

Mercurial Commitments: Minimal Assumptions and Efficient Constructions

Dario Catalano¹, Yevgeniy Dodis², and Ivan Visconti³

¹ CNRS-Ecole Normale Supérieure, Laboratoire d’Informatique, 45 Rue d’Ulm, 75230 Paris Cedex 05 - France. dario.catalano@ens.fr

² Department of Computer Science, New York University, 251 Mercer Street, New York, NY 10012, USA. dodis@cs.nyu.edu

³ Facoltà di Scienze Matematiche, Fisiche e Naturali, Università di Salerno, via S. Allende n. 2, 84081 Baronissi (SA) - Italy. visconti@unisa.it

Abstract. (Non-interactive) *Trapdoor Mercurial Commitments* (TMCs) were introduced by Chase et al. [8] and form a key building block for constructing zero-knowledge sets (introduced by Micali, Rabin and Kilian [28]). TMCs are quite similar and certainly imply ordinary (non-interactive) *trapdoor commitments* (TCs). Unlike TCs, however, they allow for some additional freedom in the way the message is opened: informally, by allowing one to claim that “if this commitment can be opened at all, then it would open to this message”. Prior to this work, it was not clear if this addition is critical or not, since all the constructions of TMCs presented in [8] and [28] used strictly stronger assumptions than TCs. We give an affirmative answer to this question, by providing simple constructions of TMCs from any trapdoor bit commitment scheme. Moreover, by plugging in various trapdoor bit commitment schemes, we get, in the trusted parameters (TP) model, *all* the efficient constructions from [28] and [8], as well as several immediate new (either generic or efficient) constructions. In particular, we get a construction of TMCs *from any one-way function* in the TP model, and, by using a special flavor of TCs, called *hybrid TCs* [6], even in the (weaker) shared random string (SRS) model.

Our results imply that (a) *mercurial commitments can be viewed as surprisingly simple variations of trapdoor commitments*; and (b) *the existence of non-interactive zero-knowledge sets is equivalent to the existence of collision-resistant hash functions*. Of independent interest, we also give a stronger and yet much simpler definition of mercurial commitments than that of [8], which is also met by our constructions in the TP model.

1 Introduction

Commitment schemes are important cryptographic primitives. They allow one party to commit to some value v so that v is kept secret from the rest of the world (this is called *hiding*), and yet everybody knows that the value v is uniquely defined at the time v was committed (this is called *binding*). In particular, binding ensures that the party cannot announce the commitment first, and then decide

later how to open it depending on the circumstances. In this sense, commitment schemes force the party to fully decide on what he is committing to.

At Eurocrypt 2005, Chase et al. [8] introduced an intriguing variant of commitments called *mercurial commitments*. The main difference comes from the fact that mercurial commitments allow for a small, and yet noticeable relaxation of the strict binding property of regular commitments. Namely, they allow for a two-stage opening protocol. In the *soft-open stage* the committer can claim that “if I committed to anything at all, then this value is m ”, while in the hard-opening stage he would indeed declare that “Yes, I really committed to the value m .” In particular, any committed value c can either be both soft- and hard-opened only to one (correct!) message m , or can be soft-opened to arbitrary messages, but then it cannot be hard-opened at all! Moreover, the committer must decide before forming the commitment which one of the two cases suits him better: to commit to only one value, or not to commit to anything at all. Although this is seemingly not much better than regular commitments, the extra freedom of the committing party comes from the fact that by showing a soft-opening of his commitment to some value m , the receivers still cannot tell if m was really committed to by c , or if c was simply a “non-commitment” to anything (and the committer might be just going around and soft-opening c to arbitrary values m'). The receivers are sure, however, that it is impossible to hard-open c to any $m' \neq m$.

Chase et al. [8] distilled the above natural primitive to abstract away a relatively complicated (but efficient!) construction of *zero-knowledge sets* by Micali et al. [28]. Such ZK sets allow one to commit to some secret set S over some universe, and then to be able to non-interactively prove statements of the form $x \in S$ and $x \notin S$, and yet no other information (which cannot be deduced from the inclusions/exclusions above) about S is leaked — not even its size! With the abstraction of mercurial commitments, Chase et al. [8] obtained an elegant and easy-to-follow general “explanation” of the construction from [28]. Namely, they showed that the construction of [28] is an instance of a general construction of ZK sets from *any* mercurial commitment scheme and any collision-resistant hash function.

PLAIN VS. TRAPDOOR MERCURIAL COMMITMENTS. We remark that to match a very strong zero-knowledge definition of ZK sets from [28], Chase et al. [8] had to require that mercurial commitments satisfy the following “equivocation” property: there exists some trapdoor information msk (ordinarily not available to anybody) which enables one to completely destroy all the binding properties of mercurial commitments. Namely, using msk one can construct fake commitments, which look just like regular commitments and yet can be soft- or hard-opened to completely arbitrary values. (This is very similar to the notion of regular *trapdoor commitments* [4], where the knowledge of the corresponding trapdoor key can enable somebody to create fake regular commitments which can be opened to any message.) As already observed by [8], this strong equivocation property does not seem to be inherent for the “plain” primitive of mercurial commitments, but they chose to insist on this extra property since it was need

for their main application. Since we believe that mercurial commitments are also interesting without equivocation, in our results we will distinguish between *plain* and *trapdoor* mercurial commitments. (Although our results described below will hold equally naturally for either case.) Indeed, we observe that one can define a weaker notion of ZK sets, which we informally call *indistinguishable sets*, which have the same functionality as ZK sets, but the privacy property is relaxed to only state that for any two sets and any sequence of inclusion/exclusion assertions which does not “separate” these sets, seeing the proofs of the corresponding assertions does not allow one to distinguish between these two sets. (This is somewhat similar to the distinction between witness indistinguishable [17] and ZK proofs [21].) And then it is easy to see that the same generic construction from [8] would give indistinguishable sets when applied to plain mercurial commitments. To summarize, we believe that both plain and trapdoor mercurial commitments are useful and deserve investigation.

MINIMAL ASSUMPTIONS FOR MERCURIAL COMMITMENTS. Having introduced a new cryptographic primitive, it is always very important to understand where it lies in the hierarchy of cryptographic assumptions. Towards this goal, [8] gave a general construction of (trapdoor) mercurial commitments in the plain model from any zero-knowledge proof system for NP. This construction, however, requires interaction. Then [8] showed a construction of non-interactive (trapdoor) mercurial commitments from non-interactive zero-knowledge proofs (NIZK) for NP, which are known, for example, to be implied by trapdoor permutations in the shared random string (SRS) model. However, this construction is mainly of theoretical interest, since it is very inefficient in practice. They also gave a more efficient (although bit-by-bit) construction of non-interactive (trapdoor) mercurial commitments from an even stronger assumption of claw-free permutations [22] in the trusted parameters (TP) model. On the other hand, [8] observed that (trapdoor) mercurial commitments are similar and trivially imply (trapdoor) regular commitments, although they pointed out some important differences as well. Thus, the following two questions were left open:

Question 1: *What minimal cryptographic and set-up assumptions are sufficient for non-interactive plain/trapdoor mercurial commitments?*

Question 2: *Can plain/trapdoor mercurial commitments be (efficiently) built from plain/trapdoor commitments?*

Our first result resolves these questions in a surprisingly simple fashion. We show a very simple and efficient construction of (bit) plain/trapdoor mercurial commitments from any bit plain/trapdoor regular commitment. The construction is a very simple generalization of the claw-free construction from [8], and since regular/trapdoor commitments are in principle equivalent to one-way functions in the SRS/TP model, we get

Theorem 1. *There exists a simple and efficient construction of (non-interactive) bit plain/trapdoor mercurial commitments from bit plain/trapdoor commitments. In particular, (non-interactive) plain/trapdoor mercurial commitments exist in the SRS/TP model if and only if one-way functions exist.*

The above result leaves open a question of basing (only trapdoor) mercurial commitments on one-way functions in the SRS model, which is a weaker set-up assumption than the TP model. Luckily, we observe that by using the same construction with a slight relaxation of trapdoor commitments, called *hybrid* trapdoor commitments, which were introduced by Catalano and Visconti [6] and shown to be equivalent to one-way functions *even in the SRS model*, we get

Theorem 2. *There exists a simple and efficient construction of (non-interactive) bit trapdoor mercurial commitments from (non-interactive) bit hybrid trapdoor commitments. In particular, non-interactive trapdoor mercurial commitments exist in the SRS model if and only if one-way functions exist.*

EFFICIENCY? Having resolved the question of feasibility, we can turn to the question of efficiency. Of course, we can plug in various efficient bit trapdoor commitment schemes to our previous construction, but this will only result in bit-by-bit constructions for long messages, which is pretty inefficient for practical use (e.g., for the ZK sets application). On the other hand, Chase et al. [8] gave two efficient constructions for long messages based on specific number-theoretic constructions (discrete log and factoring; the discrete log construction was implicit in [28]). Examining these constructions, one can see that there seems to be some kind of similarity between them, although it is not obvious exactly where this similarity comes from. Also, it is relatively hard to understand why each construction is really secure, without going into the details of the proof. Motivated by this, we ask

Question 3: *Is there an efficient and yet reasonably **general** construction of plain/trapdoor mercurial commitments, which would abstract and explain the efficient number-theoretic constructions from [8]?*

Our second result gives a surprisingly general answer to this question. Namely, we present a construction which directly transforms a plain/trapdoor bit commitment \mathcal{C} into an efficient and (typically) multi-bit plain/mercurial commitment \mathcal{C}' . Namely, we still base it on general plain/trapdoor commitment, just like in Theorem 1. However, a small catch is that we will need to assume an extra property from \mathcal{C} (see Section 2.1 for a definition of Σ -protocol):

Theorem 3. *Assume \mathcal{C} is a plain/trapdoor bit commitment which has an efficient Σ -protocol Π proving that one knows a witness d that a given (regular) commitment c can be opened to 0.¹ Then one can construct an efficient plain/trapdoor mercurial commitment \mathcal{C}' whose message space is equal to the challenge space of Π .*

Thus, to get message-efficient constructions, it will be important to design “challenge-efficient” Σ -protocols for our plain/trapdoor commitment schemes.

¹ As explained in Section 5 proving this theorem, we will need a slight extra property (*) from such Σ -protocols, but it will always hold in any practical construction we are aware of. So we omit it from this statement.

