

A preliminary version of this paper appeared in *Advances in Cryptology – CRYPTO '02*, Lecture Notes in Computer Science Vol. ?? , M. Yung ed., Springer-Verlag, 2002. This is the full version.

GQ and Schnorr Identification Schemes: Proofs of Security against Impersonation under Active and Concurrent Attacks

MIHIR BELLARE*

ADRIANA PALACIO†

April 28, 2003

Abstract

The Guillou-Quisquater (GQ) and Schnorr identification schemes are amongst the most efficient and best-known Fiat-Shamir follow-ons, but the question of whether they can be proven secure against impersonation under active attack has remained open. This paper provides such a proof for GQ based on the assumed security of RSA under one more inversion, an extension of the usual one-wayness assumption that was introduced in [5]. It also provides such a proof for the Schnorr scheme based on a corresponding discrete-log related assumption. These are the first security proofs for these schemes under assumptions related to the underlying one-way functions. Both results extend to establish security against impersonation under concurrent attack.

Keywords: Identification schemes, Guillou-Quisquater scheme, Schnorr scheme, concurrent attacks, proofs of security.

*Dept. of Computer Science & Engineering, University of California at San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. E-Mail: mihir@cs.ucsd.edu. URL: <http://www-cse.ucsd.edu/users/mihir>. Supported in part by NSF grant CCR-0098123, NSF grant ANR-0129617 and an IBM Faculty Partnership Development Award.

†Dept. of Computer Science & Engineering, University of California at San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. E-Mail: apalacio@cs.ucsd.edu. URL: <http://www-cse.ucsd.edu/users/apalacio>. Supported in part by above-mentioned grants of first author.

Contents

1	Introduction	2
1.1	Identification schemes and their security	2
1.2	The GQ scheme and our results about it	2
1.3	The Schnorr scheme and our results about it	3
1.4	Discussion and related work	4
2	Definitions	4
3	Reset lemma	6
4	Security of GQ under concurrent attack	8
5	Security of Schnorr under concurrent attack	12
	References	13
A	Proof of Theorem 5.1	15

1 Introduction

The Guillou-Quisquater (GQ) [20] and Schnorr [26] identification schemes are amongst the most efficient and best known Fiat-Shamir [16] follow-ons, but the question of whether they can be proved secure against impersonation under active attack has remained open. This paper addresses this question, as well as its extension to even stronger attacks, namely concurrent ones. We begin with some background.

1.1 Identification schemes and their security

An identification (ID) scheme enables a prover holding a secret key to identify itself to a verifier holding the corresponding public key. Fiat and Shamir (FS) [16] showed how the use of zero-knowledge techniques [19] in this area could lead to efficient schemes, paving the road for numerous successors including [20, 26], which are comparable to FS in computational cost but have much smaller key sizes.

The accepted framework for security notions for identification schemes is that of Feige, Fiat and Shamir [14]. As usual, one considers adversary goals as well as adversary capabilities, or attacks. The adversary goal is impersonation: playing the role of prover but denied the secret key, it should have negligible probability of making the verifier accept. Towards this goal, one can allow it various attacks on the honest, secret-key equipped prover which, as per [14], take place and complete before the impersonation attempt. The weakest reasonable attack is a passive attack, in which the adversary obtains transcripts of interactions between the prover and verifier. However, the attack suggested by [14] as defining the main notion of security is an active attack in which the adversary plays the role of cheating verifier, interacting with the prover numerous times before the impersonation attempt.

Security against impersonation under active attack has been the classical goal of identification schemes. However, interest has been growing in stronger attacks, namely concurrent ones. Here, the adversary would still play the role of cheating verifier prior to impersonation, but could interact with many different prover “clones” concurrently. The clones all have the same secret key but are initialized with independent coins and maintain their own state. Security against impersonation under concurrent attack implies security against impersonation under active attack.

Analyses often approach the establishment of security against impersonation via consideration of whether or not the protocol is a proof of knowledge, honest-verifier zero knowledge, witness indistinguishable [15] and so on. These auxiliary properties are important and useful tools, but not the end goal, which remains establishing security against impersonation.

1.2 The GQ scheme and our results about it

GQ is RSA based. The prover’s public key is (N, e, X) , where N is an RSA modulus, e is a prime RSA exponent, and $X \equiv x^e \pmod{N}$ where $x \in \mathbb{Z}_N^*$ is the prover’s secret key. As typical for practical ID schemes, the protocol, depicted in Figure 2, has three moves: the prover sends a “commitment,” the verifier sends a random challenge, the prover sends a “response,” and the verifier then accepts or rejects. The protocol is honest-verifier zero knowledge and a proof of knowledge of x [20], and it follows easily that it is secure against impersonation under passive attack, assuming RSA is one-way.

The main question is whether the protocol is secure against impersonation under active attack. No attack has been found. However, no proof of security has been provided either. Furthermore, it is difficult to imagine such a proof being based solely on the assumption that RSA is one-way. (The prover response is the RSA inverse of a point that is a function of the verifier challenge, giving a cheating verifier some sort of limited chosen-ciphertext attack capability, something one-wayness does not consider.) In other words, the protocol seems to be secure against impersonation under active attack, but due to properties of RSA that go beyond mere one-wayness.

The research community is well aware that RSA has important strengths beyond one-wayness, and have captured some of them with novel assumptions. Examples include the strong RSA assumption, introduced in [17, 2] and exploited in [18, 13]; the dependent-RSA assumptions [24]; and the assumption of security under one more inversion [5]. The intent, or hope, of introducing such assumptions is that they underlie not one but numerous uses or protocols. Thus our approach is to attempt to build on this existing experience, and prove security based on one of these assumptions.

We prove that the GQ identification scheme is secure against impersonation, under both active and concurrent attacks, under the assumption that RSA is secure under one more inversion. The precise statement of the result is Corollary 4.2. Let us now explain the assumption.

