by keys of length k from Theorem 3. Recall from the theorem that $\epsilon_{\mathcal{F}}$ together with our choices of $v = \log n + \log |\Sigma|$ and $t = 2n$ determine the required key length k . We treat elements of \mathcal{F} as Boolean functions on $\{1, \dots, n\} \times \Sigma$ and, for such a function f we let $f_i : \Sigma \rightarrow \{0, 1\}$ denote the function $f_i(\sigma) = f(i, \sigma)$.

We will analyze the stegosystem below in terms of the parameters $n, r, \delta, \epsilon_{\mathcal{F}}$, relegating the discussion of how these parameters determine the overall efficiency, correctness and security of the system to Section 3.3.

Key generation consists of selecting an element $f \in \mathcal{F}$. This will be facilitated by sharing a random bit string κ of length k . Alice and Bob then communicate using the algorithms SE for embedding and SD for extracting as described in Figure 1.

PROCEDURE SE :	PROCEDURE SD :
Input: Key κ , hidden text m' , history h let $m = \text{Enc}(m')$ parse m as $m = m_1 m_2 \dots m_n$ for $i = 1$ to n { $c_i = \text{rejsam}_h^{f_i}(m_i)$ set $h \leftarrow h \circ c_i$ } Output: $c_{\text{stego}} = c_1 c_2 \dots c_n \in \Sigma^n$	Input: Key κ , stegotext c_{stego} parse c_{stego} as $c = c_1 c_2 \dots c_n$ for $i = 1$ to n { set $\bar{m}_i = f_i(c_i)$ } let $\bar{m} = \bar{m}_1 \bar{m}_2 \dots \bar{m}_n$ Output: $\text{Dec}(\bar{m})$

Figure 1: Encryption and Decryption algorithms for the one-time stegosystem in the current history model.

In SE , after applying the error-correcting code E , we use $\text{rejsam}_h^{f_i}(m_i)$ to obtain an element c_i of the channel for each bit m_i of the message and update the history h . The resulting stegotext $c_1 \dots c_n$ is denoted c_{stego} . In SD , the received stegotext is parsed block by block by evaluating the key function f_i at c_i ; this results in a message bit. After performing this for each received block, a message of size n is received, which is subjected to decoding via Dec . Note that we sample at most twice from the channel for each bit we wish to send. The error-correcting code is needed to recover from the errors introduced by this process. The detailed correctness, security and overhead analysis for both models follow in the next sections.

3.1.1 Correctness

In this section we argue about the correctness of our one-time stegosystem in the current history model. We examine the minimum message length needed for achieving (ϵ, δ) -correctness for any choice of these parameters. We are particularly interested in the case when δ is small (perhaps even approaching 0 as a function of k) as the difficulty of parameter selection is amplified in this case (in contrast when δ is bounded away from 0, the cross-over probability p is bounded away from $1/2$ and thus the parameter selection is simplified).

Theorem 5. *For any $\epsilon, \delta > 0$, consider the current history model stegosystem (SK, SE, SD) of Section 3.1 under the parameter constraints $r = \Omega(1 - H(p))$ and $\epsilon_{\mathcal{F}} \leq \epsilon/2$ where $p = (1 + 2^{-\delta})/4$. Then the stegosystem is (ϵ, δ) -correct so long as the message has length*

$$\Omega(\delta^{-2} \cdot \log(\epsilon^{-1}) \log \log(\epsilon^{-1})) + 2^{O(\delta^{-4})}$$

as $\delta \rightarrow 0$ while the dependency on δ vanishes when δ is bounded away from 0.

Proof. Let us first consider the case where the function f corresponding to the shared key between the two participants is a truly random function. In this case, by Lemma 4, the underlying communication channel simulates a binary symmetric channel with cross-over probability $p = (1 + 2^{-\delta})/4$. Based on this fact and Theorem 1, the probability of error in reception would be at most $2^{-\theta n / \log n}$ for sufficiently large n . Specifically, it should hold that $n \geq n_0$ for some n_0 that satisfies $\log n_0 = \Theta(Z^{-2} \log Z^{-1})$ where $Z = 1 - H(p) - r$. Also recall that $\theta = \Theta(Z)$. Given the statement of the theorem we can postulate that $Z = \Omega(1 - H(p))$ and as a result $Z^{-1} = O((1 - H(p))^{-1})$. Observe now that the choice of $p = (1 + 2^{-\delta})/4$ implies that

$$(1 - H(p))^{-1} = O((1 - 2^{-\delta})^{-2})$$

in the light of Proposition 15. It follows that $Z^{-1} = O((1 - 2^{-\delta})^{-2})$ and thus $n_0 = 2^{O((1-2^{-\delta})^{-4})} Z^{-1}$. From this we see that the minimum message length is of the form $2^{O((1-2^{-\delta})^{-4})}$ in order to attain an error correction bound of the form $2^{-\theta n / \log n}$. To force this latter function to be below, say, $\epsilon/2$ we need to select $n / \log n = \Omega(\theta^{-1} \log(1/\epsilon))$ which implies a lower bound for n of the form $\Omega((1 - 2^{-\delta})^{-2} \cdot \log(1/\epsilon) \log \log(1/\epsilon))$. The above guarantees an error of at most $\epsilon/2$ when the function f corresponding to the shared key between the two participants is a truly random function. Now we consider the case where the selection of f is based on an ϵ -away from t -wise independent family of functions. Given that the postulated distance of our function family from truly random functions is at most $\epsilon/2$, we see that

$$\forall m \in \{0, 1\}^\ell, \Pr[SD(\kappa, SE(\kappa, m; \Theta)) \neq m \mid \kappa \leftarrow SK(1^k)] \leq \epsilon$$

which establishes the correctness of the stegosystem for messages of length ℓ that are suitably large as postulated since $(1 - 2^{-\delta})^{-2} = O(\delta^{-2})$ for small values of $\delta \leq 1$ while for larger δ it holds that $(1 - 2^{-\delta})^{-2} = O(1)$. \square