While such Σ -protocol's Π in principle (see Theorem 5 below) can always be built from one-way functions, in general this will not outperform the simple construction in Theorem 1. However, the utility of this transformation comes from the fact that all number-theoretic (trapdoor) bit commitment schemes have very efficient Σ -protocols, and usually with rich challenge spaces. Plugging in various such commitment schemes with efficient protocols, we get many efficient constructions of mercurial commitments. In particular, both the discrete log and the factoring construction of [8] become special cases of our general transformation, when applied to an appropriate trapdoor commitment scheme! And several new constructions can be obtained as well (e.g., from RSA and Paillier [32] assumptions, as well as new discrete log and factoring constructions; see Section 5.1). More generally, we also believe that our construction is much easier to understand and sheds more light onto why the previous number-theoretic constructions were built in this particular way.

SIMPLER DEFINITION. As another small contribution, by *strengthening* the definition of trapdoor mercurial commitments as compared to the definition of [8], we considerably simplified the equivocation property of mercurial commitments. Since all our constructions (with the exception of Theorem 2) satisfy the stronger definition, and it results in easier and shorter proofs, we believe our strengthening is justified and could be of independent interest.

IMPLICATION TO ZK SETS. It is known from [8] and [31] that collision-resistant hash functions (CRHF) suffices for constructing interactive ZK sets in the plain model. Chase et al. [8] also made a simple observation that ZK sets imply the existence of CRHFs and therefore interactive ZK sets in the plain model and CRHFs are equivalent. Chase et al. in [8] also show that non-interactive indistinguishable/ZK sets can be constructed from any non-interactive plain/trapdoor mercurial commitment scheme and a collision-resistant hash function (CRHF). Using Theorem 1, Theorem 2, Theorem 3, and the fact that CRHFs imply both one-way functions (and, thus, plain/trapdoor/hybrid commitments in the SRS/TP/SRS models) and efficient *plain* commitment schemes (see [12] and [25]), we immediately obtain:

Theorem 4. *The existence of ZK (and, thus, indistinguishable) sets in the SRS model is equivalent to the existence of CRHFs. Moreover, ZK sets can be efficiently constructed from CRHFs and trapdoor bit commitment schemes, while indistinguishable sets can be efficiently constructed using CRHFs alone (by also building commitments out of them). The constructions become even more efficient if the commitment scheme in question has a challenge-efficient Σ -protocol needed for Theorem 3.*

2 Definitions

2.1 Σ -Protocols

Let $\mathcal{R} = \{(x, w)\}$ be some NP-relation (i.e., it is efficiently testable to see if $(x, w) \in \mathcal{R}$ and $|w| \leq \text{poly}(|x|)$). We usually call x the input, and w — the witness

(for x). Consider a three move protocol run between a PPT prover P , with input $(x, w) \in R$, and a PPT verifier V with input x , of the following form. P chooses a random string r_p , computes $a = \text{Start}(x, w; r_p)$, and sends a to V . V then chooses a random string e (called “challenge”) from some appropriate domain E (see below) and sends it to P . Finally, P responds with $z = \text{Finish}(x, w, e; r_p)$. The verifier V then computes and returns a bit $b = \text{Check}(x, a, e, z)$. We require that **Start**, **Finish**, and **Check** be polynomial-time algorithms, and that $|e| \leq \text{poly}(|x|)$. Such a protocol (given by procedures **Start**, **Finish**, **Check**) is called a Σ -Protocol for \mathcal{R} if it satisfies the following properties, called completeness, special soundness, and special honest-verifier zero-knowledge:

- **Completeness:** If $(x, w) \in R$ then the verifier outputs $b = 1$ (with all but negligible probability).
- **Special Soundness:** There exists a PPT algorithm **Extract**, called the (knowledge) extractor, such that it is computationally infeasible to produce an input tuple (x, a, e, z, e', z') such that $e \neq e'$ both lie in the proper “challenge” domain, $\text{Check}(x, a, e, z) = \text{Check}(x, a, e', z') = 1$, and yet **Extract** (x, a, e, z, e', z') fails to output a witness w such that $(x, w) \in R$. Intuitively, if some prover can correctly respond to two different challenges e and e' on the same first flow a , then the prover must “know” a correct witness w for x (in particular, x has a witness).
- **Special HVZK:** There exists a PPT algorithm **Simul**, called the simulator, such that for any $(x, w) \in R$ and for any fixed challenge e , the following two distributions are computationally indistinguishable. The first distribution (x, a, e, z) is obtained by running an honest prover P (with some fresh randomness r_p) against a verifier whose challenge is fixed to e . The second distribution (x, a, e, z) is obtained by computing the output $(a, z) \leftarrow \text{Simul}(x, e)$ (with fresh randomness r_s). Intuitively, this says that for any a-priori fixed challenge e , it is possible to produce a protocol transcript computationally indistinguishable from an actual run with the prover (who knows w).

Since the standard zero-knowledge protocol for the Hamiltonian Cycle (see [16] and [23]) language is a (binary challenge) Σ -protocol, we get

Theorem 5 ([23],[16]). *Any NP-relation \mathcal{R} has a (binary challenge) Σ -protocol if secure commitment schemes exist (in particular, in the SRS model if one-way functions exist).*

Of course, we will see and crucially exploit the fact that many natural specific languages have much more efficient Σ -protocols. We also notice that, aside from computational efficiency, a good quality measure for a given Σ -protocol is the size of its challenge space E (the larger the better). One reason for this dependency comes because the special soundness property easily implies that if a malicious prover does not “know” a valid witness w for x , then he can succeed in fooling the verifier with probability at most (only negligibly better than) $1/|E|$. In our application, we will also see that the large size of E will also naturally translate to more efficient solutions, and we will therefore strive to use “challenge-efficient” Σ -protocols.

GENERALIZATIONS. First, we will allow \mathcal{R} to depend on some honestly generated public parameter pk (known to everybody after generation); e.g. the standard discrete-log relation would be $\mathcal{R}_{p,g}(x, w) = 1$ if and only if $x = g^w \bmod p$, where the prime p and the generator g could be randomly generated. In this case the corresponding properties of the Σ -protocol should computationally hold over the choice of such parameters. However, for one of our applications we will require an even stronger technical property. Namely, we will say that a family of relations $\{\mathcal{R}_{pk}\}$ has a Σ -protocol which is *strongly hiding w.r.t. instance generation procedure* \mathcal{P} if the special HVZK property holds even in the following experiment: \mathcal{P} produces (pk, x, w, \mathcal{I}) , where pk is the public key for \mathcal{R} , x is the input, w is the witness, and \mathcal{I} is some side information available to attacker. Then we either give to the distinguisher a tuple $(\mathcal{I}, pk, x, a, e, z)$ obtained by having the prover run the real protocol with x and w , or where (a, z) is produced by the simulator $\text{Simul}_{pk}(x, e)$. To put it differently, the side information \mathcal{I} does not help the distinguisher to break the special HVZK property. We notice that, essentially all of the practical Σ -protocols known (including all the ones we will actually consider) will satisfy the statistical HVZK property, in which case they will be strongly-hiding w.r.t. any \mathcal{P} . Also, the generic protocol from Theorem 5 will also be strongly-hiding w.r.t. any efficient procedure \mathcal{P} which only depends on the public parameters of the commitments used inside the protocol. This, once again, includes essentially all interesting procedures (including the specific one we will need later). In other words, for all practical purposes this extra property is just a technicality we need for the proof to go through.

As a second, orthogonal generalization, we can also consider “auxiliary-input” Σ -protocols, where in order to run the protocol, the prover P might need some extra information aux satisfying some property (which, presumably can be generated together with (x, w)), in addition to w . Notice, w alone is enough to allow for verification that $(x, w) \in \mathcal{R}$, so aux is only needed by the prover to fulfill his completeness requirement (in particular, the simulator does not need to know aux and special soundness and HVZK stay the same as before).

EFFICIENT Σ -PROTOCOLS. We briefly survey the following efficient Σ -protocols which we will use in the sequel. (The exact details will not be crucial for our purposes, so we will not present them here.) We notice that most of them will be unconditional: the security assumption behind the relation (such as discrete log) will be used later in the application; for example, in claiming that the hypothetical extraction of the witness contradicts the corresponding assumption.

The Schnorr Σ -protocol [34] allows one to unconditionally prove the knowledge of the discrete log in cyclic groups of prime order. A less known fact [19,, 9] is that (a slightly modified)² Schnorr protocol also works over the subgroup of quadratic residues Q_n over \mathbb{Z}_n^* , where n is the product of two safe primes. Interestingly, unlike in prime order groups, where the special soundness holds

² In particular, the prover works over the integers instead of over $\mathbb{Z}_{|Q_n|}$, since he does not know $|Q_n|$. Because of that the special HVZK guarantee is statistical here rather than perfect.

unconditionally, here it will hold computationally under the strong RSA assumption. In both of these cases the challenge space is exponential.

Very similar to Schnorr protocol, Gilliou-Quisquater (GQ) [24] protocol proves the knowledge of the e -th root over \mathbb{Z}_n^* (i.e., solution to RSA), where $\gcd(e, \varphi(n)) = 1$ and n is the product of two safe primes. Here, however, the challenge space should be smaller than the exponent e , so this protocol is challenge-efficient only if e is large (which is typically required when this protocol is used).

The Fiat-Shamir protocol is an unconditional binary-challenge Σ -protocol proving the knowledge of the square root over \mathbb{Z}_n^* , where n is the product of two primes. One way to make it challenge-efficient is to repeat it in parallel, but this is computationally inefficient. A better way is to use the elegant technique of Ong-Schnorr [30], at the expense of working over the set of quadratic residues Q_n , requiring n to be a Blum integer, and, more crucially, requiring an auxiliary witness to the prover. Namely, in order to make the challenge space to be of size 2^ℓ , the prover not only needs to know a square root of the input $x \in Q_n$, but also the 2^ℓ -root root $u \in Q_n$ of x (which is well defined when n is a Blum integer): see Lemma 3.1 in [1] explicitly stating the special soundness of this protocol. Of course, to run this protocol in practice one would first pick u and then set $w = u^{2^{\ell-1}} \bmod n$ (by repeated squaring) and $x = w^2 \bmod n$.

All the above mentioned protocols have statistical special HVZK, so they always satisfy strong-hiding. To summarize, natural relations arising from well established cryptographic assumptions have very computationally and challenge-efficient Σ -protocols.

