Provable Secure Universal Steganography of Optimal Rate

Provably Secure Steganography does not Necessarily Imply One-Way Functions

Sebastian Berndt
University of Lübeck
Ratzeburger Allee 160, 23562 Lübeck, Germany
berndt@tcs.uni-luebeck.de

Maciej Liśkiewicz
University of Lübeck
Ratzeburger Allee 160, 23562 Lübeck, Germany
liskiwi@tcs.uni-luebeck.de

ABSTRACT

We present the first complexity-theoretic secure steganographic protocol which, for any communication channel, is provably secure, reliable, and has nearly optimal bandwidth. Our system is unconditionally secure, i.e. our proof does not rely on any unproven complexity-theoretic assumption, like e.g. the existence of one-way functions. This disproves the claim that the existence of one-way functions and access to a communication channel oracle are both necessary and sufficient conditions for the existence of secure steganography, in the sense that secure and reliable steganography exists independently of the existence of one-way functions.

Keywords

Steganography; Cryptographic Primitives; Lower Bounds

1. INTRODUCTION

Digital steganography has received substantial interest in modern computer science since it allows secret communication without revealing its presence. Therefore, the investigation of steganography has recently become the subject of intensive studies, both theoretical and experimental. In this paper we provide a theoretical analysis on the existence of secure universal steganography with optimal rate – a problem which belongs to one of the most fundamental ones of this area. In the scenario considered here, we assume secret-key communication with the presence of a passive adversary and the security defined in the computational setting.

A common computational model for secret-key steganography, also used in this paper, was introduced by Hopper, Langford, and von Ahn [13, 14, 15]. Independently, Katzenbeisser and Petitcolas [16] provided a similar formulation. In this setting, a *stegosystem* is defined as a pair of probabilistic algorithms, called *encoder* and *decoder*, which share a secret-key. The aim of the encoder (often called Alice or the steganographer) is to hide a secret message in a document and to send it to the decoder (Bob) via a public *channel* $\mathcal{C}$, which is completely monitored by an *adversary* (Warden or steganalyst). The channel is modeled as a probability distribution of legal documents, called *cover-documents* and the adversary's task is to distinguish those from altered ones called *stego-documents*.

To hide a secret message m, the encoder can take sample cover-documents, based on past communication, and manipulate them to embed m. The decoder, receiving stego-documents, should be able to decode the hidden message correctly. The stegosystem is called *reliable* if the decoder succeeds with high probability. The adversary is a probabilistic algorithm with access to additional knowledge about the channel. A stegosystem is *secure* with respect to a channel $\mathcal{C}$ if no adversary of *polynomial* time complexity is able to distinguish with significant probability between cover-documents from $\mathcal{C}$ and stego-documents generated by the stegosystem's encoder. This implies in general that the distributions of cover-documents and stego-documents have to be fairly close in a complexity-theoretic sense. The *insecurity* of a stegosystem is the advantage of the best adversary to distinguish between cover- and stego-documents. Thus, a stegosystem is secure if its insecurity is sufficiently small, i.e. negligible in the security parameter κ defining the length of the shared secret-key. This work deals with the construction of unconditional secure stegosystems of high rate and with the relation between steganography and cryptography.

1.1 Steganography and Cryptography

Although there is a strong connection between these two areas, steganography is *not* cryptography. Our example below shows even more, namely that polynomial-time bounded steganography is *not* cryptography. A commonly heard argument for the premise that steganography is cryptography goes as follows:

> Let m and m' be two different secret messages and s and s' be stego-documents which embed m, resp. m'. If the distributions of s and s' are indistinguishable from the distribution of the cover-documents, then by the triangle-inequality, the distributions of s and s' are also indistinguishable. Hence, a secure stegosystem is also a secure cryptosystem.

While the argument concerning the triangle-inequality is true, one can not simply use the stegosystem as a cryptosystem, as the stegosystem needs access to samples from the channel. Arguably, the most researched channel is those of natural digital pictures (say in the JPEG-format). A typical stegosystem for this channel takes a sample picture and modifies it in a way that is not detectable. A cryptosystem that simulates this stegosystem thus needs a way to get a

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

sample picture. But the standard definition of cryptosystems does *not* assume such access and it is highly unlikely that an efficient algorithm to simulate sampling for this channel can be constructed. We will note later on that ignoring this access leads to misunderstandings, e.g. in the often cited work of Hopper, Langford and von Ahn [15].

1.2 Secure and Reliable Universal Systems

Clearly, security and reliability are necessary attributes of any reasonable stegosystem. If one forgoes one of these, steganography becomes trivial: If one is not concerned about security, we can simply send the secret message plainly over the channel. If one is not concerned about reliability, we can simply send a random sample of the channel.

The next property a reasonable stegosystem should satisfy is the versatility of the system: it should be secure and reliable with respect not only to a concrete channel $\mathcal{C}$ (like e.g. images of a specific data set) but to a broad range of channels (e.g. images of different sources and different characteristics). In this paper we study the most general type of systems, called *universal*[1]. Recall that a stegosystem is universal if the encoding method does not rely on knowledge of the distribution for the channel $\mathcal{C}$ except that its min-entropy is sufficiently large. The importance of the universality is based on the fact that typically no good description of the probability distribution on a channel is known. In this paper we assume the standard definition for security of a universal stegosystem $\mathcal{S}$, i.e. we say $\mathcal{S}$ is secure (against a chosen hiddentext attack) if it is secure with respect to any channel $\mathcal{C}$ as long as $\mathcal{C}$ does not violate the hardness construction $\mathcal{S}$ is based on [14, 15].

But, in addition to these fundamental properties, other properties are also important. To reduce the overall overhead of the protocol, the transmission *rate*, i.e. the number of bits transmitted per single stego-document, should be as high as possible. Next, the expected *query complexity* of the stegosystem, i.e. the number of cover documents the encoder needs to construct a single stego-document, should be as small as possible since sampling large numbers of documents from the communication channel is expensive in general. High query complexity causes not only difficulties in true sampling of documents but also a high total running time of the encoder and/or decoder.

Obviously, for nontrivial systems, i.e. for such of small insecurity and unreliability, there is a trade-off between these requirements, as depicted exemplary in Fig. 1. We analyze there three hypothetical universal stegosystems for cover documents of length n. As usually, we assume that $n := n(\kappa)$ is specified as a function of security parameter κ and that n is polynomially related with κ, i.e. $n(\kappa) = \mathsf{pol}(\kappa)$. To embed λ-bit secret messages per document the systems need q samples to achieve negligible insecurity and unreliability. $\mathbf{InSec}(\kappa)$ denotes in our paper insecurity over all wardens of *polynomial time complexity* and $\mathbf{UnRel}(\kappa)$ the unreliability of the system (for definitions see next section). For channels of sufficiently high entropy, $\mathcal{S}_2$ and $\mathcal{S}_3$ are scalable with respect to the rate, but $\mathcal{S}_1$ is not. System $\mathcal{S}_1$ would illustrate e.g. a *spread-spectrum* steganography: although, strictly speaking, not universal, such systems are very general. They need just one sample document to embed a secret message but their rate is very limited (see e.g. [9] for more discussion). Systems $\mathcal{S}_2$ and $\mathcal{S}_3$ achieve almost optimal rate but a drawback of $\mathcal{S}_3$ is that its query complexity grows exponentially with respect to the rate.

[1]In the literature universal stegosystems are also called "black-box".

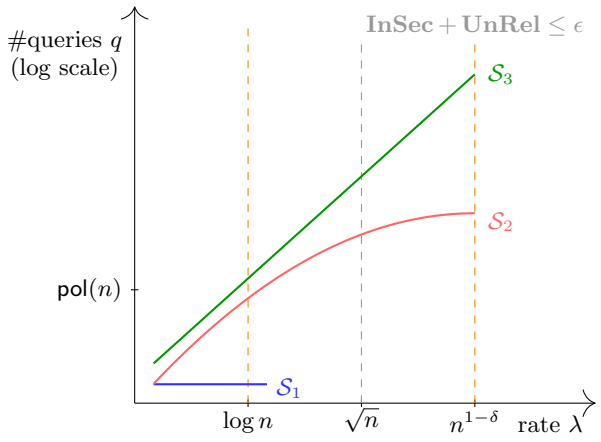

Figure 1: Dependences between rate and number of queries of three hypothetical stegosystems of small insecurity and unreliability. The systems $\mathcal{S}_2$ and $\mathcal{S}_3$ are scalable with respect to the rate, but $\mathcal{S}_1$ is not. However to increase the rate in $\mathcal{S}_3$ the number of queries increases drastically.

In this paper we investigate the problem of existence of *provably unconditionally* secure and reliably steganography, i.e. without using unproven assumptions like existence of one-way functions. The work emphasizes also on possible trade-offs between rate and query complexity for universal stegosystems. Particularly, we investigate if such scalable systems as $\mathcal{S}_2$ and $\mathcal{S}_3$ discussed above may exist. We study if it is possible to construct provable secure and reliable systems which can achieve specific rates, like e.g. square root rate, and analyze lower bounds on the query complexity for achieving this rate.

1.3 Previous Work

Secure Steganography and One-Way Functions. Beside providing a rigorous definition for computationally secure steganography, the main contribution of Hopper et al.'s work [14, 15] is demonstrating that a widely believed complexity-theoretic assumption – the existence of one-way functions – and access to a channel oracle are both necessary and sufficient conditions for the existence of secure and reliable steganography (Corollary 1 in [15]):

THEOREM 1 ([15], INFORMAL). *Relative to an oracle for channel $\mathcal{C}$, secure (and reliable) stegosystems exist if and only if one-way functions exist.*

This claim is now widely circulated in the literature. In her handbook [9] on steganography Fridrich writes: "One of the most intriguing implications of this complexity-theoretic view of steganography is the fact that secure stegosystems exist if and only if secure one-way (hash) functions exist [...]"

Unfortunately, as we show in this work the proof of this equivalence provided in [15] turns out to be incorrect in the stated form (see Section 3 for details). Moreover, we construct a universal stegosystem which is reliable and unconditionally secure, i.e. secure without any cryptographic

assumptions (see Theorem 4). This disproves the claim of Theorem 1 in the sense that secure and reliable steganography exists independently of the existence of one-way functions. However, it should be noted that our construction needs super-exponential time.