3.1.2 Security

In this section we argue about the security of our one-time stegosystem in the current history model. First, we will observe that the output of the rejection sampling function rejsam_h^f , with a truly random function f , is indistinguishable from the channel distribution C_h (this folklore result was implicit in previous work—we prove it formally below). We then show that if f is selected from a family that is $\epsilon_{\mathcal{F}}$ -away from $2n$ -wise independent, the advantage of an adversary \mathcal{A} to distinguish between the output of the steganographic embedding protocol SE and the channel C_h is bounded above by $\epsilon_{\mathcal{F}}$. Let $\mathcal{R} = \{f : \Sigma \rightarrow \{0, 1\}\}$. We will show the following:

Theorem 6. *For any $\epsilon, \delta > 0$, consider the current history stegosystem (SK, SE, SD) of Section 3.1 under the parameter constraint $\epsilon_{\mathcal{F}} \leq \epsilon$. Then the stegosystem is (ϵ, δ) -secure so long as the key has length*

$$(2 + o(1)) \left(\frac{1}{r} \cdot \ell + \log \log |\Sigma| + \log(\epsilon^{-1}) \right),$$

where ℓ is the message length and r is the rate of the error correcting code employed by the stegosystem.

Anticipating the proof of the theorem we start with some preliminary results. First, we characterize the probability distribution of the rejection sampling function:

Proposition 7. *Fix some function $f : \Sigma \rightarrow \{0, 1\}$ and channel history $h \in \Sigma^*$. The function $\text{rejsam}_h^f(m)$ is a random variable with probability distribution expressed by the following function: Let $c \in \Sigma$ and $m \in \{0, 1\}$. Let $\text{miss}_f(m) = \Pr_{c' \leftarrow C_h}[f(c') \neq m]$ and $p_c = \Pr_{c' \leftarrow C_h}[c' = c]$. Then*

$$\Pr[\text{rejsam}_h^f(m) = c] = \begin{cases} p_c \cdot (1 + \text{miss}_f(m)) & \text{if } f(c) = m, \\ p_c \cdot \text{miss}_f(m) & \text{if } f(c) \neq m. \end{cases}$$

Proof. Let c_1 and c_2 be the two (independent) samples drawn from C_h during rejection sampling. (For simplicity, we treat the process as having drawn two samples even in the case where it succeeds on the

first draw.) Note, now, that in the case where $f(c) \neq m$, the value c is the result of the rejection sampling process precisely when $f(c_1) \neq m$ and $c_2 = c$; as these samples are independent, this occurs with probability $\text{miss}_f(m) \cdot p_c$. In the case where $f(c) = m$, however, we observe c whenever $c_1 = c$ or $f(c_1) \neq m$ and $c_2 = c$. As these events are disjoint, their union occurs with probability $p_c \cdot (\text{miss}_f(m) + 1)$, as desired. \square

Lemma 8. *For any $h \in \Sigma^*$, $m \in \{0, 1\}$, the random variable $\text{rejsam}_h^f(m)$ is perfectly indistinguishable from the channel distribution C_h when f is drawn uniformly at random from the space of \mathcal{R} .*

Proof. Let f be a random function, as described in the statement of the lemma. Fixing the elements c , and m , we condition on the event E_{\neq} , that $f(c) \neq m$. In light of Proposition 7, for any f drawn under this conditioning we shall have that $\Pr[\text{rejsam}_h^f(m) = c]$ is equal to

$$\Pr_{c' \leftarrow \mathcal{C}_h} [c' = c] \cdot \text{miss}_f(m) = p_c \cdot \text{miss}_f(m) ,$$

where we have written $\text{miss}_f(m) = \Pr_{c' \leftarrow \mathcal{C}_h} [f(c') \neq m]$ and $p_c = \Pr_{c' \leftarrow \mathcal{C}_h} [c' = c]$. Conditioned on E_{\neq} , then, the probability of observing c is

$$\mathbf{E}_f [p_c \cdot \text{miss}_f(m) \mid E_{\neq}] = p_c \left(p_c + \frac{1}{2}(1 - p_c) \right) ,$$

where the above follows from the fact that in the conditional space we can expand $\text{miss}_f(m)$ as

$$\Pr_{c' \leftarrow \mathcal{C}_h} [f(c') \neq m \mid c' = c \wedge E_{\neq}] \cdot \Pr_{c' \leftarrow \mathcal{C}_h} [c' = c \mid E_{\neq}] + \Pr_{c' \leftarrow \mathcal{C}_h} [f(c') \neq m \mid c' = c \wedge E_{=}] \cdot \Pr_{c' \leftarrow \mathcal{C}_h} [c' \neq c \mid E_{\neq}] .$$