2.2 Commitments and Trapdoor Commitments

COMMITMENTS. A (non-interactive) commitment scheme consists of four efficient algorithms: $\mathcal{C} = (\text{Com-Gen}, \text{Com}, \text{Open}, \text{Ver})$. The generation algorithm $\text{Com-Gen}(1^k)$, where k is the security parameter, outputs a public commitment key pk (possibly empty, but usually consisting of public parameters for the commitment scheme). Given a message m from the associated message space \mathcal{M} (e.g., $\{0, 1\}^k$, although we will mainly concentrate on bit commitments), $\text{Com}_{pk}(m; r)$ (computed using the public key pk and additional randomness r) produces a commitment string c for the message m . We will sometimes omit r and write $c \leftarrow \text{Com}_{pk}(m)$. Similarly, the opening algorithm $\text{Open}_{pk}(m; r)$ (which is supposed to be run using the same value r as the commitment algorithm) produces a decommitment value d for c . Finally, the verification algorithm $\text{Ver}_{pk}(m, c, d)$ accepts (i.e., outputs 1) if it thinks the pair (c, d) is a valid commitment/decommitment pair for m . We require that for all $m \in \mathcal{M}$, $\text{Ver}_{pk}(m, \text{Com}_{pk}(m; r), \text{Open}_{pk}(m; r)) = 1$ holds with all but negligible probability. We remark that without loss of generality we could have assumed that the opening algorithm simply outputs its randomness r as the decommitment, and the verification algorithm simply checks if $c = \text{Com}_{pk}(m; r)$. However, we will find our more general notation more convenient for our purposes. When clear from the context, we will sometimes omit pk from our notation. Regular commitment schemes have two security properties:

Hiding: No PPT adversary (who knows pk) can distinguish the commitments to any two message of its choice: $\text{Com}_{pk}(m_1) \approx \text{Com}_{pk}(m_2)$. That is, $\text{Com}_{pk}(m)$ reveals “no information” about m .

Binding: Having the knowledge of pk , it is computationally hard for the PPT adversary \mathcal{A} to come up with c, m, d, m', d' such that (c, d) and (c, d') are valid commitment pairs for m and m' , but $m \neq m'$ (such a tuple is said to cause a *collision*). That is, \mathcal{A} cannot find a value c which it can open in two different ways.

Commitments are known to be theoretically equivalent to one-way functions [29],[26] (at least in the SRS model). However, efficient commitments can be built from collision-resistant hash functions [12],[25], and many number-theoretic assumptions (such as factoring, discrete log and RSA, and Paillier [32]; see below). In fact, most of these number-theoretic construct give a stronger kind of commitment — called trapdoor commitment — which we explain next.

TRAPDOOR COMMITMENTS. A (non-interactive) trapdoor commitment scheme consists of six efficient algorithms: $\mathcal{C} = (\text{TrCom-Gen}, \text{Com}, \text{Open}, \text{Ver}, \text{Fake}, \text{Equiv})$. The generation algorithm $\text{TrCom-Gen}(1^k)$, where k is the security parameter, outputs a public commitment key pk and a secret *trapdoor key* sk . Once pk is fixed, the meaning of Com , Open and Ver is exactly the same as for regular commitments. In particular, we will require that these algorithms satisfy the usual hiding and binding properties of the commitment schemes.

The trapdoor key sk is used in the algorithms Fake and Equiv to break the binding property of commitments. Namely, $\text{Fake}_{sk}(\cdot; r)$ (which takes no input except for randomness r) produces “fake” commitment c , initially not associated to any message m . On other other hand, for any message m , $\text{Equiv}_{sk}(m; r)$ (which is supposed to be run using the same value r as the fake commitment algorithm) produces a “fake decommitment” value d for $c = \text{Fake}_{sk}(\cdot; r)$. In particular, we require that such fake (c, d) still satisfy the verification equation: for all $m \in \mathcal{M}$, $\text{Ver}_{pk}(m, \text{Fake}_{sk}(\cdot; r), \text{Equiv}_{sk}(m; r)) = 1$ holds with all but negligible probability. Even stronger, we require that

Equivocation: for any $m \in \mathcal{M}$ (chosen by the adversary), a “true” commitment tuple $(m, \text{Com}_{pk}(m; r), \text{Open}_{pk}(m; r))$ should look computationally indistinguishable (over r) from the fake tuple $(m, \text{Fake}_{sk}(\cdot; r), \text{Equiv}_{sk}(m; r))$. More importantly, we require that these distributions should look indistinguishable even if the distinguisher knows not only the commitment key pk , but also *the trapdoor key* sk (we will explain the rational for this shortly)!

We notice that equivocation easily implies that trapdoor commitments satisfy the usual hiding property of commitments (since all commitments $\text{Com}_{pk}(m)$ are indistinguishable from a single distribution $\text{Fake}_{sk}(\cdot)$): in fact, this indistinguishability holds even if the distinguisher knows sk ! Thus, binding and equivocation are enough to argue the security of trapdoor commitment schemes.

We briefly give the rational of why we need such a strong equivocation property. This is done for the purposes of *composition*. Indeed, we would like to

argue that given several “real” pairs (c, d) , we can replace all of them by the corresponding “fake” pairs (c', d') , without anybody “noticing”. However, the standard left-to-right hybrid argument requires us to be able to generate not only the “real left-pairs” (c, d) , which we can do using pk , but also “fake right-pairs” (c', d') , and this we cannot do without the knowledge of sk . Requiring the indistinguishability to hold even with the knowledge of sk resolves this problem, and gives us all the natural composition properties.

CONSTRUCTIONS. There are many constructions of trapdoor commitments (and each of them also gives a regular commitment, of course). For example, efficient trapdoor commitments exist based on a variety of number-theoretic assumptions: factoring [27],[33], discrete log [3],[4]), RSA (combining [16],[24]), Paillier [5],[11]. In fact, some of these schemes (e.g., those based on discrete log and RSA) are special cases of a beautiful general construction by Feige and Shamir [16]. This construction *efficiently* transforms any Σ -protocol corresponding to a “hard” language in NP into a trapdoor commitment scheme. In particular, since we mentioned that all of NP has such Σ -protocols if one-way functions exists (see Theorem 5), and the latter also imply that some languages in NP are “hard” (at least, in the TP model), one can in principle construct a trapdoor commitment scheme *from any one-way function* in the TP model (see sec. 4.9.2.3 of [20]). We note that the message space for the resulting trapdoor commitment will be exactly the challenge space of the corresponding Σ -protocol, which, once again, demonstrates why we want to construct challenge-efficient Σ -protocols.³ Quite interestingly, this construction *of* trapdoor commitments will be somewhat reminiscent to our main construction *from* trapdoor commitments (possessing a certain Σ -protocols; see Section 5), although this seems to be more of a coincidence.⁴

We also mention another, less general construction [27] of trapdoor commitments from claw-free permutation pairs [22]. This construction is only efficient for bit trapdoor commitments (which, once again, are sufficient for us). Looking at various known claw-free permutation constructions (e.g., see [14] for such a list), we immediately get efficient bit trapdoor commitment constructions from various assumptions, such as the already mentioned constructions from factoring [27], Paillier [11] and the bit-version of the discrete log construction of [3],[4]. In regards to discrete log, we finally mention the following “ad-hoc” construction of trapdoor bit commitments. The public key consists of two random generators g and $h = g^x$ of some prime order q cyclic group \mathbb{G} , where the discrete log is hard (here x is a random non-zero element of \mathbb{Z}_q), while the trapdoor key is x . To commit to 0, one computes g^{r_0} (for random non-zero $r_0 \in \mathbb{Z}_q$), while to commit

³ Of course, since both of our generic mercurial commitment constructions only use *bit* commitments, even binary Σ -protocols for hard languages suffice for our purpose.

⁴ Perhaps partially explained by the fact that mercurial commitment are trapdoor commitments with several very special properties (see Section 2.3). Correspondingly, in our main construction we will need “hard” languages also satisfying some special properties. Somehow remarkably, though, these extra properties have more or less led us to trapdoor commitments themselves! See Section 5.

to 1 one similarly computes h^{r_1} . The openings are r_0 and r_1 , respectively. To break binding one needs to satisfy $g^{r_0} = h^{r_1}$, which means that one can compute $x = r_0 r_1^{-1} \bmod q$ (and this contradicts discrete log). On the other hand, if x is known, it is trivial to open a “fake” commitment h^{r_1} both to 1 (by simply presenting r_1) and to 0 (by presenting $r_1 x \bmod q$).

HYBRID TRAPDOOR COMMITMENTS. In [6] Catalano and Visconti presented the notion of *hybrid* trapdoor commitment schemes (in the context of constructing concurrent zero-knowledge proofs). Informally an hybrid trapdoor commitment scheme is a general commitment primitive that allows for two commitment parameters generation algorithms **HGen** and **HTGen**. If the commitment parameters are obtained as the output of **HGen**, then the resulting scheme is an unconditionally binding commitment scheme, while if the parameters are generated by **HTGen**, the produced scheme is actually a trapdoor commitment scheme. Moreover, no polynomially bounded adversary, taking as input only the (public) commitment parameters, should be able to tell the difference between parameters generated from **HGen** and parameters produced by **HTGen**. In [6], the authors show that 1) non-interactive hybrid trapdoor commitments can be constructed from any one-way function in the SRS model; and 2) efficient non-interactive hybrid trapdoor commitments can be constructed under standard number-theoretic assumptions in both the SRS and the TP models.

2.3 Mercurial Commitments

We now define mercurial commitments introduced by Chase et al. [8]. Our definition will be similar, but *stronger* than the definition from [8]. There are two reasons for making the change. First, all the efficient constructions in [8] and here will anyway satisfy the stronger definition. More importantly, by making our definition stronger we will also make it noticeably simpler (and shorter!) than the definition of [8]. More detailed comparison will be presented later.

