Efficient Public Key Steganography Secure Against Adaptively Chosen Stegotext Attacks

Tri Van Le[1] $^{\infty}$ and Kaoru Kurosawa[2]

[1] Department of Computer Science
Florida State University
Tallahassee, Florida 32306-4530, USA.
Email: levan@cs.fsu.edu.

[2] Department of Computer and Information Sciences
Ibaraki University 4-12-1
Nakanarusawa, Hitachi, Ibaraki 316-8511, Japan
Email: kurosawa@cis.ibaraki.ac.jp.

Abstract. We define the notion of adative chosen stegotext security. We then construct *efficient* public key steganographic schemes secure against adaptively chosen stegotext attacks, without resort to any special existence assumption such as unbiased functions. This is the first time such a construction is obtained. Not only our constructions are *secure*, but also are essentially optimal and have *no error* decoding. We achieve this by applying a primitive called $\mathcal{P}$-codes.

Keywords: foundation, steganography, public key, computational security, coding theory.

1 Introduction

Motivations. The *Prisoner's Problem* introduced by G.J. Simmons [14] and generalized by R. Anderson [1] can be stated informally as follows: Two prisoners, Alice and Bob, want to communicate to each other their secret escape plan under the surveillance of a warden, Wendy. In order to pass Wendy's censorship, Alice and Bob have to keep their communications as innocent as possible so that they will not be banned by Wendy.

Existing results. Previously, the Prisoner's Problem was considered in the secret key setting by: Cachin [3], Mittelholzer [11], Moulin and Sullivan [12], Zollner et.al. [16] in the unconditional security model; and Katzenbeisser and Petitcolas [10], Hopper et.al. [8], Reyzin and Russell [13] in the conditional security model. In this article, we consider the problem in the *public key* setting. In this setting, Craver [4] and Anderson[1] proposed several heuristic ideas to solve the problem. Katzenbeisser and Petitcolas [10] gave a formal model. Hopper and Ahn [9] constructed proven secure schemes assuming the existence of *unbiased functions* (more details on page **??**). Unbiased functions are currently required for *all public key schemes* in the literature [9], and are the main ingredients in most of other secure generic steganographic schemes [3, 8, 13]. Further, current approaches [3, 8, 9, 13] result in very low information rate for steganography in both the public key and secret key settings. There appear some other work independent of ours, which also consider adaptively chosen stegotext security [15, 2]. However, they admit the same draw back of extremely low information rate and the special assumption of unbiased functions as before. We show here our new paradigm for constructing secure steganography. Not only our construction is very efficient, it also achieves provable adaptive chosen stegotext security and do not require special assumptions.

$^{\infty}$ The author was partially supported by NSF grant 9903216.

Organization. The article is organized as follows: we describe the model in Section 2, our new primitive $\mathcal{P}$-codes in Section 3, show constructions of public key steganographic schemes and their security proofs in Section 4, and give a rate calculation for our schemes in Section 5. We conclude in Section 6.

Our Solution. We solve the steganographic problem in a novel way. At the heart of our solution are *uniquely decodable variable length coding schemes Γ, called $\mathcal{P}$-codes, with source alphabet Σ and destination alphabet C such that: if $x \in \Sigma^\infty$ is chosen uniformly randomly then $\Gamma(x) \in C^\infty$ distributes according to $\mathcal{P}$, where $\mathcal{P}$ is a given distribution over sequences of covertexts.*

Note that such a coding scheme is quite related to homophonic coding schemes [7], which are uniquely decodable variable length coding scheme Γ' with source alphabet C and destination alphabet Σ such that: if $c \in C^*$ is chosen randomly according to distribution $\mathcal{P}$ then $\Gamma'(c) \in \Sigma^*$ is a sequence of independent and uniformly random bits.

Of course, one can hope that such a homophonic coding scheme Γ' will give rise to a uniquely decodable $\mathcal{P}$-code Γ. However, this is not necessarily true because Γ' can map one-to-many, as in the case of [7]. Therefore by exchanging the encoding and decoding operations in Γ', we will obtain a non-uniquely decodable $\mathcal{P}$-coding scheme Γ'', which is not what we need.

To construct these $\mathcal{P}$-codes, we generalize a heuristic idea of Ross Anderson [1] where one can use a perfect compression scheme on the covertexts to obtain a perfectly secure steganographic scheme. Nevertheless, in practice one can never obtain a perfect encryption scheme, so we have to build our $\mathcal{P}$-coding schemes based on non-perfect compression schemes, such as arithmetic compression. The result is a coding scheme which achieve near optimal information rate, and has no error.

2 Definitions

2.1 Channel

Let C be a finite *message space*. A *channel* $\mathcal{P}$ is a probability distribution over the space C^∞ of infinite message sequences $\{(c_1, c_2, \dots) \mid c_i \in C, i \in \mathbb{N}\}$. The communication channel $\mathcal{P}$ may be *stateful*. This means that: for all $n > 0$, c_n might depend probabilistically on $c_1, \dots, c_{n-1}$. When individual messages are used to embed hiddentexts, they are called covertexts. Therefore C is also called the covertext space. Denote C^* the space of all finite message sequences $\{(c_1, \dots, c_l) \mid l \in \mathbb{N}, c_i \in C, 1 \le i \le l\}$. If $h \in C^*$ is a prefix of $s \in C^\infty$, that is $s_i = h_i$ for all $1 \le i < \ell(h)$, then we write $h \subset s$. The expression $s \in_{\mathcal{P}} C^\infty$ means that s is chosen randomly from C^∞ according to distribution $\mathcal{P}$. Denote $\mathcal{P}(c) = \Pr[c \subset s \mid s \in_{\mathcal{P}} C^\infty]$ for all $c \in C^*$.

