Fully Bideniable Public-Key Encryption

Marcel Šebek

October 23, 2013

Department of Algebra, Faculty of Mathematics and Physics,
Charles University in Prague, Czech Republic
marcel.sebek@matfyz.cz

Abstract

Bideniable encryption allows both sender and receiver in a public-key setting to simultaneously claim that a different message of their choice was transmitted, and support this claim by a good-looking encryption and key-generation randomness, respectively. A weaker version with two variants of algorithms is called flexible or multi-distributional deniability, a stronger one-algorithm version is called full deniability. Bendlin et al. (ASIACRYPT 2011) showed that certain kinds of fully-deniable schemes are constructible using corresponding flexible schemes. In this paper, we show that their construction in the bideniable case has a flaw, and we provide a fixed scheme. Our construction relies on a deterministic subset matching algorithm that assigns to each nonempty subset of a finite set its proper subset reduced by exactly one element. Although this algorithm is crucial in our construction, it may also be of independent interest.

Keywords: deniable encryption, plausible deniability, fully bideniable PKE, subset matching.

1 Introduction

Security guarantees of encryption schemes largely depend on the adversarial model. Although some basic requirements like semantic security are common for most models, there is a plenty of extensions, some of which are contradictory. An example is the commitment property – does the sender commit to her inputs, or are there alternative inputs that produce the same ciphertext? In certain applications, commitment property is desirable to obtain non-repudiation, while sometimes we need a non-committing scheme, such as to construct adaptively secure multi-party computation [CFGN96]. Non-committing encryption does not provide a procedure to obtain alternative inputs in general. Instead, a special algorithm is given that simultaneously samples a set of inputs for the regular encryption and key generation algorithm, each of which leads to the same ciphertext and public key, respectively.

Deniable encryption offers exactly the missing part. Given a ciphertext, public-key, all secret knowledge, and an alternative message, the sender and/or receiver is able to compute alternative secret knowledge (i.e., encryption algorithm randomness or secret key). The alternative secrets are required to be indistinguishable from honest secrets while delivering the alternative message.

The main motivation of deniable encryption is coercion resistance. A powerful adversary may demand secret key and encryption randomness for the intercepted communication. In various countries, this power is given to the government, but we can find a few other examples. When users run deniable scheme, they can provide a different message and alternative secrets that correspond exactly to this message. The adversary cannot distinguish honest and computed secrets.

We briefly overview basic kinds of deniable encryption. If the alternative message can be chosen after ciphertext generation, the scheme is called ad-hoc. Otherwise, if the possible alternative messages are inputs of encryption algorithm, we call the scheme plan-ahead. In this paper, we consider ad-hoc schemes only. An orthogonal property is the set of coerced parties – we distinguish sender-deniability, receiver-deniability and bideniability.

Just for ah-doc schemes, we consider full and flexible (also known as multi-distributional) deniability. Fully-deniable schemes have only one key generation and encryption algorithm. Flexible schemes have two algorithms for the coerced parties (e.g., fully sender-deniable scheme has two encryption algorithms). One algorithm is honest, and its run cannot be later faked to an alternative message. The other algorithm is dishonest, and its run can later be faked as a run of the *honest* algorithm with any alternative message.

The last property is the indistinguishability level. We consider two cases, negligible and inverse-polynomial distinguishability. In this work, we construct an inverse-polynomial scheme.

1.1 Related Work

A concept similar to deniable encryption called *plausible deniability* was studied by Beaver [Bea96]. However, all subsequent texts are based on the work by Canetti et al. [CDNO97]. They introduced fully and flexibly deniable encryption in both public and shared key setting, and sketched the idea of plan-ahead schemes. They focused on sender deniability and constructed fully $1/\kappa$-sender-deniable and flexibly sender-deniable public-key scheme, as well as a bideniable shared-key scheme. Finally, they proposed interaction as a way how to transform sender-deniability to receiver-deniability and vice versa.

Since the initial proposal of the concept, one direction of research has focused on feasibility results. O'Neill et al. [OPW11] proposed two constructions of flexibly bideniable scheme, one built from simulatable encryption [DN00] and the other from a lattice-based cryptosystem [GPV08]. Additionally, a straightforward modification of the first construction yields a flexible receiver-deniable scheme. Bendlin et al. [BNNO11] showed that fully receiver- and bideniable scheme may be obtained with inverse-polynomial distinguishability only. They also provided constructions of fully deniable schemes from flexible schemes. For sender deniability, this construction repeats the one by Canetti et al., while for receiver and bideniability, the reduction is new. Unfortunately, for bideniability, the reduction is incorrect, as we show in this paper. An unsuccessful attempt to construct interactive fully sender-deniable scheme was due to Dürmuth and Freeman [DF11], but a mistake have been found by Peikert and Waters, rendering the scheme to be inverse-polynomially indistinguishable only, like the previously known scheme. The details of the flaw can be found in the updated version of the original paper [DF11]. Dachman-Soled [DS12] showed that no black-box construction of sender-deniable scheme from simulatable encryption is possible on the negligible level. However, Sahai and Waters [SW13] leveraged indistinguishability obfuscation to obtain a non-interactive fully sender-deniable scheme.

To complete our picture, we mention some works that target efficiency and practical aspects of deniability. Klonowski et al. [KKK08] considered nested constructions, embedding of secret message into ElGamal encryption randomness, and various other things. Deniable encryption also motivated the study of deniable and steganographic filesystems [ANS98, MK99, HKX03, XHK04].

1.2 Our Contributions and Open Problems

We provide a fixed construction of fully $1/\kappa$-bideniable scheme from a flexible scheme. Since the main open problem of existence of sender-deniable scheme has recently been positively resolved [SW13], the study of feasibility of non-interactive full deniability is finished. For flexible deniability, sender [CDNO97] and bideniable [OPW11] schemes are known, while for receiver deniability, we need to slightly modify simulatable encryption based scheme [OPW11, Šeb12].

In the interactive case, sender and receiver deniability is equivalent. So, the only question that remains open is the existence of interactive fully bideniable scheme.

1.3 Organization

Section 2 fixes the notation and recalls basic facts and definitions from the theory of computational complexity, probability, and from other related fields. Section 3 introduces main types of deniable encryption, defines translucent sets, and provides a generic construction of fully bideniable scheme based on a flexible scheme or, more generally, on bitranslucent sets. Section 4 contains a construction of bideniable scheme based on a subset matching algorithm, which is the topic of Sect. 5. Finally, Appendix A shows that the original scheme [BNNO11] admits a constant advantage distinguisher.

2 Preliminaries and Notation

In the following, $\kappa \in \mathbb{N}$ will be the main security parameter. By an algorithm we mean a probabilistic Turing machine. An algorithm is feasible if its running time is polynomial in the input length. If not specified otherwise, all algorithms are required to be feasible. We use the notation $o_1, \ldots, o_k \leftarrow \mathsf{A}(i_1, \ldots, i_j; r)$ to describe interface of an algorithm A, the variables $i_1, \ldots, i_j$ are inputs, $o_1, \ldots, o_k$ are outputs and r is the randomness which may be omitted if we need not refer to it explicitly. The same notation is also used for algorithm invocation.

A function $\nu \colon \mathbb{N} \to \mathbb{R}$ is *negligible* if it vanishes faster than inverse of any polynomial, i.e., for any $n \in \mathbb{N}$ there is $k_0 \in \mathbb{N}$ so that for $k \geq k_0$ we have $\nu(k) < k^{-n}$. A function f is *overwhelming* if $1 - f$ is negligible.

2.1 Probability and Statistics

The terms random variable and probability distribution are closely related and we will often interchange them. Let A be a discrete random variable. Instead of $\Pr_{a \leftarrow A}[a \in S]$ we will often write $A(S)$, and when $S = \{s\}$, then we will use $A(s)$.

For discrete distributions A and B with domain D and for $X \subseteq D$, we define (somewhat non-standard) *X-statistical distance of A and B* as

$$\Delta_X(A,B) = \frac{1}{2}\sum_{x\in X}|A(x)-B(x)|\,.$$

We write $\Delta(A,B)$ instead of $\Delta_D(A,B)$ and call the number just *statistical distance of A and B*. For continuous distributions, analogy of the previous definition exists – it uses integral and density function instead of sum and probability mass function.