PLAIN MERCURIAL COMMITMENTS. Such commitment schemes consist of seven efficient algorithms: $\mathcal{C} = (\text{MCom-Gen}, \mathbb{H}\text{Com}, \mathbb{H}\text{Open}, \mathbb{H}\text{Ver}, \mathbb{S}\text{Com}, \mathbb{S}\text{Open}, \mathbb{S}\text{Ver})$. The first four algorithms (**MCom-Gen**, **HCom**, **HOpen**, **HVer**) follow the syntax (and the functionality!) of regular commitment schemes (see Section 2.2). Namely, generation algorithm **MCom-Gen**(1^k), where k is the security parameter, outputs a public mercurial commitment key mpk . Given a message $M \in \mathcal{M}$, the *hard-commit* algorithm $\mathbb{H}\text{Com}_{mpk}(M; R)$ produces a *hard-commitment* string C for M . We will sometimes write $C \leftarrow \mathbb{H}\text{Com}_{mpk}(M)$. Similarly, the *hard-opening* algorithm $\mathbb{H}\text{Open}_{mpk}(M; R)$ (which is supposed to be run using the same value R as the hard-commit algorithm) produces a *hard-decommitment* value π for C . Finally, the *hard-verification* algorithm $\mathbb{H}\text{Ver}_{mpk}(M, C, \pi)$ accepts (i.e., outputs 1) if it thinks π proves that C is indeed a valid hard-commitment to M . We require that for all $M \in \mathcal{M}$, $\mathbb{H}\text{Ver}_{mpk}(m, \mathbb{H}\text{Com}_{mpk}(M; R), \mathbb{H}\text{Open}_{mpk}(M; R)) = 1$ holds with all but negligible probability.

We now turn to the novel “soft algorithms”. The *soft-commit* algorithm $\mathbb{S}\text{Com}_{mpk}(\cdot; R)$ produces a *soft-commitment* string C (to no message in partic-

ular). We will sometimes write $C \leftarrow \mathbb{SCom}_{mpk}()$. The *soft-opening* algorithm $\mathbb{SOpen}_{mpk}(M, \text{flag}; R)$, where $M \in \mathcal{M}$ and $\text{flag} \in \{\mathbb{H}, \mathbb{S}\}$ now produces a *soft-decommitment* τ to M , which should say that “if the commitment produced using R can be hard-opened at all, then it would open to M ”. A bit more precisely, if $\text{flag} = \mathbb{H}$, then τ is supposed to “correspond” to the hard-commitment $C = \mathbb{HCom}_{mpk}(M; R)$, and if $\text{flag} = \mathbb{S}$, then τ is a fake soft-decommitment “corresponding” to the soft-commitment $C = \mathbb{SCom}_{mpk}(); R$. Either one of these cases is verified using the *soft-verification* algorithm $\mathbb{SVer}_{mpk}(M, C, \tau)$, which outputs 1 if it thinks that C could potentially be hard-opened to M in the future (which, intuitively, should be the case only when τ was produced from a hard-commitment). Specifically, we require that for all $M \in \mathcal{M}$, $\mathbb{SVer}_{mpk}(M, \mathbb{HCom}_{mpk}(M; R), \mathbb{SOpen}_{mpk}(M, \mathbb{H}; R)) = 1$ holds with all but negligible probability, and similarly $\mathbb{SVer}_{mpk}(M, \mathbb{SCom}_{mpk}(); R, \mathbb{SOpen}_{mpk}(M, \mathbb{S}; R)) = 1$ holds with all but negligible probability.

We notice that in many cases (including all our constructions) the soft-decommitment τ to a hard-commitment C will consist of some proper part of the hard-decommitment π , and, correspondingly, the soft-verification algorithm will perform a proper subset of the tests performed by the hard-verification algorithm. For a lack of better name, we call such natural mercurial commitments *proper*.

SECURITY. The binding property of plain mercurial commitments consists of two requirements, stating that a valid hard- or soft-opening of C to some M implies that C can not be then hard-opened to any other message $M' \neq M$:

- **Mercurial Binding:** Having the knowledge of mpk , it is computationally hard for the PPT adversary \mathcal{A} to come up with C, M, π, M', π' (resp. C, M, τ, M', π') such that π (respectively, τ) is a valid hard- (respectively soft-) decommitment of C to M and π' is a valid hard-decommitment of C to M' , but $M \neq M'$ (such a tuple is said to cause a *hard* (respectively *soft*) *collision*). That is, \mathcal{A} cannot find a value C which it can hard- or soft-open in one way and then hard-open in a different way.

We remark that for proper mercurial commitments it suffices to prove that no soft collisions can be found.

As for the analog of the hiding property, we require that not only hard-commitments to some M look indistinguishable from soft-commitments (to “nothing”), but this continues to hold even if they are both soft-opened to M (notice that by the mercurial binding property, the hard-commitment to M cannot be soft-opened to anything other than M).

- **Mercurial Hiding:** No PPT adversary (who knows mpk) can find $M \in \mathcal{M}$ for which it can distinguish a random “real” hard-commitment/soft-decommitment tuple $(M, \mathbb{HCom}_{mpk}(M; R), \mathbb{SOpen}_{mpk}(M, \mathbb{H}; R))$ from a random “fake” soft-commitment/soft-decommitment tuple $(M, \mathbb{SCom}_{mpk}(); R), \mathbb{SOpen}_{pk}(M, \mathbb{S}; R)$.

(TRAPDOOR) MERCURIAL COMMITMENTS. Such commitment schemes consist of ten efficient algorithms: $\mathcal{C} = (\text{TrMCom-Gen}, \mathbb{H}\text{Com}, \mathbb{H}\text{Open}, \mathbb{H}\text{Ver}, \mathbb{S}\text{Com}, \mathbb{S}\text{Open}, \mathbb{S}\text{Ver}, \text{MFake}, \mathbb{H}\text{Equiv}, \mathbb{S}\text{Equiv})$. The generation algorithm $\text{TrMCom-Gen}(1^k)$, where k is the security parameter, outputs a public mercurial commitment key mpk and a secret mercurial *trapdoor key* msk . Once mpk is fixed, the meaning of $\mathbb{H}\text{Com}$, $\mathbb{H}\text{Open}$, $\mathbb{H}\text{Ver}$, $\mathbb{S}\text{Com}$, $\mathbb{S}\text{Open}$ and $\mathbb{S}\text{Ver}$ is exactly the same as for plain mercurial commitments. In particular, we will require that these algorithms satisfy the usual mercurial hiding and binding properties of the plain mercurial commitment schemes.

The trapdoor key msk is used in the algorithms MFake , $\mathbb{H}\text{Equiv}$ and $\mathbb{S}\text{Equiv}$ to break the binding property of commitments. The algorithm $\text{MFake}_{msk}(\cdot; R)$ is somewhat similar in spirit to the soft-commitment algorithm $\mathbb{S}\text{Com}_{mpk}$ and produces “fake” commitment C , initially not associated to any message M . The meaning of the other two algorithms $\mathbb{H}\text{Equiv}_{msk}(M; R)$ and $\mathbb{S}\text{Equiv}_{msk}(m; R)$ is also similar to that of the corresponding algorithms $\mathbb{H}\text{Open}_{mpk}$, $\mathbb{S}\text{Open}_{mpk}$, except they *always* operate on the fake commitments C not really associated to any message. Specifically, $\mathbb{H}\text{Equiv}(M; R)$ produces a supposedly valid hard-opening π (called *hard-fake*) of the fake commitment $C = \text{MFake}(\cdot; R)$ to M , while $\mathbb{S}\text{Equiv}(M; R)$ produces a supposedly valid soft-opening τ (called *soft-fake*) of the fake commitment $C = \text{MFake}(\cdot; R)$ to M . In particular, we require that for all $M \in \mathcal{M}$, $\mathbb{H}\text{Ver}_{mpk}(M, \text{MFake}_{mpk}(\cdot; R), \mathbb{H}\text{Equiv}_{mpk}(M; R)) = 1$ holds with all but negligible probability, and similarly $\mathbb{S}\text{Ver}_{mpk}(M, \text{MFake}_{mpk}(\cdot; R), \mathbb{S}\text{Equiv}_{mpk}(M; R)) = 1$ holds with all but negligible probability. While the ability to soft-fake such bogus commitments is consistent with the previous ability of soft-opening, the ability to hard-fake them certainly contradicts the binding property that we had, and this is exactly the function of the trapdoor key msk !

Somewhat similar to the equivocation property of trapdoor commitments, we require that trapdoor mercurial commitments satisfy three related equivocation conditions. In each of them we say that no efficient distinguisher \mathcal{A} can non-negligibly tell apart the corresponding “real” from the corresponding “ideal” game, *even if it is given the trapdoor key msk at the beginning of each real or ideal game*. In the following, the value R is always random.

- **HH Equivocation:** The real game consists of \mathcal{A} choosing $M \in \mathcal{M}$ and getting back $(M, \mathbb{H}\text{Com}_{mpk}(M; R), \mathbb{H}\text{Open}_{mpk}(M; R))$; while the ideal game consists of \mathcal{A} choosing $M \in \mathcal{M}$ and getting back $(M, \text{MFake}_{msk}(\cdot; R), \mathbb{H}\text{Equiv}_{msk}(M; R))$.
- **HS Equivocation:** The real game consists of \mathcal{A} choosing $M \in \mathcal{M}$ and getting back $(M, \mathbb{H}\text{Com}_{mpk}(M; R), \mathbb{S}\text{Open}_{mpk}(M, \mathbb{S}; R))$; while the ideal game consists of \mathcal{A} choosing $M \in \mathcal{M}$ and getting back $(M, \text{MFake}_{msk}(\cdot; R), \mathbb{S}\text{Equiv}_{msk}(M; R))$.
- **SS Equivocation:** The real game consists of \mathcal{A} getting the value $C = \mathbb{S}\text{Com}_{mpk}(\cdot; R)$, then choosing $M \in \mathcal{M}$, and finally getting $\mathbb{S}\text{Open}_{mpk}(M, \mathbb{S}; R)$; while ideal game consists of \mathcal{A} getting the value $C = \text{MFake}_{mpk}(\cdot; R)$, then choosing $M \in \mathcal{M}$, and finally getting $\mathbb{S}\text{Equiv}_{mpk}(M; R)$.

Notice that similar-looking SH condition does not make sense in the real life (due to mercurial binding). Next, HS and SS Equivocations easily imply the Mercurial Hiding property, so it does not need to be checked. Also, for proper mercurial commitments it is easy to see that HH Equivocation implies HS Equivocation, so it is enough to check only HH and SS Equivocations.

RELATION TO THE ORIGINAL DEFINITION IN [8]. The main difference from [8] is in the equivocation property, which is considerably simpler to state and verify in our case. Moreover, it is also *stronger* than the definition of [8]. Essentially, the latter definition consists of playing an arbitrary composition of HH, HS and SS Equivocation games either in the real, or in the ideal world,⁵ but where the distinguisher \mathcal{A} is *not given the trapdoor key msk* . In this scenario the usual hybrid argument does not work (since \mathcal{A} cannot simulate stuff in the ideal world by himself), so one cannot reduce the composed game to the one of the three atomic HH, SE or SS games. As a result, one has to build a full-fledged simulator, and formally argue that it fools the distinguisher. In contrast, in our scenario the hybrid argument easily works, so the security of our 3 atomic games easily implies the security of the composed game *even if the distinguisher knows msk* .