Letting $E_{=}$ be the event that $f(c) = m$, we similarly compute

$$\mathbf{E}_f [p_c \cdot (1 + \text{miss}_f(m)) \mid E_{=}] = p_c \left(1 + \frac{1}{2}(1 - p_c) \right) .$$

As $\Pr[E_{=}] = \Pr[E_{\neq}] = 1/2$, we conclude that the probability of observing c is exactly

$$\frac{1}{2} \left(p_c \left(p_c + \frac{1 - p_c}{2} \right) + p_c \left(1 + \frac{1 - p_c}{2} \right) \right) = p_c ,$$

as desired. \square

Having established the behavior of the rejection sampling function when a truly random function is used, we proceed to examine the behavior of rejection sampling in our setting where the function is drawn from a function family that is $\epsilon_{\mathcal{G}}$ -away from $2n$ -wise independent. In particular we will show that the insecurity of the defined stegosystem is characterized as follows:

Proof of Theorem 6. Consider the following two games $G_1^{\mathcal{A}}$ and $G_2^{\mathcal{A}}$ that can be played with the adversary \mathcal{A} . Here $\lambda = |SE(\kappa, m; \mathcal{O})|$.

$G_1^{\mathcal{A}}(1^k)$	
1.	$\kappa \leftarrow \{0, 1\}^k$
2.	$(m^*, s) \leftarrow SA_1^{\mathcal{O}(h)}(1^\kappa, h)$, $m^* \in \{0, 1\}^\ell$
3.	$b \xleftarrow{r} \{0, 1\}$
4.	$c^* = \begin{cases} c_0, c_1, \dots, c_{\lambda-1} & \text{if } b = 0 \\ \text{from the channel with history } h & \text{if } b = 1 \end{cases}$ $c_i = \text{rejsam}_h^{f_{\kappa, i}}(m_i)$, $h = h \circ c_i$
5.	$b^* \leftarrow SA_2(c^*, s)$
6.	if $b = b^*$ then success

$G_2^A(1^k)$	
1.	$f \leftarrow \mathcal{R}$
2.	$(m^*, s) \leftarrow SA_1^{\mathcal{O}(h)}(1^\kappa, h), m^* \in \{0, 1\}^\ell$
3.	$b \xleftarrow{r} \{0, 1\}$
4.	$c^* = \begin{cases} c_0, c_1, \dots, c_{\lambda-1} & c_i = \text{rejsam}_h^{f,i}(m_i), h = h \circ c_i \\ \text{from the channel with history } h & \text{if } b = 0 \\ & \text{if } b = 1 \end{cases}$
5.	$b^* \leftarrow SA_2(c^*, s)$
6.	if $b = b^*$ then success

$$\begin{aligned} \text{Adv}_S^A(k) &= \left| \Pr[G^A(1^k) = \text{success}] - \frac{1}{2} \right| \\ &= \left| \Pr[G_1^A(1^k) = \text{success}] - \Pr[G_2^A(1^k) = \text{success}] \right| \leq \epsilon_{\mathcal{F}} \end{aligned}$$

and the theorem follows by the definition of insecurity. From Theorem 3 we get that the minimum key length required for security is $(2 + o(1))(\ell/r + \log \log |\Sigma| + \log(\epsilon^{-1}))$, where ℓ is the message length and r is the rate of the error correcting code employed by the stegosystem. \square

3.2 Adapting to the look-ahead model

In this section we note the differences in the construction for the look-ahead model from the current history model. In this model, Alice and Bob agree to communicate over a channel with distribution C_h^τ over an alphabet Σ^τ where $\tau = \delta^{-1}$. The min-entropy is now $H_\infty(C_h^{(\delta^{-1})}) \geq 1$. The binary symmetric channel cross-over probability p is no more than $3/8$. To recover from the cross-over error, they use the error-correcting code $E = (\text{Enc}, \text{Dec})$ which is an efficient $(r, 3/8, \epsilon_{\text{enc}})$ -code of length n from Theorem 1. For the look-ahead model, we record the corollary below which follows directly from Theorem 5.

Corollary 9. *For any $\epsilon, \delta > 0$, consider the look-ahead model stegosystem (SK, SE, SD) under the parameter constraints $r = \Omega(1 - H(p))$ and $\epsilon_{\mathcal{F}} \leq \epsilon/2$ where $p = 3/8$. Then the stegosystem is (ϵ, δ) -correct so long as the message has length $\Omega(\log(1/\epsilon) \log \log(1/\epsilon))$.*

In the current history model, as $\delta \rightarrow 0$, the rejection sampling procedure has a high probability of failure. This is because the overlaid binary symmetric channel cross-over probability converges to a $1/2$ very quickly since $p = (1 + 2^{-\delta})/4$. With p converging to $1/2$, the binary symmetric channel becomes informationless. Consequently, we would have to employ error correcting codes that can recover from very high error rates. This translates to a very high minimum message length requirement and explains the exponential dependence on δ^{-1} . In the look-ahead model, we amplify the entropy of the channel upto 1, thereby removing the minimum message length's exponential dependence on δ^{-1} . In either model, if we want to transmit a message of length shorter than the minimum message length, this can be accomplished by padding the original message to attain the required length. The rejection sampling procedure in the two models only differ in the size of their domain space. Observe that this does not affect the security analysis. The corollary recorded below follows directly from Theorem 6:

Corollary 10. *For any $\epsilon, \delta > 0$, consider the look-ahead stegosystem (SK, SE, SD) under the parameter constraint $\epsilon_{\mathcal{F}} \leq \epsilon$. Then the stegosystem is (ϵ, δ) -secure so long as the key has length $(2 + o(1))(\ell/r + \log \log(|\Sigma| \cdot \delta^{-1}) + \log(1/\epsilon))$, where ℓ is the message length and r is the rate of the error correcting code employed by the stegosystem.*

3.3 Putting it all together

The objective of this section to combine the results of the previous sections and illustrate the results for our stegosystem in the two channel models. As our system is built over two-sample rejection sampling, a

process that faithfully transmits each bit with cross-over probability $p = (1 + 2^{-\delta})/4$, the target rate that we may approximate is $1 - H(p)$. In the case of the look-ahead model, we have the cross-over probability $p \leq 3/8$ and the target rate that we may approximate is $1 - H(3/8)$. Indeed, as described below, the system asymptotically converges to the rate of this underlying rejection sampling channel. We remark that with sufficiently large channel entropy, there are ways for one to draw more samples during rejection sampling to reduce the error-rate without compromising security, but nevertheless this would not have any (asymptotic) bearing to our overhead objective.

Theorem 11. *For any $\epsilon, \delta > 0$, the stegosystem (SK, SE, SD) of Section 3.1 in the current history model under the parameter constraints $r = \Omega(1 - H(p))$ and $\epsilon_{\mathcal{F}} \leq \epsilon/2$ where $p = (1 + 2^{-\delta})/4$ is (ϵ, δ) -correct and (ϵ, δ) -secure so long as : the message has length $\Omega(\delta^{-2} \cdot \log(1/\epsilon) \log \log(1/\epsilon)) + 2^{O(\delta^{-2})}$ as $\delta \rightarrow 0$ while the dependency on δ vanishes for $\delta \rightarrow \infty$. When the size of the channel alphabet is polynomial in the length of the message m and $\epsilon = 2^{-|m|}$, (SK, SE, SD) has overhead $O((1 - H(p))^{-1}) = O(\delta^{-2})$ as $\delta \rightarrow 0$ while the dependency on δ vanishes for $\delta \rightarrow \infty$.*

The above theorem implies that for any fixed δ our stegosystem exhibits $O(1)$ overhead, i.e., the ratio of the key length over the message length is constant.

Theorem 12. *For any $\epsilon, \delta > 0$, the stegosystem (SK, SE, SD) of Section 3.1 in the look-ahead model under the parameter constraints $r = \Omega(1 - H(p))$ and $\epsilon_{\mathcal{F}} \leq \epsilon/2$ where $p = 3/8$ is (ϵ, δ) -correct and (ϵ, δ) -secure so long as : the message has length $\Omega(\log(1/\epsilon) \log \log(1/\epsilon))$. When the size of the channel alphabet is polynomial in the length of the message m and $\epsilon = 2^{-|m|}$, (SK, SE, SD) exhibits $O(1)$ overhead.*

4 A provably secure stegosystem for longer messages

In this section we show how to apply the “one-time” stegosystem of Section 3.1 together with a pseudorandom generator so that longer messages can be transmitted.

Definition 6. Let U_k denote the uniform distribution over $\{0, 1\}^k$. A polynomial time deterministic algorithm G is a pseudorandom generator (PRG) if the following conditions are satisfied:

Variable output For all seeds $x \in \{0, 1\}^*$ and $y \in \mathbb{N}$, $|G(x, 1^y)| = y$.

Pseudorandomness For every polynomial p the set of random variables $\{G(U_k, 1^{p(k)})\}_{k \in \mathbb{N}}$ is computationally indistinguishable from the uniform distribution $\{U_{p(k)}\}_{k \in \mathbb{N}}$.

For a PRG G and $0 < k < k'$, if A is some statistical test, we define the advantage of A over the PRG as follows:

$$\mathbf{Adv}_G^A(k, k') = \left| \Pr_{w \leftarrow G(U_k, 1^{k'})} [A(w) = 1] - \Pr_{w \leftarrow U_{k'}} [A(w) = 1] \right|.$$

The insecurity of the above PRG G against all statistical tests A computable by circuits of size $\leq P$ is then defined as

$$\mathbf{InSec}_G(k, k'; P) = \max_{A \in \mathcal{A}_P} \{\mathbf{Adv}_G^A(k, k')\}$$

where \mathcal{A}_P is the collection of statistical tests computable by circuits of size $\leq P$.

It is convenient for our application that typical PRGs have a procedure G' such that if $z = G(x, 1^y)$, it holds that $G(x, 1^{y+y'}) = G'(x, z, 1^{y'})$ (i.e., if one maintains z , one can extract the y' bits that follow the first y bits without starting from the beginning).

Consider now the following stegosystem $S' = (SK', SE', SD')$ that can be used for steganographic transmission of longer messages using the one-time stegosystem $S = (SK, SE, SD)$ defined in Section 3.1. S' can handle messages of length polynomial in the security parameter k and employs a PRG G . The two

players Alice and Bob, share a key of length k denoted by x . The function SE' is given input x and the message $m \in \{0, 1\}^\nu$ to be transmitted of length $\nu = p(k)$ for some fixed polynomial p . SE' in turn employs the PRG G to extract k' bits (it computes $\kappa = G(x, 1^{k'})$, $|\kappa| = k'$). The length k' is selected to match the number of key bits that are required to transmit the message m using the one-time stegosystem of Section 3.1. Once the key κ of length k' is produced by the PRG, the procedure SE' invokes the one-time stegosystem on input κ, m, h . The function SD' is defined in a straightforward way based on SD .