Let $\mathcal{A} = \{A_\kappa\}_{\kappa\in\mathbb{N}}$ and $\mathcal{B} = \{B_\kappa\}_{\kappa\in\mathbb{N}}$ be two ensembles of probability distributions. Let $\epsilon\colon \mathbb{N}\to\mathbb{R}$ be a function. We say that $\mathcal{A}$ and $\mathcal{B}$ are *ϵ-statistically indistinguishable* if $\Delta(A_\kappa, B_\kappa) \leq \epsilon(\kappa)$ for all sufficiently large values of κ. If ϵ is negligible, we say that $\mathcal{A}$ and $\mathcal{B}$ are just *statistically indistinguishable.*

For *discrete* distributions A and B, and for an algorithm D called distinguisher that outputs values from $\{0,1\}$, let the advantage of D with respect to A and B be

$$\mathrm{Adv}_D(A,B) = \left|\Pr_{a\leftarrow A}[D(a)=1] - \Pr_{b\leftarrow B}[D(b)=1]\right|\,.$$

Probabilities are computed over the randomness of D, and over the choice of a or b, respectively. The advantage is bounded from above by the statistical distance of A and B.

Now let us again consider the ensembles $\mathcal{A}$ and $\mathcal{B}$. We say that $\mathcal{A}$ and $\mathcal{B}$ are *ϵ-computationally indistinguishable* and denote it by $\mathcal{A} \overset{\mathrm{c}}{\approx}_\epsilon \mathcal{B}$ if for any feasible distinguisher D we have $\mathrm{Adv}_D(A_\kappa, B_\kappa) \leq \epsilon(\kappa)$ for sufficiently large values of κ. If ϵ is negligible, we say that $\mathcal{A}$ and $\mathcal{B}$ are just *computationally indistinguishable* and write $\mathcal{A} \overset{\mathrm{c}}{\approx} \mathcal{B}$.

In the rest of the paper, the ensembles are parameterized by κ implicitly. For simplicity, we will just speak about computational indistinguishability of probability distributions instead of their ensembles, and so on.

2.2 Subsets

For $n\in\mathbb{N}$ let $[n] = \{1,\dots,n\}$ and let the set of its subset be denoted by $\mathcal{P}([n])$. A subset $\mathbf{x}\subseteq[n]$ is identified with an n-bit string $x_1x_2\dots x_n \in \{0,1\}^n$ such that $x_k = 1$ iff $k\in\mathbf{x}$, and we call this string *binary representation of* $\mathbf{x}$. For $k=1,\dots,n$, let $\mathbf{e}_k = \{k\}$ be the k-th singleton. For $i=0,\dots,n$ let $L_i = \{\mathbf{x}\in\mathcal{P}([n]) \mid |\mathbf{x}| = i\}$ be the i-th layer of $\mathcal{P}([n])$. Clearly, the binary representation of an element in the i-th layer contains i ones and $n-i$ zeros, and cardinality of the i-th layer is $\binom{n}{i}$.

For convenience, we introduce an artificial layer $L_{-1} = \{\bot\}$, and we put $\mathcal{P}^*([n]) = \mathcal{P}([n])\cup L_{-1}$. To make things consistent, we set $|\bot| = -1$ and $\binom{n}{-1} = 1$.

3 Deniable Encryption

For simplicity, we consider the message space $M = \{0,1\}$. A *public-key encryption (PKE) scheme* consists of algorithms G, E, D such that

1. $pk \leftarrow \mathsf{G}(1^\kappa; sk)$ is the key-generation algorithm (randomness is identified with the secret key),

2. $c \leftarrow \mathsf{E}(pk, m; r_{\mathsf{E}})$ is the encryption algorithm,

3. $m' \leftarrow \mathsf{D}(sk, c)$ is the decryption algorithm,

4. correctness condition holds: for any $m \in M$, we have $m = m'$ with overwhelming probability (over algorithms randomness).

A PKE scheme is $\mathsf{IND\text{-}CPA}$-*secure* if $(pk, c_0) \overset{c}{\approx} (pk, c_1)$ for $c_m \leftarrow \mathsf{E}(pk, m)$.

Let $\epsilon = \epsilon(\kappa)$ be a parameter. A *flexibly ϵ-bideniable scheme* consists of algorithms G_{H}, G_{D}, E_{H}, E_{D}, D, F_{S}, F_{R} such that

1. (G, E, D) is an $\mathsf{IND\text{-}CPA}$-secure PKE scheme for $G \in \{\mathsf{G}_{\mathrm{H}}, \mathsf{G}_{\mathrm{D}}\}$ and $E \in \{\mathsf{E}_{\mathrm{H}}, \mathsf{E}_{\mathrm{D}}\}$,

2. $\tilde{r}_{\mathsf{E}} \leftarrow \mathsf{F}_{\mathrm{S}}(pk, r_{\mathsf{E}}, m, m_{\mathrm{f}})$ is the sender faking algorithm that, given inputs of dishonest encryption $c \leftarrow \mathsf{E}_{\mathrm{D}}(pk, m; r_{\mathsf{E}})$, produces an alternative randomness $\tilde{r}_{\mathsf{E}}$ such that c "looks like" it was produced honestly using $\mathsf{E}_{\mathrm{H}}(pk, m_{\mathrm{f}}; \tilde{r}_{\mathsf{E}})$,

3. $\tilde{sk} \leftarrow \mathsf{F}_{\mathrm{R}}(sk, c, m_{\mathrm{f}})$ is the receiver faking algorithm that, given $pk \leftarrow \mathsf{G}_{\mathrm{D}}(1^\kappa; sk)$ and $c \leftarrow \mathsf{E}_{\mathrm{D}}(pk, m)$, produces an alternative secret key $\tilde{sk}$ "consistent" with pk such that c "should decrypt" to m_{f},

4. quoted phrases above mean: for any $m, m_{\mathrm{f}} \in M$, we have $(pk, c, sk, r_{\mathsf{E}}) \overset{c}{\approx}_\epsilon (pk_{\mathrm{D}}, c_{\mathrm{D}}, \tilde{sk}, \tilde{r}_{\mathsf{E}})$ where

$$\begin{array}{ll} pk \leftarrow \mathsf{G}_{\mathrm{H}}(1^\kappa; sk) & pk_{\mathrm{D}} \leftarrow \mathsf{G}_{\mathrm{D}}(1^\kappa; sk_{\mathrm{D}}) \\ c \leftarrow \mathsf{E}_{\mathrm{H}}(pk, m_{\mathrm{f}}; r_{\mathsf{E}}) & c_{\mathrm{D}} \leftarrow \mathsf{E}_{\mathrm{D}}(pk_{\mathrm{D}}, m; r_{\mathsf{E}}) \\ & \tilde{r}_{\mathsf{E}} \leftarrow \mathsf{F}_{\mathrm{S}}(pk_{\mathrm{D}}, r_{\mathsf{E}}, m, m_{\mathrm{f}}) \\ & \tilde{sk} \leftarrow \mathsf{F}_{\mathrm{R}}(sk_{\mathrm{D}}, c_{\mathrm{D}}, m_{\mathrm{f}}). \end{array}$$

If ϵ is negligible, we call the scheme just *flexibly bideniable.* In addition, if the last requirement is stated just for $m = 1$ and $m_{\mathrm{f}} = 0$, we call the scheme *bitranslucent set scheme (BTS).* If ϵ, m, and m_{f} are arbitrary, $\mathsf{G}_{\mathrm{H}} = \mathsf{G}_{\mathrm{D}}$, and $\mathsf{E}_{\mathrm{H}} = \mathsf{E}_{\mathrm{D}}$, then $(\mathsf{G}_{\mathrm{D}}, \mathsf{E}_{\mathrm{D}}, \mathsf{D}, \mathsf{F}_{\mathrm{S}}, \mathsf{F}_{\mathrm{R}})$ is called *fully ϵ-bideniable* or just ϵ-bideniable.

Remark 1. If we distinguish between key-generation algorithm randomness and secret key, we can give somewhat weaker definition of deniable encryption, such that the receiver faking algorithm produces just the fake secret key. As we are interested in feasibility results, we decided to primarily state the stronger and simpler variant of the definition.

Remark 2. We modified the original definition of BTS [OPW11] slightly by dropping the support of identity-based setting. It is easy to see that a BTS fulfilling the original definition is a BTS by our definition.

We will be interested in black-box constructions of a fully $1/\kappa$-bideniable scheme from a BTS (or from a flexible scheme since it is also a BTS) that follow a general structure described below.

Construction 1. Let n be a suitable polynomial in κ.

1. Key generation algorithm samples $\mathbf{b} \in \mathcal{P}([n])$ from a given distribution $\mathcal{B}_\kappa$. It then invokes the underlying key-generation algorithms in parallel, one for each bit in the binary representation of $\mathbf{b}$, G_{H} for $b_i = 0$ and G_{D} for $b_i = 1$.