KNOWN CONSTRUCTIONS. Chase et al. [8] gave several elegant constructions of (trapdoor) mercurial commitments from the following assumptions:

- *One-way functions*. This construction works in the plain model but unfortunately is interactive. All next constructions are non-interactive.
- *Non-interactive zero-knowledge (NIZK) proofs for all of NP* [2], [15] and *unconditionally-binding commitment schemes*. However, this construction is mainly of theoretical interest, since all known NIZK constructions (especially for all of NP) are extremely inefficient. Interestingly, it also does *not* satisfy our stronger definition. However, in the sequel we will provide more general constructions (from one-way functions) which satisfy our stronger definition in the trusted parameters model and are still more efficient than this construction.⁶

⁵ There is one other, more syntactic strengthening that we had to make in order to simplify the definition. Namely, in the more general definition of [8] one could have syntactically unrelated real and ideal experiments for generating mpk , so it did not make sense to give msk to \mathcal{A} in the real game. In contrast, we insist that the public key generation even in the real world can be carried by generating both the public and the trapdoor key. While slightly more restrictive, since (1) all our efficient constructions in the trusted parameters model satisfy this restricted notion of key generation and (2) it considerably simplifies (and also *strengthens*) the definition, we feel it is very justified.

⁶ We remark, however, that the NIZK is in the SRS model, while our OWF-based construction satisfying the stronger definition will be in the TP model. We can make SRS-based constructions by either building trapdoor commitments in the SRS model (which is known how to do from one-way *permutations* or specific number-theoretic assumptions), or by using our technique from Section 4 (while reverting to the weaker definition of [8]).

- *Claw-free permutations* [22]. This construction give only bit mercurial commitment, and will be a special case of our first general construction from bit trapdoor commitments.
- *Discrete log*. This is a “distillation” of the original construction implicitly used in [28]; it supports long messages and is pretty efficient. It will be a special case of our second construction when used with the corresponding discrete-log based bit trapdoor commitment.
- *Factoring*. This is a new construction which supports long messages and is relatively efficient. It will be a special case of our second construction when used with the corresponding factoring-based bit trapdoor commitment.

IMPLICATIONS TO (TRAPDOOR) COMMITMENTS. It is simple to see that by “ignoring” all the “soft” algorithms of a secure plain/trapdoor commitment scheme, we immediately get a plain/trapdoor regular commitment scheme. (Concentrating, for example, on a slightly more complicated “trapdoor case”, $\mathbb{H}\text{Com}$ plays the role of Com , $\mathbb{H}\text{Open}$ — of Open , $\mathbb{H}\text{Ver}$ — of Ver , $\mathbb{M}\text{Fake}$ — of Fake , and $\mathbb{H}\text{Equiv}$ — of Equiv .) In the following, we show two simple constructions proving that the converse of this statement is true as well.

3 General Construction from (Trapdoor) Bit Commitments

As advocated in the introduction, we will first consider the construction of plain mercurial bit commitments from regular bit commitments, and then argue that the same construction extends to the trapdoor case as well.

BUILDING PLAIN MERCURIAL COMMITMENTS. Assume $\mathcal{C} = (\text{Com-Gen}, \text{Com}, \text{Open}, \text{Ver})$ is a regular bit commitment scheme. Define plain mercurial commitment $\mathcal{C}' = (\text{MCom-Gen}, \mathbb{H}\text{Com}, \mathbb{H}\text{Open}, \mathbb{H}\text{Ver}, \mathbb{S}\text{Com}, \mathbb{S}\text{Open}, \mathbb{S}\text{Ver})$ for a bit $b \in \{0, 1\}$ as follows (we set $\text{MCom-Gen} = \text{Com-Gen}$ and let pk be the corresponding public key):

- $\mathbb{H}\text{Com}_{pk}(b; (r_0, r_1))$: output $(c_0, c_1) = (\text{Com}_{pk}(b; r_0), \text{Com}_{pk}(1 - b; r_1))$. Notice, commitment to 0 changes its place from left to right depending on b .
- $\mathbb{H}\text{Open}_{pk}(b; (r_0, r_1))$: output $(d_0, d_1) = (\text{Open}(b; r_0), \text{Open}(1 - b; r_1))$.
- $\mathbb{H}\text{Ver}_{pk}(b, (c_0, c_1), (d_0, d_1))$: accept if and only if $\text{Ver}_{pk}(b, c_0, d_0) = \text{Ver}_{pk}(1 - b, c_1, d_1) = 1$.
- $\mathbb{S}\text{Com}_{pk}(\cdot; (r_0, r_1))$: output $(c_0, c_1) = (\text{Com}_{pk}(0; r_0), \text{Com}_{pk}(0; r_1))$.
- $\mathbb{S}\text{Open}_{pk}(b, \text{flag}; (r_0, r_1))$: irrespective of $\text{flag} \in \{\mathbb{H}, \mathbb{S}\}$, output $d = \text{Open}(0; r_b)$.
- $\mathbb{S}\text{Ver}_{pk}(b, (c_0, c_1), d)$: accept if and only if $\text{Ver}_{pk}(0, c_b, d) = 1$.

The correctness of the scheme is obvious. Intuitively, mercurial commitment to $b = 0$ looks $(0, 1)$, to $1 = (1, 0)$, and the fake — $(0, 0)$. Since the soft-opening of the hard commitment only opens the corresponding left or right 0, the fake commitment can indeed be soft-opened in both way, by honestly opening the appropriate left or right 0. On the other hand, seeing a hard-opening of some

commitment $C = (c_0, c_1)$ (to some bit b) opens to 1 one of the two regular commitments, while the subsequent soft-opening of C to $(1 - b)$ would then open this regular commitment to 0, which contradicts binding. Below, we formalize this in a straightforward manner.

Mercurial Binding. Since the mercurial commitment is proper, we only need to rule out soft collisions. For that, assume the attacker can find a soft collision. By symmetry, let us assume that 1 is the softly-opened message, and 0 is the hardly-opened one. So we denote this collision by $C = ((c_0, c_1), d_0, d_1, d'_1)$ where $\text{Ver}(0, c_0, d_0) = \text{Ver}(1, c_1, d_1) = \text{Ver}(0, c_1, d'_1) = 1$. But then c_1 can be opened to both 0 and 1, a contradiction to the binding property of \mathcal{C} .

Mercurial Hiding. Assume first $b = 0$. Then, the “real” hard-commitment/soft-decommitment tuple $(\mathbb{H}\text{Com}(0; (r_0, r_1)), \mathbb{S}\text{Open}(0, \mathbb{H}; (r_0, r_1)))$ looks like $(\text{Com}(0; r_0), \text{Com}(1; r_1), \text{Open}(0; r_0))$, while the corresponding “fake” tuple $(\text{Fake}(\cdot; (r_0, r_1)), \mathbb{S}\text{Open}(0, \mathbb{S}; (r_0, r_1)))$ looks like $(\text{Com}(0; r_0), \text{Com}(0; r_1), \text{Open}(0; r_0))$. Clearly, such distributions are indistinguishable if $\text{Com}(0)$ cannot be distinguished from $\text{Com}(1)$, which follows from the hiding property of \mathcal{C} . A similar argument holds for $b = 1$ as well.

TRAPDOOR CASE. The extension to the trapdoor case is simple as well. We now have additional algorithms **Fake** and **Equiv** for trapdoor commitments, and need to build the corresponding algorithms **MFake**, $\mathbb{H}\text{Equiv}$ and $\mathbb{S}\text{Equiv}$ for mercurial commitments.

- $\text{MFake}_{sk}(\cdot; (r_0, r_1))$: output $(\text{Fake}_{sk}(\cdot; r_0), \text{Fake}_{sk}(\cdot; r_1))$.
- $\mathbb{H}\text{Equiv}_{sk}(b; (r_0, r_1))$: output $(\text{Equiv}_{sk}(b; r_0), \text{Equiv}_{sk}(1 - b; r_1))$.
- $\mathbb{S}\text{Equiv}_{sk}(b; (r_0, r_1))$: output $\text{Equiv}_{sk}(0; r_b)$.

Correctness is obvious from definition. As for hiding, we only need to argue $\mathbb{H}\mathbb{H}$ and $\mathbb{S}\mathbb{S}$ Equivocations (since this is a proper mercurial commitment). Both are simple corollaries of the regular Equivocation properties of trapdoor commitments.

HH Equivocation. Let us assume $b = 0$, since $b = 1$ is symmetric. Then $(\mathbb{H}\text{Com}(0; (r_0, r_1)), \mathbb{H}\text{Open}(0; (r_0, r_1)))$ is equal to $D_{\text{real}} = (\text{Com}(0; r_0), \text{Com}(1; r_1), \text{Open}(0; r_0), \text{Open}(1; r_1))$, while $(\text{MFake}(\cdot; (r_0, r_1)), \mathbb{H}\text{Equiv}(0; (r_0, r_1)))$ is equal to $D_{\text{ideal}} = (\text{Fake}(\cdot; r_0), \text{Fake}(\cdot; r_1), \text{Equiv}(0; r_0), \text{Equiv}(1; r_1))$. Since r_0 and r_1 are independent, this amounts to two independent applications of the regular Equivocation property to bits 0 and 1, respectively. Notice, though, already for this simple hybrid argument *we are using the fact that the attacker knows the trapdoor key sk !* To be precise, we must first consider a hybrid distribution $D_{\text{hyb}} = (\text{Fake}(\cdot; r_0), \text{Com}(1; r_1), \text{Equiv}(0; r_0), \text{Open}(1; r_1))$, and then show $D_{\text{real}} \approx D_{\text{hyb}}$ (here we only need pk to sample $(\text{Com}(1; r_1), \text{Open}(1; r_1))$) and $D_{\text{hyb}} \approx D_{\text{ideal}}$ (here we need sk to sample $(\text{Fake}(\cdot; r_0), \text{Equiv}(0; r_0))$).