Sampler. A *sampler* S for the channel $\mathcal{P}$ is a sampling oracle such that upon a query $h \in C^*$, S randomly outputs a message $c_i \in C$ according to the marginal probability distribution $\mathcal{P}_h$:

$$\mathcal{P}_h(c_i) = \Pr[(h \| c_i) \subset s \mid h \subset s \;\wedge\; s \in_{\mathcal{P}} C^\infty],$$

where $h \| c_i$ is the concatenation of h and c_i. In general, we define $\mathcal{P}_h(c) = \Pr[(h \| c) \subset s \mid h \subset s \;\wedge\; s \in_{\mathcal{P}} C^\infty]$ for all $h \in C^*$ and $c \in C^* \cup C^\infty$. The expression $s = S(h)$ means s is the result of querying $S(h)$. Since S responses randomly, each individual query may have a different result. Finally, $x \in_R X$ means x is chosen uniformly randomly from the set X. Finite messages sequences can always be included in $\mathcal{P}$ by appending copies of a special *null* symbol to get infinite sequences.

Assumption. From now on, we assume that $\mathcal{P}$ is a channel over message space C, and that a corresponding sampler S is given. The channel $\mathcal{P}$ represents the probability distribution of an innocent communication channel; the sampler S generates covertexts according to $\mathcal{P}$, see [1, 8, 10]. Our purpose is to construct steganographic systems whose stegotext distributions are indistinguishable from $\mathcal{P}$. We also assume that the query h given to sampler S is always the history of messages communicated between Alice and Bob.

2.2 Steganographic systems

A public key *steganographic system* is specified by a pair of key spaces $\mathcal{K}_e \times \mathcal{K}_d$, and three randomized algorithms, `Setup`, `Embed`, `Extract`, that works as follows:

- `Setup`: takes a security parameter k as input, and returns system parameters `params` and a pair of keys $(e, d) \in \mathcal{K}_e \times \mathcal{K}_d$. Among other things, the system parameters `params` include a short description of a finite hiddentext space $\mathcal{M}$.
- `Embed`: takes as input a public key $e \in \mathcal{K}_e$, a hiddentext $m \in \mathcal{M}$, and returns a stegotext $s \in C$. The algorithm may query the sampler S.
- `Extract`: takes as input a secret key $d \in \mathcal{K}_d$, a stegotext $s \in C$, and returns either the symbol $\perp$ on failure, or a hiddentext $m \in \mathcal{M}$.

As usual, we require that $\texttt{Extract}(d, \cdot)$ reverses the action of $\texttt{Embed}(e, \cdot)$.

2.3 Adaptively Chosen Stegotext Attacks

The task of warden Wendy is to distinguish two cases: whether the communications between the prisoners are innocent, or contain hiddentexts. In order to detect hiddentexts, Wendy is allowed to mount chosen stegotext attacks, which are plausible in practice when Wendy has oracle access to the extraction machine and would like to detect who uses this machine to communicate steganographically.

Chosen stegotext attacks on steganographic systems are parallel to chosen ciphertext attacks on encryption systems. The only difference is in the purposes of the two attacks. In the first attack, the objective is to detect the existence of hidden messages or hiddentexts. In the second attack, the objective is to discover partial information about the content of the secret messages. Our definition of *chosen stegotext security* reflects this difference:

- In an indistinguishability under (adaptive) chosen ciphertext attack (IND-CCA), the challenger randomly chooses one of the two plaintexts submitted by the adversary and encrypts it. An encryption scheme is secure against this attack if an adversary cannot tell which plaintext was encrypted.
- In a hiding under (adaptive) chosen stegotext attack (HID-CSA), the challenger randomly flips a coin, and depending on the result decides to encrypt the submitted hiddentext or to randomly sample a cover message. A steganographic scheme is secure against this attack if an adversary cannot tell stegotexts from covertexts.

While the hiding objective of steganographic systems is substantially different from the semantic security objective of encryption systems, we shall show later that HID-CSA security implies IND-CCA.

Formally, we say that a steganographic system is secure against an adaptive chosen stegotext attack if no polynomial time adversary $\mathcal{W}$ has non-negligible advantages against the challenger in the following game:

– Setup: The challenger takes a security parameter k and runs `Setup` algorithm. It gives the resulting system parameters `params` and public key e to the adversary, and keeps the secret key d to itself.
– Phase 1: The adversary issues j queries $c_1, \ldots, c_j$ where each query c_i is a covertext in $\mathcal{C}$. The challenger responds to each query c_i by running `Extract` algorithm with input secret key d and message c_i, then sending the corresponding result of $\texttt{Extract}(d, c_i)$ back to the adversary. The queries may be chosen adaptively by the adversary.
– Challenge: The adversary stops Phase 1 when it desires, and sends a hiddentext $m \in \mathcal{M}$ to the challenger. The challenger then picks a random bit $b \in \{0, 1\}$ and does the following:
  - If $b = 0$, the challenger queries S for a covertext s, and sends $s = S(h)$ back to the adversary.
  - If $b = 1$, the challenger runs the `Embed` algorithm on public key e and plaintext m, and sends the resulting stegotext $s = \texttt{Embed}(e, m)$ back to the adversary.
– Phase 2: The adversary makes additional queries $c_{j+1}, \ldots, c_q$ where each query $c_i \neq c$ is a covertext in $\mathcal{C}$. The challenger responds as in Phase 1.
– Guess: The adversary outputs a guess $b' \in \{0, 1\}$. The adversary wins the game if $b' = b$.

