

Lossy Trapdoor Permutations with Improved Lossiness^{*}

Benedikt Auerbach¹, Eike Kiltz¹, Bertram Poettering², and Stefan Schoenen³

¹ Horst-Görtz Institute for IT Security, Ruhr-University Bochum, Bochum, Germany
{benedikt.auerbach,eike.kiltz}@rub.de

² Royal Holloway, University of London, Egham, UK
bertram.poettering@rhul.ac.uk

³ paluno – The Ruhr Institute for Software Technology, University of
Duisburg-Essen, Essen, Germany
stefan.schoenen@paluno.uni-due.de

Abstract. Lossy trapdoor functions (Peikert and Waters, STOC 2008 and SIAM J. Computing 2011) imply, via black-box transformations, a number of interesting cryptographic primitives, including chosen-ciphertext secure public-key encryption. Kiltz, O’Neill, and Smith (CRYPTO 2010) showed that the RSA trapdoor permutation is lossy under the Phi-hiding assumption, but syntactically it is not a lossy trapdoor function since it acts on \mathbb{Z}_N and not on strings. Using a domain extension technique by Freeman et al. (PKC 2010 and J. Cryptology 2013) it can be extended to a lossy trapdoor permutation, but with considerably reduced lossiness.

In this work we give new constructions of lossy trapdoor permutations from the Phi-hiding assumption, the quadratic residuosity assumption, and the decisional composite residuosity assumption, all with improved lossiness. Furthermore, we propose the first all-but-one lossy trapdoor permutation from the Phi-hiding assumption. A technical vehicle used for achieving this is a novel transform that converts trapdoor functions with index-dependent domain to trapdoor functions with fixed domain.

1 Introduction

LOSSY TRAPDOOR FUNCTIONS. Lossy trapdoor functions (LTFs) are like classic (one-way) trapdoor functions but with strengthened security properties. Instances of an LTF can be created in two computationally indistinguishable ways: An instance generated with the standard key-generation algorithm describes an injective function that can be efficiently inverted using the trapdoor; and an instance generated with the lossy key-generation algorithm describes a “lossy” function, meaning its range is considerably smaller than its domain. The *lossiness factor* $L \geq 1$, defined as the ratio of the cardinalities of domain and range,

^{*} The full version of this article can be found in the IACR eprint archive as article 2018/1183 at <https://eprint.iacr.org/2018/1183>.

measures the LTF’s quality.⁴ The larger the lossiness factor, the better the cryptographic properties of the LTF. In case the non-lossy instances define permutations, we will refer to the whole object as a lossy trapdoor permutation (LTP).

Lossy trapdoor functions were introduced by Peikert and Waters [21,22] who showed that they imply (via black-box constructions) fundamental cryptographic primitives such as classic trapdoor functions, collision-resistant hash functions, oblivious transfer, and chosen-ciphertext secure public-key encryption. Furthermore, LTFs have found various other applications, including deterministic public-key encryption [7], OAEP-based public-key encryption [17], “hedged” public-key encryption for protecting against bad randomness [2,4], security against selective opening attacks [5], efficient non-interactive string commitments [20], threshold encryption [26], correlated-product secure trapdoor functions [24], adaptive trapdoor functions [16], and many others.

LTFs WITH INDEX-DEPENDENT DOMAINS. In the original definition by Peikert and Waters, all instances of an LTF are defined over the same fixed domain $\{0, 1\}^k$. That is, the domain is independent of the specific index output by the key-generation algorithm (‘index’ is used synonym with the public key describing the instance). Subsequently, LTFs were generalized to *LTFs with index-dependent domains* [11] where the domain may depend on the function’s index. To illustrate index-dependent domains, consider the well-known RSA trapdoor permutation $f_{\text{RSA}}: \mathbb{Z}_N \rightarrow \mathbb{Z}_N; x \mapsto x^e \bmod N$. Its index consists of a modulus $N = pq$ (of fixed bit-length k) and an exponent e ; its domain is \mathbb{Z}_N , hence it is index-dependent. For $e \leq 2^{k/4}$, permutation f_{RSA} was proved to be lossy [17] with lossiness factor $L = e$ under the Phi-hiding assumption [9].⁵ Similarly, constructions of trapdoor functions based on quadratic residuosity or Paillier’s assumption yield LTPs with index-dependent domains [10,11].

As pointed out in [11], LTFs with index-dependent domains do not seem to be sufficient for constructing correlated-product secure trapdoor functions [24] or chosen-ciphertext secure public-key encryption [21]. The difficulty is that in these applications a fixed value has to be evaluated on many independently generated instances of the trapdoor function. It is therefore crucial that the domains are the same for all these instances. Furthermore, most constructions of deterministic encryption schemes (e.g., [7,3,8,18,23]) assume message distributions that do not depend on the public key and hence cannot be constructed from LTFs with index-dependent domains. Fortunately, however, LTFs with index-dependent domains turn out to be sufficient for many other applications.

⁴ The original definition of lossy trapdoor functions [21,22] measures lossiness on a logarithmic scale. That is, $\ell := \log_2(L)$ is the lossiness of the LTF and L is the lossiness factor (which we use in this work).

⁵ In brief, the Phi-hiding assumption states that (N, e) , where $N = pq$ and $e \nmid \varphi(N)$, is computationally indistinguishable from (N, e) , where $N = pq$ and $e \mid \varphi(N)$. The Phi-hiding assumption is conjectured to hold for $e \leq N^{1/4-\epsilon}$ and does not hold for $e > N^{1/4}$ (due to a Coppersmith-like attack). If $e \mid \varphi(N)$, then $f_{\text{RSA}}(x) = x^e \bmod N$ is roughly an e -to-1 function.

In [11, Section 3.2], a general domain-extension technique was (implicitly) proposed that transforms an LTF $f: \mathbb{Z}_N \rightarrow \mathbb{Z}_N$ with index-dependent domain \mathbb{Z}_N (with $2^{k-1} \leq N < 2^k$) into an LTF $f_{\text{de}}: \{0, 1\}^k \rightarrow \{0, 1\}^k$ with index-independent domain $\{0, 1\}^k$ by defining

$$f_{\text{de}}(x) := \begin{cases} f(x) & 0 \leq x < N \\ x & N \leq x < 2^k \end{cases} . \quad (1)$$

However, this transform does bad in preserving lossiness, in particular in the case where N is close to 2^{k-1} . Indeed, if the lossiness factor of f is L then the lossiness factor of f_{de} is about $L_{\text{de}} = 2 \cdot L / (L + 1) < 2$. Note that such a small lossiness factor does not even imply one-wayness, i.e., the resulting LTF is, taken by itself, essentially useless. (Based on a result by Mol and Yilek [19] it can still be used to build IND-CCA secure encryption, but with considerably worse efficiency.) In [11, Section 4.4] also an alternative domain-extension technique was sketched that can be used to construct an LTF f_{de} with index-independent domain $\{0, 1\}^{k+\log(L)}$ and lossiness factor $L_{\text{de}} \approx L$. Here, every evaluation of f_{de} requires $\log(L)$ many applications of f . For interesting values of L this is again prohibitively inefficient.

ALL-BUT-ONE LOSSY TRAPDOOR FUNCTIONS. All-but-one lossy trapdoor functions (ABO-LTFs) are a generalization of LTFs. An ABO-LTF is associated with a set $\mathcal{B}r$ of branches. The corresponding generator algorithm is invoked on input a target branch $br^* \in \mathcal{B}r$ and outputs a trapdoor and a family of functions $(f_{br})_{br \in \mathcal{B}r}$ with the property that f_{br} is injective for all $br \neq br^*$ (and can be inverted using the trapdoor), but function f_{br^*} is lossy. Moreover, the lossy branch is hidden (computationally) by the description of the function family. ABO-LTFs with just two branches are equivalent to LTFs, and, similarly to LTFs, ABO-LTFs can have index-independent or index-dependent domains. Using the techniques of Peikert and Waters [21] an ABO-LTF with exponentially large branch set can be constructed from any LTF, but the latter is required to have a sufficiently large lossiness factor L . (This transformation also works for LTFs with index-dependent domains.) Many of the mentioned applications of LTFs require in fact ABO-LTFs.

KNOWN LTFs AND ABO-LTFs. Roughly speaking, cryptographic assumptions are typically rooted in one out of three different environments: over cyclic groups, over lattices, or over RSA moduli. Over cyclic groups as well as over lattices, constructions of LTFs and ABO-LTFs are known [21]. They have index-independent domain and can be instantiated to have an arbitrarily large lossiness factor L . In the RSA setting, the situation is different.⁶ There are constructions known from the quadratic residuosity assumption [11], Paillier’s decisional composite residuosity assumption [11], and from the Phi-hiding assumption [9,17] (for a fourth one, see below). All constructions have index-dependent domains (the transform

⁶ When we say an LTF is “RSA-based” we mean it is defined in respect to some composite number $N = pq$ where p, q are primes. This shall not suggest its security relies on the RSA assumption (the hardness of computing e -th roots).

sketched above fixes this, but the results are essentially useless due to the small lossiness factor). Unfortunately, for the constructions based on the Phi-hiding assumption and the quadratic residuosity assumption the lossiness factor cannot be made arbitrarily large and, in particular, it is not sufficient to construct efficient ABO-TDFs. However, both an index-independent LTF and an ABO-LTF based on the decisional composite residuosity assumption are known [11].