The computational insecurity of the stegosystem S' is defined by adapting the definition of information-theoretic stegosystem security from Section 2.1 for the computationally bounded adversary as follows:

$$\mathbf{InSec}_{S'}(k, k'; P) = \max_{\mathcal{A} \in \mathcal{A}_P} \{\mathbf{Adv}_{S'}^{\mathcal{A}}(k, k')\},$$

this maximum taken over all adversaries \mathcal{A} , where SA_1 and SA_2 have circuit size $\leq P$ and the definition of advantage $\mathbf{Adv}_{S'}^{\mathcal{A}}(k, k')$ is obtained by suitably modifying the definition of $\mathbf{Adv}_S^{\mathcal{A}}(k)$ in Section 2.1. In particular, we define a new adversarial game $G^{\mathcal{A}}(1^k, 1^{k'})$ which proceeds as the previous game $G^{\mathcal{A}}(1^k)$ in Section 2.1 except that in this new game $G^{\mathcal{A}}(1^k, 1^{k'})$, algorithms SA_1 and SA_2 receive as input the security parameter k' and SE' invokes SE as $SE(\kappa, m; \mathcal{O})$ where $\kappa = G(x, 1^{k'})$. This matches the model of [5] which referred to such schemes as steganographically secret against chosen hidtext attacks.

Theorem 13. *The stegosystem $S' = (SK', SE', SD')$ is steganographically secret against chosen hidtext attacks. In particular employing a PRG G to transmit a message m we get*

$$\mathbf{InSec}_{S'}(k, k'; P) \leq \mathbf{InSec}_G(k, k'; P) + \mathbf{InSec}_{S'}(k')$$

where $\mathbf{InSec}_{S'}(k')$ is the information-theoretic insecurity defined in Section 2.1 and $|m| = \ell(k')$.

Performance comparison of the stegosystem S' and the Hopper, Langford, von Ahn system.

The system of Hopper et al. [5] concerns a current history model where the min-entropy of all C_h is at least 1. In this case, we may select an $(p, 3/8, \epsilon_{\text{enc}})$ -error-correcting code. Then the system of Hopper, et al. correctly decodes a given message with probability at least $1 - \epsilon_{\text{enc}}$ and makes no more than $2n$ calls to a pseudorandom function family. Were one to use the pseudorandom function family of Goldreich et al. [4], then this involves production of $\Theta(\ell \cdot \log(\ell \cdot |\Sigma|))$ pseudorandom bits, where ℓ is the message length. Of course, the security of the system depends on the security of the underlying pseudorandom generator with parameter k . On the other hand, with the same error-correcting code, the steganographic system described in this work utilizes $O(\ell + \log \log |\Sigma| + \log(1/\epsilon))$ pseudorandom bits, correctly decodes a given message with probability $1 - \epsilon$, while it possesses insecurity no more than ϵ . In order to compare the two schemes, note that by selecting $\epsilon = 2^{-k}$, both the decoding error and the security of the two systems differ by at most 2^{-k} , a negligible function in terms of the security parameter k . (Note also that pseudorandom functions utilized in the above scheme have security no better than 2^{-k} with security parameter k .) In this case, the number of pseudorandom bits used by our system is

$$(2 + o(1))(\ell + \log \log |\Sigma| + k) = \Theta(\ell + k + \log \log |\Sigma|),$$

a non-trivial improvement over the $\Theta(\ell \cdot \log(\ell \cdot |\Sigma|))$ bits of the scheme above. In the look-ahead model, the number of pseudorandom bits used by our system is $\Theta(\ell + \log \log(|\Sigma| \cdot \delta^{-1}) + k)$ as we operate on the concatenated channel $C_h^{(\delta^{-1})}$.

References

- [1] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple construction of almost k -wise independent random variables. *Random Struct. Algorithms*, 3(3):289–304, 1992.
- [2] Christian Cachin. An information-theoretic model for steganography. *Inf. Comput.*, 192(1):41–56, 2004.

- [3] G. D. Forney, Jr. *Concatenated Codes*. Research Monograph No. 37. MIT Press, 1966.
- [4] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- [5] Nicholas J. Hopper, John Langford, and Luis von Ahn. Provably secure steganography. In *CRYPTO*, pages 77–92, 2002.
- [6] Nicholas J. Hopper, Luis von Ahn, and John Langford. Provably secure steganography. *IEEE Trans. Computers*, 58(5):662–676, 2009.
- [7] Aggelos Kiayias, Yona Raekow, and Alexander Russell. Efficient steganography with provable security guarantees. In Mauro Barni, Jordi Herrera-Joancomartí, Stefan Katzenbeisser, and Fernando Pérez-González, editors, *Information Hiding*, volume 3727 of *Lecture Notes in Computer Science*, pages 118–130. Springer, 2005. ISBN 3-540-29039-7.
- [8] Thomas Mittelholzer. An information-theoretic approach to steganography and watermarking. In *Information Hiding*, pages 1–16, 1999.
- [9] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM J. Comput.*, 22(4):838–856, 1993.
- [10] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. *J. ACM*, 51(2):231–262, 2004.
- [11] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423 and 623–656, July and October, 1948.
- [12] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, Illinois, 1949.
- [13] Gustavus J. Simmons. The prisoners’ problem and the subliminal channel. In *CRYPTO*, pages 51–67, 1983.
- [14] Luis von Ahn and Nicholas J. Hopper. Public-key steganography. In *Advances in Cryptology Proceedings of Eurocrypt 04*, pages 323–341. Springer-Verlag, 2004.
- [15] Jan Zöllner, Hannes Federrath, Herbert Klimant, Andreas Pfitzmann, Rudi Piotraschke, Andreas Westfeld, Guntram Wicke, and Gritta Wolf. Modeling the security of steganographic systems. In *Information Hiding*, pages 344–354, 1998.