Such an adversary $\mathcal{W}$ is called an HID-CSA attacker. We define the adversary $\mathcal{W}$'s advantage in attacking the system as $|\Pr[b' = b] - \frac{1}{2}|$ where the probability is over the random coin tosses of both the challenger and the adversary.

We remind you that a standard IND-CCA attacker would play a different game, where at the challenge step:

– Challenge: The adversary sends a pair of plaintexts $m_0, m_1 \in \mathcal{M}$ upon which it wishes to be challenged to the challenger. The challenger then picks a random bit $b \in \{0, 1\}$, runs the encryption algorithm on public key e and plaintext m_b, and sends the resulting ciphertext $c = \texttt{Encrypt}(e, m_b)$ back to the adversary.

We note that a HID-CHA game is a restriction of the HID-CSA game where the adversary makes $q = 0$ queries.

As in IND-CCA game against an encryption system, we also define an IND-CCA game against a steganographic system. The definition is exactly the same, except with necessary changes of names: the `Encrypt` and `Decrypt` algorithms are replaced by the `Embed` and `Extract` algorithms; and the terms plaintext and ciphertext are replaced by the terms hiddentext and stegotext, respectively. Similarly, a steganographic system is called IND-CCA secure if every polynomial time adversary $\mathcal{W}$ has negligible advantages in an IND-CCA game against the steganographic system.

3 Construction of $\mathcal{P}$-Codes

A uniquely decodable coding scheme Γ is a pair consisting of a probabilistic encoding algorithm Γ_e and a deterministic decoding algorithm Γ_d such that $\forall m \in dom(\Gamma_e) : \Gamma_d(\Gamma_e(m)) = m$. In this article, we are interested in coding schemes whose source alphabet is binary, $\Sigma = \{0, 1\}$.

Definition 1. *Let $\mathcal{P}$ be a channel with message space C. A $\mathcal{P}$-code, or a $\mathcal{P}$-coding scheme, is a uniquely decodable coding scheme Γ whose encoding function $\Gamma_e : \Sigma^* \to C^*$ satisfies:*

$$\epsilon(n) = \sum_{c \in \Gamma_e(\Sigma^n)} \left| \Pr\left[\Gamma_e(x) = c \mid x \in_R \Sigma^n\right] - \mathcal{P}(c) \right|$$

is a negligible function in n. In other words, the distribution of $\Gamma_e(x)$ is statistically indistinguishable from $\mathcal{P}$ when x is chosen uniformly randomly. The function

$$e(n) = \frac{1}{n} \sum_{c \in \Gamma_e(\Sigma^n)} \mathcal{P}(c) H_{\mathcal{P}}(c)$$

is called the expansion rate of the encoding. [3]

Let $\mathcal{P}$ be a channel with sampler S. We assume here that $\mathcal{P}_h$ is polynomially sampleable, which was also assumed in [8, 9] in order to achieve proven security.[4] This is equivalent to saying that S is an efficient algorithm that given a sequence of covertexts $h = (c_1, \ldots, c_n)$ and a uniform random string $r \in_R \{0,1\}^{R_n}$, S outputs a covertext $c_{n+1} \in C$ accordingly to probability distribution $\mathcal{P}_h$. Nevertheless, we assume less that the output of S to be statistically close to $\mathcal{P}_h$.

We use algorithm S to construct a $\mathcal{P}$-coding scheme Γ. For $x = (x_1, \ldots, x_n) \in \Sigma^n$, denote $\overline{x}$ the non-negative integer number whose binary representation is x. For $0 \le a \le 2^n$, denote $\underline{a} = (a_1, \ldots, a_n)$ the binary representation of integer number a. In the following, let t be an integer parameter, h_0 is the history of all previous communicated messages between Alice and Bob. Further let us assume that the distribution $\mathcal{P}_h$ has minimum entropy bounded from below by a constant $\xi > 0$. Let G be a cryptographically secure pseudo-random generator.

Γ_1-*`Encode`*. **Input**: $z \in_R \{0,1\}^{R_n}$, $x = (x_1, \ldots, x_n) \in \Sigma^n$.
Output: $c = (c_1, \ldots, c_l) \in C^*$.

1. **let** $a = 0, b = 2^{2n}, h = \epsilon$.
2. **let** z be the seed to initialize G.

[3] Ideally, we would have used $\Pr\left[\Gamma_e(x) = c \mid x \in_R \Sigma^n\right]$ instead of $H_{\mathcal{P}}(c)$. However, the two distributions are statistically indistinguishable so this makes no real difference.
[4] Theoretically, allowing $\mathcal{P}_h$ to be non-polynomially sampleable would allow hard problems to be solvable.