As it is quite general, we describe in more detail the technique from [21] for building LTFs. Starting with an additively homomorphic encryption scheme, function indices correspond with element-wise encryptions of the identity matrix. The range of the construction consists of vectors of ciphertexts. If ElGamal encryption is used to instantiate the encryption scheme one obtains an LTF with security based on DDH. Constructions of LTFs and ABO-LTFs in the same spirit, but that achieve smaller index sizes and output lengths, are proposed in [15,6]. Using a generalization of the Goldwasser–Micali homomorphic encryption scheme [12] allows this construction, in contrast to processing the LTF input bit-by-bit, to consider input values sequences of numbers of some fixed bit-length. The construction’s security is based not only on the DDH assumption but also on the quadratic residuosity assumption for a restricted class of RSA moduli and an additional non-standard assumption, which can be removed by making further restrictions on the modulus.

While the described constructions from [21,15,6] achieve high lossiness factors, a common disadvantage is that their indices are ciphertext matrices and the function ranges are ciphertext vectors, and thus quite large. Further, [15,6] require strong hardness assumptions in a quite restricted RSA setting.

As shown in [27], collision-resistant hash functions, CPA- and CCA-secure public-key encryption, and deterministic encryption can be constructed from adversary-dependent lossy trapdoor functions and ABO-LTFs, a variant of LTFs and ABO-LTFs with relaxed security conditions. The authors give index-independent constructions of these primitives from the factoring assumption for semi-smooth RSA moduli. The proposed instantiations achieve high lossiness factors and have compact indices and ciphertexts of roughly the size of an RSA modulus.

1.1 Our Results

In this work we propose a new general domain-extension transformation that can be used to transform index-dependent LTPs into index-independent LTPs without sacrificing much lossiness. Concretely, our transformation decreases the lossiness factor by at most by a factor of 2. For the special cases of the LTP based on the Phi-hiding assumption and the LTP from [11] based on the quadratic residuosity assumption, a more refined analysis even shows that the lossiness factor effectively stays invariant. That is, ultimately we construct an LTP with index-independent domain $\{0, 1\}^k$ and lossiness factor as large as $L = 2^{k/4}$ from the Phi-hiding assumption, and an LTP with index-independent domain $\{0, 1\}^k$ and lossiness factor 2 from the quadratic residuosity assumption. In comparison, the index-independent variants obtained via the transform implicitly given in [11] would result in lossiness factors of 2 and 4/3 respectively. Furthermore, in

the full version [1] we apply our transformation to the index-dependent LTF and ABO-LTF of [11] based on the decisional composite residuosity assumption. As a result we obtain index-independent variants with slightly larger domain and lossiness factor than the index-independent constructions given in [11]. Finally we construct the first ABO-LTP from (a variant of) the Phi-hiding assumption. We highlight that in particular our Phi-hiding based construction has particularly compact indices (of the size of an RSA modulus) and range elements.

DOMAIN EXTENSION FOR LTFs WITH INDEX-DEPENDENT DOMAINS. We explain our domain extension technique for the special case of a LTF $f: \mathbb{Z}_N \rightarrow \mathbb{Z}_N$ with index-dependent domain \mathbb{Z}_N (with $2^{k-1} \leq N < 2^k$). We use a two-round construction in the spirit of Hayashi, Okamoto and Tanaka [13], who used a similar construction to extend the domain of the RSA one-way permutation. We define the function

$$f'_{\text{de}}: \{0, 1\}^k \rightarrow \{0, 1\}^k, \quad f'_{\text{de}}(x) := f_{\text{de}}(\pi(f_{\text{de}}(x))), \quad (2)$$

where f_{de} is defined in (1) and permutation $\pi: \{0, 1\}^k \rightarrow \{0, 1\}^k$ is given as $\pi(x) = x - (N - 1) \bmod 2^k$. The intuition of this construction is that the LTF f is applied to every $x \in \{0, 1\}^k$ at least once. Indeed, if f is one-way, then f'_{de} defined in (2) is one-way [13]. Our first main result states that if f is a LTF with index-dependent domain and lossiness factor L , then f'_{de} is a LTF with index-independent domain $\{0, 1\}^k$ and lossiness factor $L'_{\text{de}} = L/2$.

In the case of the RSA-based LTF f_{RSA} we can even prove that the lossiness factor of f'_{de} is completely preserved, i.e. $L'_{\text{de}} = L$. Under the Phi-hiding assumption this gives us a LTP with index-independent domain and lossiness factor as large as $k^{1/4}$. We also show how to obtain index-independent LTPs from the quadratic residuosity and the decisional composite residuosity assumption, which have a larger lossiness factor than the constructions of [11].

AN ABO-LTP IN THE RSA SETTING. Our second main result is the construction of an ABO-LTP with index-dependent domain from the Phi-hiding assumption. Our generic domain extension technique also works for ABO-LTFs, so it can be transformed into an ABO-LTP with index-independent domain $\{0, 1\}^k$.

Our construction essentially follows [16, Section 5.2] who construct an adaptive trapdoor function from the instance-independent RSA assumption, a decisional version of the RSA assumption. It makes use of a new primitive that we call *prime family generator* (PFG), an abstraction that may be of independent interest. An instance of a PFG indicates a fixed sequence of (distinct) primes e_1, \dots, e_{2^n} of some specified bit-length $l \geq n/2$. A specific programmability feature allows embedding any given prime at any given position, where the position remains hidden (computationally) from the instance. We give an information-theoretic construction of a PFG that is based on work by Cachin, Micali, and Stadler [9]. A PFG instance consists of l^2 bits and we leave it as an open problem to construct a (computationally secure) PFG with improved parameters, for example by using the PRF-based construction as implicitly in the work of Hohenberger and Waters [14].

Given a PFG we define our new RSA-setting based ABO-LTP for a branch $br \in \{0, 1\}^n$ as

$$f_{br}: \mathbb{Z}_N \rightarrow \mathbb{Z}_N; \quad f_{br}(x) := x^{e_{br}} ,$$

where e_{br} is the br -th prime of the PFG prime sequence. To prove the ABO-LTF security property we first use the Phi-hiding assumption to change the distribution of the RSA modulus N to satisfy $e^* \mid \varphi_N$, for some random prime e^* . Next, we use the PFG's programmability feature to make sure that $e_{br^*} = e^*$, meaning the function $f_{br}(\cdot)$ is injective if $br \neq br^*$ and e^* -to-1 if $br = br^*$.⁷

APPLICATIONS. Our constructions of index-independent LTFs and LTPs over domain $\{0, 1\}^k$ (and our techniques to build them) are mostly of theoretical interest with potential future applications. Whereas with our current knowledge we are not able to present a killer application, let us still discuss possible minor applications. Most importantly, correlated-product secure trapdoor functions [24] and IND-CCA secure public-key encryption [21] can be constructed from index-independent LTFs over domain $\{0, 1\}^k$. Both require the lossiness factor L to be larger than $2^{k/2}$, whereas our construction based on the Phi-hiding assumption cannot go beyond $L = 2^{k/4}$. One can still apply the amplification result by Mol and Yilek [19] to build IND-CCA secure encryption. The efficiency loss will be smaller than with the previous constructions from the Phi-hiding assumption (having lossiness factor $L \approx 2$).

2 Preliminaries

2.1 Notation

If $a, b \in \mathbb{N}, a < b$, we use notations $[a..b] = \{a, \dots, b\}$, $[b] = [1..b]$, $\llbracket a..b \rrbracket = [a..(b-1)]$, and $\llbracket b \rrbracket = [0..(b-1)]$. We say $m \in \mathbb{N}$ is an l -bit number if $m \in \llbracket 2^{l-1} .. 2^l \rrbracket$. For any set $M \subseteq \mathbb{N}$ we denote with $M_l := M \cap \llbracket 2^{l-1} .. 2^l \rrbracket$ its subset of l -bit elements. We write $\{0, 1\}^l$ for the set of strings of length l and denote the bit-wise exclusive-or operation of same-length strings with \oplus . For all $l \in \mathbb{N}$ we assume a canonic bijection $\#: \llbracket 2^l \rrbracket \rightarrow \{0, 1\}^l$ and correspondingly denote with $\#x$ the interpretation of an element x of $\llbracket 2^l \rrbracket$ as a string in $\{0, 1\}^l$. The support of a randomized algorithm A on input x , i.e., the set of values it outputs with non-zero probability, is denoted by $[A(x)]$. We annotate a disjoint union with \cup .

2.2 (All-but-one) lossy trapdoor permutations

We recall the concepts of lossy trapdoor functions and all-but-one lossy trapdoor functions as introduced by Peikert and Waters [21]. More precisely, we slightly deviate from their formalizations by restricting attention to permutations, supporting index-dependent domains [11], and considering permutations that are not perfectly correct.

⁷ In fact this requires a slightly strengthened variant of the Phi-hiding assumption where for a larger set \mathcal{E} it is known that precisely one element $e \in \mathcal{E}$ is a divisor of φ_N . We call this the *unique-divisor Phi-hiding assumption*, see Section 2.3.