Upper Bounds on the Insecurity and Unreliability. To provide secure and reliable steganography based on the existence of one-way functions, Hopper et al. construct a universal stegosystem using so called "rejecting sampling". Roughly speaking, in order to transmit a bit β, the encoder repeatedly samples from the cover-document distribution $\mathcal{C}$ until she gets a document s which is mapped to the given β by a pseudorandom function F_k indexed by the secret key k. This encoding method has been extended to embed multiple bit messages m per document in such a way that, using rejecting sampling, one searches for a document s with $F_k(s) = m$. The authors showed that the proposed stegosystem is secure and reliable and, since the existence of one-way functions implies pseudorandom functions (see [11]) they conclude the following:

THEOREM 2 ([15], INFORMAL). *There exists a universal stegosystem $\mathcal{S}(\kappa)$ with security parameter κ (describing the length of the secret key) that hides a $\lambda := \lambda(\kappa)$ bit message in a sequence of ℓ stego documents of a channel $\mathcal{C}$ and*

(a) $\mathcal{S}(\kappa)$ takes $q(\kappa) = \lambda 2^\lambda$ sample sequences, each containing ℓ cover-documents and

(b) it achieves insecurity $\mathbf{InSec}(\kappa)$ bounded from above by the term $\Phi_{\mathcal{C}}(\mathsf{pol}(\lambda 2^\lambda), \kappa)$ and unreliability $\mathbf{UnRel}(\kappa)$ bounded by $\Phi_{\mathcal{C}}(\lambda 2^\lambda, \kappa) + e^{-\lambda} + \lambda 2^{\lambda - \mathcal{H}(\mathcal{C}^\ell)}$,

where the function $\Phi_{\mathcal{C}}$ describes a term caused by the insecurity of the pseudorandom function used by the encoder and decoder and $\mathcal{H}(\mathcal{C}^\ell)$ denotes the min-entropy of ℓ consecutive documents of the channel.

Importantly, $\Phi_{\mathcal{C}}$ in the theorem above is defined relative to an oracle for channel $\mathcal{C}$. This means that an adversary that distinguishes between a randomly chosen function and a pseudorandom function indexed by a secret-key has also access to the oracle for $\mathcal{C}$ besides an access to the standard challenging oracle (for a formal definition and a discussion on this subject, see the next section). Thus, in the estimation

$$\mathbf{InSec}(\kappa) + \mathbf{UnRel}(\kappa) \;\leq\; \Phi_{\mathcal{C}}(\mathsf{pol}(\lambda 2^\lambda), \kappa) + e^{-\lambda} + \lambda 2^{\lambda - \mathcal{H}(\mathcal{C}^\ell)} \qquad (1)$$

the upper bound on the insecurity and unreliability of Hopper et al.'s system is negligible if and only if the term involving $\Phi_{\mathcal{C}}$ is negligible. We notice that, if the transmission rate exceeds the logarithm of the key length κ, then the proofs provided in [15] do not guarantee that unreliability and insecurity (recall, even against polynomial-time bounded warden) of the proposed stegosystems are negligible.

More precisely, in case the number of bits $\lambda := \lambda(\kappa)$ embedded in a single document grows asymptotically faster than $\log \kappa$, the term $\Phi_{\mathcal{C}}(\mathsf{pol}(\lambda 2^\lambda), \kappa)$ in the right hand side of Eq. (1) is not guaranteed to be negligible in κ even if the existence of pseudorandom functions is assumed. This is due to the fact that one assumes security of pseudorandom functions only against polynomial-time attacker and the term $\lambda 2^\lambda$ is super-polynomial for $\lambda \in \omega(\log \kappa)$. Thus, if a channel $\mathcal{C}$ allows to embed up to n bits per document, i.e. if its min-entropy $\mathcal{H}(\mathcal{C})$ is very high, the stegosystem of Hopper et al. is not scalable to meet optimum rate: for any $\lambda \leq n$ its query complexity is $q = \lambda 2^\lambda$ but its insecurity and unreliability is guaranteed negligible only for $\lambda \in \mathcal{O}(\log n)$ since $n(\kappa) = \mathsf{pol}(\kappa)$. We illustrate this in Fig. 2.

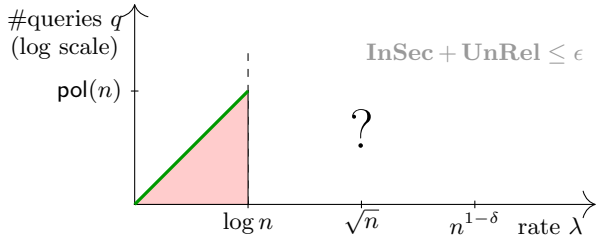

Figure 2: Known results (under cryptographic assumptions): the green line shows the dependence between the rate and number of queries to ensure negligible insecurity and unreliability of Hopper et al.'s system (Eq. (1)). This bound is sharp: any system of rate and with number of queries in the red area is insecure or unreliable (due to Eq. (2) by Dedić et al.). The situation for $\lambda \in \omega(\log n)$ has remained open so far.

Dedić et al. [6] proposed two new universal stegosystems with the upper bounds on the insecurity and unreliability similar to Eq. (1). Particularly, the bounds involve the term $\Phi_{\mathcal{C}}(\ell 2^b, \kappa)$, where b denotes the number of bits encoded per document, i.e. $b = \lambda/\ell$. Similarly to Hopper et al.'s system, if the number of bits b per document grows asymptotically faster than $\log \kappa$ then $\Phi_{\mathcal{C}}(\ell 2^b, \kappa)$ can grow faster than any negligible function, even if the encoder and decoder use pseudorandom functions.

Lower Bounds on the Insecurity and Unreliability. In [6], Dedić et al. prove (under cryptographic assumptions) the existence of channels such that the number of samples the encoder of any secure and reliable universal stegosystem must obtain from those channels is exponential in the number of bits embedded per document. In our terms, their result can be stated as:

THEOREM 3 ([6], INFORMAL). *For every universal stegosystem $\mathcal{S}(\kappa)$ which hides $\lambda := \lambda(\kappa)$ bits and takes $q := q(\kappa)$ samples per stego-document there exists a family of channels $\mathcal{C}(\kappa)$ such that*

$$\mathbf{InSec}(\kappa) + \mathbf{UnRel}(\kappa) \;\geq\; \frac{1}{2} - \frac{e \cdot q}{2^\lambda} - \Psi(q, \kappa) - o(1), \quad (2)$$

where $\mathbf{InSec}(\kappa)$ denotes the insecurity (against polynomial-time bounded wardens) and $\mathbf{UnRel}(\kappa)$ the unreliability of $\mathcal{S}(\kappa)$ on $\mathcal{C}(\kappa)$, and Ψ describes a term caused by the insecurity of the pseudorandom function used in the construction of $\mathcal{C}(\kappa)$.

They thus prove that the exponential query complexity $\lambda 2^\lambda$ of the universal systems by Hopper et al. is asymptotically optimal: indeed, if $q \in o(\lambda 2^\lambda)$ and $q \in \mathsf{pol}(\kappa)$, the right hand side of the inequality (2) goes to $1/2$. However, analogously to our discussion on the upper bound (1), we notice that the lower bound (2) is not meaningful if $q \in \omega(\mathsf{pol}(\kappa))$ (even if $q \in o(\lambda 2^\lambda)$), as the right hand side of the inequality does not necessarily need to go to 0 in this case. The red area in Fig. 2 illustrates this lower bound.

Later, Hopper et al. [15] provided another lower bound on the insecurity and unreliability. They show that for every universal stegosystem $\mathcal{S}(\kappa)$ of query complexity $q(\kappa)$ which hides $\lambda(\kappa)$ bits per document and for any κ there exists a channel such that:

$$\mathbf{InSec}(\kappa, q) + \mathbf{UnRel}(\kappa) \;\geq\; 1 - q/2^{\lambda} - 2^{-\kappa}, \qquad (3)$$

where $\mathbf{InSec}(\kappa, q)$, in contrast to $\mathbf{InSec}(\kappa)$, denotes insecurity over wardens of time complexity and size $> q(\kappa)$. Note that in the case of $\lambda \in \omega(\log n)$, bounds (2) and (3) are incomparable in the following sense. Due to (2), if in a reliable universal stegosystem $\mathcal{S}$ the number of queries q is dominated by 2^{λ} then there exists a *polynomial-time* bounded Warden whose advantage to detect $\mathcal{S}$ is big. The time complexity of the Warden must not dependend on the query complexity q of $\mathcal{S}$ but (2) needs the assumption that pseudorandom functions exist and it may be meaningless if the rate exceeds $\log n$. Bound (3) does not need any cryptographic assumption, it is meaningful for any λ but, the Warden who detects $\mathcal{S}$ needs time and size bigger than the query complexity q of $\mathcal{S}$. Thus, in cases of super-polynomial q, the Warden is not polynomial-time bounded anymore implying $\mathbf{InSec}(\kappa, q) \gg \mathbf{InSec}(\kappa)$.

1.4 Our Contribution

Thus, as shown above, if high rate is required we have no guarantee that the discussed systems are secure and reliable. And indeed, *no* secure and reliable universal stegosystem (irrespective of its query complexity) with rate larger than $\log n$ was known before, even under unproven cryptographic assumptions. Note that the secure stegosystems used in practice typically achieve a much larger rate of $\sqrt{n}$ [18]. A longstanding conjecture, the *Square Root Law of Steganographic Capacity* [8, 17, 19] deals with just this fact. It says that a rate of the form $(1-\varepsilon)\sqrt{n}$ is always achievable (not necessarily in a setting of universal steganography). We thus have the situation, that the best known theoretical rate is $\log n$, while all practical rates are of order $\sqrt{n}$.

To close this gap between theory and practice, we introduce the notion of *rate-efficiency* and analyze its impact on steganography. We say that a stegosystem is *rate-efficient*, if there is a constant $1 > \alpha > 0$ such that the number of embedded bits in a stego-document of length n is at least n^{α} (in channels of sufficiently large entropy). One of the main results of this paper is the construction of a stegosystem that is scalable with respect to the rate up to n^{α} for every $\alpha < 1$. However, to achieve this rate, an exponential number of queries is needed. On the other hand we prove that this query complexity is minimal. Thus, such a system as $\mathcal{S}_2$ analyzed in Fig.1 cannot exist and we give a complete answer to the question illustrated in Fig. 2 of determining the relationship between rate and number of queries. For an illustration of our results see Fig. 3.

Speaking more precisely: In this paper we present, up to our best knowledge, the first secret-key universal stegosystem that is provably secure (in the complexity-theoretic setting), reliable, and it has nearly optimal rate. Moreover, the security and reliability do not rely on any unproven assumptions, like the existence of one-way functions, as long as the channel does not allow one to break the hard functions the stegosystem is based on. Furthermore, no channel that is sampleable in exponential time can break these hard functions.