3. **while** $\lceil a/2^n \rceil < \lfloor b/2^n \rfloor$ **do**
   (a) **let** $c_i^* = S(h_0 \| h, G)$ for $0 \le i < t$.
   (b) Order the c_i^*'s in some fixed increasing order:
   $$c_0^* = \cdots = c_{i_1-1}^* < c_{i_1}^* = \cdots = c_{i_2-1}^* < \cdots < c_{i_{m-1}}^* = \cdots = c_{t-1}^*,$$
   where $0 = i_0 < i_1 < \cdots < i_m = t$.
   (c) **let** $0 \le j \le m-1$ be the unique j such that
   $$i_j \le \lfloor (2^n \overline{x} - a)t/(b-a) \rfloor < i_{j+1}.$$
   (d) **let** $a' = a + (b-a)i_j/t$, $b' = a + (b-a)i_{j+1}/t$.
   (e) **let** $(a, b) = (a', b')$.
   (f) **let** $h = h|c_{i_j}^*$.
4. Output $c = h$.

Everyone who is familiar with information theory will immediately realize that the above encoding resembles to the arithmetic decoding of number $\overline{x}/2^n$. In fact, the arithmetic encoding of the sequence c is exactly the number $\overline{x}/2^n$.

Each time the sender outputs a covertext $c_{i_j}^*$, the receiver will obtain some information about the message x, i.e. the receiver is able to narrow the range $[a, b]$ containing $2^n x$. The sender stops sending more covertexts until the receiver can completely determine the original value x, i.e. when the range $[a, b]$ is less than 2^n. Thus the decoding operation for the $\mathcal{P}$-coding scheme Γ follows.

Γ_1-`Decode`. **Input**: $z \in_R \{0,1\}^{R_n}$, $c = (c_1, \ldots, c_l) \in C^*$.
Output: $x = (x_1, \ldots, x_n) \in \Sigma^n$.

1. **let** $a = 0, b = 2^{2n}, h = \epsilon$.
2. **let** z be the seed to initialize G.
3. **for** *ind* **from** 0 **to** $|c| - l_0 - 1$ **do**
   (a) **let** $c_i^* = S(h_0 \| h, G)$ for $0 \le i \le t-1$.
   (b) Order the c_i^*'s in some fixed increasing order:
   $$c_0^* = \cdots = c_{i_1-1}^* < c_{i_1}^* = \cdots = c_{i_2-1}^* < \cdots < c_{i_{m-1}}^* = \cdots = c_{t-1}^*,$$
   where $0 = i_0 < i_1 < \cdots < i_m = t$.
   (c) **let** $0 \le j \le m-1$ be the unique j such that $c_{i_j}^* = c_{(l_0+1+ind)}$.
   (d) **let** $a' = a + (b-a)i_j/t$, $b' = a + (b-a)i_{j+1}/t$.
   (e) **let** $(a, b) = (a', b')$.
   (f) **let** $h = h|c_{i_j}^*$.
4. **let** $v = \lceil a/2^n \rceil$.
5. Output $x = \underline{v}$.

If x is chosen uniformly randomly from Σ^n then the correctness of our $\mathcal{P}$-coding scheme Γ is established through the following theorem.

Theorem 1. *Γ_1 described above is a $\mathcal{P}$-code.*

Proof. First, the values of $i_0, \ldots, i_t, j, a', b', h, a, b$ in the encoding are the same as in the decoding. Further, due to our choice of j, $2^n \overline{x} \in [a, b)$ is true not only before the iterations, but also after each iteration. Therefore at the end of the encoding, we obtain $\lceil a2^{-n} \rceil = \lfloor b2^{-n} \rfloor = \overline{x}$.

Because the values of a, b in encoding are the same as in decoding, this shows that the decoding operation's output is the same as the encoding operation's input x, i.e. Γ_1 is uniquely decodable. Next, we will prove that it is also a $\mathcal{P}$-code.

Indeed, let us assume temporarily that a, b were real numbers. Note that the covertexts $c_0^*, \ldots, c_{t-1}^*$ are generated independently of x, so $i_0, \ldots, i_t$ are also independent of x. By simple induction we can see that after each iteration $i \leq l-1$, the conditional probability distribution of $\overline{x}$ given the history $h = c_1 \| \ldots \| c_i$, is uniformly random over integers in the range $[a2^{-n}, b2^{-n})$. However, in our algorithms the numbers a, b are represented as integers using rounding. So the conditional distribution of $\overline{x}$ at the end of each iteration except the last one is not uniformly random, but anyway at most $4/(b-a) \leq 2^{2-n}$ from uniformly random due to rounding, and due to the fact that $b - a \geq 2^n$. Since 2^{2-n} is negligible, and our encoding operations are polynomial time, they can not distinguish a truly uniformly random $\overline{x}$ from a statistically-negligible different one. So for our analysis, we can safely assume that $\overline{x}$ is indeed uniformly random in the range $[a2^{-n}, b2^{-n})$ at the *beginning* of each iteration, including the last one.