2. There exists a deterministic algorithm computing $V\colon \mathcal{P}([n]) \to M$.

3. Encryption algorithm samples $\mathbf{a} \in \mathcal{P}([n])$ from a given distribution $\mathcal{A}_\kappa$. For each bit of $\mathbf{a}$, the underlying encryption algorithm is called, $\mathsf{E}_{\mathrm{H}}(pk_i, 0)$ for $a_i = 0$ and $\mathsf{E}_{\mathrm{D}}(pk_i, 1)$ for $a_i = 1$. The resulting ciphertext is a concatenation of the underlying ciphertexts, followed by a bit $b = m \oplus V(\mathbf{a})$.

4. Decryption algorithm calls underlying decryption algorithm for all ciphertexts parts to obtain $\mathbf{a}$, and calls $V(\mathbf{a}) \oplus b$ to get the result.

5. Faking algorithms return honest randomness when opening true message. To fake ciphertext to the opposite message, they leverage a function $F\colon \mathcal{P}([n]) \to \mathcal{P}([n])$ to compute a subset $\mathbf{f} = F(\mathbf{a})$. Sender faking algorithm invoke underlying faking algorithm for positions given by $\mathbf{f} \cap \mathbf{a}$, while receiver invoke faking for positions given by $\mathbf{f} \cap \mathbf{b}$.

We summarize construction correctness in the following proposition. We retain the same notation as above.

Proposition 1. *Let* $\mathbf{f} \cap \mathbf{a} \subseteq \mathbf{f} \cap \mathbf{b}$ *with probability at least* $1 - \epsilon_1$*, and* $V(\mathbf{a} \setminus \mathbf{f}) = 1 - V(\mathbf{a})$ *with probability at least* $1 - \epsilon_2$*. Let* $(\mathbf{a}, \mathbf{b}) \overset{c}{\approx}_{\epsilon_3} (\mathbf{a} \setminus \mathbf{f}, \mathbf{b} \setminus \mathbf{f})$ *where the distributions on the right-hand side are conditioned on* $\mathbf{f} \cap \mathbf{a} \subseteq \mathbf{f} \cap \mathbf{b}$ *and* $V(\mathbf{a} \setminus \mathbf{f}) = 1 - V(\mathbf{a})$*. Let* $\epsilon_1 + \epsilon_2 + \epsilon_3 \leq 1/\kappa - negl(\kappa)$*. Let* $\Delta(V(\mathcal{A}_\kappa), \mathsf{Ber}(1/2)) = negl(\kappa)$*. Then the Construction 1 provides a fully* $1/\kappa$*-bideniable scheme.*

Proof. Correctness and security follow easily from the corresponding properties of the underlying scheme, and from the fact that the padding bit $V(\mathcal{A}_\kappa)$ is close to uniform.

Let us prove $1/\kappa$-deniability, that is, $(pk, c, sk, r_{\mathsf{E}}) \overset{c}{\approx}_{1/\kappa} (pk_{\mathrm{D}}, c_{\mathrm{D}}, \tilde{sk}, \tilde{r}_{\mathsf{E}})$ for any m and m_{f}. By definition, the distributions are identical for $m = m_{\mathrm{f}}$, so let us assume that $m \neq m_{\mathrm{f}}$. We construct a simulator that, given the sets $\mathbf{a}$ and $\mathbf{b}$, produces $\mathcal{T} = (pk, c, sk, r_{\mathsf{E}})$ in the same way as these values would be obtained in Construction 1 when invoked for message m_{f} and when the values $\mathbf{a}$ and $\mathbf{b}$ are provided on input instead of their sampling. When $\mathbf{a} \leftarrow \mathcal{A}_\kappa$ and $\mathbf{b} \leftarrow \mathcal{B}_\kappa$, the simulator produces $\mathcal{T}$ of the honest game. On the other hand, when the simulator is supplied with $(\mathbf{a} \setminus \mathbf{f}, \mathbf{b} \setminus \mathbf{f})$, it produces a distribution close to $(pk_{\mathrm{D}}, c_{\mathrm{D}}, \tilde{sk}, \tilde{r}_{\mathsf{E}})$ with the only difference that for the coordinates $\mathbf{f}$ honest algorithms are used. We proceed by a sequence of hybrids by faking first i coordinates from $\mathbf{f}$, for $i = 0, \ldots, |\mathbf{f}|$. By assumptions on the underlying scheme, adjacent hybrids are computationally indistinguishable. The last hybrid is equivalent to dishonest game. □

4 Our Construction

We leverage generic Construction 1. One of the building blocks is a special matching algorithm provided by the following theorem. The proof is postponed to Sect. 5.

Theorem 1. *There exists a feasible deterministic algorithm* M *that assigns to each subset* $\varnothing \neq \mathbf{x} \subseteq [n]$ *a singleton* $\mathbf{e}_k \subseteq \mathbf{x}$ *with the following properties. Let* $\tilde{\mathsf{M}}(\mathbf{x}) = \mathbf{x} \setminus \mathsf{M}(\mathbf{x})$*, and* $\tilde{\mathsf{M}}(\varnothing) = \perp$*. For* $\mathbf{y} \in \mathcal{P}^*([n])$ *let* $\mathrm{p}(\mathbf{y})$ *be the number of preimages of* $\mathbf{y}$ *in the mapping* $\tilde{\mathsf{M}}$*. Then* $\mathrm{p}(\mathbf{y}_1) - \mathrm{p}(\mathbf{y}_2) \leq 1$ *for any* $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{P}^*([n])$*,* $|\mathbf{y}_1| = |\mathbf{y}_2|$*.*

For $0 \leq i \leq n$, let us define $W(i) = (n-i)/(i+1)$, and let $W(-1) = 1$. Moreover, let $\underline{W}(i) = \lfloor W(i) \rfloor$ and $\overline{W}(i) = \lceil W(i) \rceil$. Given that Theorem 1 holds, we can show by a simple counting argument that $\mathrm{p}(\mathbf{y}) \in \{\underline{W}(i), \overline{W}(i)\}$ for $-1 \leq i \leq n$ and $\mathbf{y} \in L_i$.

Next, we define a few distributions over $\mathcal{P}^*([n])$ for $0 \leq \rho \leq 1$.

- Let $\mathsf{H}_n(\rho)$ (honest) be a distribution which samples n bits independently from $\mathsf{Ber}(\rho)$, i.e., the bit is 1 with probability ρ. Probability of $\perp$ is 0.
- Let $\mathsf{PD}_n(\rho)$ (probabilistic dishonest) be a distribution which samples $\mathbf{x} \leftarrow \mathsf{H}_n(\rho)$ and subtracts uniformly chosen element from the result, or outputs $\perp$ for $\mathbf{x} = \varnothing$.
- Let $\mathsf{DD}_n(\rho)$ (deterministic dishonest) be the distribution $\tilde{\mathsf{M}}(\mathsf{H}_n(\rho))$.

We proceed with the scheme description based on generic Construction 1.

Construction 2. Let $0 < \alpha < 1/3$ be a parameter, $\beta = \alpha/(\alpha+2)$, and $n = n(\kappa)$ be a suitable polynomial specified later.

1. The mapping V returns parity of the input.
2. The distribution $\mathcal{B}_\kappa$ is $\mathsf{H}_n(1 - 1/n^\beta)$.
3. The distribution $\mathcal{A}_\kappa$ is $\mathsf{H}_n(1/n^\alpha)$.
4. Faking mapping F coincides with the algorithm M from Theorem 1.

In order to prove desired security properties, we need a few lemmas.