Lossy trapdoor permutations. Let \mathcal{X} be a domain, \mathcal{Id} a universe of function indices, and for each index $id \in \mathcal{Id}$ let $\mathcal{X}(id) \subseteq \mathcal{X}$ be a specific (sub)domain. A *lossy trapdoor permutation* (LTP) for $\mathcal{X}, \mathcal{Id}$ then consists of a trapdoor space \mathcal{Td} and three efficient algorithms $F = (\text{FGen}, \text{FEv}, \text{FInv})$ for which the following hold: Algorithm

$$\{0, 1\} \rightarrow \text{FGen} \rightarrow_{\text{s}} \mathcal{Id} \times (\mathcal{Td} \cup \{\perp\})$$

is a randomized instance generator. Its input $b \in \{0, 1\}$ specifies whether the generated instance is *injective* ($b = 1$) or *lossy* ($b = 0$). We require $[\text{FGen}(1)] \subseteq \mathcal{Id} \times \mathcal{Td}$ and $[\text{FGen}(0)] \subseteq \mathcal{Id} \times \{\perp\}$. In injective mode, if $(id, td) \in [\text{FGen}(1)]$, we refer to td as the trapdoor corresponding to id . Algorithms

$$\mathcal{Id} \times \mathcal{X} \rightarrow \text{FEv} \rightarrow \mathcal{X} \quad \text{and} \quad \mathcal{Td} \times \mathcal{X} \rightarrow \text{FInv} \rightarrow \mathcal{X}$$

are the *evaluation* and *inversion* algorithms, respectively. We require it hold $\text{FEv}(id, x) \in \mathcal{X}(id)$ for all $id \in \mathcal{Id}$ and $x \in \mathcal{X}(id)$. For correctness we further require that in injective mode the mapping $\mathcal{X}(id) \rightarrow \mathcal{X}(id)$ induced by FEv can be effectively inverted on (almost) all values if the trapdoor is known. Formally, we say that F is $(1 - \epsilon_1)$ -correct if

$$\Pr[(id, td) \leftarrow_{\text{s}} \text{FGen}(1), x \leftarrow_{\text{s}} \mathcal{X}(id), y \leftarrow \text{FEv}(id, x) : \text{FInv}(td, y) \neq x] \leq \epsilon_1 .$$

This means that for $\epsilon_1 > 0$ the function implemented by $\text{FEv}(id, \cdot)$ might technically not be a permutation. For security we require (a) that FEv lose information in lossy mode, and (b) that injective mode and lossy mode be indistinguishable. Concerning (a), we say the LTP is *L-lossy* if for all $(id, \perp) \in [\text{FGen}(0)]$ we have $|\text{FEv}(id, \mathcal{X}(id))| \leq |\mathcal{X}(id)|/L$.⁸ Concerning (b), we say the LTP is (τ, ϵ_2) -*indistinguishable* if for all τ -time distinguishers \mathcal{D} we have

$$\left| \Pr[(id, td) \leftarrow_{\text{s}} \text{FGen}(1) : \mathcal{D}(id) \Rightarrow 1] - \Pr[(id, \perp) \leftarrow_{\text{s}} \text{FGen}(0) : \mathcal{D}(id) \Rightarrow 1] \right| \leq \epsilon_2 .$$

All-but-one lossy trapdoor permutations. All-but-one LTPs are a generalization of LTPs where in addition to the universe of function indices there is a universe of branches; function FEv is lossy for one branch and injective for all others. In particular, a (regular) LTP is equivalent to an all-but-one LTP if the branch space consists of precisely two elements.

Let \mathcal{Br} be a branch space, \mathcal{X} a domain, \mathcal{Id} a universe of function indices, and for each index $id \in \mathcal{Id}$ let $\mathcal{X}(id) \subseteq \mathcal{X}$ be a specific (sub)domain. An *all-but-one lossy trapdoor permutation* (ABO-LTP) for $\mathcal{Br}, \mathcal{X}, \mathcal{Id}$ then consists of a trapdoor space \mathcal{Td} and three efficient algorithms $A = (\text{FGen}, \text{FEv}, \text{FInv})$ for which the following hold: Algorithm

$$\mathcal{Br} \rightarrow \text{FGen} \rightarrow_{\text{s}} \mathcal{Id} \times \mathcal{Td}$$

⁸ According to our definition, *L-lossiness* indicates that the size of the lossy image is by a factor L smaller than the domain. The original definition by Peikert and Waters indicates the same quantity on a logarithmic scale, i.e., they report $\log_2(L)$ instead of L .

is an instance generator such that the invocation $(id, td) \leftarrow_{\S} \text{FGen}(br)$, for a branch br , generates a function index id with trapdoor td . Similarly as for LTPs, algorithms

$$\mathcal{B}r \times \mathcal{I}d \times \mathcal{X} \rightarrow \text{FEv} \rightarrow \mathcal{X} \quad \text{and} \quad \mathcal{B}r \times \mathcal{T}d \times \mathcal{X} \rightarrow \text{FInv} \rightarrow \mathcal{X}$$

are the evaluation and inversion algorithms. We require that for all $br, br^* \in \mathcal{B}r$ and $(id, td) \in [\text{FGen}(br^*)]$ and $x \in \mathcal{X}(id)$, if $y = \text{FEv}(br, id, x)$ then $y \in \mathcal{X}(id)$. We further require that the mappings $\mathcal{X}(id) \rightarrow \mathcal{X}(id)$ induced by FEv on all branches with exception of br^* can be effectively inverted (on almost all values) if the trapdoor is known. Formally, we say that A is $(1 - \epsilon_1)$ -correct if for all $br, br^* \in \mathcal{B}r$, $br \neq br^*$, we have

$$\Pr[(id, td) \leftarrow_{\S} \text{FGen}(br^*), x \leftarrow_{\S} \mathcal{X}(id) : \text{FInv}(br, td, \text{FEv}(br, id, x)) \neq x] \leq \epsilon_1 .$$

For security we require that FEv lose information on its *lossy branch*, i.e., the branch br^* the instance was generated for. Further, it shall be unfeasible to identify the lossy branch. Concretely, we say the ABO-LTP is *L-lossy* if for all $br^* \in \mathcal{B}r$ and $(id, td) \in [\text{FGen}(br^*)]$ we have $|\text{FEv}(br^*, id, \mathcal{X}(id))| \leq |\mathcal{X}(id)|/L$, and we say it is (τ, ϵ_2) -*indistinguishable* if for all $br_0, br_1 \in \mathcal{B}r$ and all τ -time distinguishers \mathcal{D} (that may depend on br_0, br_1) we have

$$\left| \frac{\Pr[(id, td) \leftarrow_{\S} \text{FGen}(br_0) : \mathcal{D}(id) \Rightarrow 1]}{\Pr[(id, td) \leftarrow_{\S} \text{FGen}(br_1) : \mathcal{D}(id) \Rightarrow 1]} \right| \leq \epsilon_2 .$$

Index-dependent vs. index-independent LTPs/ABO-LTPs. In the above definition of LTPs, the domain $\mathcal{X}(id) \subseteq \mathcal{X}$ on which $\text{FEv}(id, \cdot)$ operates may depend on function index id . We say the LTP is index-independent if this restriction does not exist, i.e., if $\mathcal{X}(id) = \mathcal{X}$ for all id . For ABO-LTPs we say correspondingly. In later sections we show how to generically transform an index-dependent trapdoor permutation into an index-independent one.

2.3 Number theoretic assumptions

For $a, b \in \mathbb{N}, a \neq 0$, we write $a \mid b$ if a divides b , i.e., if there exists $d \in \mathbb{N}$ s.t. $b = da$. We further write $a \mid_1 b$ if a divides b exactly once, i.e., if $a \mid b \wedge a^2 \nmid b$. The greatest common divisor of a, b is denoted $\text{gcd}(a, b)$. We denote the set of prime numbers with \mathcal{P} . Recall from Section 2.1 that \mathbb{N}_l and \mathcal{P}_l denote the sets of l -bit natural and prime numbers, respectively.

If k is an even number, a product $N = pq$ is a *k-bit RSA modulus* if $N \in \mathbb{N}_k$, $p, q \in \mathcal{P}_{k/2}$, and $p \neq q$. The order of the multiplicative group \mathbb{Z}_N^* is $\varphi_N := \varphi(N) = (p-1)(q-1)$. We denote the space of k -bit RSA moduli with \mathcal{RSA}_k . If we want to restrict attention to k -bit RSA moduli that fulfill a specific condition C , we write $\mathcal{RSA}_k[C]$. The set of k -bit Blum integers, i.e., RSA moduli where the prime factors satisfy $p \equiv q \equiv 3 \pmod{4}$, is denoted by $\mathcal{BRSA}_k := \mathcal{RSA}_k[p \equiv q \equiv 3 \pmod{4}]$.