Then at the beginning of each iteration i, conditioned on the previous history $h = c_0 \| \ldots \| c_{i-1}$, $u = \lfloor (2^n \overline{x} - a)t/(b-a) \rfloor$ is a uniformly random variable on the range $[0, t-1]$, thus u is probabilistically independent of $c_0^*, \ldots, c_{t-1}^*$. Since $c_0^*, \ldots, c_{t-1}^*$ are identically distributed, c_u must also be distributed identically. Further, by definition, $i_j \leq k < i_{j+1}$, so $c_u = c_{i_j}^* = c_i$. Hence c_i distributes identically as each of $c_0^*, \ldots, c_{t-1}^*$ does. By definition of S, this distribution is $\mathcal{P}_{h_0 \| h}$, i.e. c distributes accordingly to $\mathcal{P}_{h_0}$. Since x is not truly uniformly random but rather statistically indistinguishable from uniformly random, we conclude that the output c of the encoding operation is statistically indistinguishable from $\mathcal{P}_{h_0}$. Therefore, by definition, our coding scheme is indeed a $\mathcal{P}$-code. Our coding scheme has a small overhead rate of $l_0/n = k/n\xi$. However, this overhead goes to 0 when $n > k^{1+\epsilon}$ as $n \to \infty$ and $\epsilon > 0$. Therefore our encoding is essentially optimal. See our formal proof in Section 5.

Note that in the case that $m = 0$, the encoding/decoding operations still work correctly, i.e. there are no errors. In such case, the range $[a, b)$ does not change: the encoding will output c_0^* without actually embedding any hidden information, while the decoding operation will read c_0^* without actually extracting any hidden information. This happens more often when the entropy of the cover distribution is very near zero. However, from now on we will assume that our distribution $\mathcal{P}_h$ will have minimal entropy bounded from below by a fixed constant $1 > \rho > 0$, i.e. $\forall h \in C^*, c \in C : \mathcal{P}_h(c) < \rho$. Then with overwhelming probability of at least $1 - |C|\rho^t$, we will have $m > 0$.

4 Construction of Public Key Steganographic Systems

Our purpose in this section is to construct steganographic systems based on the $\mathcal{P}$-coding scheme Γ. Using the notations from Sections 2 and 3, our construction is the following. Here, h denotes the history of previously communicated messages.

4.1 Public Key Steganographic Systems

We use the idea of Diffie-Hellman key exchange to obtain an efficient public key steganographic scheme. Denote $H_{\mathcal{P}}(c) = -\log_2(\mathcal{P}(c))$ the entropy of $c \in C^*$ according to the covertext distribution $\mathcal{P}$. We assume that there exists a constant $0 < \rho < 1$ such that:

$$\forall h \in C^*, \forall c \in C : \mathcal{P}_h(c) < \rho.$$

In other words, $\mathcal{P}_h$ has its minimum entropy bounded from below by a positive constant ($-\log_2(\rho)$). Furthermore, let $D = (\texttt{Setup}, \texttt{Sign}, \texttt{Verify})$ be a secure digital signature scheme.

S_1-***Setup.*** Call D-`Setup` to generate a key pair (k_{sig}, k_{ver}). The system parameter is a generator g of a prime order cyclic group $<g>$, whose decisional Diffie-Hellman problem is hard. Let (g, g^a, k_{ver}) be the public key of sender Alice, and (g, g^b, k'_{ver}) be the public key of receiver Bob. Let $F(X, Y)$ be a public cryptographically secure family of pseudo-random functions, indexed by variable $X \in <g>$. Let k be the security parameter and $n = O(poly(k))$. The embedding and extracting operations are as follows.

S_1-***Embed.*** **Input**: $m \in \{0,1\}^n$.
Output: $c \in C^*$.

1. Let $l = \lceil \frac{k}{\log_2 \frac{1}{\rho}} \rceil$, $h_0 = \epsilon$.
2. **for** i **from** 1 **to** l **do** $c_i = S(h_0)$; $h_0 = h_0 \| c_i$.
3. **let** $r \| z = F((g^b)^a, h_0)$.
4. **let** $m' = m \| \texttt{Sign}(k_{sig}, m)$.
5. Output $c = h_0 \| \mathit{\Gamma}_e(z, r \oplus m')$.

Note that in the call to $\mathit{\Gamma}_e(r \oplus m')$, we initialize h with h_0, instead of ϵ.

S_1-***Extract.*** **Input**: $c \in C^*$.
Output: $m \in \{0,1\}^n$.

1. Let $l = \lceil \frac{k}{\log_2 \frac{1}{\rho}} \rceil$, $c = (h_0, c')$ where $|h_0| = l$.
2. **let** $r \| z = F((g^a)^b, h_0)$.
3. **let** $m' = \mathit{\Gamma}_d(z, c') \oplus r$.
4. **if** $\texttt{Verify}(k_{ver}, m') \neq \mathit{success}$ **then** return $\perp$.
5. Parse $m' = m \| \texttt{Sign}(k_{sig}, m)$.
6. Output m.

Similarly to the construction $F((g^a)^b, \cdot)$, the secretly shared family of pseudo random function H used in $\mathit{\Gamma}_e, \mathit{\Gamma}_d$ can be constructed from a public family H' with index g^{ab}, e.g. using $H(X, Y) = H'(g^{ab}, X, Y)$.