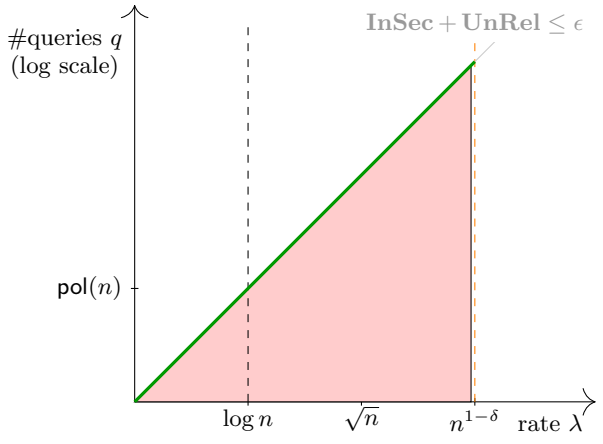

Figure 3: Our results (without any assumptions): our stegosystem achieves negligible insecurity and unreliability for the number of queries depending on the rate as shown by the green line (Eq. (4)). This bound is sharp: any system of rate and query complexity in the red area is unreliable or insecure against polynomial-time bounded wardens (Eq. (5)).

THEOREM 4 (INFORMAL). *There exists a universal stegosystem $\mathcal{S}$ that is unconditionally secure and reliable. Moreover $\mathcal{S}$ is rate-efficient.*

This stays in contrast with the claim of Hopper at al. [15] that provably secure steganography does imply existence of one-way functions (see Theorem 1). Analyzing the proofs in [15] we experienced the need for formal definitions and more careful handling of the security notion for cryptographic primitives *relative to an oracle for a channel* in order to conduct the typical cryptographic security-reductions. Otherwise, a security analysis can lead to misunderstandings and errors.

Our stegosystem is stateless and uses rejecting sampling in a similar way as the stateless secret-key construction proposed in [6] and the public-key system used in [1]. To achieve the security which does not rely on any unproven assumptions we construct pseudorandom functions of very high hardness based on constructions for sparse pseudorandom distributions due to Goldreich and Krawczyk [12]. Overall, our construction satisfies the following conditions:

THEOREM 5 (INFORMAL). *For every $1 > \alpha_1 \geq \alpha_2 > 0$, there exists a stegosystem $\mathcal{S}(\kappa)$, with security parameter κ, that has documents of length κ^{α_1}, such that for every channel $\mathcal{C}$ with min-entropy $\mathcal{H}(\mathcal{C}) > 2 \cdot \kappa^{\alpha_2}$, the system $\mathcal{S}$*

- *hides an $\ell \cdot \kappa^{\alpha_2}$ bit message, $\lambda(\kappa) = \kappa^{\alpha_2}$ bits per document sent, with $\ell \leq \mathsf{pol}(\kappa)$,*
- *takes $q(\kappa) \leq \kappa 2^{\kappa^{\alpha_2}}$ samples per κ^{α_2} bits, and*
- *has insecurity and unreliability negligible in κ (if $\mathcal{C}$ does not break the used hard functions):*

$$\mathbf{InSec}(\kappa) + \mathbf{UnRel}(\kappa) \;\leq\; \mathsf{negl}(\kappa). \qquad (4)$$

The running time of the stegosystem $\mathcal{S}$ is $2^{2^{o(\kappa)}}$ to ensure the unconditional security against polynomial wardens, but it achieves minimal query complexity. To prove that the query complexity of our construction is optimal, we show the following lower bound, improving the bounds (2) and (3):

THEOREM 6 (INFORMAL). *There exists a family of channels $\mathcal{C}(\kappa)$ such that for every universal stegosystem $\mathcal{S}(\kappa)$ where the encoder takes an expected number of $q(\kappa)$ queries, it holds that*

$$\mathbf{InSec}(\kappa) + \mathbf{UnRel}(\kappa) \;\geq\; \frac{1}{2} - \frac{e \cdot q}{2^{\lambda}} - o(1). \qquad (5)$$

In the theorems $\mathbf{InSec}(\kappa)$ denotes, like in (2), the insecurity over polynomial-time bounded wardens and $\mathbf{UnRel}(\kappa)$ denotes the unreliability of stegosystem $\mathcal{S}(\kappa)$ on $\mathcal{C}(\kappa)$.

As mentioned in the introduction, we also show that the proof of Theorem 1 in [15] is incorrect and that the theorem is wrong in the stated form. This is a direct consequence of our Theorem 5.

1.5 Related Work

The running time of the rejecting sampling algorithm was improved by Kiayias et al. in [20]. They use *t-wise independent* families of functions instead of a pseudorandom function to choose corresponding documents from the channel. As a t-wise independent family utilizes much less random bits than a pseudorandom function, this decreases the running time of the algorithm. They also provide another refinement of the rejection sampling approach by proving that a key length of $(1+o(1))\ell$ is sufficient in order to embed a message of length ℓ while achieving high security. The authors define the security of a stegosystem in an *information-theoretic* setting but they show that the results are also applicable to the computational model assuming the existence of pseudorandom functions. However, the system provided in [20] is rate-inefficient – it embeds 1 bit per document.

For some specific channels, Van Le and Kurosawa [21] proposed a stegosystem based on arithmetic coding that achieves a better rate than the universal systems proposed in [15] and [6]. However in the model assumed in [21] it is allowed that the system has access to additional knowledge of the channel. For some specific *families* of channels, like e.g. represented by Boolean functions, Liśkiewicz et al. [22] provided systems for secure and reliable steganography but, similarly to the results presented in [6, 13, 14, 15], the upper bound on insecurity of the construction in [22] is meaningless if the rate exceeds $\log n$.

Universal systems with optimum embedding rate are also a subject of intensive study using an information-theoretic modeling of secure steganography as defined in the influential work by Cachin [4]. In [24], Wang and Moulin have introduced a powerful information-theoretic framework for studying capacity in perfectly secure steganography. For relevant theoretical results and practical applications see also [2, 5, 7, 8, 23].

The paper is organized as follows. Due to space constraints, the formal definitions concerning steganography, cryptography and the security relative to a channel are skipped and we refer the reader to the literature, e.g. [6, 15]. The next section contains a short summary of our notations. Section 3 deals with the problems in the proof of Theorem 1. We analyze those problems and discuss possible solutions to them. In Section 4 two ensembles of pseudorandom functions of very high hardness are defined and in the subsequent section these ensembles are used to construct a universal stegosystem which is secure and efficient in terms of rate and reliability. This proves Theorems 4 and 5. In order to prove that our algorithm has almost optimal query complexity, Section 6 improves the lower bounds (2) and (3) of [6] and [15] that only work under certain circumstances. Our proof of Theorem 6 does without those circumstances. Finally, we discuss the conclusions one can draw from our work and possible future research directions.

2. NOTATIONS

Due to space constraints, we skip the formal security definitions and refer the reader to the literature, e.g. [6, 15].

Channel, Reliability, Security. Our stegosystems and cryptographic primitives will be sufficiently parameterized by the secret-key length κ. As usually, we will require that the ability of any polynomial-time algorithm to attack this constructions is lower than the inverse of every polynomial in κ. This is modeled by the notion of *negligible* function $\mathsf{negl}\colon \mathbb{N} \to [0,1]$, which satisfies that for every polynomial p, there is an $N_0 \in \mathbb{N}$ such that $\mathsf{negl}(N) < p(N)^{-1}$ for every $N \geq N_0$.

We denote the *min-entropy* of a probability distribution D as $\mathcal{H}(D)$. A *channel* $\mathcal{C}$ is a function, that maps for any $n \in \mathbb{N}$ (the *document length*) a sequence of documents h (the *history*) to a probability distribution $\mathcal{C}_{h,n}$ on documents of length n. A *stegosystem* $\mathcal{S} = [SE, SD]$ consists of two probabilistic algorithms SE and SD. The *encoder* SE takes a key k, a message $m \in \mathcal{M}_\kappa$ (the *message space*), and a history h and produces a sequence of ℓ documents (ℓ is the *output length* of $\mathcal{S}$), while the *decoder* SD tries to reconstruct m from those documents with the help of k and h. The *rate* b of the encoder is defined as $\log(|\mathcal{M}_\kappa|)/\ell$ and returns the number of bits that the encoder embeds into a single document. The encoder is allowed to draw $c \in \mathbb{N}$ (the *sample complexity*) samples from the channel $\mathcal{C}$. The *unreliability* $\mathbf{UnRel}_{\mathcal{S},\mathcal{C}}(\kappa)$ of the stegosystem $\mathcal{S}$ on channel $\mathcal{C}$ is the maximal probability that the decoding of a message fails. This is the same notation as we used in the introduction, where we skipped the mention of $\mathcal{C}$ and $\mathcal{S}$ to increase readability. An attacker (or *warden*) W on the stegosystem $\mathcal{S}$ on the channel $\mathcal{C}$ is allowed to give messages to a *challenging oracle* that either returns encodings of those messages or random documents. The *advantage* $\mathbf{Adv}^{\mathrm{cha}}_{W,\mathcal{S},\mathcal{C}}(\kappa)$ of W is its success probability in distinguishing those two cases. The maximal advantage of any warden W that runs in time t and makes q requests to the challenging oracle is the *insecurity* of $\mathcal{S}$ on $\mathcal{C}$ and denoted by $\mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S},\mathcal{C}}(q,t,\kappa)$. By $\mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S},\mathcal{C}}(\kappa)$ we mean the insecurity over all wardens of polynomial-time complexity. As with the unreliability, this is exactly the notion of the introduction, but in a parameterized form. The *insecurity* $\mathbf{InSec}^{\mathrm{prf}}_{F}(q,t,\kappa)$ of a *pseudorandom function* F is the maximal probability of any algorithm that runs in time t and makes q queries to an oracle, that either equals one of the pseudorandom functions or a totally random functions, to distinguish these cases.