Phi-hiding assumption. In standard RSA encryption, public exponent e is chosen constraint to $e \nmid \varphi_N$ so that the mapping $x \mapsto x^e$ is a bijection. Some applications in addition use exponents $e \mid \varphi_N$ and require that it be hard, given (N, e) , to decide whether $e \mid \varphi_N$ or $e \nmid \varphi_N$. Roughly, the Phi-hiding assumption [9,17] for a set of primes \mathcal{E} says that $N \in \mathcal{RSA}_k$ can be generated such that for uniformly picked $e \in \mathcal{E}$ the cases $N \in \mathcal{RSA}_k[e \nmid \varphi_N]$ and $N \in \mathcal{RSA}_k[e \mid \varphi_N]$ are computationally indistinguishable. Formally, we say that the (τ, ϵ) -Phi-hiding assumption holds for (k, \mathcal{E}) if for all τ -time adversaries \mathcal{D} we have

$$\left| \Pr[e \leftarrow_{\S} \mathcal{E}; (N, \varphi_N) \leftarrow_{\S} \mathcal{RSA}_k[e \mid \varphi_N] : \mathcal{D}(N, e) \Rightarrow 1] - \Pr[e \leftarrow_{\S} \mathcal{E}; (N, \varphi_N) \leftarrow_{\S} \mathcal{RSA}_k[e \nmid \varphi_N] : \mathcal{D}(N, e) \Rightarrow 1] \right| \leq \epsilon .$$

In the probability expressions we write $(N, \varphi_N) \leftarrow_{\S} \mathcal{RSA}_k[C]$ for an algorithm that generates a k -bit RSA modulus satisfying condition C , and also outputs $\varphi_N = |\mathbb{Z}_N^*|$.

In this paper we also need a variant of this assumption: An added restriction is that precisely one $e \in \mathcal{E}$ shall be a divisor of φ_N , and, as before, if e divides φ_N then at most once.⁹ This is expressed by condition

$$C(\mathcal{E}, \varphi_N, e) \quad :\iff \quad e \mid \varphi_N \wedge \gcd(\mathcal{E}, \varphi_N/e) = 1 ,$$

where the gcd term encodes that φ_N/e is relative prime to all elements of \mathcal{E} ; this in particular implies $e \mid \varphi_N$. We say the *unique-divisor* (τ, ϵ) -Phi-hiding assumption holds for (k, \mathcal{E}) if for all τ -time adversaries \mathcal{D} we have

$$\left| \Pr[e_0 \leftarrow_{\S} \mathcal{E}; (N, \varphi_N) \leftarrow_{\S} \mathcal{RSA}_k[C(\mathcal{E}, \varphi_N, e_0)] : \mathcal{D}(N, e_0) \Rightarrow 1] - \Pr[e_0, e_1 \leftarrow_{\S} \mathcal{E}; (N, \varphi_N) \leftarrow_{\S} \mathcal{RSA}_k[C(\mathcal{E}, \varphi_N, e_0)] : \mathcal{D}(N, e_1) \Rightarrow 1] \right| \leq \epsilon .$$

Quadratic residuosity assumption. Roughly, the quadratic residuosity assumption says that it is hard to distinguish quadratic residues modulo a Blum integer from quadratic non-residues that have positive Jacobi symbol.

Formally, for all $N \in \mathbb{N}$ denote with $\mathcal{QR}_N \subseteq \mathbb{Z}_N^*$ the set of quadratic residues modulo N and with $\mathcal{J}_N \subseteq \mathbb{Z}_N^*$ the set of numbers with positive Jacobi symbol. (In particular we have $\mathcal{QR}_N \subseteq \mathcal{J}_N$.) We say that the (τ, ϵ) -quadratic residuosity assumption holds for k if for all τ -time adversaries \mathcal{D} we have

$$\left| \Pr[(N, p, q) \leftarrow_{\S} \mathcal{BRSA}_k, x \leftarrow_{\S} \mathcal{QR}_N : \mathcal{D}(N, x) \Rightarrow 1] - \Pr[(N, p, q) \leftarrow_{\S} \mathcal{BRSA}_k, x \leftarrow_{\S} \mathcal{J}_N \setminus \mathcal{QR}_N : \mathcal{D}(N, x) \Rightarrow 1] \right| \leq \epsilon .$$

In the probability expressions we write $(N, p, q) \leftarrow_{\S} \mathcal{BRSA}_k$ for an algorithm that generates a k -bit Blum integer and also outputs its prime factors. Note that sampling elements of \mathcal{QR}_N and $\mathcal{J}_N \setminus \mathcal{QR}_N$ can be done efficiently if these factors are known.

⁹ While this assumption is stronger than the standard Phi-hiding assumption, we conjecture that it is rather mild (possibly in the same way as the strengthened Quadratic Residuosity assumption from [15] that is specialized towards defining the 2^k -th Power Residue symbol).

$\text{FGen}_{\text{ii}}(b)$ 00 $(id, td) \leftarrow_{\mathcal{S}} \text{FGen}(b)$ 01 Return (id, td)	$\text{FEv}_{\text{ii}}(id, x)$ 02 If $x \in \mathcal{X}(id)$: 03 $x \leftarrow \text{FEv}(id, x)$ 04 $y \leftarrow \pi_{id}(x)$ 05 If $y \in \mathcal{X}(id)$: 06 $y \leftarrow \text{FEv}(id, y)$ 07 Return y	$\text{FInv}_{\text{ii}}(td, y)$ 08 If $y \in \mathcal{X}(id)$: 09 $y \leftarrow \text{FInv}(td, y)$ 10 $x \leftarrow \pi_{id}^{-1}(y)$ 11 If $x \in \mathcal{X}(id)$: 12 $x \leftarrow \text{FInv}(td, x)$ 13 Return x
---	---	---

Fig. 1. Transformation of index-dependent LTP into index-independent LTP. To make algorithm FInv_{ii} well-defined we assume implicitly that trapdoor td contains a copy of function index id . A visualization of the construction is in Figure 2.

3 From index-dependence to index-independence

Many natural constructions of lossy trapdoor permutations are index-dependent, i.e., for each index id the function $\text{FEv}(id, \cdot)$ operates on an individual set $\mathcal{X}(id) \subseteq \mathcal{X}$. However, for applications it might be necessary that there is only one domain: $\mathcal{X}(id) = \mathcal{X}$ for all id . In this section we convert index-dependent LTPs into index-independent LTPs. Some transforms of this type have been proposed before. For instance, [11] implicitly uses the somewhat trivial approach of leaving elements in $\mathcal{X} \setminus \mathcal{X}(id)$ untouched (i.e., elements in $\mathcal{X}(id)$ are processed with the LTP, the others are passed through without modification). As discussed in the introduction, the performance of this conversion is generally rather poor: In the worst case, if $|\mathcal{X}(id)| \ll |\mathcal{X}|$, lossiness is bounded by $L = 1$.

Below we study a two-round construction that was first proposed in [13], in a different context. There, the goal was to extend the domain of the RSA trapdoor permutation; aspects of lossiness were not studied. Further, our exposition is more generic. The idea behind the transformation is to ensure that FEv is applied to every point of \mathcal{X} at least once. In both rounds the points of $\mathcal{X}(id)$ are permuted with FEv while the remaining points of \mathcal{X} stay unchanged. To achieve the property stated above, after the first round a permutation π_{id} is used to move into $\mathcal{X}(id)$ all those points that have not yet been touched by FEv .

Let $F = (\text{FGen}, \text{FEv}, \text{FInv})$ be a LTP with domain \mathcal{X} and index space \mathcal{Id} . Assume F has index-dependent domains. For all $id \in \mathcal{Id}$ write $\overline{\mathcal{X}(id)} = \mathcal{X} \setminus \mathcal{X}(id)$ and let $\pi_{id}: \mathcal{X} \rightarrow \mathcal{X}$ be an efficiently computable and efficiently invertible permutation satisfying $\pi_{id}(\overline{\mathcal{X}(id)}) \subseteq \mathcal{X}(id)$ or, equivalently, $\pi_{id}^{-1}(\overline{\mathcal{X}(id)}) \subseteq \mathcal{X}(id)$. (Note that such a π_{id} can exist only if $|\mathcal{X}(id)| \geq |\mathcal{X}|/2$ for all id .) From F and $(\pi_{id})_{id \in \mathcal{Id}}$ we construct a LTP $F_{\text{ii}} = (\text{FGen}_{\text{ii}}, \text{FEv}_{\text{ii}}, \text{FInv}_{\text{ii}})$ with index-independent domain \mathcal{X} , i.e., $\mathcal{X}(id) = \mathcal{X}$ for all id . The algorithms are specified in Fig. 1 and illustrated in Fig. 2. The analysis is in Lemma 1 (which is proved in the full version [1]).

Lemma 1. *Let F be a $(1 - \epsilon_1)$ -correct, (τ, ϵ_2) -indistinguishable L -lossy trapdoor permutation with index-dependent domain. Furthermore, let $(\pi_{id})_{id \in \mathcal{Id}}$ be a family of permutations on \mathcal{X} as described. Then F_{ii} is an $(1 - 2\epsilon_1)$ -correct, (τ, ϵ_2) -indistinguishable $L/2$ -lossy trapdoor permutation with index-independent domain \mathcal{X} . In particular, if F is 1-correct, then so is F_{ii} .*

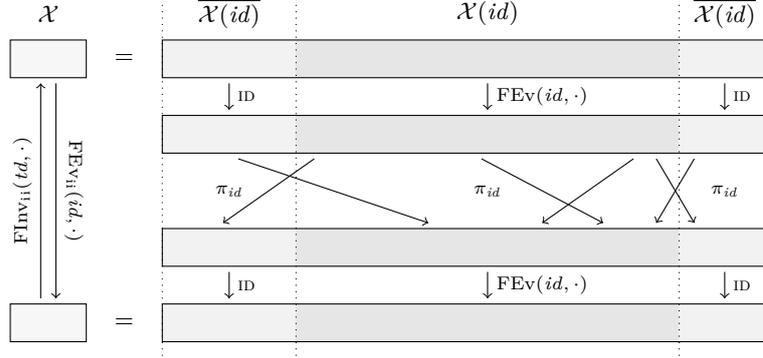


Fig. 2. Working principle of transformation of index-dependent LTP into index-independent LTP. The corresponding algorithms are in Figure 1. Note that π_{id} is chosen such that every point in \mathcal{X} is permuted by FEv at least once.