A Error-correcting codes

In this section, we provide the proof for Theorem 1 from Section 2.2.

Theorem 14. (based on Shannon [12, 11], Forney [3]) For any $p \in [1/4, 1/2]$ and any $R < 1 - H(p)$, there is an efficient (R, p, ϵ) -code for which there is $\theta \in [0, 1]$, $n_0 \in \mathbb{N}$ such that $\epsilon \leq 2^{-\theta n / \log n}$ for any $n \geq n_0$. Furthermore, it holds that $\theta^{-1} = \Theta(Z^{-1})$ and $\log n_0 = \Theta(Z^{-2} \log Z^{-1})$ as $Z \rightarrow 0$ where $Z = 1 - H(p) - R$.

Proof. We provide the details below of the classic construction for such error-correcting codes E based on concatenated codes. In the standard notation used for codes, q stands for the alphabet size, n for the block length, k for the message length (where $|C| = q^k$) and d for the minimum distance of the code. A code with the above parameters is called an $(n, k, d)_q$ code and an $[n, k, d]_q$ code if it is linear. Thus, k/n is the rate of the code and d/n is the relative distance.

Here, we take advantage of the concatenation $C_1 \cdot C_2$ of two codes C_1 and C_2 , called the “outer” and “inner” code respectively. The procedure is as follows:

$$\text{Outer Code : } C_1 = (N, K, D)_Q$$

$$\text{Inner Code : } C_2 = (n, k, d)_q$$

$$\text{Such that : } Q = q^k .$$

The alphabet of the outer code is in one-to-one correspondence with the codewords of the inner code. Given a message in Q^K , first apply C_1 onto this message to get a string $s \in Q^N$. Then, on every symbol of s , viewed as a message in q^k , apply C_2 and concatenate the results. The resulting codeword is $C_2(s_1) \cdot C_2(s_2) \cdots C_2(s_N)$ and the resulting code has $Q^K = (q^k)^K = q^{kK}$ messages. The length of the codewords is nN and the distance of concatenation is at least dD . Hence,

$$C_1 \cdot C_2 = (N, K, D)_{q^k} \cdot (n, k, d)_q \Rightarrow (nN, kK, dD)_q .$$

This operation was due to Forney [3]. We now show how we implement Forney’s code for constructing an asymptotically good code.

The inner code. In the inner code schema, we will be transmitting binary strings of length $n_1 = c \cdot \log n$ over a binary symmetric channel with cross-over probability $p \in [1/4, 1/2]$.

The encoding schema uses a set \mathcal{C} of $2^{\lfloor rn_1 \rfloor}$ random codewords drawn from $\{0, 1\}^{n_1}$ where r is a parameter to be determined later that corresponds to the rate of the inner code. These codewords are mapped arbitrarily to the elements of $\{0, 1\}^{\lfloor rn_1 \rfloor}$. The decoding procedure is a maximum likelihood decoder: given a received word, the message corresponding to the codeword closest to it, is determined and returned, with ties broken arbitrarily.

We next analyze the probability the decoding procedure fails. In what follows $e \in \{0, 1\}^{n_1}$ is selected at random such that $\Pr[e_i = 1] = p$; we would like to bound the probability $\Pr[c_i \oplus e \text{ is closest to } c_i]$ from below. Note that $\Pr[c_i \oplus e \text{ is closest to } c_i] = \Pr[d(c_i \oplus e, c_i) < d(c_i \oplus e, c_j)]$ for all c_j with $j \neq i$, where d is the Hamming distance. Without loss of generality, we can let $c_i = \mathbf{0}$ and proceed as follows. $\Pr[e \text{ is closest to } \mathbf{0}] = 1 - \Pr[\text{Some } c \in \mathcal{C} \text{ is closer to } \vec{e} \text{ than } \mathbf{0}]$. To this end, let us first proceed to upper bound the probability $\Pr[\text{Some } c \in \mathcal{C} \text{ is closer to } \vec{e} \text{ than } \vec{0}]$ which is the decoding error probability of the inner code.

