Efficient Verifiable Partially-Decryptable Commitments from Lattices and Applications

Muhammed F. Esgin[1,2], Ron Steinfeld[1], and Raymond K. Zhao[1]

[1] Faculty of Information Technology, Monash University, Clayton, Australia
[2] CSIRO's Data61, Melbourne, Australia
{Muhammed.Esgin,Ron.Steinfeld,Raymond.Zhao}@monash.edu

Abstract. We introduce *verifiable partially-decryptable commitments* (VPDC), as a building block for constructing efficient privacy-preserving protocols supporting *auditability* by a trusted party. A VPDC is an extension of a commitment along with an accompanying proof, convincing a verifier that (i) the given commitment is well-formed and (ii) a certain *part* of the committed message can be decrypted using a (secret) trapdoor known to a trusted party.

We first formalize VPDCs and then introduce a general decryption feasibility result that overcomes the challenges in *relaxed* proofs arising in the lattice setting. Our general result can be applied to a wide class of Fiat-Shamir based protocols and may be of independent interest.

Next, we show how to extend the commonly used lattice-based 'Hashed-Message Commitment' (HMC) scheme into a succinct and efficient VPDC. In particular, we devise a novel 'gadget'-based Regev-style (partial) decryption method, compatible with efficient relaxed lattice-based zero-knowledge proofs. We prove the soundness of our VPDC in the setting of adversarial proofs, where a prover tries to create a valid VPDC output that fails in decryption.

To demonstrate the effectiveness of our results, we extend a private blockchain payment protocol, MatRiCT, by Esgin et al. (ACM CCS '19) into a formally auditable construction, which we call MatRiCT-Au, with very low communication and computation overheads over MatRiCT.

Keywords: Lattice · Zero Knowledge · Verifiable Partially-Decryptable Commitment · Auditable RingCT · Accountable Ring Signature

1 Introduction

Commitment schemes and accompanying zero-knowledge proofs (ZKPs) have become crucial tools used in countless privacy-preserving protocols. For example, they are extensively used in privacy-aware blockchain applications such as Monero and Zcash cryptocurrencies to hide sensitive information such as user identities and transaction amounts. In many such privacy-preserving applications, there is a need for *auditability*, i.e., the ability of a trusted third-party to revoke the privacy or anonymity of the protocol, in order to catch or punish misbehaving entities. For instance, it is well known that the privacy features of cryptocurrencies have been exploited by cyber criminals to hide their illegal financial

activities, and some level of government oversight may be required in future to allow such activities to be traced by law authorities. Many applications where such an auditability feature is needed exist, including group signatures [5], fair exchange [1], key escrow [23] and e-voting. To enable the auditability property of the privacy protocol, we would like the protocol to use a *decryptable* commitment scheme, supporting a trapdoor decryption algorithm that enables the authority with some trapdoor to recover a message from a given commitment.[3] At the same time, to prevent malicious parties from escaping the auditability property, the protocol must support a *verifiable* decryptable commitment, which allows protocol parties to verify that a commitment is decryptable by the authority, while still hiding its contents from all other parties.

A problem similar to constructing verifiable decryptable commitments has been previously studied under the name of *verifiable encryption* [4] in the classical setting of DL-based and factoring-based public-key cryptography. The approach here is to use a public-key encryption scheme as the commitment (with the secret key known to the authority), and attach to it a zero-knowledge proof of plaintext knowledge in order to turn it into a verifiable commitment. This approach was extended to the post-quantum lattice-based setting in [18], instantiating the encryption scheme by a variant of Regev's encryption scheme [21] based on (Ring/Module)-LWE. We note that Regev's encryption scheme can also be viewed as a decryptable *short message* variant of the 'Unbounded-Message Commitment' (UMC) scheme [3] (see the full version of this paper on IACR's ePrint archive). Despite allowing Regev-style decryption, UMC also has the practical efficiency drawbacks we discuss below.

The use of verifiable Regev encryption as in [18] can result in very long commitments and communication overheads in typical applications. This is because both the randomness length and commitment length of Regev-encryption commitments have an additive term proportional to the dimension of the message vector. In typical lattice-based ZKPs such as [7, 9, 10, 12], the structure of the protocol requires the prover to send commitments to a large number of messages including masking randomness values as well as auxiliary terms, in addition to the commitment of the 'real' message which needs to be decrypted by the opening authority (e.g., the payment amount, or payer/payee identity in cryptocurrencies). The protocol also requires the prover to send masked variants of the commitment randomness. Both those factors lead to long proofs with Regev encryption commitments. To illustrate, in the MatRiCT cryptocurrency protocol of [12], an aggregated binary proof is used (see [12, Section 1.2]) to significantly reduce the proof length. In this proof, it is necessary to commit to individual bits of integers by separate ring elements so that each bit can be manipulated independently. As a result, the total message dimension in a commitment over the underlying polynomial ring R_q is in the order of several hundreds (the 'real' message is still a few hundreds dimensional). If one were to use a Regev-style

[3] We note here that our notion of a decryptable commitment is different from a *trapdoor* commitment. For a given commitment and a message, the latter allows a trapdoor holder to find a properly distributed opening randomness for the commitment.

encryption for this commitment, the commitment *alone* would cost around 100-200 KB. In comparison, this commitment costs only about 13 KB in MatRiCT thanks to the use of a *compressing* commitment.

To reduce the length of commitments/proofs, an alternative approach (used in [12]) to Regev-encryption commitments is to instead use lattice-based 'Hashed-Message Commitments' (HMC), where message hashing leads to a short commitment dimension independent of the total dimension of the committed messages. In HMC, message hashing is achieved by multiplying the (long) message vector by a *random* 'fat' (i.e., compressing) matrix and one relies on the hardness of (Ring/Module)-SIS to accomplish the binding property. However, in our context of *decryptable* commitments, the lack of a unique decryption for such HMC commitments (due to compression) makes them not directly suitable. Therefore, we study HMC in the *partial* decryption setting where the committed message has two parts: (i) a decryptable message (that contains the 'real' message the authority wants to recover), and (ii) a non-decryptable/auxiliary message (that contains other auxiliary terms that need not be recovered). This way, we can achieve both of our succinctness and (partial) decryptability goals simultaneously. We note that a straightforward combination of using UMC for the decryptable part of the message and HMC for the non-decryptable message part, although it deals with the auxiliary terms, still suffers from an overhead of at least two commitments plus the large cost of a UMC commitment. In contrast to HMC, the latter UMC commitment dimension over R_q is linear in the message dimension over R_q, which is over 100 in the context of MatRiCT discussed above.

An initial attempt to overcome the above-mentioned efficiency issues of UMC-like commitments in constructing VPDCs, was proposed in [12, Section 6.1], where a method of incorporating a lightweight Regev-style decryption trapdoor into an HMC commitment was proposed. However, although a promising direction to combine the best of both HMC and Regev encryption commitments, the work of [12] does not give a full solution to the problem, as it does not address two main technical challenges that we now explain.

Firstly, the decryption algorithm in [12] is only analyzed for *honestly-created* commitments *without* a rigorous framework. The analysis against *adversarially-created* commitments/proofs that pass the verification check, which is an important requirement in the auditability setting of VPDCs, is missing. We recall that for the underlying efficient ZKPs of opening for the HMC scheme we study in this work (see, e.g., [9,10]), the ZKP soundness only guarantees the existence of a *relaxed* commitment opening $(\boldsymbol{m}, \boldsymbol{r}, y)$ of a commitment C, satisfying the relaxed opening relation

$$(yC = \mathsf{Com}_{ck}(y\boldsymbol{m}; \boldsymbol{r})) \wedge (y \in \Delta\mathcal{C}) \wedge (\boldsymbol{m} \in \mathcal{M}), \tag{1}$$

where y is a short non-zero relaxation factor, $\Delta\mathcal{C}$ is the set of challenge differences and $\mathcal{M}$ is a public message space. Observe that the message opening $\boldsymbol{m}$ is proven to be in some set $\mathcal{M}$, which is important for our analysis, and for example, $\mathcal{M} = \{0,1\}^v$ for some $v \geq 1$ for the proof systems in [9, 10, 12]. Also, note that the relaxation factor y is *unknown* to the decryption algorithm as it is part of the prover's secret. Thus, it is not clear how one could enable such a decryption

feature in the setting of relaxed proofs as the decryptor does not even know what to decrypt exactly. The work by Lyubashevsky and Neven [18] addresses this problem in the setting of verifiable Regev encryption. Particularly in [18], it is shown that choosing a random y from the set of possible relaxation factors is in fact a good way to go, and the *expected* running time for their decryption algorithm is shown to be proportional to the number of random oracle queries made by the prover to generate the protocol transcript[4]. However, this result is specific to the Fiat-Shamir (FS) protocol[5] and the Regev-style decryption described in [18].

A second technical challenge in constructing an HMC-based partially decryptable commitment following the approach of [12] is that even if a suitable relaxation factor y is known by the decryption algorithm, decrypting the commitment with the Regev-style trapdoor key does not directly yield the decryptable message, but reveals a noisy inner-product (over the underlying polynomial ring R_q) of the message with a known *random* vector $\boldsymbol{a}$, of the form $\langle y\boldsymbol{a}, \boldsymbol{m}\rangle + e$ for some short noise term e (and y the relaxation factor). This leaves the question of how to efficiently recover the message $\boldsymbol{m}$ from this noisy information. The work of [12] addressed this issue only for *small* message spaces (and honestly generated commitments) by performing an exhaustive search over all possible messages, which is very restrictive and computationally expensive. How to make decryption work *efficiently* for *exponentially large* message spaces and guarantee the decryption soundness even against adversarially constructed commitments having such a relaxed opening has since remained unaddressed.

1.1 Our Contributions

Verifiable Partially-Decryptable Commitments. In this work, we first formalize the notion of a *Verifiable Partially-Decryptable Commitment* (VPDC), which is closely related to proofs of plaintext knowledge and verifiable encryption. In particular, a VPDC extends a commitment scheme C and a matching Non-Interactive Zero-Knowledge Proof (NIZK) Σ of opening for C by adding a trapdoor key generation algorithm CAddTd and a matching decryption algorithm CDec for C. The VPDC ensures that any valid commitment-proof pair (C, π) can be (partially) decrypted using the (secret) trapdoor td output by CAddTd.

The above notion is similar to verifiable encryption except that C is not an encryption, but rather a commitment. The differences, as pointed out in the introduction, are as follows. First, a commitment scheme in general allows for a more *succinct* encoding of a message (i.e., can be compressing unlike an encryption) and is readily compatible with many existing proof systems (see, e.g., [2, 3, 9–12]), hence has a matching NIZK already available. Second, in a VPDC, there are two message spaces: (i) a *decryptable* message space $\mathcal{D}$, whose elements can be committed *and* recovered in decryption, and (ii) an *auxiliary*

[4] We refer to [18] for methods that can be used to restrict an attacker from making a lot of random oracle queries.

[5] We call a public-coin proof made non-interactive via the Fiat-Shamir transformation as a Fiat-Shamir (FS) protocol.

message space $\mathcal{U}$, whose elements can be used to create a commitment, but are not decryptable. As a result, a VPDC eliminates the need for an additional set of requirements due to an encryption scheme, avoids potential compatibility issues and enables partial decryption while still permitting a succinct encoding of the whole message (together with additional auxiliary terms). We therefore believe VPDCs can serve as an important building block in constructing *efficient* cryptographic schemes supporting accountability, such as group signatures, fair exchange protocols, key escrow and e-voting.

Generalized analysis of decryption feasiblity for *relaxed* ZKPs. To address the first main technical challenge of handling relaxed ZKPs in decryption of VPDCs, we show how to abstract and generalize the decryption algorithm of [18] that works only for the specific Regev-based (UMC-like) encryption considered therein, to design an efficient decryption algorithm for *any* VPDC satisfying a few natural properties. In particular, the expected number of iterations until the decryption function terminates is about the number of random oracle queries made by the prover in generating the transcript to be decrypted as in [18]. Our general result is applicable to any VPDC whose underlying NIZK is derived via the Fiat-Shamir transform in the random oracle model from a Sigma protocol satisfying a variant of special soundness that is satisfied by all known instantiations of such Sigma protocols.

A novel gadget-based Regev-style decryption for HMC. Building on the above general foundations, we construct a VPDC extending one of the most commonly used lattice-based commitment schemes, namely HMC[6]. For example, the HMC scheme is an integral part of one of the most efficient post-quantum ring signatures and set membership proofs in [11], arising from [9,12], as well as sublinear-sized arithmetic circuit satisfiability proofs in [2].

In particular, to address the second main technical challenge, we introduce an HMC-compatible trapdoor decryption method that works even when the decryptable message opening is proven to be in a set of exponential size (such as 2^{256}). We analyze this method in the setting of adversarially-created VPDC outputs and provide decryption soundness guarantees. As opposed to the trapdoor decryption of [12], where the trapdoor decryption yields $\langle y\boldsymbol{a}, \boldsymbol{m}\rangle + e$ for a *random* vector $\boldsymbol{a}$, small noise e and relaxation factor y (which is hard to decrypt), our new Regev-style partial trapdoor embeds a *structured* 'gadget' vector $\bar{t}\boldsymbol{g}$ in place of $\boldsymbol{a}$ in the HMC submatrix corresponding to the decryptable message. With this, trapdoor decryption of a commitment yields $\bar{t}y\langle \boldsymbol{g}, \boldsymbol{m}\rangle + e$ for a 'large' integer $\bar{t}$, which is efficiently decryptable by exploiting the structure of the gadget vector $\bar{t}\boldsymbol{g}$ using a rounding procedure similar to standard Regev decryption. The runtime of our new trapdoor decryption is polylogarithmic in the message space size $|\mathcal{D}|$ and we prove that it works correctly even against adversarially-generated commitments and ZKPs, as long as the system modulus q is sufficiently large and the message is proven to be a part of a decryptable message space $\mathcal{D}$.

[6] This distinguishes our VPDC construction from the verifiable encryption scheme of [18], that extends a UMC-type commitment scheme.

Table 1. Comparison between MatRiCT [12] and MatRiCT-Au (this work).

Anonymity level		1/10		1/100	
# of inputs → # of outputs		1 → 2	2 → 2	1 → 2	2 → 2
Proof Size	MatRiCT [12]	93	110	103	120
	MatRiCT-Au	**96**	**113**	**106**	**123**
Spend / Verify Runtimes	MatRiCT [12]	242 / 20	375 / 23	360 / 31	610 / 40
	MatRiCT-Au	**233 / 21**	**414 / 25**	**402 / 33**	**654 / 42**
Parameters	MatRiCT [12]	PK Size: 4.36 KB		Moduli: $< 2^{53.0}$	
	MatRiCT-Au	PK Size: 4.36 KB		Moduli: $< 2^{55.3}$	

Our lightweight Regev-style 'partial trapdoor' also avoids the heavyweight machinery of 'full' lattice trapdoors a-la [19], and still supports SIS-style HMC commitment, compatible with efficient ZKP techniques used in [9, 12]. Using the 'full' trapdoors in [19] in our commitments (with ternary coordinate trapdoor vectors) requires SIS matrices with n rows and $m \geq n \log q$ columns over the underlying ring, while the 'partial trapdoor' commitments we use, $m = 2n$ columns are sufficient (still with ternary coordinate trapdoor vectors). We save a significant factor $\approx \log q$ in both public parameter length and the length of masked messages in the ZKP protocol, for the same security level.