Theorem 2. *The steganographic scheme S_1 is CHA-secure.*

Proof. By definition of the family F and the hardness of the Diffie-Hellman problem over $<g>$, we obtain that g^{ab}, and therefore r, is computationally indistinguishable from uniformly random. Thus, by definition of our $\mathcal{P}$-code, c is computationally indistinguishable from $\mathcal{P}$.

Further, since $H_{\mathcal{P}}(h_0) \geq k$, with overwhelming probability h_0 is different each time we embed. Therefore even when the embedding oracle is queried repeatedly, r still appears to the attacker as independently and uniformly random. Therefore in the attacker's view the ciphertexts obtained by him in the warmup step are independent of of the challenged ciphertext, i.e. they are useless for the attack. That means our scheme is CHA-secure.

Proposition 1. *The steganographic scheme S_1 is CSA-secure.*

Proof. Since signature scheme D is secure, i.e. unforgeable, an active adversary cannot construct a new valid covertext sequence. Therefore with overwhelming probability, all queries made to the extraction oracle will return $\perp$ at step 4 of the extraction algorithm. Therefore an active adversary obtain no more advantage than a passive one does. Since S_1 is already CHA-secure, we obtain that our scheme is CSA-secure.

Expansion Rate. The expansion rate of this scheme equals to the rate of the underlying $\mathcal{P}$-code plus the overhead in sending h_0 and a signature. Nevertheless, the overhead of h_0 and the signature, which is $O(\lceil \frac{k}{\log_2(\frac{1}{\rho})} \rceil)$, only depends on the security parameter k. Thus it diminishes when we choose n large enough so that $k = o(n)$, say $n = k \log(k)$. Therefore the expansion rate of our steganographic system is essentially that of the $\mathcal{P}$-code.

4.2 Private Key Steganographic Systems

Let G be a cryptographically secure pseudo-random generator, and k be a shared secret key. In the setup step, k is given as seed to G. The state of G is kept between calls to G. This state is usually not much more than the space for a counter, which is quite small.

S_2-***Embed***. **Input**: $m \in \Sigma^n$.
Output: $c \in C^*$.

1. **let** $r \| z = G(k)$.
2. **let** $m' = m \| \texttt{Sign}(k_{sig}, m)$.
3. Output $c = \Gamma_e(z, r \oplus m)$.

S_2-***Extract***. **Input**: $c \in C^*$.
Output: $m \in \Sigma^n$.

1. **let** $r \| z = G(k)$.
2. **let** $m' = \Gamma_e(z, r \oplus m)$.
3. **if** $\texttt{Verify}(k_{ver}, m') \neq \mathit{success}$ **then** return $\perp$.
4. Parse $m' = m \| \texttt{Sign}(k_{sig}, m)$.

Theorem 3. *The steganographic scheme S_2 is CHA-secure.*

Proof. The proof is straight-forward: z and $r \oplus m$ is computationally indistinguishable from uniformly random, so by the property of Γ_e, the output covertext sequence $c = \Gamma_e(z, r \oplus m)$ is computationally indistinguishable from $\mathcal{P}$. Further, each time the embedding operation is performed, the pseudo-random generator G changes its internal state, so its output z, r are independent of each others in the attacker's view. Consequently, the values of $z, r \oplus m$, and so do the values of $c = \Gamma_e(z, r \oplus m)$, are probabilistically independent of each others to the attacker. This means that the ciphertexts obtained by the attacker in the warmup step do not help him in the guessing step in anyway. Therefore our scheme is secure against chosen hiddentexts attack.

Proposition 2. *The steganographic scheme S_2 is CSA-secure.*

Expansion Rate. It is clear that the expansion rate of this scheme is the same as the expansion rate of the $\mathcal{P}$-code. Additionally, both sides must maintain the status of the generator G. However, this status is very small, similar to a synchronized counter used in [8]. Note that our scheme S_2 is a somewhat more efficient than S_1 because it does not have to send the preamble h_0. In the next section, we will see that they are both assymptotically optimal.

5 Essentially Optimal Rates

In this section we consider applications of our schemes in two cases: distribution $\mathcal{P}$ is given explicitly by a cumulative distribution function F, and is given implicitly by a black-box sampler S. In both cases, we show that the achieved information rate is essentially optimal.

5.1 Cumulative Distribution Function

We show here that in case we have additionally a cumulative distribution function F of the given distribution, then the construction can be much more efficient. First, let us define what a cumulative distribution function is, and then how to use this additional information to construct $\mathcal{P}$-coding schemes.

Let the message space C be ordered in some strict total order $'<'$ so that $v_0 < v_1 < \dots$ is a sorted sequence of all covertexts. A *cumulative distribution function* (CDF) for the channel $\mathcal{P}$ is a family of functions $F_h : C \to [0,1]$ such that $F_h(v) = \sum_{v'<v} \mathcal{P}_h(v')$ for all $h \in C^*$ and $v \in C$. We modify our $\mathcal{P}$-code slightly so that it can use the additional information available effectively.

Γ_2-`Encode`. **Input**: $x = (x_1, \dots, x_n) \in \Sigma^n$.
Output: $c = (c_1, \dots, c_l) \in C^*$.