Security of RSA under one more inversion, as introduced in [5], considers an adversary given input an RSA public key N, e , and access to two oracles. The *challenge oracle* takes no inputs and returns a random target point in \mathbb{Z}_N^* , chosen anew each time the oracle is invoked. The *inversion oracle* given $y \in \mathbb{Z}_N^*$ returns $y^d \bmod N$, where d is the decryption exponent corresponding to e . The assumption states that it is computationally infeasible for the adversary to output correct inverses of all the target points if the number of queries it makes to its inversion oracle is strictly less than the number of queries it makes to its challenge oracle. (When the adversary makes one challenge query and no inversion queries, this is the standard one-wayness assumption, which is why security under one more inversion is considered an extension of the standard one-wayness assumption.) This assumption was used in [5] to prove the security of Chaum’s RSA-based blind-signature scheme [12] in the random oracle model. (Our results, however, do not involve random oracles.) It was also used in [6] to prove the security of an RSA-based transitive signature scheme due to [21].

Our result is based on a relatively novel and strong assumption that should be treated with caution. But the result still has value. It reduces the security of the GQ identification scheme to a question which is solely about the security of the RSA function. Cryptanalysts need no longer attempt to attack the identification scheme, but can instead concentrate on a simply stated assumption about RSA, freeing themselves from the details of the identification model. Furthermore, our result helps clarify and unify the global picture of protocol security by showing that the properties of RSA underlying the security of the GQ identification scheme and Chaum’s RSA-based blind-signature scheme are the same. Thus our result brings the benefit we usually expect with a proof of security, namely reduction of the security of many cryptographic problems to a single number-theoretic problem. Finally, a proof under a stronger than standard assumption is better than no proof at all in the context of a problem whose provable security has remained an open question for more than ten years.

1.3 The Schnorr scheme and our results about it

The Schnorr identification scheme is discrete logarithm based. The prover’s public key is (g, X) , where g is a generator of a suitable prime-order group and $X = g^x$ where x is the prover’s secret key. The protocol, having the usual three-move format, is depicted in Figure 4. Again the protocol is honest-verifier zero knowledge and a proof of knowledge of x [26], and it follows easily that it is secure against impersonation under passive attack, assuming hardness of computation of discrete logarithms in the underlying group. (That is, one-wayness of the discrete exponentiation function.) As with GQ, the scheme appears to be secure against impersonation under active attack in the sense that no attacks are known, but proving security has remained open.

We prove that the Schnorr scheme is secure against impersonation, under both active and concurrent attacks, under the assumption that discrete exponentiation is secure under one more inversion in the underlying group. The precise statement of the result is Corollary 5.2. The assumption, also introduced in [5], is the natural analogue of the one we used for RSA. The adversary gets input the generator g . Its challenge oracle returns a random group element, and its inversion oracle computes

discrete logarithms relative to g . The assumption states that it is computationally infeasible for the adversary to output correct discrete logarithms of all the target points if the number of queries it makes to its inversion oracle is strictly less than the number of queries it makes to its challenge oracle. (When the adversary makes one challenge query and no inversion queries, this is the standard discrete logarithm assumption, meaning the standard assumption of one-wayness of the discrete exponentiation function.)

The benefits of this result are analogous to those for GQ. Although the assumption is relatively novel and strong, our result reduces the security of the Schnorr identification scheme to a question about the hardness of a number-theoretic problem, thereby freeing a cryptanalyst from consideration of attacks related to the identification problem itself.

1.4 Discussion and related work

Within the large class of FS follow-on identification schemes, proven security properties vary. Some like GQ and Schnorr did not have proofs of security against active or concurrent attacks. However, the FS scheme itself can be proven secure against impersonation under active and concurrent attacks assuming factoring is hard by exploiting its witness-indistinguishability (WI) and proof-of-knowledge (POK) properties. Okamoto’s discrete logarithm based scheme [22] is also WI and a POK, and can thus be proven secure against impersonation under active and concurrent attacks, assuming hardness of the discrete logarithm problem. Similar results hold for other schemes having the WI and POK properties. However, GQ and Schnorr are not WI, since there is only one secret key corresponding to a given public key, so these techniques do not work for them. On the other hand, they are preferable in terms of cost. Both have smaller key size than FS, and Schnorr is more efficient than Okamoto.

The so-called 2^m -th root identification scheme can be viewed as the analogue of the GQ scheme with the RSA encryption exponent e replaced by a power of two, or as a special case of the Ong-Schnorr scheme [23]. This has been proven secure against impersonation under active attack assuming factoring is hard [29, 27]. As far as we know, its security against impersonation under concurrent attack is an open question.

Shoup [28] had proved that the Schnorr scheme is secure against impersonation under active attack in the model where the attacker is restricted to be a generic algorithm, meaning one that does not exploit any special property of the encoding of group elements. Our results are in the standard and less restrictive model where the adversary is an arbitrary algorithm.

The signature schemes obtained from the GQ and Schnorr identification schemes via the Fiat-Shamir transform are already known to be provably-secure in the random oracle model assuming, respectively, the one-wayness of RSA and the hardness of the discrete logarithm problem [25], yet the security of the ID schemes against impersonation under active attack has remained open. This is not a contradiction, since the security of the signature scheme in the random oracle model relies on relatively weak security properties of the ID scheme, namely the security of the latter against impersonation under passive attack [1].

Reset attacks (where the cheating verifier can reset the internal state of prover clones with which it interacts [10, 3]) are not considered here since GQ and Schnorr, as with all proof-of-knowledge based schemes, are insecure against these attacks.

2 Definitions

The empty string is denoted ε . If x is a string then $|x|$ denotes its length, and if S is a set then $|S|$ denotes its size.

ID SCHEMES. An *identification* (ID) scheme $\mathcal{ID} = (\mathcal{K}, P, V)$ is a triple of randomized algorithms. On input security parameter $k \in \mathbb{N}$, the poly(k)-time key-generation algorithm \mathcal{K} returns a pair consisting of a public key pk and a matching secret key sk . P and V are polynomial-time algorithms that implement the prover and verifier, respectively. We require the natural correctness condition, namely that the boolean decision produced by V , in the interaction in which P has input pk, sk and V has input pk , is one with probability one. This probability is over the coin tosses of both parties. We assume that the first and last moves in the interaction always belong to the prover.