Analogously to the construction in Fig. 1 we can transform an index-dependent ABO-LTP $A = (\text{FGen}, \text{FEv}, \text{FInv})$ into an index-independent ABO-LTP $A_{ii} = (\text{FGen}, \text{FEv}_{ii}, \text{FInv}_{ii})$. Note that A_{ii} uses the same instance generator as A . Algorithms FEv_{ii} and FInv_{ii} work as their counterparts for LTPs defined in Fig. 1, the only difference being the use of the additional input br to evaluate FEv and FInv. We obtain the following.

Lemma 2. *Let A be a $(1 - \epsilon_1)$ -correct, (τ, ϵ_2) -indistinguishable L -lossy ABO-LTP with index-dependent domain. Let $(\pi_{id})_{id \in \mathcal{I}d}$ be a family of permutations on \mathcal{X} as described. Then A_{ii} is an $(1 - 2\epsilon_1)$ -correct, (τ, ϵ_2) -indistinguishable $L/2$ -lossy ABO-LTP with index-independent domain \mathcal{X} . In particular, if A is 1-correct, then so is A_{ii} .*

4 Lossy trapdoor permutations from Phi-hiding

Fix an RSA modulus N and let $e \ll \varphi_N$ be prime. We say e is *injective* for N if $e \nmid \varphi_N$ and that it is *lossy* for N if $e \mid \varphi_N$. In the injective case the mapping $E: \mathbb{Z}_N \rightarrow \mathbb{Z}_N; x \mapsto x^e$ is inverted by $D: y \mapsto y^d$, where d is such that $ed = 1 \pmod{\varphi_N}$. In the lossy case, the restriction $E|_{\mathbb{Z}_N^*}$ of E to domain \mathbb{Z}_N^* is e -to-1, i.e., we have $|E(\mathbb{Z}_N^*)|/|\mathbb{Z}_N^*| = 1/e$. The Phi-hiding assumption from Section 2.3 then precisely says that it is hard to decide whether a candidate exponent e is injective or lossy for N .

We propose two LTPs in the RSA setting, both with security based on the Phi-hiding assumption. The first construction is quite natural but has index-dependent domains. The second construction is the index-independent analogue of the first, obtained via the transformation from Section 3. Here, our contribution is establishing a better bound on the lossiness than is possible with the generic result. (Our arguments are based on structures specific to the RSA setting.)

FGen(b) 00 $e \leftarrow_{\mathcal{E}}$ 01 If $b = 0$: (lossy mode) 02 $(N, \varphi_N) \leftarrow_{\mathcal{S}} \mathcal{RSA}_k[e \mid_1 \varphi_N]$ 03 $id \leftarrow (N, e); td \leftarrow \perp$ 04 If $b = 1$: (injective mode) 05 $(N, \varphi_N) \leftarrow_{\mathcal{S}} \mathcal{RSA}_k[e \uparrow \varphi_N]$ 06 $d \leftarrow e^{-1} \bmod \varphi_N$ 07 $id \leftarrow (N, e); td \leftarrow (N, d)$ 08 Return (id, td)	FEv(id, x) 09 $(N, e) \leftarrow id$ 10 $y \leftarrow x^e \bmod N$ 11 Return y FInv(td, y) 12 $(N, d) \leftarrow td$ 13 $x \leftarrow y^d \bmod N$ 14 Return x	FEv*(id, x) 15 $(N, e) \leftarrow id$ 16 If $x = 0$: Return 0 17 If $x \notin \mathbb{Z}_N^*$: 18 $p \leftarrow \gcd(x, N)$ 19 $q \leftarrow N/p$ 20 $\varphi_N \leftarrow (p-1)(q-1)$ 21 If $e \mid \varphi_N$: Return 0 22 $y \leftarrow x^e \bmod N$ 23 Return y
---	---	---

Fig. 3. LTPs F and F* from Phi-hiding assumption (with index-dependent domains).

4.1 Index-dependent domain LTP from Phi-hiding assumption

Let k be an even number indicating a desired bit length of RSA moduli. Let \mathcal{E} be a distribution of prime numbers such that the (τ, ϵ) -Phi-hiding assumption holds for (k, \mathcal{E}) . Consider the constructions of LTPs $F = (\text{FGen}, \text{FEv}, \text{FInv})$ and $F^* = (\text{FGen}, \text{FEv}^*, \text{FInv})$ given by the algorithms in Fig. 3. Observe that condition $e \mid_1 \varphi_N$ in line 02 implies that no element of \mathcal{E} can be longer than $k/2$ bits. Further, to protect from known attacks it is necessary that $\max \mathcal{E} \leq 2^{k/4}$.

The working principle of F is as follows: Function indices id correspond with RSA parameters (N, e) . The domain corresponding to index id is $\mathcal{X}(id) = \mathbb{Z}_N$. In injective mode, (N, e) are chosen such that e is invertible modulo φ_N , i.e., such that a corresponding decryption exponent d exists. The FEv and FInv algorithms, in this case, are the standard RSA mappings $x \mapsto x^e$ and $y \mapsto y^d$ (lines 10 and 13). In lossy mode, e is a divisor of φ_N . In this case, mapping $x \mapsto x^e$ is e -to-1 for elements in \mathbb{Z}_N^* . The resulting overall lossiness (i.e., for full \mathbb{Z}_N) is analyzed in Lemma 3.

We next discuss F*. This variant achieves better lossiness by building on the fact that given an element of $\mathbb{Z}_N \setminus \mathbb{Z}_N^*$ it is possible to effectively determine whether the function index (N, e) is injective or lossy. In the first case FEv* uses the standard RSA map; in the second case elements in $\mathbb{Z}_N \setminus \mathbb{Z}_N^*$ are detected and explicitly mapped to 0. The identification of lossy indices and elements in $\mathbb{Z}_N \setminus \mathbb{Z}_N^*$ is handled in lines 17–21. Observe that the condition in line 17 can be checked efficiently.

We analyze constructions F and F* in Lemma 3 (the proof of which is in the full version [1]). While the second LTP is more complicated to implement, the achieved lossiness bound is easier to work with.

Lemma 3. *If for (k, \mathcal{E}) the (τ, ϵ) -Phi-hiding assumption holds and $L \leq \min \mathcal{E}$ is a lower bound on the elements in the support of \mathcal{E} , LTP F is a 1-correct, (τ, ϵ) -indistinguishable $(1/L + 2^{-k/2+3})^{-1}$ -lossy trapdoor function. Furthermore, LTP F* is a 1-correct (τ, ϵ) -indistinguishable L -lossy trapdoor function. Both LTPs have index-dependent domain.*

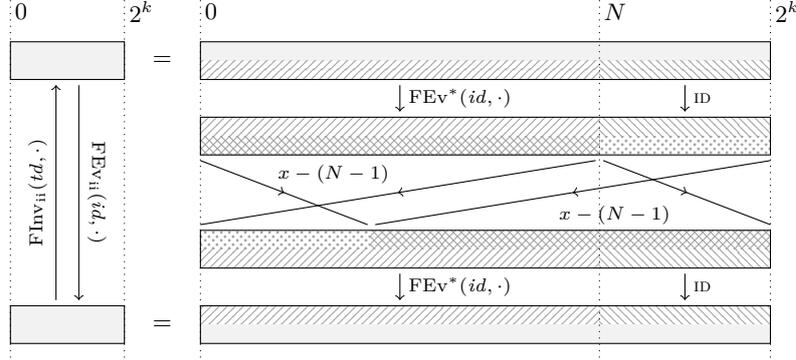


Fig. 4. Illustration of Phi-hiding based LTP F_{ii} with index-independent domain.

4.2 Index-independent domain LTP from Phi-hiding assumption

The LTP F^* from Section 4.1 has index-dependent domains: for function index $id = (N, e)$, algorithm $FEV^*(id, \cdot)$ operates on domain $\mathcal{X}(id) = \mathbb{Z}_N$. By construction we have $N \in [2^{k-1} .. 2^k]$ and thus $\mathcal{X}(id) \subseteq \mathcal{X}$ for $\mathcal{X} = [2^k]$. To obtain an LTP F_{ii} with index-independent domain $[2^k]$ we can apply to F^* the generic transform of Section 3. By Lemma 1, assuming appropriately chosen permutations (π_{id}) , if F^* is L -lossy, then F_{ii} is $L/2$ -lossy. The contribution of the current section is to show that for a specifically defined family (π_{id}) using direct (non-generic) arguments this result can be strengthened: If F^* is L -lossy, then also F_{ii} is L -lossy. In other words, there is no price to pay for switching from index-dependent domains to index-independent domains. (This holds for the lossiness; computation time might double.)