1. **let** $a = 0, b = 2^{2n}, h = \epsilon$.
2. **while** $\lceil a/2^n \rceil < \lfloor b/2^n \rfloor$ **do**
   (a) **let** $i_j = tF_{h_0\|h}(v_j)$.
   (b) **let** j be the unique integer such that
   $$i_j \le \lfloor (2^n\overline{x} - a)t/(b-a) \rfloor < i_{j+1}.$$
   (c) **let** $a' = a + (b-a)i_j/t$, $b' = a + (b-a)i_{j+1}/t$.
   (d) **let** $(a, b) = (a', b')$.
   (e) **let** $h = h|v_{i_j}$.
3. Output $c = h$.

The only difference here is that instead of calling S repeatedly to generate c_i^* $(0 \le i \le t-1)$ and then deduce i_j $(0 \le j \le m-1)$, we use $v_0, v_1, \dots \in C$ directly and let $i_j = tF_{h_0\|h}(c_j)$ for $j = 0, 1, \dots$. Note that the sorted sequence $v_0, v_1, \dots$ of all covertexts can either be given explicitly, or be given by a function $v : \mathbb{N} \to C$. In the either case, the determination of j in step $2(b)$ can be done by binary searching, thus allows large covertext space C to be used.

Γ_2-`Decode`. **Input**: $c = (c_1, \dots, c_l) \in C^*$.
Output: $x = (x_1, \dots, x_n) \in \Sigma^n$.

1. **let** $a = 0, b = 2^{2n}, h = \epsilon$.
2. **for** i **from** 1 **to** $|c|$ **do**
   (a) **let** $i_j = tF_{h_0\|h}(v_j)$.
   (b) **let** j be the unique integer such that $v_{i_j} = c_i$.
   (c) **let** $a' = a + (b-a)i_j/t$, $b' = a + (b-a)i_{j+1}/t$.
   (d) **let** $(a, b) = (a', b')$.
   (e) **let** $h = h|v_{i_j}$.
3. **let** $v = \lceil a/2^n \rceil$.
4. Output $x = \underline{v}$.

Theorem 4. *The coding scheme described above is a $\mathcal{P}$-code.*

Proof. The proof is the same, word by word, as in proof of Theorem 1, with only necessary changes of c_i^* and i_j as noted above.

Theorem 5. *The expansion rate $e(n)$ is bounded from above by $1 + \frac{1}{n}\log_2(|C|)$.*

Proof. At each iteration i, the range $[a, b]$ is reduced in size by a factor of $(b'-a')/(b-a) = (i_{j+1}-i_j)/t = F_h(v_{j+1}) - F_h(v_j) = \mathcal{P}_h(v_{i_j}) = \mathcal{P}_h(c_i)$. Further, before the last iteration $b - a \geq 2^n$, so we get:

$$\mathcal{P}(c_1\| \dots \|c_{l-1}) = \prod_{i=1}^{l-1} \mathcal{P}_{c_1\|\dots\|c_{i-1}}(c_i) \geq \frac{2^n}{2^{2n}} = 2^{-n}.$$

This means $H_{\mathcal{P}}(c_1\| \dots \|c_{l-1}) \leq n$. Summing over all $x \in \Sigma^n$ we get:

$$\sum_{c \in \Gamma_e(\Sigma^n)} \mathcal{P}(c) H_P(c) \leq n + \log_2(|C|).$$

This shows that the expansion rate $e(n)$ is bounded above by:

$$e(n) = \frac{1}{n} \sum_{c \in \Gamma_e(\Sigma^n)} \mathcal{P}(c) H_P(c) \leq 1 + \frac{\log_2(|C|)}{n}.$$

Since $\log_2(|C|)$ is a constant, we obtain that $e(n) \to 1$ when $n \to \infty$.

5.2 General Case

In this case, we know nothing about the distribution $\mathcal{P}_h$, except a given black box sampler S. We give a proof showing that our scheme is optimal.

Theorem 6. *The $\mathcal{P}$-code defined in Section 3 is essentially optimal.*

Proof. See appendix.

6 Conclusions

We have shown in this article:

- Introduction and construction of $\mathcal{P}$-codes, and their applications.
- Efficient general construction of public key steganographic schemes secure against chosen hiddentext attacks using public key exchange assuming no special conditions.
- Efficient general construction of private key steganographic schemes secure against chosen hiddentext attacks assuming the existence of a pseudo-random generator.

Our constructions are essentially optimal in many cases, and they are general constructions, producing no errors in extraction. Nevertheless, our solutions do not come for free, i.e. they require polynomially sampleable cover distributions. Readers are refered to [8, 9] for more discussions on this issue.