The following security notion uses the basic two-phase framework of [14] in which, in a first phase, the adversary attacks the secret-key equipped P , and then, in a second phase, plays the role of cheating prover, trying to make V accept. We define and prove security only for impersonation under concurrent attack, since the usual (serial) active attack [14] is a special case of a concurrent attack.

IMPERSONATION UNDER CONCURRENT ATTACK. An *imp-ca adversary* $A = (\widehat{V}, \widehat{P})$ is a pair of randomized polynomial-time algorithms, the *cheating verifier* and *cheating prover*, respectively. We consider a game having two phases. In the first phase, \mathcal{K} is run on input k to produce (pk, sk) , a random tape is chosen for \widehat{V} and it is given input pk . It then interacts concurrently with different clones of prover P , all clones having independent random tapes and being initialized with pk, sk . Specifically, we view P as a function that takes an incoming message and current state and returns an outgoing message and updated state. Cheating verifier \widehat{V} can issue a request of the form (ε, i) . As a result, a fresh random tape R_i is chosen, the initial state St_i of clone i is set to (pk, sk, R_i) , the operation $(M_{\text{out}}, St_i) \leftarrow P(\varepsilon; St_i)$ is executed, M_{out} is returned to \widehat{V} , and the updated St_i is saved as the new state of clone i . Subsequently, \widehat{V} can issue a request of the form (M, i) , in which case message M is sent to clone i , who computes $(M_{\text{out}}, St_i) \leftarrow P(M; St_i)$, returns M_{out} to \widehat{V} , and saves the updated state St_i . These requests of \widehat{V} can be arbitrarily interleaved. Eventually, \widehat{V} outputs some state information St and stops, ending the first phase. In the second phase of the game, the cheating prover \widehat{P} is initialized with St , verifier V is initialized with pk and freshly chosen coins, and \widehat{P} and V interact. We say that adversary A wins if V accepts in this interaction, and the *imp-ca advantage* of A , denoted

$$\text{Adv}_{\mathcal{ID}, A}^{\text{imp-ca}}(k)$$

is the probability that A wins, taken over the coins of \mathcal{K} , the coins of \widehat{V} , the coins of the prover clones, and the coins of V . (There is no need to give \widehat{P} separate coins, or even pk , since it can get them from \widehat{V} via St .) We say that \mathcal{ID} is *secure against impersonation under concurrent attack* (IMP-CA secure) if the function

$$\text{Adv}_{\mathcal{ID}, A}^{\text{imp-ca}}(\cdot)$$

is negligible for all imp-ca adversaries A of time complexity polynomial in the security parameter k .

We adopt the convention that the *time complexity* of imp-ca adversary A does not include the time taken by the prover clones and the verifier to compute replies to the adversary's requests. Rather we view these as oracles, each returning replies in unit time. Barring this, the time complexity of A is the execution time of the entire two-phase game, including the time taken for key generation and initializations. This convention simplifies concrete security considerations.

An active attack [14] is captured by considering cheating verifiers that interact serially, one by one, with prover clones. (This means the cheating verifier initializes a clone and finishes interacting with it before starting up another one.)

COMMENTS. We clarify that we do *not* allow reset attacks such as considered in [10, 3]: although \widehat{V} can interact concurrently and in interleaved fashion with the prover clones, the internal state of a clone progresses in a normal serial fashion and cannot be reset by \widehat{V} . Indeed, the GQ and Schnorr protocols, object of our study, are both insecure under reset attacks. We also clarify that we stay within the two-phase framework of [14] even while considering concurrent attacks, in the sense that

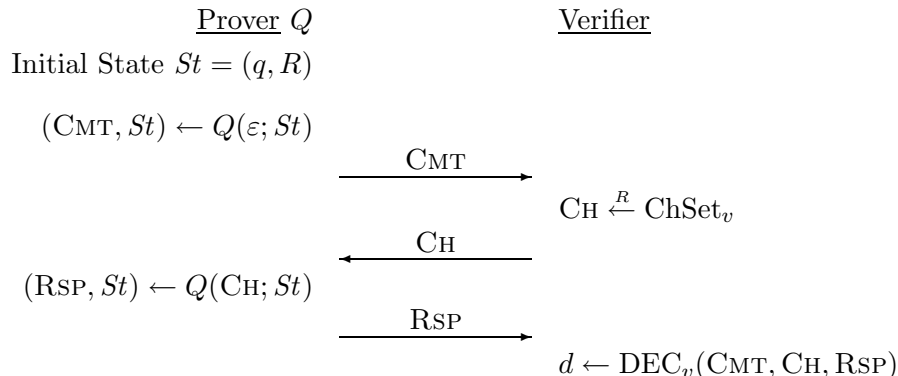


Figure 1: **A canonical protocol.** Prover Q has input q and random tape R , and maintains state St . The verifier has input v and returns boolean decision d .

the first phase (in which the adversary mounts a concurrent attack on the secret-key equipped P) is assumed to be completed before the start of the second phase (in which the adversary plays the role of cheating prover and tries to make V accept). This reflects applications such as smart card based identification for ATMs [14]. For identification over the Internet, it is more suitable to consider adversaries that can interact with the prover or prover clones even while they are interacting with the verifier in an attempt to make the latter accept. With this, one moves into the domain of authenticated key-exchange protocols which is definitionally more complex (see for example [9, 8, 30, 11]) and where identification without an associated exchange of a session-key is of little practical value.