SS Equivocation. In the real experiment, the attacker is first getting $(\text{Com}(0; r_0), \text{Com}(0; r_1))$, then he has to choose a bit b , after which he gets $\text{Open}(0; r_b)$. In the ideal game, the attacker is getting $(\text{Fake}(\cdot; r_0), \text{Fake}(\cdot; r_1))$, then he has to choose a bit b , after which he gets $\text{Equiv}(0; r_b)$. By symmetry, the choice of b does not matter here, so we can assume $b = 0$, so it suffices to argue $(\text{Com}(0; r_0), \text{Com}(0; r_1))$,

$\text{Open}(0; r_0) \approx (\text{Fake}(\cdot; r_0), \text{Fake}(\cdot; r_1), \text{Equiv}(0; r_0))$. Once again, this follows by the hybrid argument, by considering an intermediate distribution $(\text{Fake}(\cdot; r_0), \text{Com}(\cdot; r_1), \text{Equiv}(0; r_0))$ and using the fact that in the second hybrid the attacker can compute $(\text{Fake}(\cdot; r_0), \text{Open}(0; r_0))$.

COMPARISON TO [8]. The above construction is a very simple generalization of the one in [8], who used the following family of trapdoor bit commitments [27] obtained from any family of claw-free permutations [22] (f_0, f_1) . Informally, recall that these are pairs of permutations where one cannot find a “claw” (r_0, r_1) satisfying $f_0(r_0) = f_1(r_1)$; also it is assumed that there exists a trapdoor f_0^{-1} allowing one to invert f_0 (in our application, we will not need a similar trapdoor for f_1). Now, to trapdoor commit to a bit b we can sample $f_b(r_b)$ (decommitment is r_b), while the knowledge of the trapdoor f_0^{-1} provides easy fake pairs: the fake commitment $c = f_1(r_1)$ (for random r_1) can be opened to 0 by giving $r_0 = f_0^{-1}(c)$, and to 1 — by giving r_1 .

We remark, though, that the equivocality proof of our extension is indeed considerably shorter, — which is what it should be for such a simple construction! — than the corresponding proof [8]. Also, our construction implies mercurial commitments from other bit commitments which are not necessarily induced by claw-free permutations, such as the general construction of [16] from any Σ -protocol for a hard language, the factoring construction of [33], the Paillier construction of [5] or the ad hoc (g^{r_0}, h^{r_1}) -construction mentioned in Section 2.2 (and, of course, the one-way function construction from Theorem 5).

4 Mercurial Commitments from One-Way Functions in the SRS model

We now discuss a construction of a non-interactive bit trapdoor mercurial commitment scheme in the SRS model which requires only a non-interactive bit hybrid trapdoor commitment scheme as underlying building block. Since the latter can be constructed from one-way functions [6], we have that the same holds for the former. We stress that for this construction we use the original definition of mercurial commitments given in [8] where the shared random string used in the simulated game is computationally indistinguishable from the real random string. For lack of space we omit the original definitions of mercurial commitments given in [8] and of hybrid trapdoor commitments given in [6] (in particular the reader is referred to [6] for details about the constructions of hybrid trapdoor commitments).

OVERVIEW OF THE TECHNIQUE. In Section 3 we have shown a construction of mercurial commitments that can be based on the trapdoor commitment scheme proposed by Feige and Shamir [16]. This scheme needs, as common parameter, an Hamiltonian graph for which it is hard to compute an Hamiltonian cycle, but such that knowledge of a cycle allows one to equivocate. This can be realized in the trusted parameters model by generating the required instance of the Hamiltonian cycle language from the hardness of inverting a one-way function. In the

SRS model, this construction is known to work assuming that one-way *permutations* exist. In a nutshell, this is because a piece of the random string infers the computationally infeasible problem of inverting a one-way permutation (that, in turn, can be reduced to an instance of finding a cycle in an Hamiltonian graph), and in order for this to work one needs to make sure that a corresponding inverse actually exists.

In order to build a solution based on any one-way function in the SRS model, we start from the following observation. Since the possibility of computing commitments that can be equivocated is required by the simulator only (we stress that we are now using the original definition of mercurial commitments given in [8]), we could construct mercurial commitments by using regular (i.e., we do not require the equivocal property) commitments in the SRS model for the real game and trapdoor commitments in the trusted parameters model for the simulated game. In order for this idea to work, however, we need a scheme that can be used either as a trapdoor commitment in the trusted parameters model or as a regular one in the SRS model, but it is infeasible for the adversarial receiver to distinguish the two cases. In particular the trusted parameters of the simulated game must be computationally indistinguishable from a random string.

Commitment schemes realizing this requirement have been recently studied, defined and constructed by Catalano and Visconti [6] under the sole assumption that one-way functions exist.

A construction of non-interactive bit trapdoor mercurial commitments on top of a non-interactive bit hybrid trapdoor commitment scheme can be very easily obtained from the construction shown in Section 3. Indeed, it suffices to replace the algorithms of the regular non-interactive bit trapdoor commitment scheme with the corresponding ones of the hybrid one. The only additional step is that since the hybrid trapdoor commitment scheme has two algorithms for the generation of the commitment parameters, one is used in the real game and the other is used in the simulated game. In the SRS model, the former simply outputs a random string (i.e., the parties can deterministically extract the commitment parameters from a random string) while the latter outputs a random/pseudorandom string along with a trapdoor.

5 Efficient Construction from (Trapdoor) Bit Commitments with Σ -Protocols

The problem with the previous generic constructions is the fact that they only allows one to commit to one bit. Of course, we can always commit to many bits by following the “bit-by-bit” approach, but this is inefficient. Alternatively, we can try to utilize a multi-bit plain/trapdoor commitment scheme in the previous construction, but it is easy to see that the resulting length of the commitment will be linearly proportional to the number of *messages* that we want to commit to. This essentially means that setting this number to 2 — as we did in Section 3 — and doing the bit-by-bit composition is the best we can do if we try to extend the previous approach.

Instead, in this section we present our main construction which will directly transform a plain/trapdoor bit commitment \mathcal{C} into an efficient and (potentially) multi-bit plain/mercurial commitment \mathcal{C}' . However, we will need to assume an extra property from \mathcal{C} : there exists an efficient Σ -protocol Π proving that one knows a witness d that a given commitment c can be opened to 0. In this case, the message space of \mathcal{C}' will be the challenge space of the corresponding Σ -protocol. Thus, if Π will be challenge-efficient, we would get a direct, large-message mercurial commitment \mathcal{C}' .

CONSTRUCTION. Let $\mathcal{C} = (\text{Com-Gen}, \text{Com}, \text{Open}, \text{Ver})$ be a regular bit commitment scheme which has a Σ -protocol $\Pi = (\text{Start}, \text{Finish}, \text{Extract}, \text{Simul})$ for the relation (family) $\mathcal{R}_{pk} = \{(c, d) \mid \text{Ver}_{pk}(0, c, d) = 1\}$. Recall, this means that the verifier only gets a commitment c , and the prover also gets, as a witness, a valid opening d of c to 0. Also, assume \mathcal{M} is the challenge space for Π .

We then define plain mercurial commitment $\mathcal{C}' = (\text{MCom-Gen}, \mathbb{H}\text{Com}, \mathbb{H}\text{Open}, \mathbb{H}\text{Ver}, \mathbb{S}\text{Com}, \mathbb{S}\text{Open}, \mathbb{S}\text{Ver})$ for message space \mathcal{M} as follows (we set $\text{MCom-Gen} = \text{Com-Gen}$ and let pk be the corresponding public key):

- $\mathbb{H}\text{Com}_{pk}(m; (r_s, r_1))$: let $c_1 = \text{Com}_{pk}(1; r_1)$ be a commitment to 1, and $(a_1, z_1) = \text{Simul}_{pk}(c_1, m; r_s)$ be a fake first and last messages of Π which (here incorrectly) claim that c_1 is a commitment to 0 on challenge m . Output (c_1, a_1) .
- $\mathbb{H}\text{Open}_{pk}(m; (r_s, r_1))$: let $c_1 = \text{Com}_{pk}(1; r_1)$ and $(a_1, z_1) = \text{Simul}_{pk}(c_1, m; r_s)$ be as before. Set $d_1 = \text{Open}_{pk}(1; r_1)$ and output (d_1, z_1) .
- $\mathbb{H}\text{Ver}_{pk}(m, (c_1, a_1), (d_1, z_1))$: accept if and only if $\text{Ver}_{pk}(1, c_1, d_1) = 1$ (d_1 is correct decommitment to 1) and $\text{Check}_{pk}(c_1, a_1, m, z_1) = 1$ (the fake transcript on challenge m that c_1 is a commitment to 0 looks good).
- $\mathbb{S}\text{Com}_{pk}(\cdot; (r_p, r_0))$: let $c_0 = \text{Com}_{pk}(0; r_0)$ be a commitment to 0, and $d_0 = \text{Open}_{pk}(0; r_0)$ be the corresponding opening, and $a_0 = \text{Start}_{pk}(c_0, d_0; r_p)$ be a real first messages of Π which (correctly!) claims that c_0 is a commitment to 0.⁷ Output (c_0, a_0) .
- $\mathbb{S}\text{Open}_{pk}(m, \mathbb{H}; (r_s, r_1))$: let $c_1 = \text{Com}_{pk}(1; r_1)$ and $(a_1, z_1) = \text{Simul}_{pk}(c_1, m; r_s)$ be the fake transcript on challenge m that c_1 is a commitment to 0. Output z_1 .
- $\mathbb{S}\text{Open}_{pk}(m, \mathbb{S}; (r_p, r_0))$: let $c_0 = \text{Com}_{pk}(0; r_0)$, $d_0 = \text{Open}_{pk}(0; r_0)$, $a_0 = \text{Start}_{pk}(c_0, d_0; r_p)$, and $z_0 = \text{Finish}_{pk}(c_0, d_0, m; r_p)$ be the correct last flow to challenge m . Output z_0 .
- $\mathbb{S}\text{Ver}_{pk}(m, (c, a), z)$: accept if and only $\text{Check}_{pk}(c, a, m, z) = 1$ (the transcript (a, m, z) stating that c is a commitment to 0 is correct).

Intuitively, the honest hard-committer is supposed to send a commitment c to 1, but *fake* the transcript that he in fact committed to 0. On the other hand, a lying soft-committer can simply send a commitment c to 0, and now can

⁷ Notice, here the prover actually knows the value r_0 , and not just d_0 . So for efficiency reasons we might consider auxiliary-input Σ -protocols where P 's witness is actually r_0 itself. We will return to this point later.