MatRiCT-Au: Auditable RingCT based on standard lattice assumptions. As an application of our compact lattice-based VPDC, we show how it can enable an extension of the lattice-based RingCT-like private cryptocurrency protocol MatRiCT [12] easily and efficiently into an auditable variant we call MatRiCT-Au, where an auditor with access to a (secret) trapdoor can revoke the anonymity of certain users (e.g., in case of misbehaviour). The auditability feature can be optional (i.e., each user individually decides whether and by whom she wants to be audited) or enforced by a simple public check. Our construction allows *adversarially-generated* transactions to be audited, whereas, in [12], the discussion about auditability is incomplete, as the decryption method given there may fail in the adversarial transaction setting, potentially allowing adversaries to avoid auditability. Furthermore, the proposal in [12] requires an exhaustive-search-based approach while we can very efficiently run **Audit** function over a message space of size $> 2^{128}$. To analyze auditability formally in confidential transactions, we also extend the formal model for RingCT-like protocols in [12] to add the *auditability* property and prove formally that MatRiCT-Au is auditable. We compute concrete parameters for MatRiCT-Au and present implementation results[7]. Our evaluation demonstrates the practicality of MatRiCT-Au, and in particular there are very little communication and computation overheads introduced over the original MatRiCT protocol [12] as shown in Table 1 (see the full version of this paper for more run-time results).

We believe that our new techniques will find further applications in the settings where accountable anonymity is desired. Particularly, our extension of

[7] The source code of our MatRiCT-Au implementation is available at `https://gitlab.com/raykzhao/matrict_au`.

HMC into a VPDC with soundness against adversarially-created outputs extends the group (or accountable ring) signature in [12] so as to enable *efficient* anonymity revocation (i.e., opening of a group signature) against *cheating* signers. Interesting research directions from here would be, for example, to design efficient post-quantum e-voting, auction and anonymous credential schemes by exploiting the accountable anonymity provided by our VPDC.

1.2 Our Results and Techniques

A novel gadget-based Regev-style decryption for HMC. Suppose that we work over a cyclotomic ring $R_q = \mathbb{Z}_q[X]/(X^d + 1)$, and have a *binary* secret vector $\boldsymbol{b} \in \{0,1\}^v \subset R_q^v$ that forms the decryptable message to be recovered in decryption. As explained above, in a typical application protocol, we commit to this message $\boldsymbol{b}$ together with a non-decryptable message $\boldsymbol{u}$ as $C = \mathsf{Com}_{ck}(\boldsymbol{b}, \boldsymbol{u}; \boldsymbol{r})$ under some commitment randomness $\boldsymbol{r}$. The application protocol also proves knowledge of a *relaxed* opening of C (i.e., knowledge of $(y, \boldsymbol{b}', \boldsymbol{u}', \boldsymbol{r}')$ such that $yC = \mathsf{Com}_{ck}(y\boldsymbol{b}', \boldsymbol{u}'; \boldsymbol{r}')$ and $\boldsymbol{b}' \in \{0,1\}^v$). For simplicity, let us consider the case $y = 1$. After dealing with this case, we will discuss how we lift the restriction of $y = 1$ using our generalized decryption analysis results from Sec. 4.

The HMC commitment we use has the form $C = \mathsf{Com}_{ck}(\boldsymbol{b}, \boldsymbol{u}; \boldsymbol{r}) = \boldsymbol{A}\boldsymbol{r} + \boldsymbol{B}\boldsymbol{b} + \boldsymbol{C}\boldsymbol{u}$ and we recover the decrypted message as an element of R_t for some $t \geq 1$. To allow trapdoor decryption of $\boldsymbol{b}$, but not $\boldsymbol{r}$ and $\boldsymbol{u}$, our trapdoor key generation algorithm embeds a Regev-style '*gadget trapdoor*' into the last row $\boldsymbol{t}_B^\top$ of matrix $\boldsymbol{B}$ and a Regev-style '*error trapdoor*' into the last row $\boldsymbol{t}_A^\top$ (resp. $\boldsymbol{t}_C^\top$) of matrix $\boldsymbol{A}$ (resp. $\boldsymbol{C}$). That is, for the 'gadget trapdoor' matrix, we have $\boldsymbol{B} = \begin{pmatrix} \boldsymbol{B}' \\ \boldsymbol{t}_B^\top \end{pmatrix}$ with 'gadget trapdoor' row $\boldsymbol{t}_B^\top = \boldsymbol{s'}^\top \boldsymbol{B}' + \boldsymbol{e}_B^\top + \bar{t}\boldsymbol{g}^\top$, where $\bar{t} = \lfloor q/t \rfloor$, $\boldsymbol{e}_B$ is a short error, $\boldsymbol{s}'$ is a random secret, and $\boldsymbol{g}^\top$ is a 'gadget' vector with coordinates of the form $(2^i X^j)_{i<\tau, j<d}$ where $2^\tau \leq t$. While for the 'error trapdoor' matrices, we have $\boldsymbol{A} = \begin{pmatrix} \boldsymbol{A}' \\ \boldsymbol{t}_A^\top \end{pmatrix}$ and $\boldsymbol{C} = \begin{pmatrix} \boldsymbol{C}' \\ \boldsymbol{t}_C^\top \end{pmatrix}$ with 'error trapdoor' rows $\boldsymbol{t}_A^\top = \boldsymbol{s'}^\top \boldsymbol{A}' + \boldsymbol{e}_A^\top$ and $\boldsymbol{t}_C^\top = \boldsymbol{s'}^\top \boldsymbol{C}' + \boldsymbol{e}_C^\top$, where $\boldsymbol{e}_A, \boldsymbol{e}_C$ are short errors. Let $\boldsymbol{s}^\top = (-\boldsymbol{s'}^\top, 1)$ be the trapdoor. We remark that in the prior work [12], the matrix $\boldsymbol{B}$ was a random SIS matrix with no decryption trapdoor, which led to an inefficient exhaustive search decryption over the message space.

Now, it is easy to observe that $C' := \langle \boldsymbol{s}, C \rangle = e + \langle \bar{t}\boldsymbol{g}, \boldsymbol{b} \rangle$, where $e := (\langle \boldsymbol{e}_A, \boldsymbol{r} \rangle + \langle \boldsymbol{e}_B, \boldsymbol{b} \rangle + \langle \boldsymbol{e}_C, \boldsymbol{u} \rangle)$ is a small error. Thanks to the structure of the gadget vector $\bar{t}\boldsymbol{g}^\top$, the integer coefficients of $\langle \bar{t}\boldsymbol{g}, \boldsymbol{b} \rangle$ are multiples of the large integer $\bar{t}$ and encode the bits of the decryptable message $\boldsymbol{b}$ in their binary representation. Thus, $\boldsymbol{b}$ can be recovered from C' in the decryption algorithm by rounding out the small error term e to a multiple of $\bar{t}$ and performing binary decomposition, whereas the non-decryptable message/randomness $\boldsymbol{u}, \boldsymbol{r}$ only contribute to the error term e.

To apply our gadget-based Regev-style decryption for HMC to adverserially-generated commitments with a relaxed proof of opening, we apply the general

result of Theorem 1. To apply the latter theorem, we give a generalized decryption algorithm for our Regev-style HMC trapdoor and analyse (in Theorem 2 in Sec. 5.5) its correctness and soundness against (i) 'false rejection' decryption errors (where the algorithm fails to recover a decryptable message opening, even though the latter exists), as well as (ii) 'false acceptance' decryption errors (where the algorithm recovers a different decryptable message than the one in the valid opening), respectively. For (i), to recover the decryptable message $\boldsymbol{b}$ even for adverserial commitments C with a non-trivial relaxed opening $yC = \mathsf{Com}_{ck}(y\boldsymbol{b}, \boldsymbol{u}; \boldsymbol{r})$ with some short relaxation factor y, our decryption algorithm recovers $y\langle \boldsymbol{g}, \boldsymbol{b}\rangle \bmod t$ after rounding $\langle \boldsymbol{s}, yC\rangle$ to a multiple of $\bar{t}$, and we rely on invertibility of relaxation factors $y \bmod t$ to recover $\boldsymbol{b}$. For (ii), we show that a mildly larger choice of modulus q than needed for (i) guarantees that incorrect (non-unique) decryptable messages are never returned by our decryption algorithm, even with adverserial commitments/proofs and relaxation factors y.

We remark that the high-level structure of our HMC gadget-based Regev-style decryption trapdoor is similar to the full LWE inversion trapdoor of [19], but there are several important technical differences due to our HMC setting that are crucial to our scheme's efficiency and security. First, our use of the gadget during decryption is in some sense 'dual' to its use in [19]: in the LWE inversion problem considered in [19], the LWE secret $\boldsymbol{s}$ is assumed to be *uniformly random* $\bmod q$ (rather than 'short'), so that trapdoor decryption yields $\boldsymbol{c} = \boldsymbol{G}\cdot\boldsymbol{s}+\boldsymbol{e'}$ for a gadget matrix $\boldsymbol{G}$ and short error vector $\boldsymbol{e'}$. Here, to efficiently recover the 'large' coordinate secret $\boldsymbol{s}$ from $\boldsymbol{c}$, the gadget matrix $\boldsymbol{G}$ is constructed to have $\log q$ powers of 2 (up to $q/2$) along each of its *columns* so that the mapping $\boldsymbol{s} \mapsto \boldsymbol{G}\cdot\boldsymbol{s}$ effectively performs *bit decomposition* of the coordinates of $\boldsymbol{s}$. This approach *expands* the dimension of $\boldsymbol{s}$ by a factor $\log q$ to allow recovery of each bit of each coordinate of $\boldsymbol{s}$ from the corresponding row of $\boldsymbol{Gs}$. Whereas in our 'dual' HMC decryption algorithm, the decryptable message $\boldsymbol{s}$ is binary (and hence 'short'), so that when our trapdoor decryption similarly yields $\boldsymbol{c} = \boldsymbol{G}\cdot\boldsymbol{s} + \boldsymbol{e'}$, we can choose the gadget matrix $\boldsymbol{G} = \boldsymbol{g}^\top$ to have powers of 2 along its *row* so that the mapping $\boldsymbol{s} \mapsto \boldsymbol{G}\cdot\boldsymbol{s}$ performs *binary reconstruction* of integers whose bits are the coordinates of $\boldsymbol{s}$. Our approach *compresses* the dimension of $\boldsymbol{s}$ to a single element over the underlying ring, and minimises the dimension of the underlying matrices/commitments. Hence, our algorithm can also be viewed as a more efficient inversion trapdoor for LWE in 'dual' *knapsack* form ($\boldsymbol{c} = \boldsymbol{B}\boldsymbol{b}$ for 'short' $\boldsymbol{b}$ and 'fat' $\boldsymbol{B}$) rather than the more usual 'primal' form ($\boldsymbol{c} = \boldsymbol{A}\boldsymbol{s} + \boldsymbol{e}$ for 'short' $\boldsymbol{e}$ and 'tall' $\boldsymbol{A}$) addressed in [19]. A second difference from [19] is our use of *error trapdoors* for the HMC submatrices corresponding to non-decryptable message/randomness. And thirdly, as outlined above, our decryption algorithm analysis handles the adverserial commitment case with relaxed opening proofs, whereas [19] only analyses decryption for honestly created LWE samples.

MatRiCT-Au application. To show the usefulness of our novel decryption method in practice, we apply it in the setting of MatRiCT [12]. In MatRiCT, a commitment B encodes (i) an index in binary form that identifies the real user creating the transaction, and (ii) the bits of the transaction amount. Therefore, we can apply our novel decryption method to decrypt this commitment. Overall,

in addition to revoking the anonymity, we can enable an auditor to recover the hidden transaction amount. This is similar in spirit to traceable range proofs [15] (though the techniques are completely different).

Recently, a newer version of MatRiCT was published in [11]. Our techniques apply also to this newer version, called MatRiCT$^+$, and the overhead of extending MatRiCT$^+$ to support auditability is just an increase of about 20% in proof size. We discuss further details in the full version of this paper.

Organization of the paper. Section 2 covers preliminaries. We introduce the formal definitions of a VPDC in Section 3. Our generalized analysis of decryption runtime for relaxed ZKPs is introduced in Section 4. Then, in Section 5, we provide, along with the ordinary HMC scheme, the details of our new lattice-based VPDC, its decryption algorithm, and its adversarial soundness and run-time analyses. We discuss how VPDC can be used to construct MatRiCT-Au in Section 6 and, due to limited space, provide the full details relating to MatRiCT-Au in the full version of this paper on IACR's ePrint archive. Particularly, our extended formal model for RingCT-like protocols, the full description of MatRiCT-Au (including parameter setting and implementation details), and the security discussions of MatRiCT-Au are provided in the full version.

2 Preliminaries

For an odd modulus q, the ring of integers modulo q, $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$, is represented by the range $\left[-\frac{q-1}{2}, \frac{q-1}{2}\right]$. To denote column vectors and matrices, we use bold-face lower-case letters such as $\boldsymbol{x}$ and bold-face capital letters such as $\boldsymbol{V}$, respectively (hence, $\boldsymbol{x}^\top$ denotes a row vector). $(\boldsymbol{x}, \boldsymbol{y})$ is used to denote concatenation of the two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ to form a single longer vector. For a vector $\boldsymbol{x} = (x_0, \ldots, x_{n-1})$, we define the following norms $\|\boldsymbol{x}\| = \sqrt{\sum_{i=0}^{n-1} x_i^2}$, $\|\boldsymbol{x}\|_\infty = \max_i |x_i|$ and $\|\boldsymbol{x}\|_1 = \sum_{i=0}^{n-1} |x_i|$. When considering a norm of a polynomial f, we define the same norms on the coefficient vector of f. For a vector $\boldsymbol{f} = (f_0, \ldots, f_{s-1})$ of polynomials, we further define $\|\boldsymbol{f}\| = \sqrt{\sum_{i=0}^{s-1} \|f_i\|^2}$, $\|\boldsymbol{f}\|_1 = \sum_{i=0}^{s-1} \|f_i\|_1$, $\|\boldsymbol{f}\|_\infty = \max_i \|f_i\|_\infty$. The Hamming weight of the (concatenated) coefficient vector of $\boldsymbol{f}$ is denoted by $\mathsf{HW}(\boldsymbol{f})$. $\mathcal{U}(S)$ denotes uniform distribution on a set S.

Capital letters such as C denote commitments, and we write $\mathcal{S}^{d \cdot k}$ when a total of kd coefficients are sampled from a set $\mathcal{S}$ in order to generate k polynomials in $R = \mathbb{Z}[X]/(X^d + 1)$ of a power-of-2 degree d. $\mathbb{S}_\mathcal{B}$ denotes the set of polynomials in R, where each coefficient has an absolute value bounded by $\mathcal{B} \in \mathbb{Z}^+$.

2.1 Security Assumptions

In our applications, we use a commitment scheme whose security relies on the following well-known lattice problems.