Relativized Security. Clearly, the (in)security of a stegosystem $\mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S},\mathcal{C}}(\cdot,\cdot,\cdot)$ depends on the concrete channel $\mathcal{C}$, as every adversary on $\mathcal{S}$ on channel $\mathcal{C}$ has oracle access to sample documents of $\mathcal{C}$. In order to base the security of a stegosystem on the security of a cryptographic primitive, a typical reduction works along the following lines: Suppose that there is a successful warden W on the stegosystem $\mathcal{S}$; Construct an attacker A on the cryptographic primitive that simulates W on $\mathcal{S}$; Prove that the advantage of A and W is very similar. Using such reductions, it is important to note that the attacker A on the cryptographic primitives

completely simulates the warden W and the encoder of $\mathcal{S}$ (assuming a black-box access to the cryptographic primitive it is based on). As both W and $\mathcal{S}$ make calls to the sampling oracle of the channel, A also needs access to those samples. The need for this was already noted by Hopper et al. [13, 14, 15]. There are essentially two solutions to take the access to the sampling oracle into account:

- One assumes that the sampling oracle can be simulated in *polynomial time*. Hence, the simulation of W and $\mathcal{S}$ can be performed in polynomial time. As the typical requirement is that the cryptographic primitives remains secure against attackers that run in polynomial time, the security reduction remains valid.
- One assumes that the cryptographic primitive remains secure even if the attacker has access to the sampling oracle of the channel $\mathcal{C}$. One then proceeds to define *relativized* versions of the common insecurity terms, e.g. we could define the insecurity $\mathbf{InSec}^{\mathrm{prf}}_{F,\mathcal{C}}(q,t,\kappa)$ of a pseudorandom function F. It is the maximal probability of any algorithm that has access to the sampling oracle of the channel $\mathcal{C}$, runs in time t and makes q queries to an oracle, that either equals one of the pseudorandom functions or a totally random functions, to distinguish these cases. Dedić et al. [6] were the first that gave a formal definition for this, but they did not use it consistently in their work.

Note that the assumption that the sampling oracle for channel $\mathcal{C}$ can be simulated in polynomial time is quite artificial: Arguably, the single most studied channels for steganography are those containing multimedia-files such as images or videos. Typically, we do not assume that one is able to sample uniformly from the set of all valid images or videos. This rules out the first possibility. On the other hand, the second possibility is completely valid, as we have access to these channels in real-life, but suspect that this access does not break the security of cryptographic primitives. Due to this advantage, we will use the second possibility in this work.

3. COMMENTS ON THE WORK OF HOPPER, LANGFORD AND VON AHN

The attentive reader may have noticed that the commonly used formulation of Hopper, Langford and von Ahn [13, 14, 15] does not bound the running time of the stegosystem, while the time complexity of the adversary is required to be bounded by a polynomial. In addition to giving the rejection-sampling stegosystem, they also argue that one-way functions are *necessary* for steganography:

THEOREM 7 ([15], INFORMAL). *For all channels $\mathcal{C}$ it is true: if secure and reliable steganography for $\mathcal{C}$ exists then there exist one-way functions relative to an oracle for $\mathcal{C}$.*[2]

Since the proof of Theorem 7 is "black-box" with respect to the one-way functions it holds relative to the presence of the channel oracle for $\mathcal{C}$, too. Thus, Hopper et al. conclude Theorem 1 (Corollary 1 in [15]) given in the introduction.

Combining Theorem 5 with Theorem 1 one would conclude that (1) relative to an oracle for channel $\mathcal{C}$ one-way functions exist and much more startling, that (2) **one-way functions exist in the standard model**, i.e., without assuming oracle access to the channel $\mathcal{C}$. As a proof on the existence of one-way functions seems to be far away from our current knowledge, one must wonder at the validity of Theorem 1. Indeed, we found errors in the proof of Theorem 7 which consequently do not allow to conclude Theorem 1.

There are three issues concerning this proof. Firstly, the time complexity of the proposed *false entropy generator* (FEG), a kind of oracle used in the construction and the use of relativized primitives. The aim was to provide an algorithm for an FEG, assuming the existence of a stegosystem $\mathcal{S}$ that is SS-KHA-$\mathcal{D}$-$\mathcal{C}$ secure for some hidentext distribution $\mathcal{D}$ and some channel $\mathcal{C}$ (for the exact definitions, see [15]).

The proposed construction for an FEG uses, as a subroutine, the encoder of $\mathcal{S}$ having an oracle access to $\mathcal{C}$. Since no restrictions on the running time of the encoder are given, it does not follow that the obtained algorithm for FEG is bounded by a polynomial. This problem can be fixed by assuming that the stegosystem runs in polynomial time. Note, however, that making the assumption of polynomial time complexity for stegosystems, the claim of [15, Section 4.3] concerning rate-optimality is false, as the proposed system requires exponential time.

Secondly, according to the definition, an FEG is a *function* (satisfying the conditions of FEG). However, the FEG relative to an oracle $\mathcal{C}$ does not seem to be deterministic, as it sometimes returns the samples generated by the channel oracle. This does not seem to be fixable easily, but one can make use of randomized cryptographic primitives in order to give an alternative proof.

The third obstacle still remains: In order to construct a cryptographic primitive out of a stegosystem, one needs to simulate the access to the channel oracle. If this simulation can be carried out in polynomial time, the constructed primitive is indeed efficient. But, as discussed in Section 2, such an assumption is quite artificial. And indeed, if the channel oracle can not be simulated in polynomial time, the constructed cryptographic primitive is not efficient. It seems that the only remedy to this is to define another way of *relativized primitives*, where the primitive has also access to the channel oracle as in this work.

4. PSEUDORANDOM FUNCTIONS OF VERY HIGH HARDNESS

We construct two families of pseudorandom functions that are secure against adversaries of exponential running times. Our result does not rely on any unproven assumptions but to construct the family, super-exponential time is needed.

Let $P = \{P_\kappa\}_{\kappa\in\mathbb{N}}$ and $Q = \{Q_\kappa\}_{\kappa\in\mathbb{N}}$ be two ensembles of discrete probability distributions. For a probabilistic algorithm A, the *advantage of A to distinguish P and Q* is defined as

$$\mathbf{Adv}^{\mathrm{dist}}_{A,P,Q}(\kappa) = \left|\Pr[A^{P_\kappa}(1^\kappa)=1] - \Pr[A^{Q_\kappa}(1^\kappa)=1]\right|,$$

where A^D has the ability to get samples distributed accordingly to D in unit time. The insecurity of P and Q is defined as $\mathbf{InSec}^{\mathrm{dist}}_{P,Q}(q,t,\kappa) = \max_A\{\mathbf{Adv}^{\mathrm{dist}}_{A,P,Q}(\kappa)\}$, where the maximum is taken over all algorithms that make an expected number of q queries to the probability distribution and run in expected time t. In order to simplify our notation, we sometimes identify a set M with the uniform distribution

[2]Hopper et al. prove even stronger result using a weaker notion of security. Theorem 9 in [15] says that if there is a stegosystem $\mathcal{S}$ that is SS-KHA-$\mathcal{D}$-$\mathcal{C}$ secure for some hiddentext distribution $\mathcal{D}$ and some channel $\mathcal{C}$, then there exists a pseudorandom generator, relative to an oracle for $\mathcal{C}$.

on M. If M and N are two finite sets, we hence write $\mathbf{Adv}^{\mathrm{dist}}_{A,M,N}, A^M, A^N$ and $\mathbf{InSec}^{\mathrm{dist}}_{M,N}$ with the meaning that M and N are uniformly distributed.

We also introduce the relativized version $\mathbf{InSec}_{P,Q,\mathcal{C}}$ of the term $\mathbf{InSec}_{P,Q}$ and define it for parameters q, t, κ analogously as $\mathbf{InSec}^{\mathrm{prf}}_{\mathcal{F},\mathcal{C}}$.

Next we recall the following definition of the *statistical distance* D_S between two discrete probability distributions on the same domain.

DEFINITION 1 (STATISTICAL DISTANCE). *Let P and Q be two discrete probability distributions on the same domain X. The* statistical distance $D_S(P,Q)$ *between P and Q is defined as* $D_S(P,Q) = \frac{1}{2}\sum_{x\in X} |P(x) - Q(x)|$.

By the following well-known theorem (see e.g. [10]) we know that the statistical distance is a stronger measure than the computational indistinguishability.

THEOREM 8. *Let $\{P_\kappa\}_{\kappa\in\mathbb{N}}$ and $\{Q_\kappa\}_{\kappa\in\mathbb{N}}$ be two probability distribution ensembles on the same domains and let $\mathcal{C}(\kappa)$ be a channel. Then it holds that for every function t and q and all $\kappa \in \mathbb{N}$: $\mathbf{InSec}^{\mathrm{dist}}_{P,Q,\mathcal{C}}(q,t,\kappa) \le q \cdot D_S(P_\kappa, Q_\kappa)$. Particularly we have $\mathbf{InSec}^{\mathrm{dist}}_{P,Q}(q,t,\kappa) \le q \cdot D_S(P_\kappa, Q_\kappa)$.*

In our stegosystem we will apply a super-polynomial time computable pseudorandom function based on an algorithm G that also takes super-polynomial time, which is given by the following result due to Goldreich and Krawczyk [12]. In order to simplify the notation throughout this and the next section, let α_1, α_2 be constants with $1 > \alpha_1 \ge \alpha_2 > 0$ and let

$$n = \kappa^{\alpha_1}, b = \kappa^{\alpha_2}, B = 2^b \cdot b, \text{ and } N = 2^n \cdot b. \qquad (6)$$

THEOREM 9 (LEMMA 5, [12]). *Let $k(n)$ be any subexponential function in n. There are (nonpolynomial) generators which expand random strings of length n into pseudorandom string of length $k(n)$.*

This result immediately implies the following theorem:

THEOREM 10. *There is a deterministic algorithm G, that on input $x \in \{0,1\}^\kappa$ produces a string $G(x) \in \{0,1\}^N$, and a negligible function* negl *such that for every polynomial t in N, it holds $\mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^N}(1,t,\kappa) \le \mathsf{negl}(N)$. There is also another deterministic algorithm G', that on input $x \in \{0,1\}^\kappa$ produces a string $G'(x) \in \{0,1\}^B$, and a negligible function* negl' *such that for every polynomial t in B, it holds $\mathbf{InSec}^{\mathrm{dist}}_{G'(\{0,1\}^\kappa),\{0,1\}^B}(1,t,\kappa) \le \mathsf{negl}'(B)$.*

The theorem says that no polynomial time algorithm in N (recall $N = 2^n \cdot b$) can distinguish between the distribution $G(\{0,1\}^\kappa)$ and the uniform distribution on $\{0,1\}^N$. Similarly, no polynomial time algorithm in B can distinguish $G'(\{0,1\}^\kappa)$ and the uniform distribution on $\{0,1\}^B$. The running time of G and G' is exponential in N (resp. B), while the running time of the distinguisher is polynomial in N (resp. B). Note that the usual construction to obtain a pseudorandom function from a pseudorandom generator due to Goldreich, Goldwasser and Micali [11] is not suited for our situation: Its security proof relies on the ability of the attacker on the function to simulate the generator. As a simulation of the generator takes exponential time and our attackers are polynomial, we can not use this approach. Instead, we observe that the generators produce very long strings. We will interpret these strings as the table of a function.