(*honestly!*) respond to any challenge/message m that he gets subsequently, which allows him to soft-open the first flow to any message m .⁸ The binding security of this scheme comes from the fact that a hard-opening of c to 1, coupled with two soft-opening of the first flow a , must enable one to extract a legal witness, which is the hard-opening of c to 0, contradicting the binding of \mathcal{C} . Similarly, the hiding property of \mathcal{C} coupled with the zero-knowledge property of Σ -protocols imply that, without the hard-opening of c (which will tell if c is a commitment to 0 or 1), the real and fake behavior cannot be told apart. More formally,

Mercurial Binding. Since our commitment is proper, we only need to rule out soft collisions. This means that the attacker cannot find a commitment value (c, a) , a decommitment d_1 proving that c is a commitment to 1, two messages $m \neq m'$, and two valid responses z and z' claiming that c is a commitment to 0. By the special soundness of the Σ -protocol, $\text{Extract}(c, a, m, z, m', z')$ must be equal to a valid decommitment d_0 of c to 0. But then we found a way to open c to both 0 and 1 (via d_0 and d_1), contradicting the binding property of \mathcal{C} .

Mercurial Hiding. Take any message/challenge m . Then, the “real” hard-commitment/soft-decommitment tuple for m looks like is given by three values $(c = \text{Com}(1; r_1), (a, z) = \text{Simul}(c, m; r_s))$. Since our commitment is hiding, and $\text{Simul}(c, m)$ is publicly computable, we get that the above distribution is indistinguishable from $(c = \text{Com}(0; r_0), (a, z) = \text{Simul}(c, m; r_s))$. Now, since c has a proper witness $d_0 = \text{Open}(0; r_0)$, the special HVZK property of Π states that the distribution on (a, z) looks indistinguishable than the one obtained by a running a real protocol on input c , witness d_0 and challenge m . But this means that the above distribution is indistinguishable from $(c = \text{Com}(0; r_0), a = \text{Start}(c, d_0; r_p), z = \text{Finish}(c, d_0, m; r_p))$, which is exactly the triple corresponding to the “fake” soft-commitment/soft-decommitment procedures.

TRAPDOOR CASE. Recall, we now have additional algorithms **Fake** and **Equiv** for trapdoor commitments, and need to build the corresponding algorithms **MFake**, **HEquiv** and **SEquiv** for mercurial commitments. As a new technical property about the Σ -protocol, however, we will have to assume that $\Pi = \Pi_{pk}$ is strongly hiding w.r.t. a particular parameter generation procedure \mathcal{P} (see Section 2.1). The parameter generation procedure we will need generates random keys $(pk, sk) \leftarrow \text{Com-Gen}(1^k)$, picks a random r , computes $c = \text{Fake}_{sk}(\cdot; r)$, $d_0 = \text{Equiv}_{sk}(0; r)$, $d_1 = \text{Equiv}_{sk}(1; r)$, and sets the side information to (sk, d_1) , the input to be c , and the witness to be d_0 . As explained in Section 2.1, this is more of a technicality which seems to be always satisfied in any non-pathological scenario arising in practice. We call this property (*), and can now describe the claimed extension.

⁸ It might appear peculiar that we require an honest party to cook-up a fake proof in order to succeed, while having a dishonest party perform such a proof correctly! Here, however, the primitive we build legally allows a dishonest party to look “slightly like an honest party”. So the we force the honest party to do something slightly bad which might be “matched” by a good action of a dishonest party.

- $\text{MFake}_{sk}(\cdot; (r_p, r))$: let $c = \text{Fake}_{sk}(\cdot; r)$ be a fake commitment, $d_0 = \text{Equiv}_{sk}(0; r)$ be its fake opening to 0, and $a_0 = \text{Start}_{pk}(c, d_0; r_p)$ be a correct first flow of the Σ -protocol. Output (c, a_0) .
- $\mathbb{H}\text{Equiv}_{sk}(m; (r_p, r))$: let $c = \text{Fake}_{sk}(\cdot; r)$, $d_0 = \text{Equiv}_{sk}(0; r)$, and $a_0 = \text{Start}_{pk}(c, d_0; r_p)$ be as before. Compute the fake opening $d_1 = \text{Equiv}_{sk}(1; r)$ of c to 1, and the correct last message $z_0 = \text{Finish}_{pk}(c, d_0, m; r_p)$. Output (d_1, z_0) .
- $\mathbb{S}\text{Equiv}_{sk}(m; (r_p, r))$: let $c = \text{Fake}_{sk}(\cdot; r)$, $d_0 = \text{Equiv}_{sk}(0; r)$, and $a_0 = \text{Start}_{pk}(c, d_0; r_p)$ be as before. Compute the correct last message $z_0 = \text{Finish}_{pk}(c, d_0, m; r_p)$ and output z_0 .

Correctness is obvious from definition. As for hiding, we only need to argue HH and SS Equivocations (since this is a proper mercurial commitment).

HH Equivocation. Take any message m . Then $(\mathbb{H}\text{Com}(m; (r_s, r_1)), \mathbb{H}\text{Open}(m; (r_s, r_1)))$ is equal to $D_{\text{real}} = (c_1, d_1, a_1, z_1)$, where $c_1 = \text{Com}(1; r_1)$, $d_1 = \text{Open}(1; r_1)$, and $(a_1, z_1) = \text{Simul}(c_1, m; r_s)$. Since $\text{Simul}(c_1, m)$ is a public transformation, the Equivocality of \mathcal{C} implies that the above distribution is indistinguishable from $(c = \text{Fake}(\cdot; r), d_1 = \text{Equiv}(1; r), a_1, z_1)$, where $(a_1, z_1) = \text{Simul}(c, m; r_s)$. We are almost done, except we need to replace the above (a_1, z_1) by (a_0, z_0) obtained by running an honest execution of Π with witness $d_0 = \text{Equiv}(0; r)$. This is almost exactly the HVZK property, except we formally need to use the strong hiding property (*) described above. Indeed, in addition to the input c and the public parameter pk , which are allowed in the usual HVZK property, here the distinguisher also knows two extra pieces of information: the trapdoor key sk (given to him at the beginning of the game) and the fake decommitment $d_1 = \text{Equiv}(1; r)$. This is why we needed to assume that this extra information does not violate the HVZK property.

SS Equivocation. In the real soft-commit/soft-open experiment, the distinguisher (who knows sk) is first getting $c_0 = \text{Com}(0; r_0)$ and the correct first flow of the Σ -protocol showing that c_0 is indeed a commitment to 0 (using witness $d_0 = \text{Open}(0; r_0)$). He then chooses a message m , and gets a correct third flow to message m . To put differently, he simply plays the role of (malicious) verifier in the honest run of the Σ -protocol on pair (c_0, d_0) . Notice that the distinguisher's view can be perfectly simulated using some public probabilistic procedure $A_{sk}(c_0, d_0)$. Using the equivocation property of \mathcal{C} , the resulting distribution should be indistinguishable from $A_{sk}(c, d_0)$, where $c = \text{Fake}(\cdot; r)$ and $d_0 = \text{Equiv}(0; r)$. But, once again, it is easy to see that this view is exactly what the attacker gets in the ideal soft-commit/soft-open experiment.

GENERALIZATION. We already noticed in Footnote 7 that in the above definition of soft-commitment, the Prover actually knows the entire randomness r_0 and not just a witness $d_0 = \text{Open}(0; r_0)$. This, of course, is of any value only in a very few schemes where $r_0 \neq d_0$. However, it will come up in one of our examples (see Section 5.1). To accommodate this extension, we can consider Σ -protocol's where the prover needs all of r_0 for the completeness of the protocol (special soundness is still only for d_0). For plain mercurial commitments, this is all we need to change. For the trapdoor variant, however, we will need an extra property from our trapdoor commitment scheme in regards to equivocation. Namely,

in the fake commitment algorithm we need to be able to equivocate $c = \text{Fake}(\cdot; r)$ to 0 by obtaining not only a good looking value d_0 , but the entire randomness r_0 . Once this is ensured, we can easily support auxiliary input Σ -protocols.

5.1 Examples

Below we briefly give several efficient instantiations of our construction, by applying it to several efficient trapdoor commitment schemes with challenge-efficient Σ -protocols. Our examples will cover *all* the previous efficient schemes, and several more, all as part of one general framework. For each scheme we will just briefly mention which trapdoor commitment and Σ -protocol to use, since the remaining details are obvious and not very illuminating.

DISCRETE LOG CONSTRUCTION FROM [8],[28]. We will consider the ad-hoc scheme from Section 2.2, where $\text{Com}(0; r_0) = g^{r_0}$, $\text{Com}(1; r_1) = h^{r_1}$, and the trapdoor $sk = \log_g h$ (here $r_0, r_1 \neq 0$). We need a Σ -protocol to prove the knowledge of $r_0 = \log_g(c)$, where c is the claimed commitment to 0. Of course, a natural thing to do is to take Schnorr protocol, but this will result in a slightly different (but equally efficient) scheme than what we are after. Instead, we will use a bit less esthetic but equally effective Σ -protocol. In the first flow the prover sends a value $T = g^t$ (for random t), he gets challenge m , and responds with $z = (t - m)/r_0 \bmod q$ (which is defined since $r_0 \neq 0$). The verifier checks if $g^m c^z = T$ (indeed, $m + r_0 z = t$, as needed). It is simple to see that this is indeed a Σ -protocol for the knowledge of the discrete log, and that by plugging it into our construction we get exactly the discrete log construction from [8, 28].

We also remark what we could use a better known discrete-log commitment $\text{Com}(0) = g^{r_0}$, $\text{Com}(1) = hg^{r_1}$, coupled with either Schnorr Σ -protocol, or the one presented above. We will get yet another (equally efficient) solution.