Lemma 1. *Let X and Y be distributions on $A \times B$, let X_A and Y_A be corresponding marginals. Let $S \subseteq A$ be such that $X_A(S), Y_A(S) \geq 1 - \tau$ for some τ, and $X_A(a), Y_A(a) > 0$ for any $a \in S$. Let $X|a$ and $Y|a$ be conditional distributions for any $a \in S$. Then $\Delta(X,Y) \leq \Delta_S(X_A, Y_A) + \max_{a \in S} \Delta(X|a, Y|a) + \tau$.*

Proof. We split the distance into two parts to treat them separately:

$$\Delta(X,Y) = \Delta_{S \times B}(X,Y) + \Delta_{(A \setminus S) \times B}(X,Y).$$

By assumption, we have

$$\Delta_{(A \setminus S) \times B}(X,Y) \leq \frac{1}{2} \sum_{a \in A \setminus S} (X_A(a) + Y_A(a)) \leq \tau.$$

For $a \in S$ and $b \in B$, we have

$$\begin{aligned} |X(a,b) - Y(a,b)| &= |(X|a)(b) X_A(a) - (Y|a)(b) Y_A(a)| \\ &\leq (X|a)(b) \left| X_A(a) - Y_A(a) \right| + Y_A(a) \left| (X|a)(b) - (Y|a)(b) \right|. \end{aligned}$$

Therefore,

$$\begin{aligned} \Delta_{S \times B}(X,Y) &= \frac{1}{2} \sum_{a \in S} \sum_{b \in B} |X(a,b) - Y(a,b)| \\ &\leq \Delta_S(X_A, Y_A) + \max_{a \in S} \Delta(X|a, Y|a). \end{aligned}$$

□

Lemma 2. *Let $1 > \alpha > 0$. Let $\rho = 1/n^{\alpha}$, $X = V(\mathsf{H}_n(\rho))$, and $Y = \mathsf{Ber}(1/2)$. Then $\Delta(X, Y) = negl(\kappa)$.*

Proof. By definition, we have

$$\begin{aligned}
\Delta(X,Y) &= \frac{1}{2}\left(|X(0)-Y(0)| + |X(1)-Y(1)|\right) = \frac{1}{2}\left|X(0)-X(1)\right| \\
&= \frac{1}{2}\left|\sum_{i=0}^{\lfloor n/2 \rfloor}\binom{n}{2i}\rho^{2i}\left(1-\rho\right)^{n-2i} - \sum_{i=0}^{\lceil n/2 \rceil - 1}\binom{n}{2i+1}\rho^{2i+1}\left(1-\rho\right)^{n-2i-1}\right| \\
&= \frac{1}{2}\left|\sum_{i=0}^{n}\binom{n}{i}(-1)^i\rho^i\left(1-\rho\right)^{n-i}\right| = \frac{1}{2}(1-2\rho)^n \\
&= \frac{1}{2}\left(1-\frac{1}{n^{\alpha}/2}\right)^{n^{\alpha}/2\cdot 2n^{1-\alpha}} \le \frac{1}{2}e^{-2n^{1-\alpha}} = negl(\kappa).
\end{aligned}$$

□

Lemma 3. *Let $1/3 > \alpha > 0$. There exists $\epsilon > 0$, such that whenever $\mathbf{x} \in L_i$ for $n^{1-\alpha} - \delta < i < n^{1-\alpha} + \delta$ and $\delta = n^{(1-\alpha)/2+\epsilon}$, then the number of preimages $\mathrm{p}(\mathbf{x})$ lies between $n^{\alpha} - c$ and $n^{\alpha} + c$ for some constant c. For further reference, let $\mathcal{I}_{\alpha} = \{i \in \mathbb{N} \mid n^{1-\alpha} - \delta < i < n^{1-\alpha} + \delta\}$ and $\mathcal{L}_{\alpha} = \bigcup_{i \in \mathcal{I}_{\alpha}} L_i$.*

Proof. We can assume that $i \neq -1$, then $W(i) = (n-i)/(i+1) = (n+1)/(i+1) - 1$. Since we can ignore constants, it suffices to analyze $n/(i+1)$. Let $\epsilon = (1-3\alpha)/2$. Then by our assumptions, we obtain

$$\frac{n^{\alpha}}{1 + n^{-(1-\alpha)/2+\epsilon} + n^{\alpha-1}} < \frac{n}{i+1} < \frac{n^{\alpha}}{1 - n^{-(1-\alpha)/2+\epsilon} + n^{\alpha-1}}\ ,$$

which can be rewritten as

$$n^{\alpha} - \frac{n^{\epsilon-(1-3\alpha)/2} + n^{2\alpha-1}}{1 + n^{-(1-\alpha)/2+\epsilon} + n^{\alpha-1}} < \frac{n}{i+1} < n^{\alpha} + \frac{n^{\epsilon-(1-3\alpha)/2} - n^{2\alpha-1}}{1 - n^{-(1-\alpha)/2+\epsilon} + n^{\alpha-1}}\ .$$

Since $\alpha < 1$, both denominators tend to 1. Moreover, by $\alpha < 1/2$, we have $n^{2\alpha-1} \to 0$, and by the choice of ϵ, we have $n^{\epsilon-(1-3\alpha)/2} = O(1)$. □

Lemma 4. *Let $1/3 > \alpha > 0$. Let $\mathbf{x} \leftarrow \mathsf{H}_n(1/n^{\alpha})$ or $\mathbf{x} \leftarrow \mathsf{DD}_n(1/n^{\alpha})$. Then $\mathbf{x} \in \mathcal{L}_{\alpha}$ with overwhelming probability.*

Proof. For the honest distribution, we get the result by Chernoff bound. For dishonest distribution, we need to make sure that subtraction of 1 does not change anything for large values of n. That is true, because $n^{1-\alpha} \to \infty$, and the constant shifting can be hidden into n^{ϵ}. □

We need to get more insight into our distributions, so let us investigate their probability mass functions in more detail. We introduce notation for probabilities in which the parameters ρ and n are implicit (clear from the context). Let $i \in \mathbb{Z}$. Then we define

$$\begin{aligned}
H(i) &= \begin{cases} \rho^i(1-\rho)^{n-i} & 0 \le i \le n \\ 0 & \text{otherwise} \end{cases} \\
D(i) &= H(i+1)W(i) \\
\overline{D}(i) &= H(i+1)\overline{W}(i) \\
\underline{D}(i) &= H(i+1)\underline{W}(i).
\end{aligned}$$

We can see that all above defined values are 0 whenever i is outside $[-1, n]$. Let $-1 \leq i \leq n$ and $\mathbf{x} \in L_i$. Then, probability of $\mathbf{x}$ in $\mathsf{H}_n(\rho)$ is $H(i)$, probability of $\mathbf{x}$ in $\mathsf{PD}_n(\rho)$ is $D(i)$. If $\mathrm{p}(\mathbf{x}) = \underline{W}(i)$, then probability of $\mathbf{x}$ in $\mathsf{DD}_n(\rho)$ is $\underline{D}(i)$, otherwise it equals $\overline{D}(i)$.

Proposition 2. *Let $1/3 > \alpha > 0$. Then $\Delta_{\mathcal{L}_\alpha}(\mathsf{H}_n(n^{-\alpha}), \mathsf{DD}_n(n^{-\alpha})) = O(n^{-\alpha})$ where $\mathcal{L}_\alpha$ is defined in Lemma 3.*

Proof. For $-1 \leq i \leq n$, let $\overline{N}(i)$ and $\underline{N}(i)$ be the number of elements in L_i having $\overline{W}(i)$ or $\underline{W}(i)$ $\tilde{\mathsf{M}}$-preimages, respectively. Then

$$\overline{N}(i) + \underline{N}(i) = |L_i| = \binom{n}{i}. \tag{1}$$

Moreover, for $0 \leq i \leq n$ we have

$$H(i) = \left(\frac{1}{n^\alpha}\right)^i \left(1 - \frac{1}{n^\alpha}\right)^{n-i} = \left(\frac{1}{n^\alpha}\right)^n (n^\alpha - 1)^{n-i},$$

so we get $H(i) = H(i+1)(n^\alpha - 1)$ for $0 \leq i \leq n-1$. By definition and by the previous equality, our statistical distance equals

$$\begin{aligned}\Delta &= \sum_{i \in \mathcal{I}_\alpha} \underline{N}(i) \left|\underline{D}(i) - H(i)\right| + \overline{N}(i) \left|\overline{D}(i) - H(i)\right| \\ &= \sum_{i \in \mathcal{I}_\alpha} \underline{N}(i) H(i+1) \left|\underline{W}(i) - (n^\alpha - 1)\right| + \overline{N}(i) H(i+1) \left|\overline{W}(i) - (n^\alpha - 1)\right|.\end{aligned}$$

By Lemma 3, both absolute values can be bounded by a constant, say C. Hence, by (1), Δ can be bounded by

$$\sum_{i \in \mathcal{I}_\alpha} \binom{n}{i} H(i+1) C = \frac{C}{n^\alpha - 1} \sum_{i \in \mathcal{I}_\alpha} \binom{n}{i} H(i) = O(n^{-\alpha}). \qquad \square$$

Lemma 5. *For $n \in \mathbb{N}$, $i \in \{0, 1, \dots\}$, and $\beta > 0$, we have*

$$\binom{n}{i} \geq (n^\beta - 1)\binom{n}{i+1} \qquad \textit{iff} \qquad i \geq n - n^{1-\beta} - n^{-\beta}. \tag{2}$$