References

1. Ross J. Anderson and Fabien A.P. Petitcolas. On the limits of steganography. *IEEE Journal of Selected Areas in Communications*, 16(4):474–481, May 1998.
2. Michael Backes and Christian Cachin. Public key steganography with active attacks. Technical report, IACR ePrint Archive 2003/231, 2003.
3. C. Cachin. An information-theoretic model for steganography. In *Information Hiding, Second International Workshop, Proceedings (Lecture Notes in Computer Science 1525)*, pages 306–318. Springer-Verlag, 1998. Portland, Oregon, April 15–17.
4. Scott Craver. On public-key steganography in the presence of an active warden. In David Aucsmith, editor, *Information Hiding, Second International Workshop, Portland, Oregon, USA*, volume 1525 of *Lecture Notes in Computer Science*. Springer, April 14-17 1998.
5. Csiszar. The method of types. *IEEETIT: IEEE Transactions on Information Theory*, 44, 1998.
6. I. Csiszar and J. Korner. *Information theory: Coding Theory for Discrete Memoryless Systems*. Academic Press, NY, 1981.
7. G. Gurther. A universal algorithm for homophonic coding. In *Eurocrypt '88*. Springer-Verlag, 1988.
8. Nick Hopper, John Langford, and Luis von Ahn. Provably secure steganography. In Moti Young, editor, *Advances in Cryptoglogy — Crypto 2002, Proceedings*, volume 2442 of *LNCS*. Springer-Verlag, August 2002.
9. Nick Hopper and Luis von Ahn. Public key steganography. Submitted to Crypto 2003.
10. S. Katzenbeisser and F. Petitcolas. On defining security in steganographic systems, 2002.
11. Mittelholzer. An information-theoretic approach to steganography and watermarking. In A. Pfitzmann, editor, *Proceedings of Third International Workshop on Information Hiding*, volume 1768 of *LNCS*. Springer-Verlag, September 1998.
12. P. Moulin and J. O'Sullivan. Information-theoretic analysis of information hiding, 1999.
13. Leonid Reyzin and Scott Russell. More efficient provably secure steganography. Technical report, IACR ePrint Archive 2003/093, 2003.
14. G. J. Simmons. The prisoner's problem and the subliminal channel. In David Chaum, editor, *Advances in Cryptology: Proceedings of Crypto '83*, pages 51–70, New York, USA, 1984. Plenum Publishing.
15. Luis von Ahn and Nick Hopper. Public key steganography. Submitted to Eurocrypt 2004.
16. Jan Zollner, 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 Proof of Theorem 6

Proof. First, note that any steganographic scheme defined over channel $\mathcal{P}$ is indeed a $\mathcal{P}$-coding scheme. Second, the expansion rate of our steganographic schemes is essentially the expansion rate of the underlying $\mathcal{P}$-code. Hence it is enough to show that our $\mathcal{P}$-code Γ_1 is optimal.

Indeed, let $\mathit{\Gamma}'$ be any $\mathcal{P}$-coding scheme that works generically like $\mathit{\Gamma}$, i.e. $\mathit{\Gamma}'$ works on any black box S whose output has minimal entropy bounded from below (e.g. by ξ). Let t be the number of oracle calls to S by $\mathit{\Gamma}'$, and let $c^* = (c_0^*, \ldots, c_{t-1}^*)$ be the corresponding results.

Then $\mathit{\Gamma}'$ can only return one of the covers $c_0^*, \ldots, c_{t-1}^*$ as its next stegotext to be sent to the receiver. Indeed, assume otherwise that this is not the case. Then consider a black box sampler S' that output covertexts including a long random string signed with a secure digital signature. Now apply $\mathit{\Gamma}'$ to S'. If $\mathit{\Gamma}'$ outputs anything that is not in the list of covers returned by S', the output of $\mathit{\Gamma}'$ will not contain a valid signature, or otherwise by definition the digital signature would have been insecure. Now such unsigned covers is immediately detectable by a polynomial time algorithm, i.e. by checking for the signature using the corresponding public key. Therefore the output of $\mathit{\Gamma}'$ is distinguishable from the output of S'. This contradicts with our assumption that $\mathit{\Gamma}'$ is a $\mathcal{P}$-code. Since $\mathit{\Gamma}'$ cannot tell the output of S contains some sort of a digital signature or not, we conclude that $\mathit{\Gamma}'$ must always output one of the c_i^*'s as its output.

We consider two cases. First if the entropy of $\mathcal{P}_h$ is at least $(1+\epsilon)\log_2(t)$ for some fixed constant $\epsilon > 0$, then from the method of types (cf. [3, 5, 6]), we know that the c_i^*'s are distinct with overwhelming probability. Therefore $\mathit{\Gamma}_1$ achieves rate of $\log_2(t)$ bits per symbol. However, we know from previous paragraph that $\mathit{\Gamma}'$ has its rate bounded by $\log_2(t)$. Hence in this case, $\mathit{\Gamma}'$ does not do better than $\mathit{\Gamma}_1$.

In the second case, the entropy of $\mathcal{P}_h$ is at most $\log_2(t)$. In this case, method of types(cf. [3, 5, 6]) tell us that for any fixed constant $\delta > 0$ and large enought t, with overwhelming probability: the induced entropy of the view $(c_0^*, \ldots, c_{t-1}^*)$ is at least $(1-\delta)H(\mathcal{P}_h)$. Thus in this case our encoding $\mathit{\Gamma}_1$ achieves at least $(1-\delta)H(\mathcal{P}_h)$ bits per symbol. Note that the rate of the encoding $\mathit{\Gamma}'$ must be bounded from above by $(1+\delta)H(\mathcal{P}_h)$, otherwise the output of $\mathit{\Gamma}'$ will be distinguishable from $\mathcal{P}_h$ with overwhelming probability by simply estimating the entropies of the two distributions [5, 6].

We conclude that all cases, for all $\delta > 0$ our encoding $\mathit{\Gamma}_1$'s rate is within $(1-\delta)$ fraction of the best possible rate minus some negligible factor, i.e. $\mathit{\Gamma}_1$ is essentially optimal.