3 Reset lemma

We refer to a three-move protocol of the form depicted in Figure 1 as *canonical*. The prover’s first message is called its *commitment*. The verifier selects a *challenge* uniformly at random from a set ChSet_v associated to its input v , and, upon receiving a *response* RSP from the prover, applies a deterministic *decision predicate* $\text{DEC}_v(\text{CMT}, \text{CH}, \text{RSP})$ to compute a boolean decision. The verifier is said to be *represented* by the pair $(\text{ChSet}, \text{DEC})$ which, given the verifier input v , defines the challenge set and decision predicate.

A prover is identified with a function Q that given an incoming message M_{in} (this is ε when the prover is initiating the protocol) and its current state St , returns an outgoing message M_{out} and an updated state. The initial state of the prover is (q, R) , where q is an input for the prover and R is a random tape.

The following lemma, which we call the Reset Lemma, upper bounds the probability that a (cheating) prover Q can convince the verifier to accept as a function of the probability that a certain experiment based on resetting the prover yields two accepting conversation transcripts. We will use this lemma in our proofs of security of both the GQ and the Schnorr schemes at the time of exploiting their proof-of-knowledge properties. In the past such analyses were based on the techniques of [14] who considered certain “execution trees” corresponding to the interaction, and their “heavy nodes.” The Reset Lemma provides a slightly better bound, has a simple proof, and is general enough to be applicable in numerous settings, saving the need to apply the techniques of [14] from scratch in each analysis, and may thus be of independent interest. Note that the lemma makes no mention of proofs of knowledge; it is just about relating two probabilities. The formulation and proof of the lemma

generalize some analyses in [4].

Lemma 3.1 (Reset Lemma) Let Q be a prover in a canonical protocol with a verifier represented by $(\text{ChSet}, \text{DEC})$, and let q, v be inputs for the prover and verifier, respectively. Let $\text{acc}(q, v)$ be the probability that the verifier accepts in its interaction with Q , namely the probability that the following experiment returns 1:

Choose random tape R for Q ; $St \leftarrow (q, R)$; $(\text{CMT}, St) \leftarrow Q(\varepsilon; St)$
 $\text{CH} \stackrel{R}{\leftarrow} \text{ChSet}_v$; $(\text{RSP}, St) \leftarrow Q(\text{CH}; St)$; $d \leftarrow \text{DEC}_v(\text{CMT}, \text{CH}, \text{RSP})$
 Return d

Let $\text{res}(q, v)$ be the probability that the following *reset* experiment returns 1:

Choose random tape R for Q ; $St \leftarrow (q, R)$; $(\text{CMT}, St) \leftarrow Q(\varepsilon; St)$
 $\text{CH}_1 \stackrel{R}{\leftarrow} \text{ChSet}_v$; $(\text{RSP}_1, St_1) \leftarrow Q(\text{CH}_1; St)$; $d_1 \leftarrow \text{DEC}_v(\text{CMT}, \text{CH}_1, \text{RSP}_1)$
 $\text{CH}_2 \stackrel{R}{\leftarrow} \text{ChSet}_v$; $(\text{RSP}_2, St_2) \leftarrow Q(\text{CH}_2; St)$; $d_2 \leftarrow \text{DEC}_v(\text{CMT}, \text{CH}_2, \text{RSP}_2)$
 If $(d_1 = 1 \text{ AND } d_2 = 1 \text{ AND } \text{CH}_1 \neq \text{CH}_2)$ then return 1 else return 0

Then

$$\text{acc}(q, v) \leq \frac{1}{|\text{ChSet}_v|} + \sqrt{\text{res}(q, v)}. \blacksquare$$

Proof of Lemma 3.1: With q, v fixed, let r denote the length of the prover's random tape. For $R \in \{0, 1\}^r$ let $\text{CMT}(q, R)$ denote Q 's commitment when it has input q and random tape R , and let $\text{RSP}(q, R, \text{CH})$ denote the response provided by Q to verifier challenge $\text{CH} \in \text{ChSet}_v$ when Q has input q and random tape R . We define functions $X, Y: \{0, 1\}^r \rightarrow [0, 1]$ as follows. For each $R \in \{0, 1\}^r$ we let

$$X(R) = \Pr[\text{DEC}_v(\text{CMT}(q, R), \text{CH}, \text{RSP}(q, R, \text{CH})) = 1],$$

the probability being over a random choice of CH from ChSet_v . For each $R \in \{0, 1\}^r$ we let

$$Y(R) = \Pr \left[\begin{array}{l} \text{DEC}_v(\text{CMT}(q, R), \text{CH}_1, \text{RSP}(q, R, \text{CH}_1)) = 1 \text{ and} \\ \text{DEC}_v(\text{CMT}(q, R), \text{CH}_2, \text{RSP}(q, R, \text{CH}_2)) = 1 \text{ and} \\ \text{CH}_1 \neq \text{CH}_2 \end{array} \right],$$

the probability being over random and independent choices of CH_1 and CH_2 from ChSet_v . A conditioning argument shows that for any $R \in \{0, 1\}^r$

$$Y(R) \geq X(R) \cdot \left[X(R) - \frac{1}{|\text{ChSet}_v|} \right].$$

We view X, Y as random variables over the sample space $\{0, 1\}^r$ of coins of Q . Then letting $p = 1/|\text{ChSet}_v|$ and using the above we have

$$\begin{aligned} \text{res}(q, v) &= \mathbf{E}[Y] \\ &\geq \mathbf{E}[X \cdot (X - p)] \\ &= \mathbf{E}[X^2] - p \cdot \mathbf{E}[X] \\ &\geq \mathbf{E}[X]^2 - p \cdot \mathbf{E}[X] \\ &= \text{acc}(q, v)^2 - p \cdot \text{acc}(q, v). \end{aligned}$$




Figure 2: **GQ identification scheme.** Prover P has input $pk = (N, e, X)$ and $sk = (N, x)$. Verifier V has input pk .