For $-1 \leq i \leq n$, we have

$$\binom{n}{i} W(i) = \binom{n}{i+1}. \tag{3}$$

Proof. The second statement is an immediate consequence of binomial coefficient definition. For the first one, we handle the trivial case $i \geq n$ separately. To prove the remaining case, we proceed by a sequence of equivalent conditions

$$\begin{aligned}
\binom{n}{i} &\geq (n^\beta - 1)\binom{n}{i+1} \\
\frac{i+1}{n-i} &\geq n^\beta - 1 \\
\frac{1}{n^\beta - 1} &\geq \frac{n+1}{i+1} - 1 \\
\frac{n^\beta}{n^\beta - 1} &\geq \frac{n+1}{i+1} \\
\frac{i+1}{n+1} &\geq \frac{n^\beta - 1}{n^\beta} \\
i &\geq \frac{(n^\beta - 1)(n+1) - n^\beta}{n^\beta} \\
i &\geq \frac{n^{\beta+1} - n - 1}{n^\beta} \\
i &\geq n - n^{1-\beta} - n^{-\beta}.
\end{aligned}$$

□

Lemma 6. *Let i be a function of $n \in \mathbb{N}$, $i \to \infty$ as $n \to \infty$, $0 < i < n$. Then*

$$\binom{n}{i} \sim \sqrt{\frac{1}{2\pi}} \cdot \frac{n^n}{i^i(n-i)^{n-i}} \cdot \sqrt{\frac{n}{i(n-i)}} \quad .$$

Proof. Follows immediately from Stirling's approximation. □

Lemma 7. *Let $1 > \beta > 0$, $\sigma = n - n^{1-\beta} - n^{-\beta}$, $\theta = \lfloor \sigma \rfloor$, and*

$$B(n) = \frac{(n^\beta - 1)^{\theta+1} n^n}{n^{n\beta}(\theta+1)^{\theta+1}(n-\theta-1)^{n-\theta-1}} \quad .$$

Then $B(n) = O(1)$.

Proof. It is easy to see that

$$\begin{aligned}
B(n) &= \left(\frac{n - n^{1-\beta}}{\theta+1}\right)^{\theta+1} \left(\frac{n^{1-\beta}}{n-\theta-1}\right)^{n-\theta-1} \\
&\leq \left(\frac{n - n^{1-\beta}}{\sigma}\right)^{n} \left(\frac{n^{1-\beta}}{n-\sigma-2}\right)^{n^{1-\beta}} \quad .
\end{aligned}$$

We bound the factors separately. For the first, we have

$$\begin{aligned}
\left(\frac{n - n^{1-\beta}}{n - n^{1-\beta} - n^{-\beta}}\right)^{n} &= \left(1 + \frac{1}{n^{1+\beta} - n - 1}\right)^{\left(n^{1+\beta} - n - 1\right)\frac{n}{n^{1+\beta}-n-1}} \\
&\leq \exp\left(\frac{n}{n^{1+\beta} - n - 1}\right) \to 1.
\end{aligned}$$

Since $n-\sigma-2 = n^{1-\beta}-2+n^{-\beta}$, the second factor equals

$$\begin{aligned}\left(\frac{n^{1-\beta}}{n^{1-\beta}-2+n^{-\beta}}\right)^{n^{1-\beta}} &= \left(1+\frac{2-n^{-\beta}}{n^{1-\beta}-2+n^{-\beta}}\right)^{\frac{n^{1-\beta}-2+n^{-\beta}}{2-n^{-\beta}}\cdot\frac{n^{1-\beta}\left(2-n^{-\beta}\right)}{n^{1-\beta}-2+n^{-\beta}}} \\ &\le \exp\left(n^{1-\beta}\cdot\frac{2-n^{-\beta}}{n^{1-\beta}-2+n^{-\beta}}\right) \to e^2.\end{aligned}$$

□

Proposition 3. *Let $1>\beta>0$. Then $\Delta(\mathsf{H}_n(1-1/n^\beta), \mathsf{PD}_n(1-1/n^\beta)) = O(1/n^{(1-\beta)/2})$.*

Proof. In this case, for $0 \le i \le n$ we have

$$H(i) = \left(1-\frac{1}{n^\beta}\right)^i \left(\frac{1}{n^\beta}\right)^{n-i} = \left(\frac{1}{n^\beta}\right)^n \left(n^\beta-1\right)^i,$$

so we obtain $H(i+1) = H(i)(n^\beta-1)$ for $0 \le i \le n-1$. By (3) and by using the fact that $H(n+1)=0$, the analyzed statistical distance equals to

$$\begin{aligned}\Delta &= \sum_{i=-1}^{n} \binom{n}{i} |H(i)-D(i)| \\ &= \sum_{i=-1}^{n} \left|\binom{n}{i}H(i) - \binom{n}{i}H(i+1)W(i)\right| \\ &= \sum_{i=-1}^{n} \left|\binom{n}{i}H(i) - \binom{n}{i+1}H(i+1)\right| \\ &= H(0) + \sum_{i=0}^{n} \left|\binom{n}{i}H(i) - \binom{n}{i+1}H(i)(n^\beta-1)\right| \\ &= H(0) + \sum_{i=0}^{n} H(i)\left|\binom{n}{i} - \binom{n}{i+1}(n^\beta-1)\right|.\end{aligned}$$

Let $\sigma = n - n^{1-\beta} - n^{-\beta}$ and $\theta = \lfloor \sigma \rfloor$. Then by (2), we get

$$\begin{aligned}&\sum_{i=-1}^{\theta}\left(\binom{n}{i+1}H(i+1) - \binom{n}{i}H(i)\right) + \sum_{i=\theta+1}^{n}\left(\binom{n}{i}H(i) - \binom{n}{i+1}H(i+1)\right) \\ &\quad = 2\binom{n}{\theta+1}H(\theta+1).\end{aligned}$$

Stirling's approximation of binomial coefficient (Lemma 6) gives us

$$\Delta \sim \sqrt{\frac{2}{\pi}} \cdot \sqrt{\frac{n}{(\theta+1)(n-\theta-1)}} \cdot \frac{(n^\beta-1)^{\theta+1}n^n}{n^{n\beta}(\theta+1)^{\theta+1}(n-\theta-1)^{n-\theta-1}} \ .$$

Finally, by Lemma 7, we have

$$\Delta = O\left(\sqrt{\frac{n}{\theta+1}} \cdot \sqrt{\frac{1}{\lceil n^{1-\beta}+n^{-\beta}\rceil}}\right) = O(n^{(\beta-1)/2}).$$

□

Theorem 2. *There exists a suitable polynomial $n = n(\kappa) = O(\kappa^{1/\beta})$, so that the scheme described in Construction 2 is fully $1/\kappa$-bideniable.*

Proof. We use Proposition 1. It is immediately seen that $\epsilon_2 = negl(\kappa)$. By Lemma 2, we have $\Delta(V(\mathsf{H}_n(\rho)), \mathsf{Ber}(1/2)) = negl(\kappa)$. We show that $\epsilon_1 + negl(\kappa) < 1/(2\kappa)$ for sufficiently large values of κ. Since $\mathbf{f} \subseteq \mathbf{a}$ holds always, the condition $\mathbf{f} \cap \mathbf{a} \subseteq \mathbf{f} \cap \mathbf{b}$ is equivalent to $\mathbf{f} \subseteq \mathbf{b}$. By definition, $|\mathbf{f}| = 1$ with overwhelming probability, and $\mathbf{f}$ is sampled independently on $\mathbf{b}$. Therefore $\mathbf{f} \nsubseteq \mathbf{b}$ with probability $n^{-\beta} = O(\kappa^{-1})$ which can be bounded by $1/(2\kappa) - negl(\kappa)$ by a wise choice of coefficient in $n = O(\kappa^{1/\beta})$.