Definition 1 (M-SIS$_{n,m,q,\beta_{\text{SIS}}}$). *Given $\boldsymbol{A} \leftarrow R_q^{n \times m}$ sampled uniformly at random, the Module-SIS (M-SIS) problem asks to find a* short $\boldsymbol{x} \in R_q^m$ *such that $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{0}$ over R_q and $0 < \|\boldsymbol{x}\| \leq \beta_{\text{SIS}}$.*

Definition 2 (M-LWE$_{n,m,q,\mathcal{B}}$). *The Module-LWE (M-LWE) problem asks to distinguish between the following two cases: (i)* $(\boldsymbol{A}, \boldsymbol{As}+\boldsymbol{e})$ *for* $\boldsymbol{A} \leftarrow R_q^{m\times n}$*, a secret vector* $\boldsymbol{s} \leftarrow \mathbb{S}_{\mathcal{B}}^n$ *and an error vector* $\boldsymbol{e} \leftarrow \mathbb{S}_{\mathcal{B}}^m$*, and (ii)* $(\boldsymbol{A}, \boldsymbol{t})$ *for* $\boldsymbol{A} \leftarrow R_q^{m\times n}$ *and* $\boldsymbol{t} \leftarrow R_q^m$*.*

It is known that the secret $\boldsymbol{s}$ can equivalently be sampled from $\mathcal{U}(R_q^n)$.

2.2 Zero-Knowledge Proofs

A Relaxed NIZK $\Sigma = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ for relation R_σ and its relaxed counterpart R'_σ with $\mathsf{R}_\sigma \subseteq \mathsf{R}'_\sigma$ (parameterized by a common reference string σ) and their corresponding languages $L_\sigma = \{u : \exists r \text{ s.t. } (u,r) \in \mathsf{R}_\sigma\}$ and $L'_\sigma = \{u : \exists r \text{ s.t. } (u,r) \in \mathsf{R}'_\sigma\}$ respectively, consists of the following algorithms (here, u denotes a language member and r denotes a witness):

$\sigma \leftarrow \mathsf{K}(1^\lambda)$: is the PPT common reference string generation algorithm of Σ that outputs a common reference string σ.

$\pi \leftarrow \mathsf{P}^{\mathcal{H}}(\sigma, u, r)$: is the PPT prover algorithm of Σ that, given a common reference string σ, access to a random oracle $\mathcal{H}$ and a language member u and a witness r with $(u,r) \in \mathsf{R}_\sigma$, outputs a proof π.

$0/1 \leftarrow \mathsf{V}^{\mathcal{H}}(\sigma, u, \pi)$: is the PPT verification algorithm of Σ that, given a common reference string σ, access to a random oracle $\mathcal{H}$ and a language member u and proof π, outputs 0 (invalid) or 1 (valid).

We remark that our lattice-based constructions regard the commitment key as part of the CRS σ (a similar issue arises in both DL-based Pedersen and lattice-based commitments). We refer to the full version of this paper for the standard definitions of completeness, soundness and zero-knowledge for NIZK proofs.

Our VDPC construction is based on a NIZK obtained using the Fiat-Shamir (FS) transform [13] applied to an interactive Zero-Knowledge Sigma protocol $\Sigma_{\mathsf{I}} = (\mathsf{K}_{\mathsf{I}}, \mathsf{P}_{\mathsf{I}}, \mathsf{V}_{\mathsf{I}})$ for relations $\mathsf{R}_\sigma, \mathsf{R}'_\sigma$ (parameterised by a common reference string σ) with a challenge space $\mathcal{C}$ and public-private inputs (u, r) with same notations for relations as above. We refer to the full version of this paper for the standard definitions of completeness, special soundness and honest-verifier zero-knowledge for Sigma protocols. The FS heuristic transforms Σ_{I} into a NIZK using a random oracle $\mathcal{H}$, by letting the prove algorithm compute the verifier's challenge from the common reference string σ, public input u, and commitment message w, setting $x = \mathcal{H}(\sigma, u, w)$.

2.3 Commitment Schemes

A commitment scheme $\mathsf{C} = (\mathsf{CKeygen}, \mathsf{Commit}, \mathsf{COpen})$ consists of three algorithms:

$pp = (ck, \mathcal{M}, \mathcal{R}) \leftarrow \mathsf{CKeygen}(1^\lambda)$: is a PPT key generation algorithm returning pp containing a commitment key ck and descriptions of message space $\mathcal{M}$ and randomness space $\mathcal{R}$. Note pp is an implicit input to the remaining algorithms.

$(C, \mathsf{o}) \leftarrow \mathsf{Commit}(\mathsf{m})$: is a PPT commitment algorithm which for message $\mathsf{m} \in \mathcal{M}$, outputs a commitment C to m together with an opening o.

$0/1 \leftarrow \mathsf{COpen}(C, \mathsf{o})$: is a deterministic poly-time opening algorithm that given commitment C and opening o, checks whether o is a valid opening of C.

An opening o of a commitment is a tuple containing a message m, randomness r, and possibly also relaxation factors used by the opening algorithm (e.g., the relaxation factor y used in the lattice-based HMC commmitment in Sec. 5.1). We write $\mathsf{m}(\mathsf{o})$ to denote the message part of opening o. We refer to the full version of this paper for standard definitions of correctness, hiding and binding properties of commitment schemes.

3 VPDC: Verifiable Partially-Decryptable Commitments

A VPDC is an extension of two building blocks: (1) a (non-decryptable) commitment scheme C, and (2) a NIZK relaxed proof of opening protocol Σ for C. The VPDC adds a new trapdoor key generation algorithm CAddTd to embed a hidden partial decryption trapdoor td in the commitment key of C, such that with this trapdoor, efficient partial decryption of commitments accompanied by a valid relaxed proof of opening is possible, using the VPDC's partial decryption algorithm CDec. In particular, for VPDC, we view the commitment scheme's message space $\mathcal{M}$ as the product of two sets $\mathcal{D}$ and $\mathcal{U}$, where $\mathcal{D}$ is the *decryptable* message space and $\mathcal{U}$ is the *auxiliary* message space. For a commitment opening o, we let $\mu(\mathsf{o})$ denote the decryptable message part of o.

Formally, a Verifiable Partially-Decryptable Commitment scheme $\mathsf{VPDC} = (\mathsf{C}, \Sigma, \mathsf{CAddTd}, \mathsf{CDec})$ consists of a (non-decryptable) commitment scheme $\mathsf{C} = (\mathsf{CKeygen}, \mathsf{Commit}, \mathsf{COpen})$ with message space $\mathcal{M} = \mathcal{D} \times \mathcal{U}$ (the decryptable message space $\mathcal{D}$ and auxiliary message space $\mathcal{U}$ respectively), and a *matching* NIZK relaxed proof of opening protocol $\Sigma = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ for C, a trapdoor key generation algorithm CAddTd and a partial decryption algorithm CDec.

We say that the underlying NIZK Σ is a *matching* NIZK relaxed proof of opening for C if:

- On input 1^λ, the CRS generation algorithm K returns a CRS of the form $\sigma = (pp, \sigma')$, where $pp = (ck, \mathcal{M}, \mathcal{R})$ is C's public parameters $pp \leftarrow \mathsf{CKeygen}(1^\lambda)$. (i.e. Σ has pp in its CRS).
- Σ satisfies the standard completeness, soundness and zero-knowledge properties with respect to the following commitment opening relations $\mathsf{R}_{\mathsf{C},pp} \subseteq \mathsf{R}'_{\mathsf{C},pp}$ (parameterised by the commitment key pp from the CRS):

$$\mathsf{R}_{\mathsf{C},pp} = \{(C, \mathsf{o}) : \exists(\mathsf{m}, r) \in (\mathcal{M} \times \mathcal{R}) \text{ with } (C, \mathsf{o}) = \mathsf{Commit}(m; r)\}$$

and

$$\mathsf{R}'_{\mathsf{C},pp} \subseteq \mathsf{R}^{\mathsf{COpen}}_{\mathsf{C},pp} := \{(C, \mathsf{o}) : \mathsf{COpen}(C, \mathsf{o}) = 1\}\,.$$

In addition to the algorithms $(\mathsf{CKeygen}, \mathsf{Commit}, \mathsf{COpen})$ of C and the algorithms $(\mathsf{K}, \mathsf{P}, \mathsf{V})$ of Σ, VPDC adds two new algorithms to enable decryptability, with the following syntax:

$(ck^{\mathsf{td}}, \mathsf{td}) \leftarrow \mathsf{CAddTd}(ck, \mathcal{D}, \mathcal{U})$: a PPT algorithm that on input a commitment key ck and a description of the decryptable and auxiliary message spaces $\mathcal{D}$ and $\mathcal{U}$ such that $\mathcal{M} = \mathcal{D} \times \mathcal{U}$, outputs a 'trapdoored' commitment key ck^{td} and a partial decryption trapdoor td.

$\mu' \leftarrow \mathsf{CDec}_{\mathsf{td}}(C, \pi)$: is a probabilistic algorithm that on input a commitment C with a corresponding proof π and a trapdoor td, outputs a message $\mu' \in \mathcal{D}$.

We now list several additional properties for a VPDC, all of which are enjoyed by our construction:

Succinctness: The bit length of the commitment should depend only poly-logarithmically on the bit length of the auxiliary message.[8]

Additive Homomorphism: The commitment message and randomness spaces are subsets of modules with operations $(+, \cdot)$ over some underlying scalar ring $\mathfrak{R}$, the commitment space is a subset of a module with operations $(\oplus, \otimes)$ over $\mathfrak{R}$, and there exists a set $S \subseteq \mathfrak{R}$ of scalars, such that for all messages $\mathsf{m}_1, \mathsf{m}_2 \in \mathcal{M}$, randomness $r_1, r_2 \in \mathcal{R}$ and scalar $\alpha \in S$, we have $C = \alpha \otimes C_1 \oplus C_2$ for $(C, \cdot) := \mathsf{Commit}(\alpha \cdot \mathsf{m}_1 + \mathsf{m}_2; \alpha \cdot r_1 + r_2)$, $(C_1, \cdot) := \mathsf{Commit}(\mathsf{m}_1; r_1)$ and $(C_2, \cdot) := \mathsf{Commit}(\mathsf{m}_2; r_2)$.

Small Integer Decryptable Message Space: The decryptable message space $\mathcal{D} \subset \mathfrak{R}^v$ is of the form $\mathcal{D} := Z_{\mathcal{B}}^v$, where $Z_{\mathcal{B}} \subseteq \mathbb{Z}$ is a set of *integers* of small maximum absolute value $\mathcal{B} = \lambda^{o(1)}$, and v is the decryptable message dimension over the underlying scalar ring $\mathfrak{R}$.

The succinctness property is essential for the efficient application of our VPDC in ZKPs. For our concrete VPDC construction, $\mathcal{U}$ is a much bigger set than $\mathcal{D}$. Thus, one can commit to auxiliary terms together with the target message to be decrypted under a single succinct commitment to save significant communication thanks to the commitment's compression feature (which is not available in encryption-based commitments). Here, we stress that *partial* decryption (as opposed to *full* decryption) is an important feature, not a drawback. If we required full decryption, then we would not be able to achieve succinctness. Similarly, the additively homomorphic property is needed to support efficient (e.g., 'Schnorr-like' [22]) ZKPs that rely on this property. Note that such a homomorphism is needed also for the non-decryptable message parts. This precludes a simple VPDC solution that would commit by hashing the auxiliary message part with a non-homomorphic collision-resistant hash function. The 'Small Integer Decryptable Message Space' property is required to efficiently support certain classes of ZK proofs needed in applications, such as the binary proofs, range proofs and 1-out-of-N proofs in [9, 12, 14]. Here, the fact that the message coordinates are integers, rather than general ring elements, allows for independent manipulation of the committed decryptable message coordinates (e.g., for computing an integer vector inner-product or independent evaluation of a quadratic function on all message coordinates) as needed in the verification of such ZK proofs. Their smallness bound (size $\mathcal{B} = \lambda^{o(1)}$) allows the length of such proofs to be kept short.

[8] Note that succinctness cannot be achieved for the *decryptable* message.

It is important to note that our VPDC model allows all of the following three properties *together* within the same environment:

- one can commit to any message in $\mathcal{M}$, where decryption is (computationally) infeasible. Such commitments are simply ordinary commitments and in this case, the commitment key ck should be used.
- one can commit to any message μ in $\mathcal{D}$ together with an auxiliary message in $\mathcal{U}$, where recovery of μ is possible using td. In this case, the commitment key ck^{td} should be used.
- one can commit to any message in $\mathcal{M}$, where decryption is not necessarily needed. Here, both commitment keys ck or ck^{td} can be used and commitments created this way are easily compatible with the rest of the protocol.

We require that C satisfies the standard correctness, hiding and binding properties of commitment schemes. Furthermore, we recall that NIZK proof Σ is required to be a matching NIZK for C (see above), and so satisfies the completeness, (relaxed) soundness and zero-knowledge properties for relations $(\mathsf{R}_{\mathsf{C},pp}, \mathsf{R}'_{\mathsf{C},pp})$ defined above.

In addition, as a partially-decryptable extension of a given commitment scheme C and matching ZK proof Σ, we would like the VPDC's trapdoor key generation algorithm for C to preserve the functionality and security properties of C and Σ. Accordingly, we say that C (resp. Σ) satisfies the *VPDC trapdoor key variants* of correctness, hiding, and binding properties for C (respectively, the trapdoor key variants of completeness, (relaxed) soundness and zero-knowledge for Σ) if the properties are still satisfied when the commitment key generation calls $pp = (ck, \mathcal{M}, \mathcal{R}) \leftarrow \mathsf{CKeygen}$ in C (resp. its call in K of Σ) are followed by the trapdoor commitment key generation calls $(ck^{\mathsf{td}}, \mathsf{td}) \leftarrow \mathsf{CAddTd}(ck)$ and $pp' = (ck^{\mathsf{td}}, \mathcal{M}, \mathcal{R})$ replaces pp. For ease of reference, we define from hereon $(ck^{\mathsf{td}}, \mathsf{td}, \mathcal{M}, \mathcal{R}) \leftarrow \mathsf{CKeygenTd}(1^\lambda)$ as the function that runs $(ck, \mathcal{M}, \mathcal{R}) \leftarrow \mathsf{CKeygen}(1^\lambda)$ and $(ck^{\mathsf{td}}, \mathsf{td}) \leftarrow \mathsf{CAddTd}(ck)$, and returns $(ck^{\mathsf{td}}, \mathsf{td}, \mathcal{M}, \mathcal{R})$. The following commitment Key Indistinguishability property for a VPDC suffices for this purpose (see Proposition 1).

Key Indistinguishability. A VPDC scheme is said to satisfy *key indistinguishability* if any PPT adversary $\mathcal{A}$ wins the following game with probability $1/2 + \mathrm{negl}(\lambda)$:

1. $pp_0 = (ck, \mathcal{M}, \mathcal{R}) \leftarrow \mathsf{CKeygen}(1^\lambda)$,
2. $pp_1 \leftarrow (ck^{\mathsf{td}}, \mathcal{M}, \mathcal{R})$, where $(ck^{\mathsf{td}}, \mathsf{td}) \leftarrow \mathsf{CAddTd}(ck)$.
3. $b \xleftarrow{\$} \{0, 1\}$,
4. $b' \leftarrow \mathcal{A}(pp_b)$.
5. $\mathcal{A}$ wins the game if $b' = b$.

The following proposition is immediate from the fact that the trapdoor key td does not appear in the view of the adversary in the security games defining the trapdoor key variants of the C and Σ properties. Therefore, by key indistinguishability, any attack against the VPDC trapdoor key variant properties of C (resp. Σ) would imply a corresponding attack contradicting the assumed (non trapdoor key variant) property of C (resp. Σ).