For a bit string $\omega = \omega_1\omega_2\ldots$ of length $2^X \cdot Y$, for some positive integers X and Y, let the function $F_\omega\colon \{0,1\}^X \to \{0,1\}^Y$ be defined as

$$F_\omega(z) = \omega_{i_z \cdot Y}\omega_{i_z \cdot Y+1}\ldots\omega_{(i_z+1)\cdot Y-1},$$

if z is the binary representation of the number i_z. For example, the bit string $\omega = 0111101100010110$ corresponds to the function $F_\omega\colon \{0,1\}^3 \to \{0,1\}^2$ with e.g. $F_\omega(000) = 01$, $F_\omega(001) = 11$, and $F_\omega(111) = 10$.

Moreover, let F_G denote the function ensemble $F_G := \{F_{G(x)}\}_{x\in\{0,1\}^\kappa}$ and $F_{G'} := \{F_{G'(x)}\}_{x\in\{0,1\}^\kappa}$. The definition of F_ω implies a bijection between $\{0,1\}^{2^X\cdot Y}$ and the set of all function from $\{0,1\}^X \to \{0,1\}^Y$, which we will denote by $\boldsymbol{F}_{X,Y}$. The following theorem shows that F_G is not distinguishable from $\boldsymbol{F}_{n,b}$ by any algorithm with time complexity $\mathsf{pol}(N)$ and $F_{G'}$ is not distinguishable from $\boldsymbol{F}_{b,b}$ by any algorithm with time complexity $\mathsf{pol}(B)$.

THEOREM 11. *For all functions q, t of κ, we have:* $\mathbf{InSec}^{\mathrm{prf}}_{F_G}(q,t,\kappa) \le \mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^N}(1, N\cdot q + t, \kappa)$ *and* $\mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}(q,t,\kappa) \le \mathbf{InSec}^{\mathrm{dist}}_{G'(\{0,1\}^\kappa),\{0,1\}^B}(1, B\cdot q + t, \kappa)$.

The proof of the theorem relies simply on the fact that any polynomial-time adversary on F_G has only access to an excerpt of size $\mathrm{poly}(\kappa)$. Theorem 10 states that even access to the whole string of length $N \gg \mathrm{poly}(\kappa)$ does not help an adversary. The advantage of any adversary is thus only negligible.

PROOF. We only prove the theorem for F_G, as the proof for $F_{G'}$ is analogous.

Let A be any algorithm trying to distinguish between F_G and $\boldsymbol{F}_{n,b}$ with running time t by making q queries to the function oracle. The algorithm A has access to a function oracle f, which is either uniformly chosen from $\boldsymbol{F}_{n,b}$ or equal to $F_{G(x)}$ for a certain $x \in \{0,1\}^\kappa$. We will now construct a distinguisher Dist for G, such that

$$\left|\Pr[\mathrm{Dist}^{G(\{0,1\}^\kappa)}(1^\kappa) = 1] - \Pr[\mathrm{Dist}^{\{0,1\}^N}(1^\kappa) = 1]\right| = \left|\Pr_{x\leftarrow\{0,1\}^\kappa}[A^{F_{G(x)}(\cdot)}(1^\kappa) = 1] - \Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[A^{f(\cdot)}(1^\kappa) = 1]\right|.$$

The distinguisher Dist makes a single query to its distribution oracle and receives a bit string $\omega \in \{0,1\}^N$, which is either a random string or produced by $G(x)$. Whenever A makes a query z to its function oracle, Dist returns $F_\omega(z)$. In the end, Dist returns the same value as A. We thus have

$$\Pr_{x\leftarrow\{0,1\}^\kappa}[A^{F_{G(x)}(\cdot)}(1^\kappa) = 1] = \Pr[\mathrm{Dist}^{G(\{0,1\}^\kappa)}(1^\kappa) = 1]$$

and because of the bijection between $\boldsymbol{F}_{n,b}$ and $\{0,1\}^N$, we have

$$\Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[A^{f(\cdot)}(1^\kappa) = 1] = \Pr[\mathrm{Dist}^{\{0,1\}^N}(1^\kappa) = 1].$$

The computation of $F_\omega(z)$ takes time $\mathcal{O}(N)$. As the running time of A is bounded by t and A makes at most q queries to its function oracle, the running time of Dist is bounded by $\mathcal{O}(N)\cdot q + t$. The distinguisher Dist performs 1 query. □

5. RATE-EFFICIENT STEGANOGRAPHY

In this section we prove that there exists secure, reliable and rate-efficient steganography. Our result does not rely on any unproven assumption.

We will use the function families $F_G, F_{G'}$ of the previous section in the stegosystem of Backes and Cachin [1] to construct a universal stegosystem, which is unconditionally secure. As in the previous section, let α_1, α_2 be constants with $1 > \alpha_1 \geq \alpha_2 > 0$ and let n, b, N, B be as defined in eq. (6) in Section 4.

The rejecting sampling stegosystem $\mathcal{S} = [SE, SD]$ is described in Algorithm 1 and in Algorithm 2.

Algorithm 1: SE

In: key $k \in \{0,1\}^\kappa$, message $m \in \{0,1\}^{\ell \cdot b}$, history h
let $f := F_{G(k)}(\cdot)$; parse m into $m_1 m_2 \dots m_\ell$ with $|m_j| = b$;
for $j = 1$ *to* ℓ **do**
  let $i := 0$; sample $x \leftarrow \mathcal{C}_{h,n(\kappa)}$;
  while $f(x) \neq m_j$ *and* $i < \kappa \cdot 2^b$ **do**
    sample $x \leftarrow \mathcal{C}_{h,n(\kappa)}$; let $i := i + 1$;
  let $x_j := x$; append x_j to h;
return $x_1 x_2 \dots x_\ell$

Algorithm 2: SD

In: key $k \in \{0,1\}^\kappa$, documents $x_1 x_2 \dots x_\ell$
let $f := F_{G(k)}(\cdot)$;
return $f(x_1) f(x_2) \dots f(x_\ell)$

Backes and Cachin [1] prove that this stegosystem is secure against polynomial-time CHA-Wardens as long as the family of functions used is pseudorandom and as long as the number of bits embedded in a single document is at most $\log \kappa$. We will expand this result and prove that one can embed up to $o(\kappa)$ bits into a single document.

For any function $\hat{f} \in \boldsymbol{F}_{n,b}$, denote by $SE_{\hat{f}}(k, m, h)$ the run of SE if we replace $f = F_{G(k)}$ by $\hat{f}$ and by $SE_{\boldsymbol{F}_{n,b}}(k, m, h)$ the output distribution of SE_f, if f is chosen at random from $\boldsymbol{F}_{n,b}$. The following result is also due to Backes and Cachin and shows that $SE_{\boldsymbol{F}_{n,b}}$ and $\mathcal{C}$ are statistically close.

THEOREM 12 (PROPOSITION 1 IN [1]). [3] *If κ is a sufficiently large key-length, then there exists a constant $\eta < 1$ such that $D_S(P, \mathcal{C}_{h,n}) \leq \ell \cdot \left(2^{b - \mathcal{H}(\mathcal{C}_{h,n})} + \eta^{2^b \cdot \kappa}\right)$, where P is the probability distribution generated by $SE_{\boldsymbol{F}_{n,b}}(\cdot, m, h)$ upon random choice of m and fixed choice of h.*

As m is chosen by the Warden, we need to "randomize" the message embedded by SE. We will thus not embed m, but CTR\$$(m, k)$ with the meaning of:

The output length of CTR\$$(m, k)$ with $|m| = \ell \cdot b$ is $(\ell+1) \cdot b$. The insecurity of this construction can be reduced to the insecurity of $F_{G'}$ by the well known result of Bellare et al. [3]. As their reduction is black-box, the result also holds in a relativized setting.

[3]The exact wording of this Proposition and a corresponding proof can be found as Proposition 7 in the full version available under the link: `www.zurich.ibm.com/~cca/papers/pkstego.pdf`

Algorithm 3: CTR\$

In: message $m \in \{0,1\}^{\ell \cdot b}$, key $k \in \{0,1\}^\kappa$
sample y from $\{0,1\}^b$; let $g = F_{G'(k)}(\cdot)$;
let
  $r = g(y) || g((y+1) \bmod 2^b) || \dots || g((y+\ell-1) \bmod 2^b)$;
return $(r \oplus m) || y$

THEOREM 13. *The probability that a probabilistic algorithm A, that has access to the channel $\mathcal{C}$, with running time t, that makes q samples to an oracle, which on input $m \in \{0,1\}^{\ell \cdot b}$ either (a) produces random strings of length $(\ell+1) \cdot b$ or (b) CTR\$$(m, k)$ for a uniformly chosen k, can distinguish between the two cases (a) and (b) is bounded by $2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'},\mathcal{C}}(\kappa) + \frac{\ell \cdot b \cdot (q-1)}{b \cdot 2^b}$.*

This operation can be simply inverted with the knowledge of k and we will denote this operation by CTR\$$^{-1}$. We hence use the stegosystem $\mathcal{S}' = [SE', SD']$ which uses CTR\$ as

$$SE'(m, k, h) = SE(\text{CTR\$}(m, k), k, h)$$

and

$$SD'(x_1 \dots x_\ell, k) = \text{CTR\$}^{-1}(SD(x_1 \dots x_\ell, k), k).$$

By using the families $F_G, F_{G'}$, we prove that every channel with sufficient min-entropy that is sampleable in exponential time has a secure stegosystem. This analysis resembles the analysis in [1], but spells out the relation of $\mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S},\mathcal{C}}$ and $\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}$ respectively $\mathbf{InSec}^{\mathrm{prf}}_{F_G}$.