As a first step we identify a family (π_{id}) of permutations on \mathcal{X} that suits the conditions of the transform from Section 3, namely $\pi_{id}(\mathcal{X}(id)) \subseteq \mathcal{X}(id)$ for all $id \in \mathcal{Id}$. Hence let $\mathcal{X} = [2^k]$ and $(N, e) = id \in \mathcal{Id}$, where $N \in [2^{k-1} .. 2^k]$. We define $\pi_{id}: \mathcal{X} \rightarrow \mathcal{X}; x \mapsto x - (N-1) \bmod 2^k$. Then π_{id} is a permutation on \mathcal{X} and we have $N \leq x < 2^k \Rightarrow 1 \leq \pi_{id}(x) \leq 2^k - N < N$ (the last inequality follows from $2^{k-1} \leq N < 2^k$); this establishes $\pi_{id}(\mathcal{X}(id)) \subseteq \mathcal{X}(id)$. We illustrate the transform from Fig. 2 in conjunction with this family of bijections (π_{id}) in Fig. 4. In the following, we first state the generic result obtained by applying Lemma 1 to this setup. We then give the one established directly. The proof of Lemma 4 is in the full version [1].

Corollary 1. *Let \mathcal{E} be a prime distribution and $L \leq \min \mathcal{E}$. Further, let $F^* = (FGen, FEV^*, FInv)$ be the L -lossy LTP defined in Fig. 3, $(\pi_{id})_{id \in \mathcal{Id}}$ the permutation family defined above, and F_{ii} the conversion of F^* via Fig. 1. If for (k, \mathcal{E}) the (τ, ϵ) -Phi-hiding assumption holds, then F_{ii} is a 1-correct, (τ, ϵ) -indistinguishable $L/2$ -lossy trapdoor function with index-independent domain $\mathcal{X} = [2^k]$.*

Lemma 4. *Let \mathcal{E} be a prime distribution and $L \leq \min \mathcal{E}$. Further, let $F^* = (\text{FGen}, \text{FEv}^*, \text{FInv})$ be the L -lossy LTF defined in Fig. 3, $(\pi_{id})_{id \in \mathcal{I}d}$ the permutation family defined above, and F_{ii} the conversion of F^* via Fig. 1. If for (k, \mathcal{E}) the (τ, ϵ) -Phi-hiding assumption holds, then F_{ii} is a 1-correct, (τ, ϵ) -indistinguishable $L/(1 + 2^{-k/2})$ -lossy trapdoor function with index-independent domain $\mathcal{X} = \llbracket 2^k \rrbracket$.*

5 Lossy trapdoor permutations from Quadratic residuosity assumption

In this section we recall the index-dependent lossy trapdoor function F of [10] based on the quadratic residuosity assumption and show how the transform of Section 3 can be used to obtain an index-independent variant F_{ii} . Since F has a lossiness factor of 2, using the generic bound is of no use in this case. However, by exploiting the algebraic structure of the construction we are able to establish that F_{ii} has essentially the same lossiness factor as F . This improves on the index-independent variant given in [10], which achieves a lossiness factor of $4/3$.

5.1 Index-dependent domain LTP from Quadratic residuosity

Let p, q be primes of bit length $k/2$ satisfying $p \equiv 3 \pmod{4}$ and $q \equiv 3 \pmod{4}$. Consider the functions $j_N : \mathbb{Z} \rightarrow \{0, 1\}$ and $h_N : \mathbb{Z} \rightarrow \{0, 1\}$ defined by

$$j_N(x) = \begin{cases} 0, & \text{if } x \in \mathcal{J}_N \cup (\mathbb{Z}_N \setminus \mathbb{Z}_N^*) \\ 1, & \text{if } x \in \mathbb{Z}_N^* \setminus \mathcal{J}_N \end{cases}$$

$$h_N(x) = \begin{cases} 0, & \text{if } x \leq N/2 \\ 1, & \text{if } x > N/2 \end{cases}.$$

Note that both j_N and h_N can be efficiently computed given N . Let $d_j, d_h \in \{0, 1\}$. Then —as pointed out in [10]— for each $y \in \mathcal{QR}_N$ exactly one of the four solutions of the equation $x^2 = y \pmod{N}$ satisfies $j_N(x) = d_j$ and $h_N(x) = d_h$. We denote this square root of y by R_{d_j, d_h} . Furthermore for every $y \in \mathbb{Z}_N \setminus \mathbb{Z}_N^*$ with $y \in \mathcal{QR}_p \vee y \in \mathcal{QR}_q$ the equation $y = x^2 \pmod{N}$ has exactly two solutions —one being the negative of the other. Hence both solutions satisfy $j_N(x) = 0$ and for $d_h \in \{0, 1\}$ exactly one of the solutions satisfies $h_N(x) = d_h$. Analogous to the situation above we denote this solution by R_{0, d_h} . In [10] the authors construct a lossy trapdoor permutation with index-dependent domain \mathbb{Z}_N . The LTP's algorithms are depicted in Fig. 5. The idea of the construction is to map elements $x \in \mathbb{Z}_N$ to x^2 , which is afterwards multiplied by some appropriately chosen group elements, which allow to reconstruct x^2 as well as both $j_N(x) =: d_j$ and $h_N(x) =: d_h$. Then the LTF can be inverted by computing R_{d_j, d_h} .

Lemma 5 ([10]). *Let $F = (\text{FGen}, \text{FEv}, \text{FInv})$ the LTP of Fig. 5. If the (τ, ϵ) -Quadratic residuosity assumption holds for k , then F is an (τ', ϵ) -indistinguishable 2-lossy trapdoor function with index-dependent domain $\mathcal{X}((N, r, s)) = \mathbb{Z}_N \subseteq \llbracket 2^k \rrbracket = \mathcal{X}$, where $\tau' \approx \tau$.*

<pre> FGen(b) 00 (N, p, q) $\leftarrow_{\S} \mathcal{BRSA}_k$ 01 $r \leftarrow_{\S} \mathbb{Z}_N^* \setminus \mathcal{J}_N$ 02 If $b = 0$: (lossy mode) 03 $s \leftarrow_{\S} \mathcal{QR}_N$ 04 $id \leftarrow (N, r, s)$; $td \leftarrow \perp$ 05 If $b = 1$: (injective mode) 06 $s \leftarrow_{\S} \mathcal{J}_N \setminus \mathcal{QR}_N$ 07 $id \leftarrow (N, r, s)$ 08 $td \leftarrow (N, p, q, r, s)$ 09 Return (id, td) </pre>	<pre> FEv(id, x) 10 (N, r, s) $\leftarrow id$ 11 $d_j \leftarrow j_N(x)$ 12 $d_h \leftarrow h_N(x)$ 13 $y \leftarrow x^2 r^t s^\tau$ 14 Return y </pre>	<pre> FInv(td, y) 15 If $y = 0$: return 0 16 (N, p, q, r, s) $\leftarrow td$ 17 $d_j \leftarrow j_N(y)$ 18 $y' \leftarrow yr^{-d_j}$ 19 If $y' \in \mathbb{Z}_N \setminus \mathbb{Z}_N^*$: 20 If $y' \notin \mathcal{QR}_p \wedge y' \notin \mathcal{QR}_q$: 21 $d_h \leftarrow 1$ 22 Else: $d_h \leftarrow 0$ 23 Elseif $y' \in \mathcal{QR}_N$: $d_h \leftarrow 0$ 24 Else: $d_h \leftarrow 1$ 25 $y'' \leftarrow y' s^{-d_h}$ 26 $x \leftarrow R_{d_j, d_h}(y'')$ 27 Return x </pre>
--	--	---

Fig. 5. LTP F from Quadratic residuosity assumption (with index-dependent domains).

5.2 Index-independent domain LTP from Quadratic residuosity

In [10] the authors propose to modify the LTP F of Section 5.1 in the following way to obtain an LTP F_{ii}^* with index-independent domain $\llbracket 2^k \rrbracket$. This is done by letting $F_{ii}^*|_{\mathcal{X}(id)}(id, \cdot) := F(id, \cdot)$ and $F_{ii}^*|_{\overline{\mathcal{X}(id)}}(id, \cdot) := \text{ID}$. The resulting LTP F_{ii}^* is a 4/3-lossy trapdoor function. In this section we show that using our transformation of Section 3 with an appropriate permutation yields an index-independent LTP based on the quadratic residuosity assumption having essentially the same lossiness factor as the underlying LTP F.

To be able to use our transformation we need a family of permutations on \mathcal{X} that suits the conditions of Section 3. Since—as in the construction based on the Phi-hiding assumption—the index-dependent domain $\mathcal{X}(id)$ for some $id = (N, r, s)$ is \mathbb{Z}_N , we are able to use the same family of permutations. Hence for $id = (N, r, s) \in \mathcal{Id}$ define $\pi_{id}: \mathcal{X} \rightarrow \mathcal{X}; x \mapsto x - (N - 1) \bmod 2^k$. Then π_{id} is a permutation on \mathcal{X} and as in Section 4.2 we obtain $\pi_{id}(\mathcal{X}(id)) \subseteq \mathcal{X}(id)$.

Note that applying Lemma 1 would only yield a bound of $2/2 = 1$ on the lossiness factor of the transformed LTP, which is of no use. However, we are able to establish a desirable result directly using techniques similar to the ones used in the proof of Lemma 4. The proof of Lemma 6 is in the full version [1].