Proposition 1. *If a VPDC scheme* $\mathsf{VPDC} = (\mathsf{C}, \Sigma, \mathsf{CAddTd}, \mathsf{CDec})$ *satisfies key indistinguishability, then* C *(resp.* Σ*) satisfies the* VPDC *trapdoor key variants of correctness, hiding, and binding properties for* C *(respectively, the* VPDC *trapdoor key variants of completeness, (relaxed) soundness and zero-knowledge for* Σ*).*

In some applications, it is desirable to strengthen the binding requirement for the VDPC so it holds even against attackers that are given the partial decryption trapdoor key td (e.g. in our blockchain application as we do not want auditors to create fake proofs). We call this requirement *trapdoor-binding*.

Trapdoor-Binding. A VPDC is *(computationally) trapdoor-binding* if, for $(pp, \mathsf{td}) \leftarrow \mathsf{CKeygenTd}(1^\lambda)$, the following probability (over the randomness of PPT $\mathcal{A}$ and $\mathsf{CKeygen}$) is negligible

$$\Pr[(C, \mathsf{o}, \mathsf{o}') \leftarrow \mathcal{A}(pp, \mathsf{td}) : \mathsf{m}(\mathsf{o}) \neq \mathsf{m}'(\mathsf{o}) \wedge \mathsf{COpen}(C, \mathsf{o}) = \mathsf{COpen}(C, \mathsf{o}') = 1].$$

We capture the decryptability requirements for VPDC by the *Decryption Soundness* and *Decryption Feasibility* properties defined as follows.

Decryption Soundness. A VPDC scheme is said to satisfy *Decryption Soundness* if any PPT adversary wins the following `Exp:Soundness` game with $\mathrm{negl}(\lambda)$ probability.

1. $P := (ck^{\mathsf{td}}, \mathsf{td}, \mathcal{M}, \mathcal{R}) \leftarrow \mathsf{CKeygenTd}(1^\lambda)$
2. $(C, \pi) \leftarrow \mathcal{A}(P)$,
3. $b \leftarrow \mathsf{V}_{ck^{\mathsf{td}}}(C, \pi)$,
4. $\mu' \leftarrow \mathsf{CDec}_{\mathsf{td}}(C, \pi)$.

$\mathcal{A}$ wins the game if $b = 1$ and one of the following conditions holds

(i) There exists no opening o such that $\mathsf{COpen}(C, \mathsf{o}) = 1$, or

(ii) There exists an opening o such that $\mathsf{COpen}(C, \mathsf{o}) = 1$ and $\mu(\mathsf{o}) \neq \mu'$.

Decryption Feasibility. A VPDC scheme is said to satisfy *Decryption Feasibility* if, for any $\alpha \geq 1$ and any PPT adversary $\mathcal{A}$, if $b = 1$ in Step 3 of game `Exp:Soundness` above, the running time of $\mathsf{CDec}_{\mathsf{td}}(C, \pi)$ in Step 4 of game `Exp:Soundness` is at most $\alpha \cdot T_{\mathcal{A}} \cdot \mathrm{poly}(\lambda)$, except with probability $\leq \frac{1}{\alpha} + \mathrm{poly}\left(\frac{T_{\mathcal{A}}}{2^\lambda}\right)$, where $T_{\mathcal{A}}$ is the runtime of $\mathcal{A}$.

Remark 1 (Decryption Soundness). The decryption soundness property captures the informal requirement that it should be infeasible for an attacker to output a maliciously-created commitment and proof (C, π) that passes the V verification check, but where C cannot be decrypted into the correct decryptable message μ using the trapdoor decryption algorithm $\mathsf{CDec}_{\mathsf{td}}$. The latter may occur either because of the non-existence of a decryptable message opening (case i), or because of the existence multiple decryptable message openings that may cause a 'false accept' decryption error (case ii).

Remark 2 (Decryption Feasibility). The decryption feasibility requirement captures the property that the decryption algorithm does not run for 'too long'. Here, a too long decryption time corresponds to exceeding the attacker runtime by a super-polynomial factor. Jumping ahead to our construction, similarly

to [18], this attacker time corresponds to the number of queries made by the attacker to a certain random oracle. Such a runtime as given in our decryption feasibility definition (arising from our results generalizing those of [18]) is currently the best one can achieve for relaxed proofs, where the relaxation factor is unknown to the decryptor.

There are examples of VPDC-like constructions in the literature satisfying some but not all of our desired properties. For example, the proofs of plaintext knowledge in general are an example, where the commitment is an encryption scheme and $\mathcal{U} = \emptyset$. For a lattice-based construction, one may see the discussions in [18, Section 3.3] and [7]. The "extractable" commitment scheme in [12] is another (weaker) example, where the decryption soundness holds only against *honest* provers and decryption runtime is linear in $|\mathcal{D}|$. The main motivation for our VPDC notion is that we want the additional properties of succinctness (i.e., C should be compressing, which is not possible for encryption) and decryption feasibility and soundness against *cheating* provers. These properties are achieved by our concrete construction $\mathsf{VPDC}_{\mathsf{HMC}}$ (Section 5).

4 Generalized Decryption Feasibility for Relaxed Proofs

In this section, we study the decryption algorithms for relaxed NIZK proofs and show a general result on the Partial-Decryption Feasibility of any VPDC in which the underlying NIZK Σ is derived from a suitable interactive Sigma protocol Σ_{I} using the Fiat-Shamir (FS) transform. Our result generalizes previous results of [18]. The discussion is kept abstract in this section to preserve the generality of our results. Our concrete lattice-based instantiation of decryption algorithms is given in the next section.

The questions we focus on are as follows. If one is given a valid transcript $\mathsf{tr} = (C, w, x, z)$ for Σ and a trapdoor td that enables recovering a message from a *well-formed* commitment of the form $\bar{x}C$ for an *unknown* relaxed opening factor $\bar{x}$ (also known as a relaxation factor) and a known commitment C, how should one precisely design the overall decryption algorithm? Moreover, what is the *expected* number of iterations until the decryption algorithm terminates?

To answer these questions, we prove the general result in Theorem 1 below for the generic decryption algorithm given in Algorithm 1. This approach first allows us to put our decryption methodology into a general framework. Then, we identify the connections between the components of Algorithm 1 that must be satisfied so that the decryption runs in polynomial time. As the result can be applied to *any* suitable functions F, Rec, V' in Algorithm 1, we can use this result to analyze the run-time of different decryption methods. Besides the Partial-Decryption Feasibility, there is, of course, also the Partial-Decryption Soundness aspect that depends on the concrete instantiation of CDec, which will be analyzed in the next section.

Our Partial-Decryption Feasibility result applies to VPDCs in which the underlying NIZK Σ is derived via the FS transform from an interactive Sigma protocol with the following mild variant of the special soundness property, that

we call *existential special-soundness*. This property relaxes the standard PPT efficiency requirement for extractor $\mathcal{E}$, but requires that the extracted witness contains a component (relaxation factor $\bar{x} = x - x'$ in Schnorr-like protocols) depending on only the two input transcript challenges via a poly-time computable function F (the latter syntactic requirement is used in our decryption compatibility definition below). Therefore, the existential special-soundness is directly implied by the standard special-soundness property for a large class of known Schnorr-like Sigma protocols, in which $\mathcal{E}$ is efficient, and $F(x, x') = x - x'$.

Definition 3 (Existential Special-Soundness). *We say that a Sigma protocol* $\Sigma_{\mathsf{I}} = (\mathsf{K}_{\mathsf{I}}, \mathsf{P}_{\mathsf{I}}, \mathsf{V}_{\mathsf{I}})$ *for relations* $\mathsf{R}_{\mathsf{C},pp}, \mathsf{R}'_{\mathsf{C},pp}$ *(parameterised by a common reference string pp) with a challenge space* $\mathcal{C}$ *and public-private inputs* (C, o)*, satisfies* existential special-soundness *if the following holds.*

- ***Existential special-soundness:*** *There exists an* extractor $\mathcal{E}$ *and a deterministic poly-time algorithm* F *such that, given* $(pp, \sigma') \leftarrow \mathsf{K}_{\mathsf{I}}(1^\lambda)$ *and two accepting protocol transcripts* $\mathsf{tr} = (C, w, x, z)$ *and* $\mathsf{tr}' = (C, w, x', z')$ *with* $x \neq x'$*, computes an extracted witness of the form* $\bar{\mathsf{o}} = (\bar{x}, \bar{\mathsf{o}}')$ *where* $\bar{x} = F(x, x')$*, satisfying* $(C, \bar{\mathsf{o}}) \in \mathsf{R}'_{\mathsf{C},pp}$ *with probability* $1 - \mathrm{negl}(\lambda)$ *over the choice of* σ*.*

The following definition captures the properties of a decryption algorithm CDec that are sufficient to ensure it terminates in feasible time, if the underlying Sigma protocol Σ_{I} satisfies existential special-soundness.

Definition 4 (Compatible CDec). *Let* $\mathsf{VPDC} = (\mathsf{C}, \Sigma, \mathsf{CAddTd}, \mathsf{CDec})$ *with* Σ *a matching NIZK relaxed proof of opening for* C*, and* Σ *is obtained from a Sigma protocol* Σ_{I} *using the Fiat-Shamir transform.*

We say that CDec *is* compatible *with* Σ *if it satisfies the following properties:*

$\mathbf{P}_1$: CDec *has a structure as in Algorithm 1, where* Rec *is a PPT algorithm.*

$\mathbf{P}_2$: *On input* td *(generated by running* $(ck^{\mathsf{td}}, \mathsf{td}, \mathcal{M}, \mathcal{R}) \leftarrow \mathsf{CKeygenTd}(1^\lambda)$*) and any* tr*, if, for some loop iteration, Step 3 of* CDec *computes a "good"* $\bar{x}$ *(such that there exists an opening of the form* $\bar{\mathsf{o}} = (\bar{x}, \bar{\mathsf{o}}')$ *satisfying* $(C, \bar{\mathsf{o}}) \in \mathsf{R}'_{\mathsf{C},pp}$*), then* CDec *terminates in this loop iteration, i.e.* Rec *recovers a message* m' *deemed "valid" by* $V'(m') = 1$*.*

The function Rec in Alg. 1 is a procedure that recovers the message from a well-formed commitment and will be instantiated depending on our decryption method. One may imagine it being similar to the decryption of Regev encryption scheme. However, the Rec function always returns a message m' that may simply be useless. Therefore, there is an additional check V' to make sure that the given message is "valid" (where "valid" is protocol-dependent).

Theorem 1 below shows that the only task required to use our results is to design a compatible decryption algorithm as in Def. 4 (as existential special-soundness is implied by special-soundness). In essence, this task itself reduces to making sure that the message recovery algorithm Rec returns a "valid" message from any given "good" $\bar{x}$ for the relaxed relation $\mathsf{R}'_{\mathsf{C},pp}$. In the next section, we will show how our VPDC allows the recovery of the *same* message used to create the proof transcript.

Algorithm 1 $\mathsf{CDec}_{\mathsf{td}}(\mathsf{tr})$

INPUT: $\mathsf{tr} = (C, \pi = (w, x, z))$ Σ_{I} protocol transcript; td trapdoor
OUTPUT: (m', x') such that $V'(m') = 1$ for some validity check V'

loop
  $x' \xleftarrow{\$} \mathcal{C}$ ▷ Choose a random challenge
  $\bar{x} = F(x, x')$ ▷ $F(x, x') = x - x'$ for 2-sound FS proofs
  $m' = \mathsf{Rec}(\bar{x}, \mathsf{tr}, \mathsf{td})$ ▷ Tries to decrypt a well-formed commitment
  if $V'(m') = 1$ **then** ▷ Check if the recovered message is "valid"
    return (m', x')
  end if
end loop

To make it easier to read Theorem 1, let us interpret it in the case of 'Schnorr-like' FS proofs that work as follows. For a homomorphic commitment Com (we use additive notation as it is the case in the lattice setting), let $C = \mathsf{Com}(r)$ be an input commitment whose opening the prover wants to prove knowledge of. The prover computes a 'masking' commitment $w = \mathsf{Com}(\rho)$ for some masking value ρ. Then, she computes a challenge $x \leftarrow \mathcal{H}(pp, C, w)$, followed by a response $z = \rho + x \cdot r$, where r is the prover's witness. The verification V in this case checks $w + x \cdot C \stackrel{?}{=} \mathsf{Com}(z)$. It is easy to see from here that this proof has the '2-soundness' property, i.e., a *knowledge extractor* can extract a (relaxed) opening of C given two 'rewinded' accepting transcripts with distinct challenges. In particular, given accepting (C, w, x, z) and (C, w, x', z') with $x \neq x'$, we have $\bar{x}C = \mathsf{Com}(\bar{z})$ for $\bar{x} := x - x'$ and $\bar{z} := z - z'$. Therefore, the concrete functions in this case are $F(x, x') = x - x'$ and $(C, (\bar{x}, \bar{z})) \in \mathsf{R}'_{\mathsf{C},pp}$ iff $\bar{x}C = \mathsf{Com}(\bar{z})$, i.e., $\mathsf{R}'_{\mathsf{C},pp}$ in Theorem 1 corresponds to the *relaxed* commitment opening relation $\mathsf{COpen}(C, (\bar{x}, \bar{m}, \bar{r}))$ where $\bar{z} = (\bar{m}, \bar{r})$. Here, $(\bar{x}, \bar{z})$ serves as an *extracted* witness/opening for C. It is easy to see that the existential special-soundness property follows from the special-soundness of the 'Schnorr-like' protocol. Although in the setting of the Schnorr proof of knowledge of discrete-log [22], one may further recover an *exact* opening of u by computing $\bar{z}/\bar{x}$, this approach does not work in the lattice variants [16,17] as $\bar{z}/\bar{x}$ must be *short* (relative to the system modulus q), which cannot be guaranteed unless some costly measures are implemented[9].

Theorem 1. *Let* $\mathsf{VPDC} = (\mathsf{C}, \Sigma, \mathsf{CAddTd}, \mathsf{CDec})$ *with* Σ *a matching NIZK relaxed proof of opening for* C*, and* Σ *is obtained from a Sigma protocol* Σ_{I} *using the Fiat-Shamir transform with random oracle* $\mathcal{H} : \{0,1\}^* \to \mathcal{C}$*.*

If Σ_{I} *satisfies Existential Special Soundness (Def. 3),* CDec *is compatible with* Σ *(Def. 4) and* $|\mathcal{C}| \geq 2^{\lambda}$*, then* VPDC *satisfies Decryption Feasibility.*

Concretely, let $\hat{H}$ *and* $\hat{D}$ *be the random coins of* $\mathcal{H}$ *and* CDec*, respectively, and* T *be the number of loop iterations in the execution of* CDec *in Step 6 of game* `Exp:Soundness` *when* $b = 1$*. Then, for any* $\mathcal{A}$ *that makes at most* $q_H - 1$

[9] For example, *exact* lattice proofs (see, e.g., [7]) require around 40-50 KB in comparison to 2-3 KB for *relaxed* lattice proofs.

queries to $\mathcal{H}$ and any positive α,

$$\Pr_{\hat{H},\hat{D}}\left[T \geq \alpha \cdot q_H\right] \leq \frac{1}{\alpha} + 2 \cdot \sqrt{\frac{q_H}{\alpha \cdot |\mathcal{C}|}} + \frac{q_H}{|\mathcal{C}|}. \tag{2}$$

Proof (Theorem 1). The proof follows essentially the same blueprint as in the proof of [18, Lemma 3.2], but we show precisely where the properties in the theorem statement are needed.