In the fourth line above, we used Jensen’s inequality¹ applied to the convex function $f(x) = x^2$. Using the above we have

$$\left(\text{acc}(q, v) - \frac{p}{2}\right)^2 = \text{acc}(q, v)^2 - p \cdot \text{acc}(q, v) + \frac{p^2}{4} \leq \text{res}(q, v) + \frac{p^2}{4}.$$

Taking the square-root of both sides of the above, and using the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for all real numbers $a, b \geq 0$, we get

$$\text{acc}(q, v) - \frac{p}{2} \leq \sqrt{\text{res}(q, v) + \frac{p^2}{4}} \leq \sqrt{\text{res}(q, v)} + \sqrt{\frac{p^2}{4}} = \sqrt{\text{res}(q, v)} + \frac{p}{2}.$$

Re-arranging terms gives us the desired conclusion. ■

4 Security of GQ under concurrent attack

A randomized, $\text{poly}(k)$ -time algorithm \mathcal{K}_{rsa} is said to be a *prime-exponent RSA key generator* if on input security parameter $k \in \mathbb{N}$, its output is a triple (N, e, d) where N is the product of two distinct primes, $2^{k-1} \leq N < 2^k$ (N is k bits long), $e < \varphi(N)$ is an odd prime, $\text{gcd}(d, \varphi(N)) = 1$, and $ed \equiv 1 \pmod{\varphi(N)}$. We do not pin down any specific such generator. Rather it is a parameter of the GQ identification scheme, and security is proved based on an assumption about it.

GQ IDENTIFICATION SCHEME. Let \mathcal{K}_{rsa} be a prime-exponent RSA key generator and let $l: \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial-time computable, polynomially-bounded function such that $2^{l(k)} < e$ for any e output by \mathcal{K}_{rsa} on input k . The *GQ identification scheme* associated to \mathcal{K}_{rsa} and challenge length l is the ID scheme whose constituent algorithms are depicted in Figure 2. The prover’s commitment is a random element $Y \in \mathbb{Z}_N^*$. For any verifier input $pk = (N, e, X)$, $\text{ChSet}_{pk} = \{0, 1\}^{l(k)}$. A challenge $c \in \text{ChSet}_{pk}$ is interpreted in the natural way as an integer in the set $\{0, \dots, 2^{l(k)} - 1\}$ in the ensuing computations. Due to the assumption that $2^{l(k)} < e$, the challenge is in \mathbb{Z}_e . The verifier’s decision predicate $\text{DEC}_{pk}(Y, c, z)$ evaluates to 1 if and only if z is the RSA-inverse of $YX^c \bmod N$.

RSA ASSUMPTION. We recall the notion of security under one more inversion (omi) [5]. An *rsa-omi* adversary is a randomized, polynomial-time algorithm I that gets input N, e and has access to

¹ Jensen’s inequality states that if f is a convex function and X is a random variable, then $\mathbf{E}[f(X)] \geq f(\mathbf{E}[X])$.

two oracles. The first is an RSA-inversion oracle $(\cdot)^d \bmod N$ that given $Y \in \mathbb{Z}_N^*$ returns $Y^d \bmod N$. The second is a challenge oracle that, each time it is invoked (it takes no inputs), returns a random challenge point $W \in \mathbb{Z}_N^*$. The game considered is to run $\mathcal{K}_{\text{rsa}}(k)$ to get N, e, d and then run $I(N, e)$ with its oracles. Let W_1, \dots, W_n denote the challenges returned by I 's challenge oracle. We say that I wins if its output is a sequence of points $w_1, \dots, w_n \in \mathbb{Z}_N^*$ satisfying $w_i \equiv W_i^d \pmod{N}$ —meaning I inverts all the challenge points—and also the number of queries made by I to its RSA-inversion oracle is *strictly less than* n . The *rsa-omi advantage* of I , denoted $\mathbf{Adv}_{\mathcal{K}_{\text{rsa}}, I}^{\text{rsa-omi}}(k)$, is the probability that I wins, taken over the coins of \mathcal{K}_{rsa} , the coins of I , and the coins used by the challenge oracle across its invocations. We say that \mathcal{K}_{rsa} is OMI secure if the function $\mathbf{Adv}_{\mathcal{K}_{\text{rsa}}, I}^{\text{rsa-omi}}(\cdot)$ is negligible for any rsa-omi adversary I of time complexity polynomial in k .

We adopt the convention that the *time complexity* of an rsa-omi adversary I is the execution time of the entire game, including the time taken for key generation and one time unit for each reply to an oracle query. (The time taken by the oracles to compute replies to the adversary's queries is not included.)

RESULT. The following theorem shows that the advantage of any imp-ca attacker against the GQ scheme can be upper bounded via the advantage of a related rsa-omi adversary and a function of the challenge length. The theorem shows the concrete security of the reduction.

Theorem 4.1 *Let $\mathcal{ID} = (\mathcal{K}, P, V)$ be the GQ identification scheme associated to prime-exponent RSA key generator \mathcal{K}_{rsa} and challenge length l . Let $A = (\widehat{V}, \widehat{P})$ be an imp-ca adversary of time complexity $t(\cdot)$ attacking \mathcal{ID} . Then there exists an rsa-omi adversary I attacking \mathcal{K}_{rsa} such that for every k*

$$\mathbf{Adv}_{\mathcal{ID}, A}^{\text{imp-ca}}(k) \leq 2^{-l(k)} + \sqrt{\mathbf{Adv}_{\mathcal{K}_{\text{rsa}}, I}^{\text{rsa-omi}}(k)}. \quad (1)$$

Furthermore, the time complexity of I is $2t(k) + O(k^4 + (n(k) + 1) \cdot l(k) \cdot k^2)$, where $n(k)$ is the number of prover clones with which \widehat{V} interacts. \blacksquare

Based on this theorem, which we will prove later, we can easily provide the following security result for the GQ scheme. In this result, we assume that the challenge length l is super-logarithmic in the security parameter, which means that $2^{-l(\cdot)}$ is negligible. This assumption is necessary, since otherwise the GQ scheme can be broken merely by guessing the verifier's challenge.