Lemma 6. *Let $F = (\text{FGen}, \text{FEv}, \text{FInv})$ be the 1-correct, 2-lossy LTP defined in Fig. 5, $(\pi_{id})_{id \in \mathcal{Id}}$ the permutation family defined above, and F_{ii} the transformation of F via Fig. 1. If the (τ, ϵ) -Quadratic residuosity assumption holds for k , F_{ii} is a 1-correct (τ', ϵ) -indistinguishable $2/(1 + 2^{-k/2})$ -lossy trapdoor function with index-independent domain $\mathcal{X} = \llbracket 2^k \rrbracket$, where $\tau' \approx \tau$.*

6 Prime family generators

In Section 7 we construct all-but-one lossy trapdoor permutations from the unique divisor Phi-hiding assumption. As a building block we use prime family

generators, a tool that deterministically derives prime numbers from a randomly picked seed. While this concept already appeared in [9], we need a variant of the tool with different functionality and security properties. Below, we first define syntax and functionality of prime family generators, and then give a construction based on polynomial evaluation.

Let $\mathcal{Q} \subseteq \mathcal{P}$ be a finite set of prime numbers and let $L \leq |\mathcal{Q}|$. For (\mathcal{Q}, L) , any instance of a *prime family generator* (PFG) indicates a sequence of distinct primes $q_1, \dots, q_L \in \mathcal{Q}$. A specific programmability feature allows for embedding any given prime at any given position. Formally, an (ϵ_1, ϵ_2) -PFG for (\mathcal{Q}, L) consists of a seed space $\mathcal{S}d$ and three algorithms PGen, PGet, PProg such that

$$\text{PGen} \rightarrow_{\mathfrak{s}} \mathcal{S}d \quad \text{and} \quad \mathcal{S}d \times [L] \rightarrow \text{PGet} \rightarrow \mathcal{Q} \quad \text{and} \quad [L] \times \mathcal{Q} \rightarrow \text{PProg} \rightarrow_{\mathfrak{s}} \mathcal{S}d .$$

For functionality we demand (a) programmability: for all $i \in [L]$ we require

$$\Pr[q \leftarrow_{\mathfrak{s}} \mathcal{Q}; sd \leftarrow_{\mathfrak{s}} \text{PProg}(i, q) : \text{PGet}(sd, i) \neq q] \leq \epsilon_1 .$$

(b) distinctness of outputs: for all $i \in [L]$ we require

$$\Pr[sd \leftarrow_{\mathfrak{s}} \text{PGen} : \exists j \in [L], i \neq j : \text{PGet}(sd, i) = \text{PGet}(sd, j)] \leq \epsilon_2 .$$

For security we require perfectly indistinguishable programmability: We demand that for all $i \in [L]$ and every distinguisher \mathcal{D} (running in arbitrary time) we have

$$\left| \Pr[sd \leftarrow_{\mathfrak{s}} \text{PGen} : \mathcal{D}(sd) \Rightarrow 1] - \Pr[q \leftarrow_{\mathfrak{s}} \mathcal{Q}; sd \leftarrow_{\mathfrak{s}} \text{PProg}(i, q) : \mathcal{D}(sd) \Rightarrow 1] \right| = 0 .$$

6.1 Construction based on polynomial evaluation

The PFG we construct here outputs (l -bit) primes from $\mathcal{Q} = \mathcal{P}_l$. While the construction is similar to one by [9], their PFG would also output primes shorter than l bits. Further, our analysis of probabilities is different, for being tailored towards our application: the construction of ABO-LTPs.

Concretely, for a set of chosen parameters $l, n, d, \lambda \in \mathbb{N}$ we construct a $(2^{-(\lambda+1)}, 2^{-\lambda})$ -PFG for $\mathcal{Q} = \mathcal{P}_l$ and $L = 2^n$. The construction is based on a family $\{F_{sd}\}$ of d -wise independent hash functions and, roughly, works as follows (see Fig. 6). The PFG's seed space $\mathcal{S}d$ is equal to $\{F_{sd}\}$'s key space. For $sd \in \mathcal{S}d$ and $i \in [2^n]$, natural numbers are generated by evaluating F_{sd} at up to $d/2$ distinct points. $\text{PGet}(sd, i)$'s output is the first prime found. Since numbers of bit length l are tested for primality, the prime number theorem guarantees that PGen will succeed in finding a prime on average after roughly l attempts. Furthermore, if d is chosen large enough finding a prime in this way will succeed except with some negligible error probability. Concretely, we instantiate $\{F_{sd}\}$ with polynomial evaluation of degree d over the field $\text{GF}(2^{l-1})$. Programming a prime q into a particular point i is done by sampling a sequence of d -many values a_j in the image of F_{sd} . Then —if existent— the first prime in this sequence is replaced by q . By polynomial interpolation it is possible to find a seed

<pre> PGen 00 $sd \leftarrow_{\mathcal{S}} \{0, 1\}^{d(l-1)}$ 01 Return sd PGet(sd, i) 02 For $j \leftarrow 1$ to $d/2$: 03 $q \leftarrow F_{sd}(\#i\ \#j)$ 04 If $q + 2^{l-1} \in \mathcal{P}_l$: 05 Return $q + 2^{l-1}$ 06 Return \perp </pre>	<pre> PProg(i, q) 07 $(a_1, \dots, a_d) \leftarrow_{\mathcal{S}} \llbracket 2^{l-1} \dots 2^l \rrbracket^d$ 08 Find smallest j with $a_j \in \mathcal{P}_l$ 09 $a_j \leftarrow q$ 10 $sd \leftarrow \text{FindC}(i, a_1 - 2^{l-1}, \dots, a_d - 2^{l-1})$ 11 Return sd </pre>
---	--

Fig. 6. PFG based on polynomial evaluation

sd such that F_{sd} evaluated at the j 'th point equals a_j . The technical challenge is to prove that if q was a uniformly distributed prime then the resulting seed sd has the correct distribution and, furthermore, with high probability satisfies $q = \text{PGet}(sd, i)$.

We now specify the construction in detail. We start by imposing necessary restrictions on its parameters. Let $l, n, d, \lambda \in \mathbb{N}$ with d even and $l \geq 25$ such that

$$n \leq l - \lambda - \log_2(l) - 2 \quad (3)$$

$$2l(\lambda + 1)/\log_2(e) \leq d < 2^{l-1-n} \quad (4)$$

where e is Euler's number. The first inequality ensures the probability of two primes sampled uniformly from \mathcal{P}_l colliding is small, the second inequality makes sure d on one hand is large enough that PGet finds a prime with high probability and on the other hand small enough, that numbers smaller than d can be encoded with few bits. Note that for $l = \mathcal{O}(\lambda)$ and $n = l/2$ equation (3) will typically be fulfilled and results in d being of order $\mathcal{O}(\lambda^2)$. The family of hash functions used in our construction is defined as follows. For $sd \in \text{GF}(2^{l-1})^d$ let

$$F_{sd}: \{0, 1\}^{l-1} \rightarrow \llbracket 2^{l-1} \rrbracket; x \mapsto \sum_{k=0}^{d-1} sd_k x^k .$$

Here x is interpreted as element of $\text{GF}(2^{l-1})$. Note that the function family $(F_{sd})_{sd \in \mathcal{S}d}$ is a d -wise independent hash function [25]. Finally, we define an algorithm FindC as follows. FindC receives as input a tuple (i, a_1, \dots, a_d) , where $i \in \llbracket 2^n \rrbracket$ and $a_1, \dots, a_d \in \llbracket 2^{l-1} \rrbracket$. It then uses Lagrange interpolation to find $sd_0, \dots, sd_{d-1} \in \text{GF}(2^{l-1})$ such that $F_{sd}(\#i\|\#j) = a_j$ for all $j \in [d]$, where $sd := (sd_0, \dots, sd_{d-1})$ (see Section 2.1 for the $\#$ notation). Here we assume $\#j \in \{0, 1\}^{l-1-n}$, which is possible since by equation 4 we have $j \leq d < 2^{l-1-n}$. FindC's output is sd . Note that for every $i \in \llbracket 2^n \rrbracket$ the function implemented by $\text{FindC}(i, \cdot)$ is a bijection between $\llbracket 2^{l-1} \rrbracket^d$ and $\mathcal{S}d$. The description of the PFG P may be found in Fig. 6.

Note that in the definition we formally do not allow PGet to return elements that are not in \mathcal{Q} . However, P returns \perp if after d tests no prime has been found. This issue could be solved by letting PGet return some fixed prime $q \in \mathcal{Q}$ in this case. We obtain the following result (the proof of which is in the full version [1]).

<pre> FGen(br^*) 00 $sd \leftarrow_{\mathcal{S}} \text{PGen}$ 01 $e^* \leftarrow_{\mathcal{S}} \text{PGet}(sd, br^*)$ 02 $(N, \varphi_N) \leftarrow_{\mathcal{S}} \mathcal{RSA}_k[C(e^*, \varphi_N)]$ 03 $id \leftarrow (N, sd)$ 04 $td \leftarrow (N, sd, \varphi_N)$ 05 Return (id, td) </pre>	<pre> FEv(br, id, x) 06 If $x = 0$: Return 0 07 $(N, sd) \leftarrow id$ 08 $e \leftarrow \text{PGet}(sd, br)$ 09 If $x \notin \mathbb{Z}_N^*$: 10 $p \leftarrow \gcd(x, N)$ 11 $q \leftarrow N/p$ 12 $\varphi_N \leftarrow (p-1)(q-1)$ 13 If $e \mid \varphi_N$: Return 0 14 $y \leftarrow x^e \bmod N$ 15 Return y </pre>	<pre> FInv(br, td, y) 16 $(N, sd, \varphi_N) \leftarrow td$ 17 $e \leftarrow \text{PGet}(sd, br)$ 18 If $\gcd(e, \varphi_N) \neq 1$: 19 Return \perp 20 $d \leftarrow e^{-1} \bmod \varphi_N$ 21 $x \leftarrow y^d \bmod N$ 22 Return x </pre>
---	---	--

Fig. 7. ABO from Phi-hiding assumption. $C(e^*, \varphi_N)$ denotes the condition defined in Section 2.3.