Let pp be some public parameters. For a given $\mathsf{tr} = (C, w, x, z)$ of Σ_{I}, define the set of "good" challenges $\mathcal{G}_{\mathsf{tr}}$ as follows

$$\mathcal{G}_{\mathsf{tr}} = \{\, x' \in \mathcal{C} \;:\; \exists z' \;:\; \mathsf{V}_{\mathsf{I}}(C, w, x', z') = 1 \,\}. \tag{3}$$

Here, V_{I} denotes the verification algorithm of Σ_{I}. Let G be the event that $\mathcal{A}$ produces tr with $|\mathcal{G}_{\mathsf{tr}}| > f$ for $f = \left\lceil \sqrt{\frac{|\mathcal{C}|}{\alpha q_H}} \right\rceil$.

Claim: For any valid $\mathsf{tr} = (C, w, x, z)$, if CDec chooses x' with $x' \in \mathcal{G}_{\mathsf{tr}} \setminus \{x\}$, then CDec terminates.

The claim follows from the following facts. If the assumption of the claim holds, then there exist (C, w, x, z) (as the input) and (C, w, x', z') such that $\mathsf{V}_{\mathsf{I}}(C, w, x, z) = \mathsf{V}_{\mathsf{I}}(C, w, x', z') = 1$ by the definition of $\mathcal{G}_{\mathsf{tr}}$. Then, by the existential special soundness of Σ_{I}, there exists $(\bar{x}, \bar{\mathsf{o}}')$ such that $\bar{x} = F(x, x')$ and $(C, (\bar{x}, \mathsf{o}')) \in \mathsf{R}'_{\mathsf{C},pp}$. Now, by the property $\mathbf{P}_2$ of CDec, the claim follows.

As a result, the probability that CDec terminates in one iteration is at least $\frac{|\mathcal{G}_{\mathsf{tr}}|-1}{|\mathcal{C}|}$. Therefore, we have

$$\mathrm{Exp}_{\hat{D}}\left[\; T \mid \mathcal{A}^{\mathcal{H}} \text{ outputs } \mathsf{tr} \;\right] \leq \frac{|\mathcal{C}|}{|\mathcal{G}_{\mathsf{tr}}| - 1}, \quad \text{and also} \tag{4}$$

$$\mathrm{Exp}_{\hat{D}}\left[\; T \mid \mathcal{A}^{\mathcal{H}} \text{ outputs } \mathsf{tr} \;\wedge\; G \;\right] \leq \frac{|\mathcal{C}|}{f}. \tag{5}$$

We say that "$\mathcal{A}^{\mathcal{H}}$ outputs tr_i" if $\mathcal{A}^{\mathcal{H}}$ outputs $\mathsf{tr} = (C, w, x, z)$ such that the output of $\mathcal{A}$'s i-th random oracle query is x. As in [18], without loss of generality, we consider an adversary $\mathcal{A}$ that (1) makes q_H random oracle queries, (2) uses one of the random oracle outputs in his output transcript and (3) only makes random oracle queries for transcripts tr_i with $|\mathcal{G}_{\mathsf{tr}_i}| > f$ as we are conditioning on G. Then, similar to [18], we have the following

$$\mathrm{Exp}_{\hat{H},\hat{D}}\left[T \mid G\right] = \sum_{i=1}^{q_H} \Pr_{\hat{H}}\left[\mathcal{A}^{\mathcal{H}} \text{ outputs } \mathsf{tr}_i \mid G\right] \mathrm{Exp}_{\hat{D}}\left[T \mid \mathcal{A}^{\mathcal{H}} \text{ outputs } \mathsf{tr}_i \;\wedge\; G\right]. \tag{6}$$

For each random oracle query made by $\mathcal{A}$ for a transcript tr_i, the probability (over $\hat{H}$) that $\mathcal{A}$ outputs tr_i is at most the probability that the random oracle query output is in $\mathcal{G}_{\mathsf{tr}_i}$, as otherwise there exists no response z such that $\mathsf{V}_{\mathsf{I}}(C, w, x, z) = 1$. Therefore, each tr_i can be output with probability at most $\frac{|\mathcal{G}_{\mathsf{tr}_i}|}{|\mathcal{C}|}$. Then, using this fact and (4), we get

$$\mathrm{Exp}_{\hat{H},\hat{D}}\left[\; T \mid G \;\right] \leq \sum_{i=1}^{q_H} \frac{|\mathcal{G}_{\mathsf{tr}_i}|}{|\mathcal{C}|} \frac{|\mathcal{C}|}{|\mathcal{G}_{\mathsf{tr}_i}| - 1} \leq q_H \cdot \max_{i=1,\ldots,q_H} \left(\frac{|\mathcal{G}_{\mathsf{tr}_i}|}{|\mathcal{G}_{\mathsf{tr}_i}| - 1} \right) \leq \frac{q_H(f+1)}{f}.$$

For any random oracle query, the probability that $\mathcal{A}$ outputs a transcript with $|\mathcal{G}_{\mathsf{tr}}| \leq f$ is at most $f/|\mathcal{C}|$. Therefore, we have

$$\Pr_{\hat{H},\hat{D}}[\,\neg G\,] \leq \frac{f \cdot q_H}{|\mathcal{C}|}. \tag{7}$$

Using now Markov's inequality and (7), we get

$$\begin{aligned}\Pr_{\hat{H},\hat{D}}[T \geq \alpha q_H] &= \Pr_{\hat{H},\hat{D}}[T \geq \alpha q_H \mid G\,] \Pr_{\hat{H},\hat{D}}[\,G\,] + \Pr_{\hat{H},\hat{D}}[T \geq \alpha q_H \mid \neg G\,] \Pr_{\hat{H},\hat{D}}[\,\neg G\,] \\ &\leq \frac{\mathrm{Exp}_{\hat{H},\hat{D}}[\,T \mid G\,]}{\alpha \cdot q_H} + \Pr_{\hat{H},\hat{D}}[\,\neg G\,] \leq \frac{f+1}{\alpha \cdot f} + \frac{f \cdot q_H}{|\mathcal{C}|} = \frac{1}{\alpha} \cdot \left(1 + \frac{1}{f}\right) + \frac{f \cdot q_H}{|\mathcal{C}|}.\end{aligned}$$

Plugging in the value of $f = \left\lceil \sqrt{\frac{|\mathcal{C}|}{\alpha q_H}} \right\rceil$ proves the result. □

5 HMC-based VPDC from Lattices

5.1 Instantiation of (ordinary) HMC

We start by describing the (ordinary) Hashed-Message Commitment (HMC) scheme C underlying our lattice-based VPDC. Let n, m, q be positive integers with $m > n$. If we want to commit to a v_1-dimensional 'real' message over R_q for $v_1 \geq 1$ together with a v_2-dimensional auxiliary message for $v_2 \geq 0$, then HMC is instantiated as follows.

- $\mathsf{CKeygen}(1^\lambda)$: Sample $\boldsymbol{A} \leftarrow R_q^{n \times m}$, $\boldsymbol{B} \leftarrow R_q^{n \times v_1}$ and $\boldsymbol{C} \leftarrow R_q^{n \times v_2}$. Output $ck = \boldsymbol{G} = [\,\boldsymbol{A} \,\|\, \boldsymbol{B} \,\|\, \boldsymbol{C}\,] \in R_q^{n \times (m+v_1+v_2)}$, message space $\mathcal{M} = \mathcal{D} \times \mathcal{U}$ with $\mathcal{D} := \mathbb{S}_\alpha^{v_1}$ and $\mathcal{U} := \mathbb{S}_\beta^{v_2}$, and $\mathcal{R} := \mathbb{S}_\mathcal{B}^m$ for some $\alpha, \beta, \mathcal{B} \geq 1$.
- $\mathsf{Commit}_{ck}(\boldsymbol{m}, \boldsymbol{u})$: Sample $\boldsymbol{r} \leftarrow \mathbb{S}_\mathcal{B}^m$. Output C and $\mathsf{o} = (1, \boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r})$, where

$$C = \mathsf{Com}_{ck}(\boldsymbol{m}, \boldsymbol{u};\, \boldsymbol{r}) = \boldsymbol{G} \cdot (\boldsymbol{r}, \boldsymbol{m}, \boldsymbol{u})^\top = \boldsymbol{A} \cdot \boldsymbol{r} + \boldsymbol{B} \cdot \boldsymbol{m} + \boldsymbol{C} \cdot \boldsymbol{u}.$$

- $\mathsf{COpen}_{ck}(C, (y, \boldsymbol{m}', \boldsymbol{u}', \boldsymbol{r}'))$: If $yC = \mathsf{Com}_{ck}(y\boldsymbol{m}', \boldsymbol{u}';\, \boldsymbol{r}')$, $\|(\boldsymbol{r}', y\boldsymbol{m}', \boldsymbol{u}')\| \leq \gamma_{\mathrm{com}}$, and $\dim(\boldsymbol{m}') = v_1$, $\dim(\boldsymbol{u}') = v_2$ and $\dim(\boldsymbol{r}') = m$ over R_q, return 1. Otherwise, return 0.

One can easily observe that HMC is additively homomorphic. Moreover, note that the opening algorithm is relaxed, where an additional *relaxation factor* $y \in R_q$ is involved. This relaxation is needed to obtain *efficient* lattice-based ZKPs. For classical commitment schemes such as Pedersen commitment, the relaxation factor is always 1. The same is true for *honestly-created* lattice-based commitments. However, *efficient* lattice-based ZKPs do not always prove that this is the case. For example, for the *exact* proof of knowledge of a commitment opening with $y = 1$ as in [7], the proof length is more than 40 KB while a *relaxed* variant leads to a proof length of only a few KBs. Therefore, the relaxation factor can be a non-trivial value when created by a cheating prover (that still succeeds

in the ZKP verification). For HMC used within our VPDC below, we say that the trapdoor-binding property is *satisfied w.r.t. to the same relaxation factor* if the relaxation factors in o and o' in the binding definition in Section 3 are restricted to be the same. This same-relaxation trapdoor-binding property is sufficient for our applications as well as many prior ones, e.g., [9,10,12], and can be based on a harder variant of the MSIS problem than the general trapdoor-binding property (see Lemma 1).

We remark that in some applications, the COpen algorithm checks a slightly different relation than above, of the form $yC = \mathsf{Com}_{ck}(\boldsymbol{m}', \boldsymbol{u}'; \boldsymbol{r}')$, where the relaxation factor y does not multiply the decryptable message. However, in this paper, we need the stronger variant in the above definition. Despite the relaxation factor, HMC defined above is still (computationally) binding and hiding as we will discuss in Lemma 1.

5.2 Instantiation of NIZK

Our lattice-based VPDC can be instantiated with any suitable Schnorr-like lattice-based relaxed NIZK proofs of opening for HMC commitments derived from a Sigma protocol via the Fiat-Shamir transform, such as the one-shot relaxed binary proof protocols in [9,11,12]. For compatibility with such protocols, we define the challenge space

$$\mathcal{C}^d_{w,p} = \{x \in \mathbb{Z}[X] : \deg(x) < d \wedge \ \mathsf{HW}(x) = w \wedge \ \|x\|_\infty \leq p\}.$$

The same set is also defined in [9] and $|\mathcal{C}^d_{w,p}| = \binom{d}{w}(2p)^w$. Thus, given d, it is easy to set (w, p) such that $|\mathcal{C}^d_{w,p}| > 2^{256}$. Throughout the manuscript, we assume that (d, w, p) is set so that $|\mathcal{C}^d_{w,p}|$ is exponentially large. We also let $\Delta\mathcal{C}^d_{w,p}$ denote the set of differences of challenges in $\mathcal{C}^d_{w,p}$ except for the zero element. We design our VPDC to work with the following definition of relaxed "well-formedness" of a commitment relation. We refer the reader to Lemma 2 in the Sec. 6 for a concrete example of such a relaxed NIZK protocol Σ in our cryptocurrency protocol application.

Definition 5 (γ-valid commitment opening relation $\mathsf{R}'_{\mathsf{C},pp}$). *We say that $\mathsf{o} := (y, (\boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r}))$ is a γ-valid opening of a commitment C with a decryptable message space $\mathcal{D}$, denoted by $(C, \mathsf{o}) \in \mathsf{R}'_{\mathsf{C},pp}$, if the following holds:*

- $y \in \Delta\mathcal{C}^d_{w,p}$,
- $\boldsymbol{m} \in \mathcal{D}$,
- $yC = \mathsf{Com}_{ck}(y\boldsymbol{m}, \boldsymbol{u}; \boldsymbol{r})$,
- $\|(y\boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r})\| \leq \gamma$ *for* $\gamma \in \mathbb{R}^+$,
- $\dim(\boldsymbol{m}) = v_1$, $\dim(\boldsymbol{u}) = v_2$ *and* $\dim(\boldsymbol{r}) = m$ *over* R_q.

The above relation definition is very similar to COpen except that we additionally have the first two requirements. For our proof systems (as in Schnorr-like proofs), the commitment w of the Fiat-Shamir protocol as given in Section 4 is uniquely determined by the rest of the proof output. Hence, it need not be included in the non-interactive proof transcript and therefore its notation is omitted in the rest of the paper.

5.3 VPDC Trapdoor for HMC

Now, we present our gadget-based Regev-style VPDC trapdoor algorithm CAddTd for our lattice-based VPDC based on the HMC commitment described in Sec. 5.1.