Corollary 4.2 *If prime-exponent RSA key generator \mathcal{K}_{rsa} is OMI secure and challenge length l satisfies $l(k) = \omega(\log(k))$, then the GQ identification scheme associated to \mathcal{K}_{rsa} and l is secure against impersonation under both active and concurrent attacks. \blacksquare*

Proof of Corollary 4.2: Let A be an imp-ca adversary of polynomial time complexity attacking \mathcal{ID} . Then the rsa-omi adversary given by Theorem 4.1 also has polynomial time complexity. The assumption that \mathcal{K}_{rsa} is OMI secure implies that $\mathbf{Adv}_{\mathcal{K}_{\text{rsa}}, I}^{\text{rsa-omi}}(\cdot)$ is negligible, and the condition on the challenge length implies that $2^{-l(\cdot)}$ is negligible. Equation (1) then implies that $\mathbf{Adv}_{\mathcal{ID}, A}^{\text{imp-ca}}(\cdot)$ is negligible. This shows that any imp-ca adversary of polynomial time complexity attacking the scheme has a negligible advantage. Since an active attack is a particular case of a concurrent attack, the conclusion holds. \blacksquare

We proceed to prove Theorem 4.1.

Proof of Theorem 4.1: We assume wlog that \widehat{V} never repeats a request. Fix $k \in \mathbb{N}$ and let (N, e, d) be an output of \mathcal{K}_{rsa} running on input k . Adversary I has access to an RSA-inversion oracle $(\cdot)^d \bmod N$ and a challenge oracle \mathcal{O}_N that takes no inputs and returns a random challenge point

Adversary $I^{(\cdot)^d \bmod N, \mathcal{O}_N}(N, e)$

Make a query to \mathcal{O}_N and let W_0 be the response ; $pk \leftarrow (N, e, W_0)$

Choose a random tape R for \widehat{V} ; Initialize \widehat{V} with (pk, R) ; $n \leftarrow 0$

Run \widehat{V} answering its requests as follows:

When \widehat{V} issues a request of the form (ε, i) do

$n \leftarrow n + 1$; Make a query to \mathcal{O}_N , let W_i be the response and return W_i to \widehat{V}

When \widehat{V} issues a request of the form (c, i) , where $c \in \{0, 1\}^{l(k)}$, do

$c_i \leftarrow c$; Make query $W_i W_0^{c_i}$ to $(\cdot)^d \bmod N$, let z_i be the response and return z_i to \widehat{V}

Until \widehat{V} outputs state information St and stops

$St \leftarrow (St, \varepsilon)$; $(Y, St) \leftarrow \widehat{P}(\varepsilon; St)$

$CH_1 \xleftarrow{R} \{0, 1\}^{l(k)}$; $(RSP_1, \overline{St}) \leftarrow \widehat{P}(CH_1; St)$

If $RSP_1^e \equiv Y W_0^{CH_1} \pmod{N}$ then $d_1 \leftarrow 1$ else $d_1 \leftarrow 0$

$CH_2 \xleftarrow{R} \{0, 1\}^{l(k)}$; $(RSP_2, \overline{St}) \leftarrow \widehat{P}(CH_2; St)$

If $RSP_2^e \equiv Y W_0^{CH_2} \pmod{N}$ then $d_2 \leftarrow 1$ else $d_2 \leftarrow 0$

If $(d_1 = 1 \text{ AND } d_2 = 1 \text{ AND } CH_1 \neq CH_2)$ then

$z \leftarrow RSP_1 \cdot RSP_2^{-1} \pmod{N}$; $(\overline{d}, a, b) \leftarrow \text{EGCD}(e, CH_1 - CH_2)$

$w_0 \leftarrow W_0^a z^b \pmod{N}$; For $i = 1$ to n do $w_i \leftarrow z_i w_0^{-c_i} \pmod{N}$

Return w_0, w_1, \dots, w_n

else Return \perp

Figure 3: rsa-omi adversary I for the proof of Theorem 4.1.

$W \in \mathbb{Z}_N^*$ each time it is invoked. The adversary's goal is to invert all the challenges returned by \mathcal{O}_N , while making fewer queries to its RSA-inversion oracle than the number of such challenges.

A detailed description of the adversary is in Figure 3. It first queries its challenge oracle to obtain a random element $W_0 \in \mathbb{Z}_N^*$ and uses it to create a public key pk for the imp-ca adversary A . It then uses A to achieve its goal by running \widehat{V} and playing the role of the prover clones to answer its requests. In response to a request of the form (ε, i) , I queries its challenge oracle \mathcal{O}_N and returns the answer W_i to \widehat{V} . By the definition of prover P , from \widehat{V} 's perspective, this is equivalent to picking a random tape R_i for prover clone i , initializing clone i with state pk, R_i , computing clone i 's commitment W_i , and returning the commitment to \widehat{V} . I is not in possession of the secret key $sk = (N, W_0^d \bmod N)$ corresponding to pk , which the prover clones would use to respond to \widehat{V} 's requests of the form (c, i) , where $c \in \{0, 1\}^{l(k)}$, but it compensates using its access to the RSA-inversion oracle to answer these requests. Specifically, in response to request (c, i) , I makes the query $W_i W_0^c$ to its inversion oracle and returns the answer z_i to \widehat{V} . Since $z_i = (W_i W_0^c)^d \bmod N = W_i^d (W_0^d)^c \bmod N$, this is exactly the response that clone i would return to \widehat{V} . Hence I simulates the behavior of the prover clones perfectly.