Let $|B_w(x)|$ be the set of words y with $d(x, y) \leq w$. The set $B_w(x)$ is called the ball with radius w and center x . $|B_w(x)| = \sum_{0 \leq k \leq w} \binom{n_1}{k}$. Let α be a constant with $1/2 > \alpha > p$.

$$\begin{aligned} & \Pr[\text{Some } c \in \mathcal{C} \text{ is closer to } \vec{e} \text{ than } \mathbf{0}] \\ &= \sum_{i=0}^{n_1} p^i (1-p)^{n_1-i} \binom{n_1}{i} \Pr[\text{Some } c \in B_i(\vec{e}) \mid \vec{e}] \\ &\leq \Pr[|e| \geq \alpha n_1] + \Pr[\text{Some } c \in B_{\alpha n_1}(\mathbf{0})] \\ &\leq \Pr[|e| \geq \alpha n_1] + \frac{2^{\lfloor rn_1 \rfloor} \cdot |B_{\alpha n_1}(\mathbf{0})|}{2^{n_1}} . \end{aligned}$$

Here, $|B_i(x)|$ is the set of words y with $d(x, y) \leq i$. Then, for $0 \leq \alpha \leq 1/2$,

$$|B_{\alpha n_1}(x)| = \sum_{0 \leq k \leq \alpha n_1} \binom{n_1}{k} \leq 2^{n_1 H(\alpha)} .$$

Thus,

$$\Pr[\text{Some } c \in \mathcal{C} \text{ is closer to } \vec{e} \text{ than } \mathbf{0}] \leq \Pr[|e| \geq \alpha n_1] + 2^{(r-1+H(\alpha))n_1} < e^{\frac{-n_1(\alpha-p)^2}{3p}} + 2^{-(1-r-H(\alpha))n_1} .$$

Here the final inequality follows from the Chernoff bound.¹ The above indicates that as long as we choose α, r such that $\alpha > p$ and $r < 1 - H(\alpha)$ it holds that the error probability of decoding in the inner code is at most ϵ_1 provided that

$$n_1 \geq \max \left\{ 3p \ln(2\epsilon_1^{-1})(\alpha - p)^{-2}, \log(2\epsilon_1^{-1})(1 - H(\alpha) - r)^{-1} \right\}.$$

The outer code. We next describe the outer code that is a Reed-Solomon code over the binary extension field $GF(2^{n_1})$. The code has length $n_2 < 2^{n_1}$ symbols and is of rate κ . We can correct a number of errors up to a rate of $(1 - \kappa)/2$. We want to ensure that we can correct the message with probability $1 - \epsilon$. Let u be the number of errors of the inner code. The expected value of u is $\epsilon_1 n_2$. We would like to bound the probability $p' = \Pr[u > (1 - \kappa)n_2/2]$. Applying Chernoff bound and setting $\zeta = \epsilon_1^{-1}(1 - \kappa)/2 - 1$ we have $p' < e^{-\epsilon_1 n_2 (\epsilon_1^{-1}(1 - \kappa)/2 - 1)^2/3} \leq \epsilon$ provided that

$$n_2 \geq \ln(\epsilon^{-1})\epsilon_1^{-1} \cdot (\epsilon_1^{-1}(1 - \kappa)/2 - 1)^{-2}.$$

We may decode the Reed-Solomon code using the Berlekamp-Welch algorithm and in the event that the error rate is greater than $(1 - \kappa)/2$, the output of the decoding procedure will be a failure symbol \perp .

The concatenation construction. The objective is to transmit with a given transmission rate $R < 1 - H(p)$ for a transmission length of n bits while having a decoding error probability of $\epsilon = e^{-\theta n / \log n}$ for suitable θ (that is independent of n). The recovery of the $R \cdot n$ bits of message should be achieved in time polynomial in n . Since we employ a concatenated code with inner-code rate r and the outer code rate κ , we can achieve our objective provided that $r = 1 - H(p) - f_1$, $\kappa = 1 - f_2$ as long as $f_1 + f_2 \leq Z$ where $Z = 1 - H(p) - R$ for some suitably chosen values $f_1, f_2 \in [0, 1]$.

To account for the fact that the decoding of the inner code is done in a brute-force manner we need that $n_1 \leq c_1 \cdot 1/r \log n$ for some constant c_1 (as in this case we maintain polynomial-time dependency in n in terms of running time).

Let $g_1 < f_1$ be some constant; we set $\alpha = H^{-1}(H(p) + g_1)$. Observe that $\alpha > p$ and $\alpha < 1/2$, i.e., this is a valid choice for the selection of α in the inner code design. Furthermore, we have that

$$1 - H(\alpha) - r = 1 - H(p) - g_1 - (1 - H(p) - f_1) = f_1 - g_1 \text{ and } \alpha - p \geq \frac{g_1}{2}.$$

This latter inequality follows from the fact that $H(p) + z \geq H(p + z/2)$ for all $p \in [1/4, 1/2]$ and $z \in [0, 1]$. Based on the above we may restate the bounds on n_1 as follows:

$$n_1 \geq \max \left\{ 5 \cdot \log(2\epsilon_1^{-1})g_1^{-2}, \log(2\epsilon_1^{-1})(f_1 - g_1)^{-1} \right\}.$$

We set $g_1 = f_1/2$ and $f_1 = Z/2$ and we obtain the requirement that $n_1 = \Omega(Z^{-2} \log \epsilon_1^{-1})$. Next observe that $(1 - \kappa)/2\epsilon_1 - 1 = f_2/(2\epsilon_1) - 1$, by setting $f_2 = 2\epsilon_1(1 + g_2)$ for some constant g_2 , the bound on n_2 can be expressed as

$$n_2 \geq \ln(\epsilon^{-1})\epsilon_1^{-1}g_2^{-2}.$$

Note that assuming $f_2 = Z/2$ we obtain that $g_2 = Z/(4\epsilon_1) - 1$, by setting $\epsilon_1 = Z/5$ we obtain that $g_2 = 1/4$. This results in the bound $n_2 \geq 80 \cdot Z^{-1} \ln(\epsilon^{-1})$, i.e., $n_2 = \Omega(Z^{-1} \log(\epsilon^{-1}))$. Combining the above results with the fact that $n_1 = O(\log n)$, $n = n_1 n_2$ we obtain that the error probability is $2^{-\theta n / \log n}$ where $\theta = \Theta(Z)$ assuming that $n \geq n_0$ where n_0 is such that $\log n_0 = \Omega(Z^{-2} \log Z^{-1})$. \square