Our trapdoor $\boldsymbol{s}$ is designed to allow partial decryption of the latter HMC commitment, i.e., to recover the decryptable binary message $\boldsymbol{m} \in \{0,1\}^{v_1}$ from the commitment $C = \mathsf{Com}_{ck}(\boldsymbol{m}, \boldsymbol{u}; \boldsymbol{r}) = \boldsymbol{G} \cdot \begin{pmatrix} \boldsymbol{r} \\ \boldsymbol{m} \\ \boldsymbol{u} \end{pmatrix}$, where $\boldsymbol{G} = [\, \boldsymbol{A} \,\|\, \boldsymbol{B} \,\|\, \boldsymbol{C} \,] \in R_q^{n \times (m+v_1+v_2)}$ is the commitment key matrix, and $\boldsymbol{r}$ and $\boldsymbol{u}$ are the short non-decryptable commitment randomness and auxiliary message, respectively. Our trapdoor $\boldsymbol{s}$ is embedded into the matrix $\boldsymbol{G}$ such that $\boldsymbol{s}^\top \cdot \boldsymbol{G} \approx [\boldsymbol{0}, \bar{t}\boldsymbol{g}^\top, \boldsymbol{0}] \in R_q^{m+v_1+v_2}$ (with the approximate equality up to a 'short' error vector) where $\bar{t}\boldsymbol{g}^\top$ (with $\bar{t} = \lfloor q/t \rfloor$) is a large gadget vector of the form $\bar{t} \cdot (1, 2, 2^2, \ldots, 2^{\tau-1}, X, X \cdot 2, \ldots, X \cdot 2^{\tau-1}, \ldots)$. This means that VPDC partial decryption of the commitment C can be carried out by computing $\boldsymbol{s}^\top C \approx \bar{t}\boldsymbol{g}^\top \boldsymbol{m}$. The (approximately) 0 entries in $\boldsymbol{s}^\top \cdot \boldsymbol{G}$ annihilate the non-decryptable $\boldsymbol{r}$ and $\boldsymbol{u}$ vectors in decryption (these vectors only contribute to the short error terms in the approximate equality), whereas the gadget vector $\bar{t}\boldsymbol{g}^\top$ entry of $\boldsymbol{s}^\top \cdot \boldsymbol{G}$ 'selects' the decryptable message $\boldsymbol{m}$ and compresses its dimension by reconstructing and packing groups of τ bits in $\boldsymbol{m}$ into the integer coefficients of $1, X, X^2, \ldots$ in the ring element $\boldsymbol{g}^\top \boldsymbol{m}$. To achieve the desired trapdoor condition $\boldsymbol{s}^\top [\, \boldsymbol{A} \,\|\, \boldsymbol{B} \,\|\, \boldsymbol{C} \,] \approx [\boldsymbol{0}, \bar{t}\boldsymbol{g}^\top, \boldsymbol{0}]$, for the 'selection' gadget, we embed a Regev-style LWE decryption '*gadget trapdoor*' into the last row $\boldsymbol{t}_B^\top$ of matrix $\boldsymbol{B}$, setting $\boldsymbol{t}_B^\top \approx \boldsymbol{s'}^\top \boldsymbol{B'} + \bar{t}\boldsymbol{g}^\top$, where $\boldsymbol{B'}$ consists of the top $n-1$ rows of $\boldsymbol{B}$ and $\boldsymbol{s'} \in R_q^{n-1}$ is random, and we use the form $\boldsymbol{s} = (-\boldsymbol{s'}, 1)$ for the trapdoor. For the annihilating 0 entries of $\boldsymbol{s}^\top \cdot \boldsymbol{G}$, we embed a Regev-style '*error trapdoor*' into the last rows $\boldsymbol{t}_A^\top$ (resp. $\boldsymbol{t}_C^\top$) of matrices $\boldsymbol{A}$ (resp. $\boldsymbol{C}$), setting them to $\approx \boldsymbol{s'}^\top \boldsymbol{A'}$ (resp. $\approx \boldsymbol{s'}^\top \boldsymbol{C'}$), where $\boldsymbol{A'}$ (resp. $\boldsymbol{C'}$) denote the top $n-1$ rows of $\boldsymbol{A}$ (resp. $\boldsymbol{C}$). Due to the errors in the above approximate equalities, to an attacker not knowing the secret trapdoor $\boldsymbol{s'}$, the trapdoor rows of matrix $\boldsymbol{G}$ are indistinguishable from uniformly random rows, assuming the hardness of the rank-$(n-1)$ M-LWE problem with respect to the secret $\boldsymbol{s'}$.

We now summarise our new gadget-based Regev-style HMC VPDC construction and start with instantiating CAddTd.

- $\mathsf{CAddTd}(ck)$: Let $ck = [\, \boldsymbol{A} \,\|\, \boldsymbol{B} \,\|\, \boldsymbol{C} \,] \in R_q^{n \times (m+v_1+v_2)}$ where $\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A'} \\ \boldsymbol{a}^\top \end{bmatrix}$ for $\boldsymbol{A'} \in R_q^{(n-1) \times m}$ and $\boldsymbol{a} \in R_q^m$, $\boldsymbol{B} = \begin{bmatrix} \boldsymbol{B'} \\ \boldsymbol{b}^\top \end{bmatrix}$ for $\boldsymbol{B'} \in R_q^{(n-1) \times v_1}$ and $\boldsymbol{b} \in R_q^{v_1}$, and $\boldsymbol{C} = \begin{bmatrix} \boldsymbol{C'} \\ \boldsymbol{c}^\top \end{bmatrix}$ for $\boldsymbol{C'} \in R_q^{(n-1) \times v_2}$ and $\boldsymbol{c} \in R_q^{v_2}$. Sample $\boldsymbol{s'} \leftarrow R_q^{n-1}$, $\boldsymbol{e}_0 \leftarrow \mathbb{S}_{\mathcal{B}_e}^m$, $\boldsymbol{e}_1 \leftarrow \mathbb{S}_{\mathcal{B}_e}^{v_1}$, and $\boldsymbol{e}_2 \leftarrow \mathbb{S}_{\mathcal{B}_e}^{v_2}$, and set $\boldsymbol{A}^{\mathsf{td}} = \begin{bmatrix} \boldsymbol{A'} \\ \boldsymbol{t}_0^\top \end{bmatrix}$, $\boldsymbol{B}^{\mathsf{td}} = \begin{bmatrix} \boldsymbol{B'} \\ \boldsymbol{t}_1^\top \end{bmatrix}$, $\boldsymbol{C}^{\mathsf{td}} = \begin{bmatrix} \boldsymbol{C'} \\ \boldsymbol{t}_2^\top \end{bmatrix}$ where $\boldsymbol{t}_0 = \boldsymbol{A'}^\top \boldsymbol{s'} + \boldsymbol{e}_0$, $\boldsymbol{t}_1 = \boldsymbol{B'}^\top \boldsymbol{s'} + \boldsymbol{e}_1 + \bar{t}\boldsymbol{g}$ and $\boldsymbol{t}_2 = \boldsymbol{C'}^\top \boldsymbol{s'} + \boldsymbol{e}_2$, with $\boldsymbol{g}^\top := (2^0 X^0, \ldots, 2^{\tau-1} X^0, 2^0 X^1, \ldots, 2^{\tau-1} X^1, \ldots, 2^0 X^{d'}, \ldots, 2^{\ell-1} X^{d'}) \in R^{v_1}$,

$\tau := \lceil \frac{v_1}{d} \rceil$, $d' := \lfloor \frac{v_1}{\tau} \rfloor$ (note that $d' \leq d$) and $\ell := v_1 \bmod \tau \in \{0, \ldots, \tau - 1\}$.
Output $(ck^{\mathsf{td}}, \mathsf{td}) = ([\, \boldsymbol{A}^{\mathsf{td}} \,\|\, \boldsymbol{B}^{\mathsf{td}} \,\|\, \boldsymbol{C}^{\mathsf{td}} \,], \boldsymbol{s})$ where $\boldsymbol{s} = \begin{pmatrix} -\boldsymbol{s}' \\ 1 \end{pmatrix}$.

The following lemma follows from the hiding/binding properties of standard HMC commitments, and the M-LWE based key indistinguishability property of CAddTd. The proof is given in the full version of this paper.

Lemma 1. *Let the ring R_q split into s fields $\mathbb{F}_{p_1}, \ldots, \mathbb{F}_{p_s}$ with $p = \min\{p_1, \ldots, p_s\}$. If $\frac{n \cdot s}{p^{m-n+1}}$ is negligible, then HMC under 'trapdoored' commitment keys as output by CAddTd defined above is*

- *correct if $\gamma_{com} \geq \sqrt{\mathcal{B}^2 md + (\gamma_y \alpha d)^2 v_1 + \beta^2 v_2 d}$,*
- *computationally trapdoor γ_{com}-binding with respect to the same relaxation factor (resp. γ_{com}-binding) if M-SIS$_{n-1,m+v_1+v_2,q,2\gamma_{\mathrm{com}}}$ is hard (resp. if M-SIS$_{n-1,m+v_1+v_2,q,2\sqrt{d}\gamma_Y \cdot \gamma_{\mathrm{com}}}$ is hard, where $\gamma_Y := max_{y \in Y} \|y\|$, and Y is the set of valid relaxation factors accepted by COpen; for our VPDC, it suffices to use $Y := \Delta\mathcal{C}^d_{w,p}$ as in Def. 5).*
- *computationally hiding if M-LWE$_{m-n,m,q,\mathcal{B}}$ and M-LWE$_{n-1,m+v_1+v_2,q,\mathcal{B}}$ problems are hard.*

Additionally, if M-LWE$_{m-n,m,q,\mathcal{B}}$ and M-LWE$_{n-1,m+v_1+v_2,q,\mathcal{B}}$ problems are hard, any commitment vector is computationally indistinguishable from a uniformly random element in R_q^n.

Note that there are two main differences in Lemma 1 compared to the assumptions required for standard HMC (see [12, Lemma 2.3], [6, Lemma 3.4]): (i) the module rank of M-SIS is reduced by 1 (from n to $n-1$), and (ii) the hardness of M-LWE$_{n-1,m+v_1+v_2,q,\mathcal{B}}$ is additionally required. As mentioned before, binding w.r.t. the same relaxation factor is sufficient for many applications (including ours) since the reduction creates a challenge commitment with a known *exact* opening (i.e., $y = 1$) and recovers another *relaxed* opening by rewinding the adversary. The former exact opening can be multiplied by the relaxation factor of the latter to solve an M-SIS problem.

5.4 Gadget-based Regev-style Decryption for HMC

We now present the decryption algorithm CDecGR for our lattice-based VPDC. When a commitment key with a trapdoor is used to generate a proof, the ZKPs we use prove knowledge of an opening $(y, \boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r})$ of a commitment C such that

$$yC = \mathsf{Com}_{ck^{\mathsf{td}}}(y\boldsymbol{m}, \boldsymbol{u};\, \boldsymbol{r}) = \boldsymbol{A}^{\mathsf{td}}\boldsymbol{r} + \boldsymbol{B}^{\mathsf{td}} y\boldsymbol{m} + \boldsymbol{C}^{\mathsf{td}}\boldsymbol{u}. \tag{8}$$

Note that the opening message is also multiplied by the relaxation factor y. From here, we can try to eliminate the randomness $\boldsymbol{r}$ and the auxiliary message $\boldsymbol{u}$ by multiplying both sides by the secret trapdoor $\boldsymbol{s}$. However, the decryptor does not know what y is. For an honest user, we simply have $y = 1$, but for adversarially-generated proofs, that may not be the case. Thankfully, we can

Algorithm 2 $\mathsf{CDecGR}(C, x, \mathsf{td}, v_1)$

INPUT: a commitment $C \in R_q^n$; a challenge $x \in \mathcal{C}_{w,p}^d$; trapdoor $\mathsf{td} = \boldsymbol{s} \in R_q^n$; the dimension v_1 such that $\mathcal{D} = \{0,1\}^{v_1}$
OUTPUT: $(\boldsymbol{m}', x') \in \mathcal{D} \times \mathcal{C}_{w,p}^d$

loop
  $x' \leftarrow \mathcal{C}_{w,p}^d$
  $y' = x - x'$ ▷ $y' = 1$ is assumed to be tried first
  $C' = \langle \boldsymbol{s}, y'C \rangle$
  $C'' = \mathsf{Rnd}_{\bar{t}}(C')$ where $\bar{t} = \lfloor q/t \rfloor$
  $\bar{m}' = (\bar{t})^{-1} \cdot C'' \in R$ ▷ Note that C'' is a multiple of $\bar{t}$ in R
  $m'' = (y')^{-1} \cdot \bar{m}' \in R_t$ ▷ If y' is not invertible in R_t, restart from Step 2
  $\boldsymbol{m}' = \mathsf{BD}_{\tau,v_1}(m'')$
  $e' = C' - C''$
  if $(\|e'\|_\infty < \|e\|_{bnd,\infty})$ **and** $(m'' \in [0, \ldots, 2^\tau - 1]^d)$ **then**
    return $(\boldsymbol{m}', x')$
  end if
end loop

use our new results from Section 4 to overcome this problem. Let us first present the full procedure in Algorithm 2. In this algorithm, the decrypted message is encoded as an element of R_t for some positive integer t, and we define the integer $\bar{t} := \lfloor q/t \rfloor$. We also use the following two functions. The function $\mathsf{BD}_{\tau,v_1}(m'')$ performs bit decomposition of the coefficients of the R_t-encoded message $m'' = m''_0 + m''_1 X + \cdots + m''_{d-1} X^{d-1}$ and returns the resulting binary vector message $\boldsymbol{m}' = (m'_0, \ldots, m'_{v_1-1}) \in \{0,1\}^{v_1}$. Namely, for $j \in \{0, \ldots, v_1 - 1\}$, it sets m'_j to the k-th bit of the coefficient $m''_{\lfloor j/\tau \rfloor}$ where $k := j \bmod \tau$. The function $\mathsf{Rnd}_{\bar{t}}(C')$ rounds each coefficient of $C' \in R$ to the nearest integer multiple of $\bar{t}$.

As mentioned in Sec. 4, an important task is to prove that the message returned by the decryption algorithm (Alg. 2) is "valid". We prove this in Theorem 2 below so that, for a commitment C with a valid NIZK relaxed proof of opening and a sufficiently large q, the message output by Alg. 2 is the same as the one used to generate the commitment C. In the theorem below, we show the decryption feasibility (which relies on the results from Sec. 4) and also the decryption soundness of our construction.

Theorem 2 (HMC Decryption). *Let* $\mathsf{VPDC}_{\mathsf{HMC}} = (\mathsf{C}, \Sigma, \mathsf{CAddTd}, \mathsf{CDecGR})$ *denote our lattice-based VPDC construction with HMC commitment scheme* C*,* Σ *a matching NIZK relaxed proof of* γ*-valid opening relation* $\mathsf{R}'_{\mathsf{C},pp}$ *as in Def. 5, with* $\mathcal{D} := \{0,1\}^{v_1}$*. Suppose that* Σ *is obtained from a Sigma protocol* Σ_{I} *using the Fiat-Shamir transform with random oracle* $\mathcal{H} : \{0,1\}^* \to \mathcal{C}$*,* Σ_{I} *satisfies Existential Special Soundness (Def. 3), that for any fixed* $x \in \mathcal{C}_{w,p}^d$*,* $x - x' \in \Delta\mathcal{C}_{w,p}^d$ *is invertible in* R_t *except with negligible probability* p_{ni} *over the uniformly random choice of* $x' \in \mathcal{C}_{w,p}^d$ *and* $|\mathcal{C}_{w,p}^d| \geq 2^\lambda$*.*

For an adversary $\mathcal{A}$ against soundness game $\texttt{Exp:Soundness}$ *making $q_H - 1$ queries to its random oracle, let*

$$\|e\|_{bnd,\infty} := \sqrt{(m+v_1+v_2)d}\mathcal{B}_e\gamma + 2pw(2^\tau - 1) + t/2. \tag{9}$$

Suppose that $t \geq 2^\tau$ and

$$\bar{t} := \lfloor q/t \rfloor > 4pw\|e\|_{bnd,\infty} + t(1/2 + 2pw). \tag{10}$$

Then the following holds:

1. ***Decryption Feasibility:*** *The scheme $\mathsf{VPDC_{HMC}}$ satisfies Decryption Feasibility. Concretely, the number of iterations T of the loop over x' in CDecGR is upper bounded by $\alpha \cdot q_H$, except with probability at most $\frac{1}{\alpha} + 2\sqrt{\frac{q_H}{\alpha \cdot |\mathcal{C}^d_{w,p}|}} + \frac{q_H}{|\mathcal{C}^d_{w,p}|} + \alpha q_H p_{\text{ni}}$.*
2. ***Decryption soundness:*** *The scheme $\mathsf{VPDC_{HMC}}$ satisfies Decryption Soundness. Concretely, we have*

$$Adv_{\texttt{Exp:Soundness}}(\mathcal{A}) \leq q_H/|\mathcal{C}^d_{w,p}|.$$

Proof (Thm. 2). **Decryption Feasibility:** To prove the run-time claim, we apply Theorem 1. For this, we need to show that the assumptions of Theorem 1 are satisfied. First, by our assumption on Σ_{I}, the existential special-soundness property is satisfied. For the compatibility of CDecGR and Σ_{I}, the property $\mathbf{P}_1$ is satisfied by structure of the algorithm CDecGR. For the property $\mathbf{P}_2$, we show that if y' in some iteration of the loop in Step 3 of Algorithm CDecGR is 'good' in the sense that y' is invertible in R_t and there exists a γ-valid opening $(y', \boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r})$ of C as in Def. 5, then decryption will terminate and return $\boldsymbol{m}' = \boldsymbol{m}$. Let us denote by E_0 the bad event that y' is not invertible in R_t. We first observe that E_0 occurs with negligible probability, i.e. $\Pr[E_0] \leq \alpha q_H p_{\text{ni}}$ over at most αq_H iterations of CDecGR, since at each iteration $y' = x - x'$ where x' sampled uniformly from $\mathcal{C}^d_{w,p}$ independently of x. Now we show that $\mathbf{P}_2$ holds if E_0 does not occur. Indeed, by γ-validity of $(y', \boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r})$, we have $y'C = \boldsymbol{A}^{\mathsf{td}}\boldsymbol{r} + \boldsymbol{B}^{\mathsf{td}}y'\boldsymbol{m} + \boldsymbol{C}^{\mathsf{td}}\boldsymbol{u}$ and $\boldsymbol{m} \in \{0,1\}^{v_1}$.