If $n(k)$ is the number of prover clones with which \widehat{V} interacts, when \widehat{V} stops I has made $n(k)$ queries to its RSA-inversion oracle and it needs to invert $n(k) + 1$ challenge points. It cannot use the inversion oracle to obtain the desired inverses. Instead, I attempts to extract from \widehat{P} , initialized with the output of \widehat{V} , the RSA-inverse of challenge W_0 . It can then use this value to compute the inverse of each of the other challenge points. To do so, I runs \widehat{P} obtaining its commitment, selects a challenge uniformly at random from $\{0, 1\}^{l(k)}$, runs \widehat{P} to obtain its response to this challenge, and evaluates

the verifier's decision predicate. It then selects another random challenge, re-runs \widehat{P} (with the same state as before) to obtain its response to the new challenge, and evaluates the verifier's decision predicate. If the decision predicate evaluates to 1, meaning \widehat{P} makes the verifier accept, on both accounts and the challenges are different, then I extracts the inverse of W_0 as follows. It computes the quotient mod N of the cheating prover's responses to the challenges and sets z to this value. We observe that $z^e \equiv W_0^{\text{CH}_1 - \text{CH}_2} \pmod{N}$. Then I uses the routine EGCD, which implements the extended Euclid algorithm, to compute (\bar{d}, a, b) , where $\bar{d} = \gcd(e, \text{CH}_1 - \text{CH}_2)$ and $a, b \in \mathbb{Z}$ are such that $ae + b(\text{CH}_1 - \text{CH}_2) = \bar{d}$. By the assumptions that e is prime and $2^{l(k)} < e$ (which implies $\text{CH}_1, \text{CH}_2 \in \mathbb{Z}_e$), $\bar{d} = 1$. Hence $ae + b(\text{CH}_1 - \text{CH}_2) = 1$. Therefore, modulo N we have

$$W_0 \equiv W_0^{ae} W_0^{b(\text{CH}_1 - \text{CH}_2)} \equiv W_0^{ae} (W_0^{\text{CH}_1 - \text{CH}_2})^b \equiv W_0^{ae} (z^e)^b \equiv (W_0^a z^b)^e.$$

This shows that $w_0 = W_0^a z^b \pmod{N}$ is the RSA-inverse of W_0 . For $i = 1, \dots, n(k)$, I computes the inverse of the i -th challenge point as $w_i = z_i w_0^{-c_i} \pmod{N}$. To prove that this computation yields the desired RSA-inverse, we show that $w_i^e \equiv W_i \pmod{N}$. Since z_i is the inverse of $W_i W_0^{c_i}$ and w_0 is the inverse of W_0 ,

$$w_i^e \equiv (z_i w_0^{-c_i})^e \equiv z_i^e (w_0^e)^{-c_i} \equiv W_i W_0^{c_i} W_0^{-c_i} \equiv W_i \pmod{N}.$$

If the decision predicate does not evaluate to 1 on both occasions or the challenges coincide, then I fails. Therefore, I wins if and only if $d_1 = 1$, $d_2 = 1$ and $\text{CH}_1 \neq \text{CH}_2$. We proceed to relate the probability of this event with the imp-ca advantage of adversary A .

We observe that pk has the same distribution as in the two-phase game that defines a concurrent attack. Since I simulates the environment provided to \widehat{V} in that game perfectly, \widehat{V} behaves as it does when performing a concurrent attack against \mathcal{ID} , and \widehat{P} is given state information with the same distribution as in that case. Therefore, the probability that $d_1 = 1$ is exactly $\text{Adv}_{\mathcal{ID}, A}^{\text{imp-ca}}(k)$.

Let $\text{acc}(St, pk)$ denote the probability that $d_1 = 1$ when the public key created by I is pk and the output of \widehat{V} is St . (This probability is over the choice of challenge CH_1 .) Let $\text{res}(St, pk)$ denote the probability that $d_1 = 1$, $d_2 = 1$ and $\text{CH}_1 \neq \text{CH}_2$ when the public key created by I is pk and the output of \widehat{V} is St . (The probability here is over the choice of challenges CH_1 and CH_2 .) Then, if $\mathbf{E}[\cdot]$ denotes the expectation of random variable \cdot over the choice of pk and St , the probability that $d_1 = 1$ is $\mathbf{E}[\text{acc}(St, pk)]$, and the probability that I wins is $\mathbf{E}[\text{res}(St, pk)]$. Applying the Reset Lemma to \widehat{P} with input St and verifier V with input pk , where the latter is implemented by I , we have

$$\text{acc}(St, pk) \leq 2^{-l(k)} + \sqrt{\text{res}(St, pk)}.$$

We obtain Equation (1) as follows.

$$\begin{aligned} \text{Adv}_{\mathcal{ID}, A}^{\text{imp-ca}}(k) &= \mathbf{E}[\text{acc}(St, pk)] \\ &\leq \mathbf{E}\left[2^{-l(k)} + \sqrt{\text{res}(St, pk)}\right] \\ &= 2^{-l(k)} + \mathbf{E}\left[\sqrt{\text{res}(St, pk)}\right] \\ &\leq 2^{-l(k)} + \sqrt{\mathbf{E}[\text{res}(St, pk)]} \\ &= 2^{-l(k)} + \sqrt{\text{Adv}_{\mathcal{K}_{\text{rsa}}, I}^{\text{rsa-omi}}(k)}, \end{aligned}$$

where the last inequality follows from Jensen's inequality² applied to the concave function $f(x) = \sqrt{x}$.

To complete the proof of Theorem 4.1, it remains to justify the claim about the time complexity of adversary I . Consider the game that defines the rsa-omi advantage of I . Our conventions for

² Jensen's inequality states that if f is a concave function and X is a random variable, then $\mathbf{E}[f(X)] \leq f(\mathbf{E}[X])$.