¹Consider e_1, \dots, e_N random variables taking values in $\{0, 1\}$ with $\Pr[e_i = 1] = p$ for $i = 1, \dots, N$. It holds that

$$\Pr \left[\sum_{i=1}^N e_i > (1 + \zeta)\mu \right] < e^{-\mu\zeta^2/3}$$

for any $0 < \zeta < 1$ where $\mu = \sum_{i=1}^N \Pr[e_i = 1]$.

B Lower bound on rate given upper bound on error

Proposition 15. *Let $0 \leq \tau < 1/4$ be a constant. Let $R' = 1 - H(1/2 - \tau)$. Then, $R' \geq \tau^2$. Here, $H(\cdot)$ is the Shannon entropy function.*

Proof. We want to lower bound the rate $R' = 1 - H(1/2 - \tau)$. To this end, let us upper bound $H(1/2 - \tau)$.

From the definition of Shannon entropy,

$$H\left(\frac{1}{2} - \tau\right) = - \left[\left(\frac{1}{2} - \tau\right) \log_2 \left(\frac{1}{2} - \tau\right) + \left(\frac{1}{2} + \tau\right) \log_2 \left(\frac{1}{2} + \tau\right) \right].$$

Rewriting the log terms as below,

$$\begin{aligned} \log_2 \left(\frac{1}{2} - \tau\right) &= \log_2 (1 + (-2 \cdot \tau)) - 1, \\ \log_2 \left(\frac{1}{2} + \tau\right) &= \log_2 (1 + (2 \cdot \tau)) - 1, \end{aligned}$$

we get,

$$\begin{aligned} H\left(\frac{1}{2} - \tau\right) &= - \left[\left(\frac{1}{2} - \tau\right) \log_2 \left(\frac{1}{2} - \tau\right) + \left(\frac{1}{2} + \tau\right) \log_2 \left(\frac{1}{2} + \tau\right) \right] \\ &= - \left[\left(\frac{1}{2} - \tau\right) \cdot \log_2 (1 + (-2\tau)) - \left(\frac{1}{2} - \tau\right) \right] - \left[\left(\frac{1}{2} + \tau\right) \cdot \log_2 (1 + (2\tau)) - \left(\frac{1}{2} + \tau\right) \right] \\ &= - \left[\left(\frac{1}{2} - \tau\right) \cdot \log_2(e) \cdot \ln (1 + (-2\tau)) + \left(\frac{1}{2} + \tau\right) \cdot \log_2(e) \cdot \ln (1 + (2\tau)) - 1 \right]. \end{aligned}$$

We now lower bound the terms $\ln (1 + (-2\tau))$ and $\ln (1 + (2\tau))$ using the natural logarithm power series:

$$\ln(1+x) = x - \frac{1}{2}x^2 + \frac{1}{3}x^3 - \frac{1}{4}x^4 \dots \quad (-1 < x \leq 1).$$

In our case, we have $0 \leq \tau < 1/4$. So, $-1/2 < -2\tau \leq 0$ and $0 \leq 2\tau < 1/2$.

When $-1/2 < x \leq 0$,

$$\ln(1+x) = x - \frac{1}{2}x^2 + \frac{1}{3}x^3 - \frac{1}{4}x^4 \geq x - \frac{1}{2}x^2(1+x+\dots) \geq x - \frac{1}{2}x^2 \cdot \frac{1}{1-x} \geq x - x^2.$$

When $0 \leq x < 1/2$,

$$\ln(1+x) = x - \frac{1}{2}x^2 + \frac{1}{3}x^3 - \frac{1}{4}x^4 \geq x - \frac{1}{2}x^2.$$

Hence, we have

$$\ln(1+x) \geq \begin{cases} x - \frac{x^2}{2} \cdot \frac{1}{(1-x)} \geq x - x^2 & -\frac{1}{2} < x \leq 0, \\ x - \frac{1}{2}x^2 & 0 \leq x < \frac{1}{2}, \end{cases}$$

and

$$\begin{aligned} H\left(\frac{1}{2} - \tau\right) &= - \left[\left(\frac{1}{2} - \tau\right) \cdot \log_2(e) \cdot \ln (1 + (-2\tau)) \right] - \left[\left(\frac{1}{2} + \tau\right) \cdot \log_2(e) \cdot \ln (1 + (2\tau)) - 1 \right] \\ &\leq - \left[\left(\frac{1}{2} - \tau\right) \cdot \log_2(e) \cdot (-2\tau - 4\tau^2) \right] - \left[\left(\frac{1}{2} + \tau\right) \cdot \log_2(e) \cdot (2\tau - 2\tau^2) - 1 \right] \\ &= 1 - \log_2(e) \cdot (\tau^2 + 2 \cdot \tau^3) \leq 1 - \log_2(e) \cdot \tau^2 \leq 1 - 1.44\tau^2. \end{aligned}$$

This gives us,

$$R' = 1 - H\left(\frac{1}{2} - \tau\right) \geq 1 - (1 - 1.44\tau^2) \geq \tau^2.$$

□