Multiplying the γ-valid relation by $\boldsymbol{s}^\top$, defining $\langle \boldsymbol{e}_0, \boldsymbol{r}\rangle + \langle \boldsymbol{e}_1, y'\boldsymbol{m}\rangle + \langle \boldsymbol{e}_2, \boldsymbol{u}\rangle := e$, and using $\boldsymbol{s}^\top \cdot \boldsymbol{A}^{\mathsf{td}} = \boldsymbol{e}_0^\top$, $\boldsymbol{s}^\top \cdot \boldsymbol{B}^{\mathsf{td}} = \bar{t}\boldsymbol{g}^\top + \boldsymbol{e}_1^\top$ and $\boldsymbol{s}^\top \cdot \boldsymbol{C}^{\mathsf{td}} = \boldsymbol{e}_2^\top$, we have $\langle \boldsymbol{s}, y'C\rangle = \bar{t}y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle + e$ over R_q. Writing $y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle = (y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle \bmod t) + t\lfloor \frac{y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle}{t} \rceil$, we get the following equality over R_q:

$$\langle \boldsymbol{s}, y'C\rangle = \bar{t}(y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle \bmod t) + e + \tilde{e}, \tag{11}$$

where $\tilde{e} := \bar{t}t \cdot \lfloor \frac{y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle}{t} \rceil \bmod q$. We have $\|\bar{t}(y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle \bmod t)\|_\infty \leq \lfloor q/t \rfloor (t-1)/2 \leq q/2 - \bar{t}/2$. Hence, if $\|e + \tilde{e}\|_\infty < \bar{t}/2$, there is no wraparound mod q on the right hand side of (11), and since $\bar{t}y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle$ is a multiple of $\bar{t}$ in R, the rounded polynomial $C'' = \mathsf{Rnd}_{\bar{t}}(C')$ (recall $C' = \langle \boldsymbol{s}, y'C\rangle \bmod q$) will be equal to $\bar{t}(y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle \bmod t)$ and decryption will succeed and return $\boldsymbol{m}$. It remains to show that $\|e + \tilde{e}\|_\infty < \bar{t}/2$. By the Schwartz inequality, $\|e\|_\infty \leq$

$\|(\boldsymbol{e}_0, \boldsymbol{e}_1, \boldsymbol{e}_2)\| \cdot \|(\boldsymbol{r}, y\boldsymbol{m}, \boldsymbol{u})\| \leq \sqrt{(m + v_1 + v_2)d}\mathcal{B}_{\mathrm{e}}\gamma$ using $\|(\boldsymbol{r}, y\boldsymbol{m}, \boldsymbol{u})\| \leq \gamma$ by γ-validity. Also, writing $\bar{t} = q/t - \epsilon$ for $0 \leq \epsilon < 1$, we have $\|\tilde{e}\|_\infty = \|(q/t - \epsilon)t \cdot \lfloor \frac{y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle}{t} \rceil \bmod q\|_\infty = \|\epsilon t \cdot \lfloor \frac{y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle}{t} \rceil\|_\infty \leq \|t \lfloor \frac{y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle}{t} \rceil\|_\infty \leq t/2 + \|y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle\|_\infty$, and $\|y'\langle \boldsymbol{g}, \boldsymbol{m}\rangle\|_\infty \leq \|y'\|_1 \|\langle \boldsymbol{g}, \boldsymbol{m}\rangle\|_\infty \leq (2pw)(2^\tau - 1)$ using $\|y'\|_1 \leq 2pw$ and $\|\langle \boldsymbol{g}, \boldsymbol{m}\rangle\|_\infty \leq 2^\tau - 1$ since $\boldsymbol{m} \in \{0,1\}^{v_1}$. Overall, we have $\|e + \tilde{e}\|_\infty \leq \|e\|_\infty + \|\tilde{e}\|_\infty \leq \sqrt{(m + v_1 + v_2)d}\mathcal{B}_{\mathrm{e}}\gamma + (2pw)(2^\tau - 1) + t/2 := \|e\|_{bnd,\infty}$, which is less than $\bar{t}/2$ by condition (10), as required.

Decryption Soundness: To show the decryption soundness claim, let E_1 denote the event that $\mathcal{A}$ wins and case (i) in `Exp:Soundness` occurs, i.e., $\mathsf{V}(C, x, \boldsymbol{z}) = 1$ but a γ-valid opening $(C, y, (\boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r}))$ of C with $y = x - x''$ and $x'' \in \mathcal{C}^d_{w,p}$ does *not* exist. Similarly, let E_2 be the event that $\mathcal{A}$ wins and case (ii) in `Exp:Soundness` occurs, i.e., $\mathsf{V}(C, x, \boldsymbol{z}) = 1$ and a γ-valid opening $(C, y, (\boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r}))$ of C with $y = x - x''$ and $x'' \in \mathcal{C}^d_{w,p}$ exists, but CDecGR returns the wrong message $\boldsymbol{m'} \neq \boldsymbol{m}$. We show that $\Pr[E_1] + \Pr[E_2] \leq \frac{q_H}{|\mathcal{C}^d_{w,p}|}$.

We first claim that $\Pr[E_1] \leq \frac{q_H}{|\mathcal{C}^d_{w,p}|}$. Indeed, for each $\mathcal{H}$-query of $\mathcal{A}$ of the form $(pp, C, \cdot)$, we have that for any query answer $x' \neq x$, there does not exist a $\boldsymbol{z}'$ such that $\mathsf{V}(C, x', \boldsymbol{z}') = 1$ (otherwise, by existential special-soundness of the protocol $\boldsymbol{\Sigma}_{\mathsf{I}}$, a γ-valid opening $(C, y, (\boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r}))$ of C with $y = x - x'$ would exist, a contradiction with E_1). It follows that $\Pr[E_1]$ is upper bounded by the probability that $\mathcal{A}$ receives the special challenge x for which a $\boldsymbol{z}$ exists in one of the $\leq q_H$ queries to $\mathcal{H}$. Since the special challenge is returned with probability $1/|\mathcal{C}^d_{w,p}|$ in each query, the claimed bound on $\Pr[E_1]$ follows.

Next, we claim that $\Pr[E_2] = 0$. On the one hand, if E_2 occurs, then the existence of the γ-valid opening $(C, y, (\boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r}))$ of C with $y = x - x''$ means that $yC = \boldsymbol{A}^{\mathsf{td}}\boldsymbol{r} + \boldsymbol{B}^{\mathsf{td}} y\boldsymbol{m} + \boldsymbol{C}^{\mathsf{td}}\boldsymbol{u}$. Similarly to (11), multiplying the latter by $\boldsymbol{s}^\top$ gives us the following relation over R_q:

$$\langle \boldsymbol{s}, yC\rangle = \bar{t}(y\langle \boldsymbol{g}, \boldsymbol{m}\rangle \bmod t) + e + \tilde{e}, \tag{12}$$

where $e := \langle \boldsymbol{e}_0, \boldsymbol{r}\rangle + \langle \boldsymbol{e}_1, y\boldsymbol{m}\rangle + \langle \boldsymbol{e}_2, \boldsymbol{u}\rangle$ and $\tilde{e} := \bar{t}t \cdot \lfloor \frac{y\langle \boldsymbol{g}, \boldsymbol{m}\rangle}{t} \rceil \bmod q$. The same bound $\|e\|_{bnd,\infty}$ on $\|e + \tilde{e}\|$ applies by the same argument as in the run-time proof, based on Schwartz inequality. On the other hand, let $y' = x - x'$ be the value chosen in the iteration of the loop in CDecGR for which the message $\boldsymbol{m'}$ is returned. Then $\bar{m}' = \bar{t}(y'\langle \boldsymbol{g}, \boldsymbol{m'}\rangle \bmod t) \in R$, and we get the following relation over R_q:

$$\langle \boldsymbol{s}, yC\rangle = \bar{t}(y'\langle \boldsymbol{g}, \boldsymbol{m'}\rangle \bmod t) + e', \tag{13}$$

where $\|e'\|_\infty < \|e\|_{bnd,\infty}$ by the decryption check of CDecGR.

We now multiply (12) by y' and subtract (13) multiplied by y. Let $b_1 := y'(y\langle \boldsymbol{g}, \boldsymbol{m}\rangle \bmod t) \in R$, $b_2 := y(y'\langle \boldsymbol{g}, \boldsymbol{m'}\rangle \bmod t) \in R$. Note that $b_1 - b_2 = y'y\langle \boldsymbol{g}, \boldsymbol{m} - \boldsymbol{m'}\rangle \bmod t$. Writing $b_1 - b_2 = (y'y\langle \boldsymbol{g}, \boldsymbol{m} - \boldsymbol{m'}\rangle \bmod t) + t\lfloor \frac{b_1 - b_2}{t} \rceil$ gives the following relation over R_q:

$$\bar{t}(y'y\langle \boldsymbol{g}, \boldsymbol{m} - \boldsymbol{m'}\rangle \bmod t) = y'(e + \tilde{e}) - ye' - \tilde{e}', \tag{14}$$

where $\tilde{e}' := \bar{t}t\lfloor\frac{b_1-b_2}{t}\rceil \bmod q$. We claim that the relation (14) leads to a contradiction, so that $\Pr[E_2] = 0$. To see this, first observe that the relation actually holds over R, not just R_q. Indeed, there is no wraparound mod q in the left hand side of (14), since the left hand side norm is at most $\|\lfloor q/t \rfloor \cdot t/2\|_\infty < q/2$. Since $\boldsymbol{m} - \boldsymbol{m'} \neq 0$, and $\boldsymbol{m} - \boldsymbol{m}' \in \{-1, 0, 1\}^{v_1}$, we have $\|\langle \boldsymbol{g}, \boldsymbol{m} - \boldsymbol{m}' \rangle\|_\infty \leq 2^\tau - 1 < t$ so $\langle \boldsymbol{g}, \boldsymbol{m} - \boldsymbol{m}' \rangle \neq 0 \bmod t$. Note that y is non-zero in R (since γ-validity of $(C, y, (\boldsymbol{m}, \boldsymbol{u}, \boldsymbol{r}))$ implies $y \in \Delta\mathcal{C}^d_{w,p}$) and y' is also non-zero in R (since it caused termination of CDecGR and hence is invertible in R_t) and R is an integral domain, the left hand side of (14) is a non-zero multiple of $\bar{t}$ in R. But the norm of the error terms on the right-hand side of (14) is bounded as follows. First, $\|y'(e + \tilde{e}) - ye'\|_\infty < \|y'\|_1\|e + \tilde{e}\|_\infty + \|y\|_1\|e'\|_\infty \leq 4pw\|e\|_{bnd,\infty}$ using $\|y'\|_1 \leq 2pw$, $\|y\|_1 \leq 2pw$, $\|e + \tilde{e}\|_\infty < \|e\|_{bnd,\infty}$ and $\|e'\|_\infty < \|e\|_{bnd,\infty}$. Also, $\|\tilde{e}'\|_\infty = \|\bar{t}t\lfloor\frac{b_1-b_2}{t}\rceil \bmod q\|_\infty \leq t\lfloor\frac{b_1-b_2}{t}\rceil \leq t(\frac{1}{2}+2pw)$ using $|\bar{t}t \mod q| < t$, and $\|\lfloor\frac{b_1-b_2}{t}\rceil\|_\infty \leq \frac{1}{2} + 2pw$ using $\|b_1\|_\infty \leq \|y'\|_1\|y\langle \boldsymbol{g}, \boldsymbol{m}\rangle \bmod t\|_\infty \leq (2pw)(t/2) \leq pw$, and similarly, $\|b_2\|_\infty \leq pw$. Overall, the norm of the right-hand side of (14) is bounded as $\|y'(e + \tilde{e}) - ye' - \tilde{e}'\|_\infty < 4pw\|e\|_{bnd,\infty} + t(\frac{1}{2} + 2pw)$, which is smaller than $\bar{t}$ by condition (10). So, the left-hand side cannot be a non-zero multiple of $\bar{t}$ in R, implying the claimed contradiction. This completes the proof that $\Pr[E_2] = 0$ and the claimed soundness bound. □

5.5 Generalized Decryption

Our gadget-based Regev-style decryption trapdoor presented in the previous section, which handles a binary decryptable message space $\mathcal{D} = \{0, 1\}^{v_1}$, can be readily generalised to handle more general decryptable message spaces $\mathcal{D} = (\{0, \ldots, \beta-1\}[X]^{<\delta})^{v_1}$ whose coordinates are polynomials in the ring R of degree $< \delta$ with β-ary coefficients for some positive integers $\delta, \beta > 1$. This generalisation can naturally be achieved via the appropriate generalisation of the reconstruction gadget vector $\boldsymbol{g}$, by setting

$$\boldsymbol{g}^\top := (\beta^0 X^0, \ldots, \beta^{\tau-1} X^0, \beta^0 X^\delta, \ldots, \beta^{\tau-1} X^\delta, \ldots, \beta^0 X^{d'\delta}, \ldots, \beta^{\ell-1} X^{d'\delta}) \in R^{v_1},$$

with $\tau := \left\lceil \frac{v_1}{\lfloor d/\delta \rfloor} \right\rceil$, $d' := \lfloor \frac{v_1}{\tau} \rfloor \leq \lfloor d/\delta \rfloor$. The decryption soundness result, Thm 2, directly extends to this generalised case with the term 2^τ replaced by β^τ.

5.6 Succinctness of Our HMC-based VPDC

An HMC commitment C as defined in Section 5.1 costs $nd \log q$ bits, i.e.,

$$\mathsf{bitlen}(C) = nd \log q. \tag{15}$$

We show that we can choose parameters such that succinctness of the VPDC is satisfied, i.e., $\mathsf{bitlen}(C) = \log^{O(1)}(\mathsf{bitlen}(\boldsymbol{u}))$, where $\mathsf{bitlen}(\boldsymbol{u})$ is the bit length of *honestly generated* auxiliary messages, assuming the following very mild assumptions: (i) $v_1/d = O(\log \lambda)$ and (ii) $\gamma = (\lambda\|\boldsymbol{u}\|)^{O(1)}$, a condition that is typically satisfied by the soundness extractor of the associated ZKP. The auxiliary message

space is defined as $\mathcal{U} := \mathbb{S}_{\mathcal{B}}^{v_2}$ with $\mathcal{B} = \lambda^{O(1)}$. Then $\mathsf{bitlen}(\boldsymbol{u}) := dv_2 \log(2\mathcal{B})$. Set $d = \Theta(\lambda)$, $v_1 = \Theta(\lambda)$ (with $v_1/d = O(\log \lambda)$), $v_2 = \lambda^{O(1)}$, $p = O(1)$, $w = \Theta(\lambda)$ (so that $|\mathcal{C}_{w,p}^{d}| \geq (d/w-1)^w \geq 2^{\lambda}$), $\mathcal{B}_e = \Theta(1)$ and $n, m = \Theta(\log \gamma)$. As a result, we have $t = O(2^{\tau}) = O(2^{v_1/d}) = \lambda^{O(1)}$.