Finally, we show that $\epsilon_3 + negl(\kappa) < 1/(2\kappa)$. We are going to apply Lemma 1 with $A = B = [n]$. The distributions X and Y are those of $(\mathbf{a}, \mathbf{b})$ and $(\mathbf{a} \setminus \mathbf{f}, \mathbf{b} \setminus \mathbf{f})$, respectively. We put $S = \mathcal{L}_\alpha$. By Lemma 4, both $\mathbf{a} \in S$ and $\mathbf{a} \setminus \mathbf{f} \in S$ with overwhelming probability. By Proposition 2, we have $\Delta(\mathbf{a}, \mathbf{a} \setminus \mathbf{f}) = O(n^{-\alpha})$. Conditioning on the value of $\mathbf{a}$ (or $\mathbf{a} \setminus \mathbf{f}$ in the second case), we can split the set of coordinates $[n]$ into two groups based on the preimages of $\mathbf{a}$. The first group Q contains coordinates not used for $\tilde{\mathsf{M}}$-preimages, i.e., $i \in Q$ whenever $\tilde{\mathsf{M}}(\mathbf{a} \cup \mathbf{e}_i) \neq \mathbf{a}$ or $\tilde{\mathsf{M}}((\mathbf{a} \setminus \mathbf{f}) \cup \mathbf{e}_i) \neq \mathbf{a} \setminus \mathbf{f}$, respectively. The second group R contains the remaining coordinates. It is easy to see that conditioning on $\mathbf{a}$, the distribution of $\mathbf{b}$ is the product of distributions of substrings $\mathbf{b}_Q$ and $\mathbf{b}_R$, and the same decomposition works for $\mathbf{b} \setminus \mathbf{f}$ (when conditioning on $\mathbf{a} \setminus \mathbf{f}$). Moreover, we see that the distributions of $\mathbf{b}_Q$ and $(\mathbf{b} \setminus \mathbf{f})_Q$ are exactly the same since no preimages use coordinates in Q. By Lemma 4 and Lemma 3, $|R| \geq n^\alpha - c$ for some constant c. Then, we apply Proposition 3 to $\mathbf{b}_R$ and $(\mathbf{b} \setminus \mathbf{f})_R$ to compute distance of these substrings as $|R|^{(\beta-1)/2} \leq (n^\alpha - c)^{(\beta-1)/2} = O(n^{\alpha(\beta-1)/2}) = O(n^{-\beta})$, which can again be bounded by $1/(2\kappa) - negl(\kappa)$ by a wise choice of $n(\kappa)$. □

5 Subset Matching Algorithm

In this section, we prove Theorem 1. To do that, we construct the algorithm M and prove its properties.

Let L_i be the destination layer, i.e., we map an element $\mathbf{x} \in L_{i+1}$ to an element $\mathbf{y} \in L_i$. We start with an informal description of the algorithm for the case $\underline{W}(i) = \overline{W}(i)$, i.e., when the size of upper (source) layer is divisible by the size of lower (target) layer. We assign to each coordinate its weight, -1 for coordinate containing bit 0 and $\underline{W}(i)$ for bit 1. We define weight of a range of coordinates to be sum of weights of individual coordinates. Then, we define min-weight of a coordinate as the minimum of weights of ranges starting by the given coordinate and ending on some higher coordinate. The algorithm outputs the lowest coordinate with min-weight at least 1.

The situation is slightly more complicated when $\overline{W}(i) = \underline{W}(i) + 1$. Then, some 1-coordinates get $\underline{W}(i)$ weight and some $\overline{W}(i)$, so that the weight of the whole string is 1. The ordering of weights is fixed, e.g., some number of lower 1-coordinates always get $\underline{W}(i)$ and the remaining 1-coordinates get $\overline{W}(i)$. We remark that our choice of weight ordering is arbitrary, but it is sufficient to prove just one case for our purposes.

Now we proceed with a rigorous treatment of the algorithm. Let $1 \leq k \leq n$ be a coordinate. For $\mathbf{x} \in \mathcal{P}([n])$, let us denote the number of ones on lower positions by $\rho(\mathbf{x}, k) = |\{m < k \mid x_m = 1\}|$. For $i \geq 0$, let

$$\begin{aligned}\overline{C}(i) &= (n-i) - (i+1)\underline{W}(i)\\ \underline{C}(i) &= (i+1) - \overline{C}(i).\end{aligned}$$

For $k \leq l \leq n$, we define *k, l-weight of* $\mathbf{x}$ to be

$$\mathrm{w}^{(i)}(\mathbf{x}, k, l) = \begin{cases} -1 & k = l \wedge x_k = 0 \\ \overline{W}(i) & k = l \wedge x_k = 1 \wedge \rho(\mathbf{x}, k) < \overline{C}(i) \\ \underline{W}(i) & k = l \wedge x_k = 1 \wedge \rho(\mathbf{x}, k) \geq \overline{C}(i) \\ \sum_{j=k}^{l} \mathrm{w}^{(i)}(\mathbf{x}, j, j) & k < l, \end{cases}$$

and we define *k-th min-weight of* $\mathbf{x}$ as $\mu^{(i)}(\mathbf{x}, k) = \min_{l=k}^{n} \mathrm{w}^{(i)}(\mathbf{x}, k, l)$.

Lemma 8. *Let $\mathbf{y} \subseteq \mathbf{x}$, $1 \leq k \leq l \leq n$, and $\mathbf{x} \cap [k-1] = \mathbf{y} \cap [k-1]$. Then $\mathrm{w}^{(i)}(\mathbf{y}, k, l) \leq \mathrm{w}^{(i)}(\mathbf{x}, k, l)$.*

Proof. The assumption $\mathbf{x} \cap [k-1] = \mathbf{y} \cap [k-1]$ implies that $\rho(\mathbf{x}, k) = \rho(\mathbf{y}, k)$. The rest follows immediately from the definition. □

Lemma 9. *For $i \geq 0$ we have $\overline{W}(i)\overline{C}(i) + \underline{W}(i)\underline{C}(i) = n - i$.*

Proof. By definition, $\overline{W}(i)\overline{C}(i) + \underline{W}(i)\underline{C}(i) = \overline{C}(i)(\overline{W}(i) - \underline{W}(i)) + (i+1)\underline{W}(i)$. The result then follows by considering the cases $\overline{W}(i) = \underline{W}(i)$ and $\overline{W}(i) = \underline{W}(i) + 1$ separately. □

Lemma 10. *For any $\mathbf{x} \in L_{i+1}$, we have $\mathrm{w}^{(i)}(\mathbf{x}, 1, n) = 1$.*

Proof. By definition, $\mathrm{w}^{(i)}(\mathbf{x}, 1, n) = (-1)(n - i - 1) + \overline{C}(i)\overline{W}(i) + \underline{C}(i)\underline{W}(i)$, which equals 1 by Lemma 9. □

Definition 1. Let M be an algorithm that for input $\mathbf{x} \in \mathcal{P}([n])$) does the following:

1. Finds i such that $\mathbf{x} \in L_{i+1}$.
2. If $i = -1$, fails.
3. Finds the lowest coordinate k such that $\mu^{(i)}(\mathbf{x}, k) \geq 1$.
4. Returns $\mathbf{e}_k$.

We claim that the algorithm M is the desired one from Theorem 1.

Lemma 11. *Let $\mathbf{x} \in L_{i+1}$ for $i \geq 0$. Then $\mathsf{M}(\mathbf{x})$ is well-defined, i.e., there is a coordinate k for which $\mu^{(i)}(\mathbf{x}, k) \geq 1$.*

Proof. We know that $\mathrm{w}^{(i)}(\mathbf{x}, 1, n) \geq 1$, by Lemma 10. If $\mu^{(i)}(\mathbf{x}, 1) \geq 1$, we are done. Otherwise, there is a coordinate $l_1 \geq 1$ such that $\mathrm{w}^{(i)}(\mathbf{x}, 1, l_1) < 1$. If $\mu^{(i)}(\mathbf{x}, l_1 + 1) \geq 1$, we are done, in the other case, we repeat the previous step to obtain a coordinate $l_2 > l_1$, and so on. Eventually, this process stops with a coordinate l_j such that $\mu^{(i)}(\mathbf{x}, l_j + 1) \geq 1$. □

Definition 2. Let $\mathbf{y} \in L_i$ and let k be a coordinate for which $y_k = 0$. Then k is called *suitable for* $\mathbf{y}$ if for any lower coordinate $l < k$ we have $\mathrm{w}^{(i)}(\mathbf{y}, l, k-1) < 1$. Otherwise, k is called *unsuitable for* $\mathbf{y}$.

Lemma 12. *Let k be unsuitable for $\mathbf{y}$ and let $\mathbf{x} = \mathbf{y} \cup \mathbf{e}_k$. Then $\mathsf{M}(\mathbf{x}) \neq \mathbf{e}_k$.*

Proof. By assumption, there is a coordinate $l < k$ for which $\mathrm{w}^{(i)}(\mathbf{y}, l, k-1) \geq 1$. Take l to be maximal. We claim that $\mathrm{w}^{(i)}(\mathbf{y}, l, m) \geq 1$ whenever $l \leq m < k-1$. If not, then $\mathrm{w}^{(i)}(\mathbf{y}, l, m) < 1$ for some $m < k-1$, and the assumption implies $\mathrm{w}^{(i)}(\mathbf{y}, m+1, k-1) \geq 1$. But that contradicts the maximality of l.