FACTORIZING CONSTRUCTION FROM [8]. This scheme will use the generalization of our construction to use auxiliary inputs, as explained earlier. Let us start with a well-known factoring-based trapdoor bit commitment from a corresponding claw-free permutation pair: the public parameter is a random square U , and $\text{Com}(0; r_0) = r_0^2 \bmod n$, $\text{Com}(1; r_1) = Ur_1^2 \bmod n$ (the trapdoor is the square root of U). Here we need a Σ -protocol for the knowledge of the square root. As we mentioned in Section 2.1, using Fiat-Shamir protocol [18] is not communication- or challenge-efficient. Instead, we use the auxiliary input Ong-Schnorr protocol [30]. For that one need to know 2^ℓ -th square root of $\text{Com}(0)$, so we modify $\text{Com}(0; r_0) = r_0^{2^\ell} \bmod n$ (but leave $\text{Com}(1; r_1) = Ur_1^2 \bmod n$). We notice, that although the decommitment to 0 is “only” the square root $d_0 = r_0^{2^{\ell-1}}$, and not r_0 itself, the fake commitment should enable us to extract (using sk) the 2^ℓ -th root from c_0 , and not just a mere square root. Of course, this is easy to achieve by defining $\text{Fake}(\cdot; r) = r^{2^\ell}$, and “fully opening” it to 0 by giving r , and to 1 — by giving $r^{2^{\ell-1}}/\sqrt{U}$. With these changes we get *precisely* the factoring construction from [8]. We also notice that by using a different claw-free permutation $(r_0^2, 4r_1^2)$ [22] defined over the so called Williams integers, we can slightly simplify the scheme and set $U = 4$.

NEW RSA-BASED CONSTRUCTION. Here we could use the RSA-based trapdoor commitment $\text{Com}(0; r_0) = r_0^e \bmod n$, $\text{Com}(1; r_1) = yr_1^e \bmod n$, where y is a public parameter, whose e -th root is the trapdoor key. Here we simply need the Σ -protocol proof of knowledge of the e -th root, which is just the GQ protocol [24]. To have the protocol to be challenge-efficient, though, we will need to use a relatively large e .

ALTERNATIVE FACTORING CONSTRUCTION. We can use the following factoring-based commitment of [33] (slightly modified for easier Σ -protocols and specialized to bits). The public key is $n = pq$, where $p = 2p' + 1$, $q = 2q' + 1$ are safe primes, and all the operations are performed in the subgroup Q_n of quadratic residues whose generator g is also part of parameters. Notice, $|Q_n| = p'q'$. Let C be a large enough constant (anything larger than n will do). Then $\text{Com}(0; r_0) = g^{C+r_0} \bmod n$, $\text{Com}(1; r_1) = g^{r_1} \bmod n$ (here r_0, r_1 are random from 0 to n (which is statistically close to $\varphi(n)$, which is the “true range” we are aiming for). The trapdoor is the value $|Q_n| = p'q'$. In this case the Σ -protocol we need to again the one of knowledge of discrete-log, but in the groups of unknown order. As mentioned before, such (computationally sound) protocol is given by [19,, 9].

PAILLIER-BASED SCHEME. Finally, we mention another trapdoor commitment based on the hardness of finding n -th roots over \mathbb{Z}_{n^2} (where n is the the product of two safe primes, for simplicity), which is implicit in [11]. Here the public parameters will include a generator g in the subgroup S of n -th powers in $\mathbb{Z}_{n^2}^*$, and the n -th root u of g will be the trapdoor. Next, $\text{Com}(0; r_0) = r_0^n \bmod n^2$, $\text{Com}(1; r_1) = gr_1^n \bmod n^2$ (here $r_0, r_1 \in \mathbb{Z}_n^*$). This scheme is perfectly hiding and computationally binding assuming it is hard to take n -th root over \mathbb{Z}_{n^2} , and could be viewed as yet another claw-free based construction. The Σ -protocol for commitment to 0 is simply the Σ -protocol for knowing the n -th root. This protocol is very similar to the GQ protocol and is formally analyzed by [10].

6 Concluding Remarks

We believe that our results elucidate the notion of mercurial commitments, put them in their place on the map of cryptographic assumptions, and better explain the rational following the previous constructions of [28],[8]. We hope that mercurial commitments will find more interesting applications in the future.

This paper joins two independent papers that can be found at [7],[13].

Acknowledgments. The second author would like to thank Leonid Reyzin for several insightful conversations about mercurial commitments, and Tal Malkin for giving a talk inspiring this research. The work of the first and third authors has been supported in part by the European Commission through the IST Programme under Contract IST-2002-507932 ECRYPT. The work of the second author was supported in part through NSF Career Award CCR-0133806 and NSF grant CCR-0311095. The work of the third author is also supported in part through the FP6 program under contract FP6-1596 AEOLUS.

References

1. M. Abdalla and L. Reyzin. A new forward-secure digital signature scheme. In T. Okamoto, editor, *Advances in Cryptology—ASIACRYPT 2000*, volume 1976 of *Lecture Notes in Computer Science*, pages 116–129, Kyoto, Japan, 3–7 Dec. 2000. Springer-Verlag. Full version available from the Cryptology ePrint Archive, record 2000/002, <http://eprint.iacr.org/>.
2. Manuel Blum, Alfredo De Santis, Silvio Micali, and Guiseppe Persiano. Non-interactive zero-knowledge. *SIAM Journal of Computing*, 20(6), 1991.
3. Joan Boyar, S. A. Kurtz, Mark W. Krentel. A Discrete Logarithm Implementation of Perfect Zero-Knowledge Blobs. In *J. of Cryptology*, 2(2):63–76, 1990.
4. G. Brassard, D. Chaum, and C. Crépeau. Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37(2):156–189, Oct. 1988.
5. Dario Catalano, Rosario Gennaro, Nick Howgrave-Graham, Phong Q. Nguyen. Paillier’s cryptosystem revisited. In *ACM Conference on Computer and Communications Security 2001*, pp. 206–214.
6. D. Catalano and I. Visconti. Hybrid Trapdoor Commitments and Their Applications. In *32nd International Colloquium on Automata, Languages, and Programming (ICALP 05)*, volume 3580 of *Lecture Notes in Computer Science*, pages 298–310. Springer-Verlag, 2005.
7. D. Catalano and I. Visconti. Non-Interactive Mercurial Commitments from One-Way Functions. Cryptology ePrint Archive, 2005.
8. Melissa Chase, Alexander Healy, Anna Lysyanskaya, Tal Malkin and Leonid Reyzin. Mercurial Commitments with Applications to Zero-Knowledge Sets. In *Proc. of EUROCRYPT*, pp. 422–439, 2005.
9. Ivan Damgård, Eiichiro Fujisaki. A Statistically-Hiding Integer Commitment Scheme Based on Groups with Hidden Order. In *ASIACRYPT 2002*, pp. 125–142.
10. Ivan Damgård, Mats Jurik. A Generalisation, a Simplification and Some Applications of Paillier’s Probabilistic Public-Key System. *Public Key Cryptography 2001*, pp. 119–136.
11. I. Damgård and J. B. Nielsen. Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In M. Yung, editor, *Advances in Cryptology—CRYPTO 2002*, Lecture Notes in Computer Science. Springer-Verlag, 18–22 Aug. 2002.
12. I. B. Damgård, T. P. Pedersen, and B. Pfitzmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. *Journal of Cryptology*, 10(3):163–194, Summer 1997.
13. Y. Dodis. Minimal Assumptions for Efficient Mercurial Commitments. Cryptology ePrint Archive, Report 2005/438.
14. Y. Dodis and L. Reyzin. On the power of claw-free permutations. In *Conference on Security in Communication Networks*, 2002.
15. U. Feige, D. Lapidot, and A. Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM J. Computing*, 29(1), 1999.
16. U. Feige and A. Shamir. Zero knowledge proofs of knowledge in two rounds. In G. Brassard, editor, *Advances in Cryptology—CRYPTO ’89*, volume 435 of *Lecture Notes in Computer Science*, pages 526–545. Springer-Verlag, 1990, 20–24 Aug. 1989.
17. U. Feige and A. Shamir. Witness indistinguishability and witness hiding protocols. In *Proceedings of the Twenty Second Annual ACM Symposium on Theory of Computing*, pages 416–426, Baltimore, Maryland, 14–16 May 1990.

18. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *Advances in Cryptology—CRYPTO '86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer-Verlag, 1987, 11–15 Aug. 1986.
19. E. Fujisaki and T. Okamoto. Statistical zero knowledge protocols to prove modular polynomial relations. In B. S. Kaliski Jr., editor, *Advances in Cryptology—CRYPTO '97*, volume 1294 of *Lecture Notes in Computer Science*, pages 16–30. Springer-Verlag, 17–21 Aug. 1997.
20. O. Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, 2001.
21. S. Goldwasser, S. Micali, and C. Rackoff. Knowledge complexity of interactive proofs. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, pages 291–304, Providence, Rhode Island, 6–8 May 1985.
22. S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Computing*, 17(2), 1988.
23. O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 38(1):691–729, 1991.
24. L. C. Guillou and J.-J. Quisquater. A “paradoxical” identity-based signature scheme resulting from zero-knowledge. In S. Goldwasser, editor, *Advances in Cryptology—CRYPTO '88*, volume 403 of *Lecture Notes in Computer Science*, pages 216–231. Springer-Verlag, 1990, 21–25 Aug. 1988.
25. S. Halevi and S. Micali. Practical and provably-secure commitment schemes from collision-free hashing. In N. Kobitz, editor, *Advances in Cryptology—CRYPTO '96*, volume 1109 of *Lecture Notes in Computer Science*, pages 201–215. Springer-Verlag, 18–22 Aug. 1996.
26. J. Håstad, R. Impagliazzo, L. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Computing*, 28(4), 1999.
27. H. Krawczyk and T. Rabin. Chameleon signatures. In *Network and Distributed System Security Symposium*, pages 143–154. The Internet Society, 2000.
28. Silvio Micali, Michael Rabin, and Joe Kilian. Zero-knowledge sets. In *Proc. 44th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2003.
29. Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):51–158, 1991.
30. H. Ong and C. P. Schnorr. Fast signature generation with a Fiat Shamir-like scheme. In I. B. Damgård, editor, *Advances in Cryptology—EUROCRYPT 90*, volume 473 of *Lecture Notes in Computer Science*, pages 432–440. Springer-Verlag, 1991, 21–24 May 1990.
31. R. Ostrovsky, C. Rackoff, and A. Smith. Efficient consistency proofs for generalized queries on a committed database. In *31st International Colloquium on Automata, Languages, and Programming (ICALP 04)*, volume 3142 of *Lecture Notes in Computer Science*, pages 1041–1053. Springer-Verlag, 2004.
32. P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In J. Stern, editor, *Advances in Cryptology—EUROCRYPT '99*, volume 1592 of *Lecture Notes in Computer Science*. Springer-Verlag, 2–6 May 1999.
33. A. Shamir and Y. Tauman. Improved online/offline signature schemes. In J. Kilian, editor, *Advances in Cryptology—CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 355–367. Springer-Verlag, 19–23 Aug. 2001.
34. C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.