THEOREM 14. *The rejection sampling stegosystem $\mathcal{S}'$ on document length $n(\kappa) = n$ with the message space $\mathcal{M}_\kappa = \mathcal{U}(\{0,1\}^{\ell \cdot b})$ satisfies for every polynomials q, t in κ the following properties relative to channel $\mathcal{C}$:*

1. $\mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S}',\mathcal{C}}(q, t, \kappa) \leq$
   $\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left(q(\ell+1)2^b\kappa\ ,\ q(\ell+1)2^b\kappa + t\ ,\ \kappa\right) +$
   $\mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}\left(\ell+1\ ,\ (\ell+1)^2\ ,\ \kappa\right) +$
   $q(\ell+1)\left(2^{b-\mathcal{H}(\mathcal{C}_n)} + \eta^{2^b\kappa}\right) + \frac{(\ell+1)^2}{2^b} +$
   $2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'},\mathcal{C}}(q, t, \kappa) + \frac{\ell \cdot b \cdot (q-1)}{b \cdot 2^b}$ *for a constant $\eta < 1$,*
2. $\mathbf{UnRel}_{\mathcal{S}',\mathcal{C}}(\kappa) \leq$
   $\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left((\ell+1)2^b\kappa\ ,\ (\ell+1)2^b\kappa\ ,\ \kappa\right) +$
   $(\ell+1) \cdot \exp(-\kappa) + (\ell+1)^2 \cdot \kappa^2 \cdot 2^{2b - \mathcal{H}(\mathcal{C}_n)}$.

PROOF. In order to bound the insecurity of the stegosystem, we construct for every warden W with running time t that makes q queries to its challenging oracle on the stegosystem $\mathcal{S}'$ with respect to the channel $\mathcal{C}$ an attacker A on the function family F_G such that

$$\left| \Pr_{k \leftarrow \{0,1\}^\kappa}[W^{\mathcal{C}, SE^{\mathcal{C}}(k,\cdot,\cdot)}(1^\kappa) = 1] - \Pr[W^{\mathcal{C}, OC(\cdot,\cdot)}(1^\kappa) = 1] \right| \leq$$

$$\mathbf{Adv}^{\mathrm{prf}}_{A,F_G,(\mathcal{C})}(A) + q \cdot (\ell+1) \cdot \left(2^{b-\mathcal{H}(\mathcal{C}_n)} + \eta^{2^b \cdot \kappa}\right) + \frac{(\ell+1)^2}{2^b}$$
$$+ \mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}\left(\ell+1, (\ell+1)^2, \kappa\right) + 2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'},\mathcal{C}}(q, t, \kappa)$$
$$+ \frac{\ell \cdot b \cdot (q-1)}{b \cdot 2^b}.$$

This yields the security of the stegosystem. Let W be any such warden on the stegosystem $\mathcal{S}$ with respect to the channel

$\mathcal{C}$. The attacker A has access to a function oracle f, which is either uniformly chosen from $\boldsymbol{F}_{n,b}$ or equal to $F_{G(k)}$ for a certain $k \in \{0,1\}^\kappa$. The attacker A simulates the warden W. Whenever W makes a query to the channel-oracle, A uses its channel-oracle to produce such a sample. Whenever W makes a query (m,h) to the challenging oracle, A uses the encoding algorithm $SE_f(k,m,h)$. The attacker A then returns the same result as W. If $f = F_{G(k)}$, the attacker A simply simulates the run of W against the stegosystem, i.e.,

$$\Pr_{k\leftarrow\{0,1\}^\kappa}[A^{F_{G(k)}(\cdot)}(1^\kappa)=1] =$$
$$\Pr_{k\leftarrow\{0,1\}^\kappa}[W^{\mathcal{C},SE^{\mathcal{C}}(k,\cdot,\cdot)}(1^\kappa)=1].$$

If f is truly randomly and $m = m_1 m_2 \dots m_\ell$ is a message of length $\ell \cdot b$ such that $m_i \neq m_j$ for every $i \neq j$, we can think of $SE_f(m,k,h)$ as ℓ-fold product of the probability distribution $SE_f(m_i,k,h)$, where f is chosen randomly for every i.

The output of $SE_f(m_i,k,h)$ is nearly identical to the channel (see Theorem 12), if the corresponding message of length b is also chosen uniformly. Theorem 13 implies that for W, the difference between the behavior of $SE_f(m_i,k,h)$ on a uniformly chosen message m_i or an m_i generated by CTR\$ is bounded by $2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'},\mathcal{C}}(q,t,\kappa) + \frac{\ell\cdot b\cdot(q-1)}{b\cdot 2^b}$.

As we do not give m to SE, but rather the message $m' = \mathrm{CTR\$}(m,k) = m'_1 m'_2 \dots m'_{\ell+1}$, the probability that there are $i \neq j$ such that $m'_i = m'_j$ is at most

$$\mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}(\ell+1,(\ell+1)^2,\kappa) + \frac{(\ell+1)^2}{2^b},$$

by constructing an attacker on $F_{G'}$ which guesses values $x_1,\dots,x_{\ell+1}$ and tests, whether $f(x_1), f(x_2), \dots, f(x_{\ell+1})$ are pairwise different.

As statistical distance is stronger than computational indistinguishability (see Theorem 8), we thus know that there is a constant $\eta < 1$ such that

$$\left|\Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[W^{\mathcal{C},SE_f}(1^\kappa)=1] - \Pr[W^{\mathcal{C},OC(\cdot,\cdot)}(1^\kappa)=1]\right| \leq$$
$$q\cdot(\ell+1)\left(2^{b-\mathcal{H}(\mathcal{C}_n)} + \eta^{2^b\cdot\kappa}\right) + 2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}(q,t,\kappa) +$$
$$\frac{\ell\cdot b\cdot(q-1)}{b\cdot 2^b} + \mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}(\ell+1,(\ell+1)^2,\kappa) + \frac{(\ell+1)^2}{2^b},$$

as W makes at most q calls to its challenging oracle. This concludes the statement concerning the advantage of A. The simulation of each call to SE_f can be carried out in time $\mathcal{O}((\ell+1)\cdot 2^b\cdot\kappa)$ if one has access to the channel oracle. The number of calls to the function oracle f is bounded by $q\cdot(\ell+1)\cdot 2^b\cdot\kappa$ The running time of A is thus at most $q\cdot\mathcal{O}((\ell+1)\cdot 2^b\cdot\kappa)+t$ and the number of queries of A is at most $q\cdot(\ell+1)\cdot 2^b\cdot\kappa$.

Concerning the reliability, we construct for every message m and every legal history h a different attacker $A_{m,h}$ against F_G. The attacker $A_{m,h}$ with function oracle f computes $m' = SD_f(k, SE_f(k,m,h))$ and returns 1 if $m = m'$. If $f = F_{G(k)}$, we have

$$\Pr_{k\leftarrow\{0,1\}^\kappa}[A^{F_{G(k)}(\cdot)}_{m,h}(1^\kappa)=1] =$$
$$\Pr_{k\leftarrow\{0,1\}^\kappa}[m \neq SD(k, SE(k,m,h))].$$

If f is a totally random function from $\boldsymbol{F}_{n,b}$ and all samples $x_1, x_2, \dots$ taken from the channel oracle $\mathcal{C}$ are different, the probabilities $\Pr[f(x_i) = m_j]$ are independent, as we can assume that a new random function is evaluated on each x_i. Denote the event that all of the x_i are pairwise different with $\overline{\mathrm{Coll}}$. The probability that none of this samples evaluates to m is then bounded by

$$\Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[m \neq SD(SE_f(m,h)) \mid \overline{\mathrm{Coll}}] \leq$$
$$\sum_{j=1}^{\ell+1}\prod_{i=1}^{2^b\cdot\kappa}\Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[f(x_i)\neq m_j] \leq$$
$$(\ell+1)\cdot\left(1-\frac{1}{2^b}\right)^{2^b\cdot\kappa} \leq (\ell+1)\cdot\exp(-\kappa). \qquad (*)$$

By definition, the maximal probability of any element from the channel is bounded from above by $2^{-\mathcal{H}(\mathcal{C}_n)}$. The probability that $x_i = x_{i'}$ for $i \neq i'$ is thus bounded by $2^{-\mathcal{H}(\mathcal{C}_n)}$. Hence

$$\Pr[\mathrm{Coll}] \leq ((\ell+1)\cdot\kappa\cdot 2^b)^2\cdot 2^{-\mathcal{H}(\mathcal{C}_n)} =$$
$$(\ell+1)^2\cdot\kappa^2\cdot 2^{2b-\mathcal{H}(\mathcal{C}_n)}. \qquad (**)$$

If $p = \Pr[\mathrm{Coll}]$, we can combine (*) and (**) and conclude

$$\Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[m \neq SD(SE_f(m,h))] =$$
$$\Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[m \neq SD(SE_f(m,h)) \mid \overline{\mathrm{Coll}}]\cdot(1-p)$$
$$+ \Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[m \neq SD(SE_f(m,h)) \mid \mathrm{Coll}]\cdot p \leq$$
$$(\ell+1)\cdot\exp(-\kappa) + (\ell+1)^2\cdot\kappa^2\cdot 2^{2b-\mathcal{H}(\mathcal{C}_n)}.$$

We thus have

$$\left|\Pr_{k\leftarrow\{0,1\}^\kappa}[A^{F_{G(k)}(\cdot)}_{m,h}(1^\kappa)=1] - \Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[A^{f(\cdot)}_{m,h}(1^\kappa)=1]\right| =$$
$$\left|\Pr_{k\leftarrow\{0,1\}^\kappa}[m \neq SD(SE(k,m,h))] - \Pr_{f\leftarrow\boldsymbol{F}_{n,b}}[A^{f(\cdot)}_{m}(1^\kappa)=1]\right| \geq$$
$$\Pr_{k\leftarrow\{0,1\}^\kappa}[m \neq SD(SE(k,m,h))] -$$
$$(\ell+1)\cdot\exp(-\kappa) - (\ell+1)^2\cdot\kappa^2\cdot 2^{2b-\mathcal{H}(\mathcal{C}_n)}.$$

We can thus conclude

$$\mathbf{UnRel}_{\mathcal{S}',\mathcal{C}}(\kappa) \leq$$
$$\mathbf{InSec}^{\mathrm{prf}}_{F_G}\left((\ell+1)\cdot 2^b\cdot\kappa, t_C\cdot(\ell+1)\cdot 2^b\cdot\kappa, \kappa\right)$$
$$+ (\ell+1)\cdot\exp(-\kappa) + (\ell+1)^2\cdot\kappa^2\cdot 2^{2b-\mathcal{H}(\mathcal{C}_n)}.$$

The simulation of the call to SE_f can be carried out in time $\mathcal{O}((\ell+1)\cdot 2^b\cdot\kappa)$ with $2^b\cdot\kappa$ calls to the function oracle f. The running time of A_m is thus at most $\mathcal{O}((\ell+1)\cdot 2^b\cdot\kappa)+t$ and the number of queries of A is at most $(\ell+1)\cdot 2^b\cdot\kappa$. □