By Lemma 8, we have $\mathrm{w}^{(i)}(\mathbf{x}, l, m) \geq \mathrm{w}^{(i)}(\mathbf{y}, l, m)$ for $m \geq l$. In combination with the previous inequalities, we get $\mathrm{w}^{(i)}(\mathbf{x}, l, m) \geq 1$ for $l \leq m < k$. Now, if $\mu^{(i)}(\mathbf{x}, k) \geq 1$ (so $\mathbf{x}$ is a possible preimage), then it is also $\mu^{(i)}(\mathbf{x}, l) \geq 1$, and there is a lower coordinate that could be used instead. □

Lemma 13. *Let $\mathbf{y} \in L_i$. Then the number of suitable coordinates for $\mathbf{y}$ is bounded from below by $\underline{W}(i)$.*

Proof. Each coordinate containing one makes unsuitable at most $\overline{W}(i)$ or $\underline{W}(i)$ zero coordinates, depending on the position. We have $\overline{C}(i)$ coordinates of the first type and $\underline{C}(i) - 1$ coordinates of the second type. The number of zero coordinates is $n - i$. Thus, the number of suitable coordinates is bounded from below by

$$n - i - \overline{C}(i)\overline{W}(i) - (\underline{C}(i) - 1)\underline{W}(i) = n - i - \overline{C}(i)\overline{W}(i) - \underline{C}(i)\underline{W}(i) + \underline{W}(i)$$

which equals $\underline{W}(i)$ by Lemma 9. □

Lemma 14. *Let k be one of the $\underline{W}(i)$ uppermost suitable coordinates for $\mathbf{y} \in L_i$. Then for $\mathbf{x} = \mathbf{y} \cup \mathbf{e}_k$, we have $\mathsf{M}(\mathbf{x}) = \mathbf{e}_k$.*

Proof. Since k is suitable, it suffices to show that $\mu^{(i)}(\mathbf{x}, k) \geq 1$. To obtain a contradiction, let us assume that for some $m \geq k$ we have $\mathrm{w}^{(i)}(\mathbf{x}, k, m) < 1$. Since $x_k = 1$ and $y_k = 0$, we have $\mathrm{w}^{(i)}(\mathbf{x}, k, k) \geq \underline{W}(i)$ and $\mathrm{w}^{(i)}(\mathbf{y}, k, m) < -\underline{W}(i)$. Take the lowest possible coordinates $k = k_0 < k_1 < \cdots < k_{\underline{W}(i)} \leq m$ such that $\mathrm{w}^{(i)}(\mathbf{y}, k, k_j) < -j$. Clearly, such coordinates exist. We will show that each k_j is suitable for $\mathbf{y}$, yielding a contradiction with lemma assumption. By the minimality of k_j, we have $\mathrm{w}^{(i)}(\mathbf{y}, k, l) \geq -j$ for $k \leq l < k_j$, so $\mathrm{w}^{(i)}(\mathbf{y}, l+1, k_j) < 0$. But that means $\mathrm{w}^{(i)}(\mathbf{y}, r, k_j - 1) < 1$ for any $k < r < k_j$. The inequality also holds for $r = k$, by assumption, and for $r < k$, by suitability of k. Thus, each k_j is suitable for $\mathbf{y}$. □

Lemma 15. *Let k be a suitable coordinate for $\mathbf{y} \in L_i$, so that there are at least $\overline{W}(i)$ higher suitable coordinates. Let $\mathbf{x} = \mathbf{y} \cup \mathbf{e}_k$. Then $\mathsf{M}(\mathbf{x}) \neq \mathbf{e}_k$.*

Proof. Let $k = k_0 < k_1 < \cdots < k_{\overline{W}(i)}$ be suitable coordinates for $\mathbf{y}$. Since $y_{k_j} = 0$ for all j, we have $\mathrm{w}^{(i)}(\mathbf{y}, k_j, k_j) = -1$. Whenever $k_{j-1} < k_j - 1$, we know that $\mathrm{w}^{(i)}(\mathbf{y}, k_{j-1}+1, k_j - 1) < 1$ by suitability of k_j. To put everything together, we have $\mathrm{w}^{(i)}(\mathbf{y}, k, k_{\overline{W}(i)}) \leq -\overline{W}(i) - 1$. By definition, that implies $\mathrm{w}^{(i)}(\mathbf{x}, k, k_{\overline{W}(i)}) \leq 0$, so $\mu^{(i)}(\mathbf{x}, k) < 1$. □

Proof of Theorem 1. Let $\mathbf{y} \in L_i$ and we count the number of preimages in $\tilde{\mathsf{M}}$. Each preimage uses distinct coordinate k for which $y_k = 0$. By Lemma 12, we can ignore unsuitable coordinates. By Lemma 13 and Lemma 14, the number of preimages is at least $\underline{W}(i)$. By Lemma 15, we cannot have more than $\overline{W}(i)$ preimages. □

6 Conclusion

This work fixes an important part of the study of deniable encryption feasibility. However, some problems in this areas still remain open, as well as the question of efficiency of the existing schemes.

Acknowledgements The author would like to thank to Michal Hojsík for review of the paper, to Petr Nižnanský for his help with solving a paradoxical statement, and to Jiří Tůma for discussions about various aspects of deniable encryption.

References

[ANS98] Ross J. Anderson, Roger M. Needham, and Adi Shamir. The steganographic file system. In *Information Hiding*, pages 73–82, Portland, Oregon, April 1998.

[Bea96] Donald Rozinak Beaver. Plausible deniability. In *PRAGOCRYPT*, pages 272–288, Prague, Czech Republic, September 1996.

[BNNO11] Rikke Bendlin, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. Lower and upper bounds for deniable public-key encryption. In *ASIACRYPT*, pages 125–142, Seoul, Korea, December 2011.

[CDNO97] Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. Deniable encryption. In *CRYPTO*, pages 90–104, Santa Barbara, California, August 1997.

[CFGN96] Ran Canetti, Uriel Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *STOC*, pages 639–648, Philadelphia, Pennsylvania, May 1996.

[DF11] Markus Dürmuth and David Mandell Freeman. Deniable encryption with negligible detection probability: An interactive construction. In *EUROCRYPT*, pages 610–626, Tallinn, Estonia, May 2011. An updated version in Cryptology ePrint Archive, Report 2011/066.

[DN00] Ivan Damgård and Jesper Buus Nielsen. Improved non-committing encryption schemes based on a general complexity assumption. In *CRYPTO*, pages 432–450, Santa Barbara, California, August 2000.

[DS12] Dana Dachman-Soled. On the impossibility of sender-deniable public key encryption. Cryptology ePrint Archive, Report 2012/727, 2012. `http://eprint.iacr.org/`.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, Victoria, British Columbia, Canada, May 2008.

[HKX03] HweeHwa Pang, Kian-Lee Tan, and Xuan Zhou. StegFS: A steganographic file system. In *ICDE*, pages 657–667, Bangalore, India, March 2003.

[KKK08] Marek Klonowski, Przemysław Kubiak, and Mirosław Kutyłowski. Practical deniable encryption. In *SOFSEM*, pages 599–609, Nový Smokovec, Slovakia, January 2008.

[MK99] Andrew D. McDonald and Markus G. Kuhn. StegFS: A steganographic file system for Linux. In *Information Hiding*, pages 462–477, Dresden, Germany, September 1999.

[OPW11] Adam O'Neill, Chris Peikert, and Brent Waters. Bi-deniable public-key encryption. In *CRYPTO*, pages 525–542, Santa Barbara, California, August 2011.

[Šeb12] Marcel Šebek. Deniable encryption. Master's thesis, Charles University in Prague, 2012. `http://www.karlin.mff.cuni.cz/~sebek/research/deniable-encryption-thesis.pdf`.

[SW13] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. Cryptology ePrint Archive, Report 2013/454, 2013. `http://eprint.iacr.org/`.

[XHK04] Xuan Zhou, HweeHwa Pang, and Kian-Lee Tan. Hiding data accesses in steganographic file system. In *ICDE*, pages 572–583, Boston, Massachusetts, March 2004.

A Failure of the Original Scheme

In this section, we describe the scheme proposed by Bendlin et al. [BNNO11], and we show that it admits a constant-advantage distinguisher. We made a few non-substantial changes so that the description fits generic Construction 1.