In general, we require two conditions to be satisfied: decryption soundness requirements and M-SIS security requirements (note that M-LWE security affects the number of columns of the commitment matrix, and thus not the commitment size). Let us analyze these two aspects.

(1) Partial-Decryption Soundness and Feasibility (Thm. 2):

$$q > \Omega(tpw\gamma B_e \sqrt{d(m+v_1+v_2)} + t^2 pw) = \Omega(\lambda^{O(1)} \gamma \log \gamma). \tag{16}$$

(2) M-SIS security (Lemma 1): The hardness of $\text{M-SIS}_{n,m,q,\beta_{\text{SIS}}}$ requires (see [9, Section 1.2]):

$$nd \log q \geq \Omega(\lambda \log^2(\beta_{\text{SIS}})) = \Omega(\lambda \log^2(\gamma)). \tag{17}$$

Note that $\beta_{\text{SIS}} = 2\gamma$ as given in Lemma 1. We can satisfy both conditions with some $\log q = \Theta(\log(\lambda\gamma))$ (ignoring $\log\log\gamma$). With this choice, we get commitment length $\mathsf{bitlen}(C) = nd \log q = \Theta(\lambda \log^2(\gamma))$. To show it is polylog in $\mathsf{bitlen}(\boldsymbol{u})$, it suffices to show that $\log \gamma = \log^{O(1)}(\mathsf{bitlen}(\boldsymbol{u}))$. Assuming that $\gamma = (\lambda \|\boldsymbol{u}\|)^{O(1)}$, we have $\log \gamma \leq O(\log(\lambda\|\boldsymbol{u}\|)) = O(\log \lambda + \log(dv_2) + \log \mathcal{B}) = \log^{O(1)}(\mathsf{bitlen}(\boldsymbol{u}))$, as required.

6 Extending MatRiCT to Auditable Setting

Having dealt with the core task of constructing and analysing a VPDC, we now explain how our VPDC construction can be applied to extend a privacy-preserving confidential transaction blockchain protocol MatRiCT [12] to the auditable setting. Unlike the auditability feature of the original MatRiCT protocol, where auditing may fail against adversarially-created transactions, our auditable MatRiCT variant, called MatRiCT-Au, takes advantage of our VPDC to efficiently provide auditability soundness guarantees against adverserial transactions. More specifically, we show that only minor modifications to MatRiCT are sufficient to add the auditability feature. As the whole MatRiCT protocol is quite involved, in this section, we only briefly review MatRiCT, focusing on the specific parts of the protocol which we modify.

MatRiCT follows the blueprint of RingCT-like [20] private blockchain payment protocols, in which there are two main entities: (i) *spenders/payers*, who create transactions together with a proof of validity, and (ii) *verifiers*, who check that the proof and transaction is valid. The goal of the private payment protocol is to enable users to conduct transactions on blockchain while hiding sensitive transaction information such as the payer/payee identities and the payment amount. Once such information is concealed from the verifiers, it gets harder to validate transactions as we cannot, for example, simply check that the transaction amount is positive and the total balance of the transaction is zero (i.e., the

amount spent equals the amount received). To this end, the payers create a NIZK proof showing that they are not creating an invalid transaction, for example, by proving that the balance is preserved.

To hide the payer identity, MatRiCT makes use of a 1-out-of-N NIZK proof (or a ring signature), where the identity of the real payer is hidden within a set of possible payers. This involves committing to the unary representation of an index $\ell \in [0, N-1]$. To hide the payment amount, a commitment to the bits of the transaction amount is used. In fact, the bits representing the user index and those representing the transaction amount are committed in a single commitment. To enable an authority to recover these two critical data pieces, we apply our new VPDC from Section 5 for this commitment, so that the VPDC decryption algorithm recovers (i) the real payer's index among N users, and (ii) the transaction amount. Let us investigate more details.

In MatRiCT, an HMC commitment $B = \mathsf{Com}_{ck}(\boldsymbol{b}, \boldsymbol{u}; \boldsymbol{r})$ is computed, where $\boldsymbol{b}$ is a binary vector over $R_{\hat{q}}$ for some $\hat{q} \in \mathbb{Z}^+$ (i.e., $\boldsymbol{b} = (b_0, b_1, \ldots)$ such that $b_i \in \{0,1\} \subset R_{\hat{q}}$), $\boldsymbol{u}$ is some short auxiliary message and $\boldsymbol{r}$ is some short randomness. Here, the binary vector $\boldsymbol{b}$ is comprised of three components: (i) the unary representation of an index ℓ that identifies the real payer (i.e., spender) index among N parties in a ring signature or a 1-out-of-N proof, (ii) the bits in the binary representation of all output amounts, and (iii) the bits in the so-called "corrector values". Our target is to recover the first two components in decryption so that the authority can learn the two hidden data pieces mentioned above. Note that the payer indeed proves in zero-knowledge that she owns the ℓ-th public key and that certain bits (with known indices) in $\boldsymbol{b}$ construct the output coins. Hence, recovering $\boldsymbol{b}$ guarantees that the real payer index and the output amounts (and thus the transaction amount) are revealed.

As it is expensive to perform an *exact* binary proof on B, MatRiCT performs a *relaxed* binary proof on B. Let us recall a simplified version of the relation proven in MatRiCT for the commitment B, which also applies to our variant MatRiCT-Au.

Lemma 2. *Assume that $\hat{q}$ is sufficiently large and that HMC is γ_{bin}-binding for some γ_{bin} that depends on the system parameters. For an input commitment $B \in R_{\hat{q}}^{\hat{n}}$ and a commitment key $ck = \hat{\boldsymbol{G}} = [\,\boldsymbol{A} \,\|\, \boldsymbol{B} \,\|\, \boldsymbol{C}\,]$ defined over $R_{\hat{q}}$, our binary proof proves knowledge of $(y, \boldsymbol{b}, \hat{\boldsymbol{c}}, \hat{\boldsymbol{r}})$ such that*

- $y \in \Delta\mathcal{C}_{w,p}^{d}$, $\hat{\boldsymbol{c}} \in R_{\hat{q}}^{v_2}$ *and* $\hat{\boldsymbol{r}} \in R_{\hat{q}}^{\hat{m}}$ *for some* $v_2, \hat{m} \geq 1$,
- $yB = \mathsf{Com}_{ck}(y\boldsymbol{b}, \hat{\boldsymbol{c}}; \hat{\boldsymbol{r}}) = \boldsymbol{A}\hat{\boldsymbol{r}} + \boldsymbol{B}y\boldsymbol{b} + \boldsymbol{C}\hat{\boldsymbol{c}}$,
- *All coordinates b_i of $\boldsymbol{b}$ are in $\{0,1\}$, i.e., $\boldsymbol{b} \in \{0,1\}^{v_1} \subset R_{\hat{q}}^{v_1}$, where $v_1 = k\beta + Sr + \lceil \log(M+S-1) \rceil (r-1)$ for parameters k, β satisfying $N = \beta^k$, M, S denoting the number of input/output accounts and r denoting the bit length of each amount,*
- $\|(y\boldsymbol{b}, \hat{\boldsymbol{c}}, \hat{\boldsymbol{r}})\| \leq \gamma_{\text{B}}$ *for some $\gamma_{\text{B}} \in \mathbb{R}^+$ with $\gamma_{\text{B}} < \hat{q}$.*

The above relation is effectively what we study in Def. 5 with $\mathcal{D} = \{0,1\}^{v_1}$. So, the extension we need to make over MatRiCT is to let the spender use the VPDC from Section 5 for the commitment B. In particular, the spender just

needs to update $ck = \hat{\boldsymbol{G}}$ with a 'trapdoored' commitment key ck^{td} generated by CAddTd. This simply means that the spender replaces the last row of $ck = \hat{\boldsymbol{G}}$ with trapdoor rows published by an authority. This way, the authority in possession of the corresponding trapdoor td can execute CDecGR (Alg. 2) to recover the vector $\boldsymbol{b}$, which in turn reveals the real payer index $\ell \in [0, N-1]$ and the transaction amount, which is equal to the sum of the output amounts. In particular, since the commitment algorithm for our VPDC remains exactly the same as the standard HMC commitment algorithm used in MatRiCT, our VPDC can be directly plugged in and used with the same efficient NIZK proof of transaction well-formedness used in MatRiCT. For the invertibility of challenge differences in R_t as required in Theorem 2, we use the results of [8, 11].

For the concrete parameter setting in MatRiCT with $N = 100$, we have $\dim(\boldsymbol{b}) = v_1 = 291$ as $(k, \beta) = (1, N)$ and $(M, S, r) = (1, 2, 64)$ with the first $k\beta = 100$ bits having exactly a single '1'. As a result, there are more than 2^{191} possibilities for $\boldsymbol{b}$ (i.e., $|\mathcal{D}| > 2^{191}$). Hence, it is infeasible to do an exhaustive search over $\mathcal{D}$ (as done in [12]) for decryption. As our new decryption's runtime is *polylogarithmic* in $|\mathcal{D}|$, we can efficiently execute it. In particular, as we discuss in the full version of this paper, to guarantee auditability soundness against adversarially-created commitments based on our VPDC security bounds while maintaining the same security level as MatRiCT against best-known lattice attacks, we only need to increase (i) the system modulus $\hat{q}$ to a 55-bit value from a 53-bit value and (ii) the commitment matrix dimensions slightly. As shown in the full version of this paper, the decryption runs very fast despite the exponentially large message space.

It is important to note here that MatRiCT and MatRiCT-Au crucially relies on an *aggregate* binary proof for compactness, where many messages are committed together inside a single commitment B. Therefore, the additional succinctness feature of VPDC plays an important role. If one were to replace this HMC commitment with an encryption (or an encryption-like commitment as in [3]), the proof/commitment length would significantly increase (as also discussed in the introduction) due to the large input message dimension (several hundreds) over the polynomial ring $R_{\hat{q}}$.

In the full version of this paper, we describe MatRiCT-Au in full details, and show that addition of a trapdoor as in CAddTd is effectively the only modification required over MatRiCT. We instantiated MatRiCT-Au concretely and implemented it in C/C++ (see the full version). We compare MatRiCT and MatRiCT-Au in Table 1. Our results show that the overhead of MatRiCT-Au over MatRiCT is very small in both communication and computation.

References

1. Bao, F., Deng, R.H., Mao, W.: Efficient and practical fair exchange protocols with off-line TTP. In: IEEE S&P. pp. 77–85. IEEE Computer Society (1998)
2. Baum, C., Bootle, J., Cerulli, A., del Pino, R., Groth, J., Lyubashevsky, V.: Sublinear lattice-based zero-knowledge arguments for arithmetic circuits. In: CRYPTO (2). Lecture Notes in Computer Science, vol. 10992, pp. 669–699. Springer (2018)

3. Baum, C., Damgård, I., Lyubashevsky, V., Oechsner, S., Peikert, C.: More efficient commitments from structured lattice assumptions. In: SCN. pp. 368–385. Springer (2018)
4. Camenisch, J., Shoup, V.: Practical verifiable encryption and decryption of discrete logarithms. In: CRYPTO. LNCS, vol. 2729, pp. 126–144. Springer (2003)
5. Chaum, D., Van Heyst, E.: Group signatures. In: EUROCRYPT. pp. 257–265. Springer (1991)
6. Esgin, M.F.: Practice-Oriented Techniques in Lattice-Based Cryptography. Ph.D. thesis, Monash University (5 2020). https://doi.org/10.26180/5eb8f525b3562
7. Esgin, M.F., Nguyen, N.K., Seiler, G.: Practical exact proofs from lattices: New techniques to exploit fully-splitting rings. In: ASIACRYPT (2). LNCS, vol. 12492, pp. 259–288. Springer (2020)
8. Esgin, M.F., Steinfeld, R., Liu, D., Ruj, S.: Efficient hybrid exact/relaxed lattice proofs and applications to rounding and vrfs. Cryptology ePrint Archive, Report 2022 (2022)
9. Esgin, M.F., Steinfeld, R., Liu, J.K., Liu, D.: Lattice-based zero-knowledge proofs: New techniques for shorter and faster constructions and applications. In: CRYPTO (1). pp. 115–146. LNCS, Springer (2019), (Full version at `ia.cr/2019/445`)
10. Esgin, M.F., Steinfeld, R., Sakzad, A., Liu, J.K., Liu, D.: Short lattice-based one-out-of-many proofs and applications to ring signatures. In: ACNS. pp. 67–88. LNCS, Springer (2019), (Full version at `ia.cr/2018/773`)
11. Esgin, M.F., Steinfeld, R., Zhao, R.K.: MatRiCT$^+$: More efficient post-quantum private blockchain payments. Cryptology ePrint Archive, Report 2021/545 (2021), `ia.cr/2021/545` (to appear at IEEE S&P 2022)
12. Esgin, M.F., Zhao, R.K., Steinfeld, R., Liu, J.K., Liu, D.: MatRiCT: Efficient, scalable and post-quantum blockchain confidential transactions protocol. In: ACM CCS. pp. 567–584. CCS '19, ACM (2019), (Full version at `ia.cr/2019/1287`)
13. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: CRYPTO. pp. 186–194. Springer (1986)
14. Groth, J., Kohlweiss, M.: One-out-of-many proofs: Or how to leak a secret and spend a coin. In: EUROCRYPT. pp. 253–280. Springer (2015)
15. Li, W., Wang, Y., Chen, L., Lai, X., Zhang, X., Xin, J.: Fully auditable privacy-preserving cryptocurrency against malicious auditors (2019), `ia.cr/2019/925`
16. Lyubashevsky, V.: Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In: ASIACRYPT. pp. 598–616. Springer (2009)
17. Lyubashevsky, V.: Lattice signatures without trapdoors. In: EUROCRYPT. pp. 738–755. Springer (2012), (Full version)
18. Lyubashevsky, V., Neven, G.: One-shot verifiable encryption from lattices. In: EUROCRYPT. pp. 293–323. Springer (2017), (Full version at `ia.cr/2017/122`)
19. Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: EUROCRYPT. pp. 700–718. LNCS, Springer (2012)
20. Noether, S.: Ring signature confidential transactions for monero. Cryptology ePrint Archive, Report 2015/1098 (2015), `ia.cr/2015/1098`
21. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6), 34:1–34:40 (2009), preliminary version in STOC 2005
22. Schnorr, C.: Efficient identification and signatures for smart cards. In: CRYPTO. Lecture Notes in Computer Science, vol. 435, pp. 239–252. Springer (1989)
23. Young, A.L., Yung, M.: Auto-recoverable auto-certifiable cryptosystems. In: EUROCRYPT. LNCS, vol. 1403, pp. 17–31. Springer (1998)