Theorem 1. *Let $l, n, d, \lambda \in \mathbb{N}$ as above. Then $P = (\text{PGen}, \text{PGet}, \text{PProg})$ from Fig. 6 is a $(2^{-(\lambda+1)}, 2^{-\lambda})$ -PPFG for $(\mathcal{P}_l, 2^n)$ with seed space $\mathcal{Sd} = \text{GF}(2^{l-1})^d$.*

7 ABO-LTP with index-independent domain from unique-divisor Phi-hiding

We use a prime family generator (for instance the one from Section 6) to construct an ABO-LTP with index-independent domain, which can be shown secure under the unique-divisor Phi-hiding assumption. The construction resembles [16, Section 5.2] who build an adaptive trapdoor function. As a starting point we first specify an ABO-LTP A having index-dependent domains. Using the transform from Section 3, A can be made index-independent. Due to the result of Lemma 4 the transformed ABO-LTP has essentially the same lossiness factor as A .

Index-dependent ABO-LTP from unique-divisor Phi-hiding. Let $n, l \in \mathbb{N}$. Consider the ABO-LTP defined in Fig. 7. We obtain the following result (the proof of which is in the full version [1]).

Lemma 7. *Let $n, l \in \mathbb{N}$ and let $P = (\text{PGen}, \text{PGet}, \text{PProg})$ be a (ϵ_1, ϵ_2) -PPFG for $(\mathcal{P}_l, 2^n)$. Consider $A = (\text{FGen}, \text{FEv}, \text{FInv})$ from Fig. 7. If the unique-divisor (τ, ϵ) -Phi-hiding assumption holds for (k, \mathcal{P}_l) , A is a $(1 - \epsilon_2)$ -correct, L -lossy, $(\tau', 2(\epsilon + \epsilon_1)/(1 - \epsilon_1))$ -indistinguishable ABO-LTP with index-dependent domain $\mathcal{X} = \llbracket 2^k \rrbracket$, where $L = 2^{l-1}$ and $\tau' \approx \tau$. Further, A has branching set $\mathcal{Br} = \llbracket 2^n \rrbracket$.*

Index-independent ABO-LTP from unique-divisor Phi-hiding. Using the technique from Section 3 it is possible to transform A into an index-independent ABO-LTP A_{ii} . Using the improved bound on the lossiness from Lemma 4 we obtain the following.

Corollary 2. *Let $n, l, k \in \mathbb{N}$ and $A = (\text{FGen}, \text{FEv}, \text{FInv})$ be the ABO defined in Fig. 7. Further, for $(N, sd) = id \in \mathcal{Id}$ let π_{id} the permutation*

$$\pi_{id}: \llbracket 2^k \rrbracket \rightarrow \llbracket 2^k \rrbracket; x \mapsto (x - N + 1) \bmod N .$$

Let A_{ii} be the conversion of A via Fig. 1. If the unique-divisor (τ, ϵ) -Phi-hiding assumption holds for (k, \mathcal{P}_1) , then A_{ii} is a $(1 - 2\epsilon_2)$ -correct, L -lossy, $(\tau', 2(\epsilon + \epsilon_1)/(1 - \epsilon_1))$ -indistinguishable index-independent ABO-LTP with domain $\llbracket 2^k \rrbracket$ and branching set $\llbracket 2^m \rrbracket$, where $L = 2^{l-1}/(1 + 2^{-k/2})$ and $\tau' \approx \tau$.

Acknowledgments

Benedikt Auerbach was supported in part by the NRW Research Training Group SecHuman and by ERC Project ERCC (FP7/615074). Bertram Poettering conducted part of the research at Ruhr University Bochum, supported by ERC Project ERCC (FP7/615074). Eike Kiltz was supported in part by ERC Project ERCC (FP7/615074) and by DFG SPP 1736 Big Data.

References

1. Auerbach, B., Kiltz, E., Poettering, B., Schoenen, S.: Lossy trapdoor permutations with improved lossiness. Cryptology ePrint Archive, Report 2018/1183 (2018), <https://eprint.iacr.org/2018/1183>
2. Bellare, M., Brakerski, Z., Naor, M., Ristenpart, T., Segev, G., Shacham, H., Yilek, S.: Hedged public-key encryption: How to protect against bad randomness. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 232–249. Springer, Heidelberg (Dec 2009)
3. Bellare, M., Fischlin, M., O’Neill, A., Ristenpart, T.: Deterministic encryption: Definitional equivalences and constructions without random oracles. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 360–378. Springer, Heidelberg (Aug 2008)
4. Bellare, M., Hoang, V.T.: Resisting randomness subversion: Fast deterministic and hedged public-key encryption in the standard model. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 627–656. Springer, Heidelberg (Apr 2015)
5. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and impossibility results for encryption and commitment secure under selective opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (Apr 2009)
6. Benhamouda, F., Herranz, J., Joye, M., Libert, B.: Efficient cryptosystems from 2^k -th power residue symbols. Journal of Cryptology 30(2), 519–549 (Apr 2017)
7. Boldyreva, A., Fehr, S., O’Neill, A.: On notions of security for deterministic encryption, and efficient constructions without random oracles. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 335–359. Springer, Heidelberg (Aug 2008)
8. Brakerski, Z., Segev, G.: Better security for deterministic public-key encryption: The auxiliary-input setting. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 543–560. Springer, Heidelberg (Aug 2011)
9. Cachin, C., Micali, S., Stadler, M.: Computationally private information retrieval with polylogarithmic communication. In: Stern, J. (ed.) EUROCRYPT’99. LNCS, vol. 1592, pp. 402–414. Springer, Heidelberg (May 1999)
10. Freeman, D.M., Goldreich, O., Kiltz, E., Rosen, A., Segev, G.: More constructions of lossy and correlation-secure trapdoor functions. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 279–295. Springer, Heidelberg (May 2010)

11. Freeman, D.M., Goldreich, O., Kiltz, E., Rosen, A., Segev, G.: More constructions of lossy and correlation-secure trapdoor functions. *Journal of Cryptology* 26(1), 39–74 (Jan 2013)
12. Goldwasser, S., Micali, S.: Probabilistic encryption. *Journal of Computer and System Sciences* 28(2), 270–299 (1984)
13. Hayashi, R., Okamoto, T., Tanaka, K.: An RSA family of trap-door permutations with a common domain and its applications. In: Bao, F., Deng, R., Zhou, J. (eds.) PKC 2004. LNCS, vol. 2947, pp. 291–304. Springer, Heidelberg (Mar 2004)
14. Hohenberger, S., Waters, B.: Short and stateless signatures from the RSA assumption. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 654–670. Springer, Heidelberg (Aug 2009)
15. Joye, M., Libert, B.: Efficient cryptosystems from 2^k -th power residue symbols. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 76–92. Springer, Heidelberg (May 2013)
16. Kiltz, E., Mohassel, P., O’Neill, A.: Adaptive trapdoor functions and chosen-ciphertext security. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 673–692. Springer, Heidelberg (May / Jun 2010)
17. Kiltz, E., O’Neill, A., Smith, A.: Instantiability of RSA-OAEP under chosen-plaintext attack. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 295–313. Springer, Heidelberg (Aug 2010)
18. Mironov, I., Pandey, O., Reingold, O., Segev, G.: Incremental deterministic public-key encryption. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 628–644. Springer, Heidelberg (Apr 2012)
19. Mol, P., Yilek, S.: Chosen-ciphertext security from slightly lossy trapdoor functions. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 296–311. Springer, Heidelberg (May 2010)
20. Nishimaki, R., Fujisaki, E., Tanaka, K.: Efficient non-interactive universally composable string-commitment schemes. In: Pieprzyk, J., Zhang, F. (eds.) ProvSec 2009. LNCS, vol. 5848, pp. 3–18. Springer, Heidelberg (Nov 2009)
21. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 187–196. ACM Press (May 2008)
22. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. *SIAM J. Comput.* 40(6), 1803–1844 (2011), <http://dx.doi.org/10.1137/080733954>
23. Raghunathan, A., Segev, G., Vadhan, S.P.: Deterministic public-key encryption for adaptively chosen plaintext distributions. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 93–110. Springer, Heidelberg (May 2013)
24. Rosen, A., Segev, G.: Chosen-ciphertext security via correlated products. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 419–436. Springer, Heidelberg (Mar 2009)
25. Shoup, V.: *A Computational Introduction to Number Theory and Algebra*. Cambridge University Press (2005)
26. Xie, X., Xue, R., Zhang, R.: Efficient threshold encryption from lossy trapdoor functions. In: Yang, B.Y. (ed.) Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011. pp. 163–178. Springer, Heidelberg (Nov / Dec 2011)
27. Yamakawa, T., Yamada, S., Hanaoka, G., Kunihiro, N.: Adversary-dependent lossy trapdoor function from hardness of factoring semi-smooth RSA subgroup moduli. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part II. LNCS, vol. 9815, pp. 3–32. Springer, Heidelberg (Aug 2016)