Construction 3. Let $n = \kappa^5$ and $r = \kappa^4$. Let us consider the following algorithms.

1. For $0 \le \gamma \le 1$, let us define an auxiliary sampling algorithm $N_\gamma\colon \mathcal{P}([r]) \to \mathcal{P}([n])$, $\mathbf{x} \mapsto \mathbf{y}$, so that if $j \notin \mathbf{x}$, then $\{(j-1)\kappa + 1, \dots, j\kappa\} \cap \mathbf{y} = \emptyset$, and if $j \in \mathbf{x}$, then for $i = 1, \dots, \kappa$, the relation $(j-1)\kappa + i \in \mathbf{y}$ holds with probability γ, and this value is sampled independently for each j and i.

2. To sample $\mathbf{b} \leftarrow \mathcal{B}_\kappa$, one first samples $\mathbf{b}' \in \mathcal{P}([r])$, so that bits of $\mathbf{b}'$ are sampled independently, and each equals 1 with probability $1 - 1/\kappa^2$. Then $\mathbf{b} \leftarrow N_1(\mathbf{b}')$.

3. To sample $\mathbf{a} \leftarrow \mathcal{A}_\kappa$, one first samples $\mathbf{a}' \in \mathcal{P}([r])$ uniformly, so that $|\mathbf{a}'| \bmod 2 = m$, and puts $\mathbf{a} \leftarrow N_{1/2}(\mathbf{a}')$.

4. The mapping V splits the input string into κ-bit blocks, let $C \in \{0, \dots, \kappa^4\}$ be the number of non-zero ones. Then V returns $C \bmod 2$, i.e., V calculates parity of non-zero κ-blocks in the input string.

5. Faking mapping F splits the input into κ-bit blocks, finds the block with lexicographically lowest non-zero value, and outputs the set containing all coordinates determining this block. Note that this block is determined uniquely with overwhelming probability.

The failing assumption of Proposition 1 is that $(\mathbf{a}, \mathbf{b}) \overset{\mathrm{c}}{\approx}_{\epsilon_3} (\mathbf{a} \setminus \mathbf{f}, \mathbf{b} \setminus \mathbf{f})$. In fact, this indistinguishability condition would pass if the adversary is given just $(\mathbf{a}', \mathbf{b})$ (or its corresponding fake counterpart). However, that would disallow sender and receiver to agree on the common faking block index.

We claim that the distributions of $\mathbf{a}$ and $\mathbf{a} \setminus \mathbf{f}$ are distinguishable with advantage that tends to $1/e$, and there is an efficient distinguisher. We sketch the main ideas first. We can convert each κ-bit block of the bitstring into a real number belonging to the interval $[0, 1]$. Then, we forget about order of these numbers. Thus, the distinguisher is asked to tell apart the following cases:

- A uniformly random set of n real numbers from $[0,1]$ is returned.
- A uniformly random set of $n+1$ real numbers from $[0,1]$ is sampled, the smallest one is dropped, and the resulting set is returned.

By means of order statistics, it can be shown that a maximum-likelihood distinguisher compares the smallest number with $1/(n+1)$, and if it is smaller, it answers that the first case is true.

We proceed by a formal argument. Let us split $\mathbf{a}$ into κ-bit strings and convert each of these strings into its corresponding natural number. Denote the set of these numbers by N_0 and let $N = N_0 \setminus \{0\}$. Similarly, $\mathbf{a} \setminus \mathbf{f}$ corresponds to the set N' which can be obtained from N by dropping the lowest element (in the natural ordering).

Up to a negligible statistical distance, N can be equivalently produced by first sampling the number of non-zero elements k from binomial distribution of length r and parameter $1/2$, and then sampling k natural numbers uniformly independently from the set $\{1,\ldots,\kappa\}$. By Chernoff bound, $m = |N| \geq \lambda r$ for some constant $\lambda > 0$. Let us condition on the value of m.

To simplify our analysis, we describe the situation using continuous distributions and real numbers. Let R_m be a set of m numbers sampled uniformly from the unit interval $[0,1]$. Let R'_m be a set of r numbers obtained by sampling from R_{m+1} and dropping the lowest value. Since the probability that two elements of R_m are equal is zero, we can consider the following probability space:

$$M_m = \{x_1,\ldots,x_m \mid 0 < x_1 < \cdots < x_m < 1\}.$$

Its volume is

$$\int_{M_m} 1 \,\mathrm{d}x_1 \ldots \mathrm{d}x_m = \int_0^1 \left(\cdots \left(\int_0^{x_3} \left(\int_0^{x_2} 1 \,\mathrm{d}x_1 \right) \mathrm{d}x_2 \right) \cdots \right) \mathrm{d}x_m = \frac{1}{m!}.$$

Therefore, density of the distribution of R_m is $m!$, and cumulative distribution function is

$$H_m(y_1,\ldots,y_m) = \int_{([0,y_1]\times\cdots\times[0,y_m])\cap M_m} m! \,\mathrm{d}x_1 \ldots \mathrm{d}x_m.$$

It is easy to see that CDF of the distribution of R'_m is

$$\begin{aligned} D_m(y_1,\ldots,y_m) &= H_{m+1}(1,y_1,\ldots,y_m) \\ &= \int_{([0,1]\times[0,y_1]\times\cdots\times[0,y_m])\cap M_{m+1}} (m+1)! \,\mathrm{d}x_1 \ldots \mathrm{d}x_{m+1} \\ &= \int_{([0,y_1]\times\cdots\times[0,y_m])\cap M_m} (m+1)! \left(\int_0^{x_2} 1 \,\mathrm{d}x_1 \right) \mathrm{d}x_2 \ldots \mathrm{d}x_{m+1} \\ &= \int_{([0,y_1]\times\cdots\times[0,y_m])\cap M_m} x_2 (m+1)! \,\mathrm{d}x_2 \ldots \mathrm{d}x_{m+1}, \end{aligned}$$

so its density is $x_1(m+1)!$ (when we shift down the variables back to $x_1,\ldots,x_m$). Before we state the main result, we need a lemma.

Lemma 16. *Let $\mathcal{F}$ be the set of all measurable functions $f\colon [0,1] \to [0,1]$, and let σ be a mapping on this set that assigns to a function $f \in \mathcal{F}$ the function $\int_x^1 f(y)\,\mathrm{d}y$. Then for any $j \in \mathbb{N}_0$, we have $\sigma^j(1) = (1-x)^j/j!$.*

Proof. For $j = 0$, that is just the definition. By induction, we have

$$\int_x^1 \frac{(1-y)^j}{j!}\,\mathrm{d}y = -\frac{(1-y)^{j+1}}{(j+1)!}\bigg|_x^1 = \frac{(1-x)^{j+1}}{(j+1)!}. \qquad \square$$

Theorem 3. *Statistical distance Δ of distributions given by cumulative distribution functions H_m and D_m tends to $1/e$ as m grows to infinity. A distinguisher realizing this distance compares first coordinate of the sample with $1/(m+1)$, and if the value is lower, answers that the distribution is H_m.*

Proof. We compute the statistical distance. In the last equality, we use Lemma 16:

$$\begin{aligned}
2\Delta &= \int_{M_m} m! \cdot |1 - x_1(m+1)|\,\mathrm{d}x_1 \ldots \mathrm{d}x_m = \\
&= \int_0^1 m! \cdot |1 - x_1(m+1)| \left(\int_{[x_1,1]^{m-1} \cap M_{m-1}} 1\,\mathrm{d}x_2 \ldots dx_m \right) \mathrm{d}x_1 \\
&= \int_0^1 m! \cdot |1 - x_1(m+1)| \frac{(1-x_1)^{m-1}}{(m-1)!}\,\mathrm{d}x_1.
\end{aligned}$$

We put $\delta(x_1) = m(1 - x_1(m+1))(1-x_1)^{m-1}$ to obtain

$$\begin{aligned}
2\Delta &= \int_0^{1/(m+1)} \delta(x_1)\,\mathrm{d}x_1 - \int_{1/(m+1)}^1 \delta(x_1)\,\mathrm{d}x_1 \\
&= m x_1 (1-x_1)^m \big|_0^{1/(m+1)} - m x_1 (1-x_1)^m \big|_{1/(m+1)}^1 \\
&= 2\left(1 - \frac{1}{m+1}\right)^{m+1} \to 2/e.
\end{aligned}$$

The statement about distinguisher is clear from the fact that $\delta(x) > 0$ if and only if $x < 1/(m+1)$. □