By combining Theorem 10, Theorem 11 and Theorem 14 together, we can conclude the existence of a secure black-box stegosystem (see Theorem 4 and 5 for an informal statement) and in particular:

THEOREM 15. *Let $\mathcal{C}$ be a channel and let α_1, α_2 be constants with $1 > \alpha_1 \geq \alpha_2 > 0$. Furthermore, let negl_G and*

$\mathsf{negl}_{G'}$ *be two negligible functions such that for every polynomial* t*, it holds that*

$$\mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^N,\mathcal{C}}(1,t,\kappa) \leq \mathsf{negl}_G(N),$$
$$\mathbf{InSec}^{\mathrm{dist}}_{G'(\{0,1\}^\kappa),\{0,1\}^B,\mathcal{C}}(1,t,\kappa) \leq \mathsf{negl}_{G'}(B).$$

Let $n(\kappa) = \kappa^{\alpha_1}$ *be the document length and the message space* $\mathcal{M}_\kappa = \mathcal{U}(\{0,1\}^{\ell\cdot b})$ *be with* $1 \leq b \leq \kappa^{\alpha_2}$*. If* $\mathcal{H}(\mathcal{C}_{n(\kappa)}) > 2b$ *and* $(\ell+1)\cdot b \leq \mathsf{pol}(\kappa)$*, then* $\mathcal{S}'$ *is a secure, reliable and rate-efficient stegosystem on* $\mathcal{C}$*.*

Proof. Recall that $N = 2^n \cdot b$ and $B = 2^b \cdot b$. Assume W is a Warden with $\mathbf{Adv}^{\mathrm{cha}}_{\mathcal{C},\mathcal{S},W}(\kappa) = \mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S},\mathcal{C}}(q,t,\kappa)$. Theorem 14 then implies that

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S},\mathcal{C}}(q,t,\kappa) \leq\\ &\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left(q(\ell+1)2^b\kappa\ ,\ (\ell+1)2^b\kappa + t\ ,\ \kappa\right) +\\ &q\cdot(\ell+1)\left(2^{b-\mathcal{H}(\mathcal{C}_n)} + \eta^{2^{b\cdot\kappa}}\right) + \frac{(\ell+1)^2}{2^b} + \frac{\ell\cdot b\cdot(q-1)}{b\cdot 2^b} +\\ &\mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}\left(\ell+1,(\ell+1)^2,\kappa\right) + 2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'},\mathcal{C}}(q,t,\kappa)\end{aligned}$$

for a constant $\eta < 1$. Note that the terms $q\cdot(\ell+1)\cdot\eta^{2^{b\cdot\kappa}}$, $\frac{(\ell+1)^2}{2^b}$, $\frac{\ell\cdot b\cdot(q-1)}{b\cdot 2^b}$ and $q\cdot(\ell+1)\cdot 2^{b-\mathcal{H}(\mathcal{C}_n)}$ are negligible in κ as $\mathcal{C}$ has sufficient min-entropy. There is thus a negligible function negl such that

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S},\mathcal{C}}(q,t,\kappa) \leq\\ &\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left(q(\ell+1)2^b\kappa\ ,\ q(\ell+1)\cdot 2^b\kappa + t\ ,\ \kappa\right) +\\ &\quad\mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}\left(\ell+1,(\ell+1)^2,\kappa\right) + 2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'},\mathcal{C}}(q,t\cdot,\kappa) +\\ &\quad\mathsf{negl}(\kappa).\end{aligned}$$

By using Theorem 11, we have

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left(q\cdot(\ell+1)\cdot 2^b\cdot\kappa, q\cdot\cdot(\ell+1)\cdot 2^b\cdot\kappa + t,\kappa\right) \leq\\ &\mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^N,\mathcal{C}}\big(1\ ,\ Nq(\ell+1)2^b\kappa +\\ &\qquad\qquad q(\ell+1)2^b\kappa + t\ ,\ \kappa\big).\end{aligned}$$

As $q,t,\ell,b \leq \mathsf{pol}(\kappa)$, we can bound this term by

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^N,\mathcal{C}}\big(1\ ,\ Nq(\ell+1)2^b\kappa +\\ &\qquad\qquad q(\ell+1)2^b\kappa + t\ ,\ \kappa\big) \leq\\ &\mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^{2^{n\cdot b}},\mathcal{C}}\left(1,\mathsf{pol}(N),\kappa\right).\end{aligned}$$

This insecurity is negligible by the assumption and there is thus a negligible function negl' such that

$$\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left(q\cdot(\ell+1)\cdot 2^b\cdot\kappa, q\cdot(\ell+1)\cdot 2^b\cdot\kappa + t,\kappa\right) \leq \mathsf{negl}'(\kappa).$$

Furthermore, Theorem 11 also implies

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}\left(\ell+1,(\ell+1)^2,\kappa\right) + 2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'},\mathcal{C}}(q,t,\kappa) \leq\\ &\mathbf{InSec}^{\mathrm{dist}}_{G'(\{0,1\}^\kappa),\{0,1\}^B,\mathcal{C}}\left(1,B\cdot(\ell+1)+(\ell+1)^2,\kappa\right)\\ &\quad + 2\,\mathbf{InSec}^{\mathrm{dist}}_{G'(\{0,1\}^\kappa),\{0,1\}^B,\mathcal{C}}\left(1,B\cdot q+t,\kappa\right)\end{aligned}$$

As $q,t,\ell,b \leq \mathsf{pol}(\kappa)$, we can bound this term by

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{dist}}_{G'(\{0,1\}^\kappa),\{0,1\}^B,\mathcal{C}}\left(1,B\cdot(\ell+1)+(\ell+1)^2,\kappa\right)\\ &\quad + \mathbf{InSec}^{\mathrm{dist}}_{G'(\{0,1\}^\kappa),\{0,1\}^B,\mathcal{C}}\left(1,B\cdot q+t,\kappa\right) \leq\\ &3\,\mathbf{InSec}^{\mathrm{dist}}_{G'(\{0,1\}^\kappa),\{0,1\}^B,\mathcal{C}}(1,\mathsf{pol}(N),\kappa)\end{aligned}$$

This insecurity is negligible in κ by assumption and there is thus a negligible function negl'' such that

$$\mathbf{InSec}^{\mathrm{prf}}_{F_{G'},\mathcal{C}}\left(\ell+1,(\ell+1)^2,\kappa\right) + 2\,\mathbf{InSec}^{\mathrm{prf}}_{F_{G'}}(q,t,\kappa) \leq \mathsf{negl}''(\kappa).$$

In conclusion, we have

$$\mathbf{InSec}^{\mathrm{cha}}_{\mathcal{S},\mathcal{C}}(q,t,\kappa) \leq \mathsf{negl}(\kappa) + \mathsf{negl}'(\kappa) + \mathsf{negl}''(\kappa).$$

The stegosystem $\mathcal{S}$ is thus secure on $\mathcal{C}$.

Concerning the unreliability, we can proceed similarly. Theorem 14 implies that

$$\begin{aligned}&\mathbf{UnRel}_{\mathcal{S},\mathcal{C}}(\kappa) \leq\\ &\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left((\ell+1)\cdot 2^b\cdot\kappa,(\ell+1)\cdot 2^b\cdot\kappa,\kappa\right) +\\ &\quad(\ell+1)\cdot\exp(-\kappa) + (\ell+1)^2\cdot\kappa^2\cdot 2^{2b-\mathcal{H}(\mathcal{C}_n)}.\end{aligned}$$

Due to sufficient min-entropy of $\mathcal{C}$ and the fact, that $\ell, b \leq \mathsf{pol}(\kappa)$, there is a negligible function negl such that

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left((\ell+1)\cdot 2^b\cdot\kappa,(\ell+1)\cdot 2^b\cdot\kappa,\kappa\right) +\\ &\quad(\ell+1)\cdot\exp(-\kappa) + (\ell+1)^2\cdot\kappa^2\cdot 2^{2b-\mathcal{H}(\mathcal{C}_n)} \leq\\ &\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left((\ell+1)\cdot 2^b\cdot\kappa,(\ell+1)\cdot 2^b\cdot\kappa,\kappa\right) + \mathsf{negl}(\kappa).\end{aligned}$$

As above, Theorem 11 shows that

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{prf}}_{F_G,\mathcal{C}}\left((\ell+1)\cdot 2^b\cdot\kappa,(\ell+1)\cdot 2^b\cdot\kappa,\kappa\right) + \mathsf{negl}(\kappa) \leq\\ &\mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^N,\mathcal{C}}\big(1\ ,\ N(\ell+1)2^b\kappa +\\ &\qquad\qquad(\ell+1)2^b\kappa\ ,\ \kappa\big) + \mathsf{negl}(\kappa).\end{aligned}$$

As $\ell, b \leq \mathsf{pol}(\kappa)$, this is bounded by

$$\begin{aligned}&\mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^N,\mathcal{C}}\big(1\ ,\ N(\ell+1)2^b\kappa +\\ &\qquad\qquad(\ell+1)2^b\kappa\ ,\ \kappa\big) + \mathsf{negl}(\kappa) \leq\\ &\mathbf{InSec}^{\mathrm{dist}}_{G(\{0,1\}^\kappa),\{0,1\}^N,\mathcal{C}}\left(1,\mathsf{pol}(N),\kappa\right) + \mathsf{negl}(\kappa).\end{aligned}$$

The insecurity is negligible in N by assumption and there is thus a negligible function negl' such that

$$\mathbf{UnRel}_{\mathcal{S},\mathcal{C}}(\kappa) \leq \mathsf{negl}'(\kappa).$$

The stegosystem $\mathcal{S}$ is thus reliable on $\mathcal{C}$.

As we embed $b \leq \kappa^{\alpha_2}$ bits into a single document, the transmission rate $b(\kappa)$ is equal to b. As $\mathcal{H}(\mathcal{C}_n) \leq \kappa^{\alpha_1}$, the stegosystem $\mathcal{S}$ is rate-efficient on $\mathcal{C}$, as α_1, α_2 are constants and as long as b is large enough. □

Note that the precondition concerning the negligible functions $\mathsf{negl}_G, \mathsf{negl}_{G'}$ is always fulfilled, if the channel oracle can be simulated in time poly(N). This is due to Theorem 11, that states the security of the pseudorandom function.

Dedić et al. [6] introduced the family $\mathcal{F}$ of *pseudorandom flat-h channels* and proved that every time- and rate-efficient stegosystem for $\mathcal{F}$ is either insecure or unreliable (under the cryptographic assumption that efficient pseudorandom functions exist). The running-time of every rate-efficient, secure and reliable stegosystem for $\mathcal{F}$ is thus at least super-polynomial. To the best of our knowledge, there is no such stegosystem known for $\mathcal{F}$. We therefore give the first secure, reliable and rate-efficient stegosystem for this family of channels:

COROLLARY 1. *The stegosystem $\mathcal{S}'$ is rate-efficient, secure and reliable on the family of pseudorandom flat-h channels.*

6. UNCONDITIONAL LOWER BOUND

In order to give an unconditional lower bound, we make use of a lower bound by Dedić et al. [6]. By providing the warden W with an efficient test, whether a document belongs to the support of the channel, they prove:

THEOREM 16 ([6]). *For every universal (not necessarily of polynomial-time complexity) stegosystem $\mathcal{S}(\kappa)$ which hides $\lambda := \lambda(\kappa)$ bits and takes $q := q(\kappa)$ samples per stego-document there exists a family of channels $\mathcal{C}(\kappa)$ such that*

$$\mathbf{InSec}^*(\kappa) + \mathbf{UnRel}(\kappa) \;\geq\; \frac{1}{2} - \frac{e \cdot q}{2^\lambda} - o(1),$$

where $\mathbf{InSec}^*$ *denotes the insecurity against polynomial wardens with an auxiliary* oracle for testing membership *in the support of $\mathcal{C}(\kappa)$, and* $\mathbf{UnRel}$ *denotes the unreliability of $\mathcal{S}(\kappa)$ on $\mathcal{C}(\kappa)$.*

Dedić et al. then argue that the assumption that a warden has an oracle for membership-testing is not feasible, if the channel is chosen completely random. By making use of the fact that the warden can choose a history, while the stegoencoder can not, we will show how an *efficient* warden is able to test membership of a completely random channel.

Let $\boldsymbol{S}_n$ be the set of all subsets of $\{0,1\}^n$ of cardinality $n/2$. For $S \in \boldsymbol{S}_n$, let $\mathcal{C}_S$ be the following channel, where $\vec{1}$ denotes the vector of length n that contains a 1 at every position:

- $\mathcal{C}(S)_{\varnothing,n}$ is the uniform distribution on $\{0,1\}^n$.
- $\mathcal{C}(S)_{\vec{1}||d,n}$ is the uniform distribution on all strings in $\{0,1\}^n$ that start with 1, if $d \in S$ or the uniform distribution on all strings in $\{0,1\}^n$ that start with 0, if $d \notin S$ (i.e. the first position indicates the membership of d in S).
- $\mathcal{C}(S)_{h,n}$ is the uniform distribution on S for all other histories.

The warden W for the family $\{\mathcal{C}(S) \mid S \in \boldsymbol{S}_n\}_{n\in\mathbb{N}}$ now works as follows: It randomly chooses a history $h \leftarrow \{0,1\}^n \setminus \{\vec{1}\}$ and $m = 00\ldots0$ – a message containing ℓ 0-bits – and gets the results $d_1, d_2, \ldots, d_\ell$ from the challenging oracle on h and m. For $i \in \{1,\ldots,\ell\}$, it takes a sample $s_i \leftarrow \mathcal{C}(S)_{\vec{1}||d_i,n}$. If every s_i starts with 1, the warden returns "Non-Stego" and else "Stego". The warden W is thus able to test membership in S efficiently by making use of the channel. Note that the stegoencoder can not make use of these capabilities of $\mathcal{C}(S)$ as it can only make queries to $\mathcal{C}_{h,n}$, where h does not start with $\vec{1}$. We use here the definition for channel access as in [15], which assumes that the encoder has an access to the marginal channel distributions $\mathcal{C}_h$ for the histories h started with adversarially chosen prefixes.

We can thus efficiently simulate an oracle for membership-testing and Theorem 16 thus implies (the formal statement of Theorem 6 in Introduction):

THEOREM 17. *For every universal (not necessarily polynomial-time bounded) stegosystem $\mathcal{S}(\kappa)$ which hides $\lambda := \lambda(\kappa)$ bits and takes $q := q(\kappa)$ samples per stego-document there exists a family of channels $\mathcal{C}(\kappa)$ such that*

$$\mathbf{InSec}(\kappa) + \mathbf{UnRel}(\kappa) \;\geq\; \frac{1}{2} - \frac{e \cdot q}{2^\lambda} - o(1),$$

where $\mathbf{InSec}$ *denotes the insecurity against polynomial wardens and* $\mathbf{UnRel}$ *denotes the unreliability of $\mathcal{S}(\kappa)$ on $\mathcal{C}(\kappa)$.*

Note that in contrast to Theorem 3, no cryptographic assumption is necessary and in contrast to Theorem 16, no membership-oracle is necessary. Our lower bound thus holds unconditionally. Furthermore, this lower bound holds even when the running time of the stegosystem $\mathcal{S}$ is much larger (say 2^{2^κ}) than the running time of W (say $\mathsf{pol}(\kappa)$).

Note that this method only works because of the asymmetry between Alice and Warden: While Warden has oracle-access for all possible histories, Alice can only use the history chosen by Warden.

7. CONCLUSIONS AND FURTHER WORK

We first gave a secure, reliable and rate-efficient stegosystem by using pseudorandom functions of very high hardness. The running time of the stegosystem is roughly $2^{2^{o(\kappa)}}$. The work of Dedić et al. [6] gives the best known lower bound of a running time of $\omega(\mathsf{pol}(\kappa))$ for any universal secure, reliable stegosystem (under cryptographic assumptions and of the rate $\omega(\log\kappa)$). We proved that by making use of the imbalance between encoder and warden, this lower bound also holds without any assumption and for any rate-efficient stegosystem. This immediately gives rise to the question, whether one can shrink the gap between $2^{2^{o(\kappa)}}$ and $\omega(\mathsf{pol}(\kappa))$, by either giving a more efficient stegosystem or by the construction of more difficult channels. If the requirements of a universal stegosystem seem too strict, one can proceed similarly to Liśkiewicz et al. [22] and try to find secure, reliable, and rate-efficient stegosystems for a large family $\mathcal{F}$ of channels. It is still open how complex $\mathcal{F}$ can be for secure, reliable, efficient and rate-efficient stegosystems.

We also showed that the common phrase "Steganography is Cryptography" is provably wrong as the communication channel is a very important part of the steganographic setting. We would hope that this motivates other authors to conduct theoretical and practical research on this fascinating topic.

8. REFERENCES

[1] M. Backes and C. Cachin. Public-key steganography with active attacks. In *Theory of Cryptography*, pages 210–226. Springer, 2005.

[2] F. Balado and D. Haughton. Optimum perfect universal steganography of finite memoryless sources. *CoRR*, abs/1410.2659, 2014.

[3] M. Bellare, A. Desai, E. Jokipii, and P. Rogaway. A concrete security treatment of symmetric encryption. In *Proc. FOCS, 1997*, pages 394–403. IEEE, 1997.

[4] C. Cachin. An information-theoretic model for steganography. *Information and Computation*, 192(1):41–56, 2004.

[5] P. Comesaña and F. Pérez-González. On the capacity of stegosystems. In *Proceedings of the 9th workshop on Multimedia & security*, pages 15–24. ACM, 2007.

[6] N. Dedić, G. Itkis, L. Reyzin, and S. Russell. Upper and lower bounds on black-box steganography. *Journal of cryptology*, 22(3):365–394, 2009.

[7] T. Filler and J. Fridrich. Fisher information determines capacity of ε-secure steganography. In *Information Hiding*, pages 31–47. Springer, 2009.

[8] T. Filler, A. D. Ker, and J. J. Fridrich. The square root law of steganographic capacity for markov covers. In *Media Forensics and Security I, part of the IS&T-SPIE Electronic Imaging Symposium, San Jose, CA, USA, January 19, 2009, Proceedings*, page 725408, 2009.

[9] J. Fridrich. *Steganography in digital media: principles, algorithms, and applications*. Cambridge University Press, 2009.

[10] O. Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.

[11] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.

[12] O. Goldreich and H. Krawczykt. Sparse pseudorandom distributions. *Random Structures and Algorithms*, 3(2), 1992.

[13] N. J. Hopper. Toward a theory of steganography. Technical report, Technical Report CMU-CS-04-157, Carnegie Mellon Univ., 2004.

[14] N. J. Hopper, J. Langford, and L. Ahn. Provably secure steganography. In *Proc. CRYPTO 2002*, volume 2442 of *LNCS*, pages 77–92. Springer Berlin Heidelberg, 2002.

[15] N. J. Hopper, L. von Ahn, and J. Langford. Provably secure steganography. *Computers, IEEE Transactions on*, 58(5):662–676, 2009.

[16] S. Katzenbeisser and F. A. Petitcolas. Defining security in steganographic systems. In *Electronic Imaging 2002*, pages 50–56. SPIE, 2002.

[17] A. D. Ker. The square root law in stegosystems with imperfect information. In *Information Hiding - 12th International Conference, IH 2010, Calgary, AB, Canada, June 28-30, 2010, Revised Selected Papers*, pages 145–160, 2010.

[18] A. D. Ker, P. Bas, R. Böhme, R. Cogranne, S. Craver, T. Filler, J. Fridrich, and T. Pevnỳ. Moving steganography and steganalysis from the laboratory into the real world. In *Proceedings of the first ACM workshop on Information hiding and multimedia security*, pages 45–58. ACM, 2013.

[19] A. D. Ker, T. Pevný, J. Kodovský, and J. J. Fridrich. The square root law of steganographic capacity. In *Proceedings of the 10th workshop on Multimedia & Security, MM&Sec 2008, Oxford, UK, September 22-23, 2008*, pages 107–116, 2008.

[20] A. Kiayias, Y. Raekow, A. Russell, and N. Shashidhar. A one-time stegosystem and applications to efficient covert communication. *J. Cryptology*, 27(1):23–44, 2014.

[21] T. V. Le and K. Kurosawa. Bandwidth optimal steganography secure against adaptive chosen stegotext attacks. In *Proc. Information Hiding, 2006*, pages 297–313, 2006.

[22] M. Liśkiewicz, R. Reischuk, and U. Wölfel. Grey-box steganography. *Theoretical Computer Science*, 505:27–41, 2013.

[23] B. Ryabko and D. Ryabko. Constructing perfect steganographic systems. *Information and Computation*, 209(9):1223–1230, 2011.

[24] Y. Wang and P. Moulin. Perfectly secure steganography: Capacity, error exponents, and code constructions. *Information Theory, IEEE Transactions on*, 54(6):2706–2722, 2008.