Figure 4: **Schnorr identification scheme.** Prover P has input $pk = (q, g, X)$ and $sk = (q, x)$. Verifier V has input pk .

measuring time complexity imply that the cost of all the steps of this game before the execution of the final “If” in the algorithm of adversary I is at most $2t(k)$ plus the cost of evaluating the verifier’s decision predicate twice. The latter involves computing two exponentiations of $|e|$ -bit exponents and two exponentiations of $l(k)$ -bit exponents. Since e is at most k bits long, this is $O(k^3 + l(k) \cdot k^2)$. We now calculate the cost of the remaining operations performed by I . The computation of the quotient mod N of the cheating prover’s responses has cost $O(k^2)$. The extended Euclid algorithm runs in time the product of the lengths of its inputs. Hence the cost of computing (\bar{d}, a, b) is $O(|e| \cdot |\text{CH}_1 - \text{CH}_2|)$, which is $O(k^2)$ because $\text{CH}_1, \text{CH}_2 \in \mathbb{Z}_e$ and $|e| \leq k$. The lengths of a and b cannot exceed the running time of EGCD, and they are exponents in the computation of w_0 . Therefore, the cost of this computation is $O(k^2 \cdot k^2) = O(k^4)$. The “For” loop has cost $n(k) \cdot O(l(k) \cdot k^2)$. The time complexity of I is then $2t(k) + O(k^4 + (n(k) + 1) \cdot l(k) \cdot k^2)$. ■

5 Security of Schnorr under concurrent attack

A randomized, $\text{poly}(k)$ -time algorithm \mathcal{K}_{dl} is said to be a *discrete logarithm parameter generator* if given security parameter $k \in \mathbb{N}$, it outputs a pair (q, g) where q is a prime such that $q | p - 1$ for a prime p with $2^{k-1} \leq p < 2^k$ (p is k bits long), and g is a generator of G_q , a subgroup of \mathbb{Z}_p^* of order q . As before, we do not pin down any specific such generator. The generator is a parameter of the Schnorr scheme, and security is proved based on an assumption about it.

SCHNORR IDENTIFICATION SCHEME. Let \mathcal{K}_{dl} be a discrete logarithm parameter generator and let $l: \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial-time computable, polynomially-bounded function such that $2^{l(k)} < q$ for any q output by \mathcal{K}_{dl} on input k . The *Schnorr identification scheme* associated to \mathcal{K}_{dl} and challenge length l is the ID scheme whose constituent algorithms are depicted in Figure 4. The prover’s commitment is a random element $Y \in G_q$. For any verifier input $pk = (q, g, X)$, $\text{ChSet}_{pk} = \{0, 1\}^{l(k)}$. A challenge $c \in \text{ChSet}_{pk}$ is interpreted as an integer in the set $\{0, \dots, 2^{l(k)} - 1\}$ in the ensuing computations. The assumption that $2^{l(k)} < q$ implies that the challenge is in \mathbb{Z}_q . The verifier’s decision predicate $\text{DEC}_{pk}(Y, c, z)$ evaluates to 1 if and only if z is the discrete log of YX^c .

DL ASSUMPTION. We recall the notion of security under one more discrete logarithm (omdl) [5]. An *omdl* adversary is a randomized, polynomial-time algorithm I that gets input q, g and has access to two oracles. The first is a discrete log oracle $\text{DLog}_{G_q, g}(\cdot)$ that given $Y \in G_q$ returns $y \in \mathbb{Z}_q$ such that

$g^y = Y$. The second is a challenge oracle that, each time it is invoked (it takes no inputs), returns a random challenge point $W \in G_q$. The game considered is to run $\mathcal{K}_{\text{dl}}(k)$ to get q, g and then run $I(q, g)$ with its oracles. Let W_1, \dots, W_n denote the challenges returned by I 's challenge oracle. We say that I wins if its output is a sequence of points $w_1, \dots, w_n \in Z_q$ satisfying $g^{w_i} = W_i$ —meaning I finds the discrete log of all the challenge points—and also the number of queries made by I to its discrete log oracle is *strictly less than* n . The *omdl advantage* of I , denoted $\mathbf{Adv}_{\mathcal{K}_{\text{dl}}, I}^{\text{omdl}}(k)$, is the probability that I wins, taken over the coins of \mathcal{K}_{dl} , the coins of I , and the coins used by the challenge oracle across its invocations. We say that \mathcal{K}_{dl} is OMDL secure if the function $\mathbf{Adv}_{\mathcal{K}_{\text{dl}}, I}^{\text{omdl}}(\cdot)$ is negligible for any omdl adversary I of time complexity polynomial in k .

We adopt the same convention regarding time complexity as in the case of an rsa-omi adversary.

RESULT. The following theorem guarantees that the advantage of any imp-ca adversary attacking the Schnorr scheme can be upper bounded via the advantage of a related omdl adversary and a function of the challenge length.

Theorem 5.1 *Let $\mathcal{ID} = (\mathcal{K}, P, V)$ be the Schnorr identification scheme associated to discrete logarithm parameter generator \mathcal{K}_{dl} and challenge length l . Let $A = (\widehat{V}, \widehat{P})$ be an imp-ca adversary of time complexity $t(\cdot)$ attacking \mathcal{ID} . Then there exists an omdl adversary I attacking \mathcal{K}_{dl} such that for every k*

$$\mathbf{Adv}_{\mathcal{ID}, A}^{\text{imp-ca}}(k) \leq 2^{-l(k)} + \sqrt{\mathbf{Adv}_{\mathcal{K}_{\text{dl}}, I}^{\text{omdl}}(k)}. \quad (2)$$

Furthermore, the time complexity of I is $2t(k) + O(k^3 + (l(k) + n(k)) \cdot k^2)$, where $n(k)$ is the number of prover clones with which \widehat{V} interacts. ■

The proof of Theorem 5.1 is in Appendix A. This theorem implies the following security result for the Schnorr scheme.

Corollary 5.2 *If discrete logarithm parameter generator \mathcal{K}_{dl} is OMDL secure and challenge length l satisfies $l(k) = \omega(\log(k))$, then the Schnorr identification scheme associated to \mathcal{K}_{dl} and l is secure against impersonation under both active and concurrent attacks. ■*

As in the case of the GQ scheme, the assumption that the challenge length l is super-logarithmic in the security parameter is necessary since otherwise the Schnorr scheme can be broken by guessing the verifier's challenge. The proof of this corollary is completely analogous to the proof of Corollary 4.2.

References

- [1] M. ABDALLA, J. AN, M. BELLARE AND C. NAMPREMPRE. From identification to signatures via the Fiat-Shamir Transform: Minimizing assumptions for security and forward-security. *Advances in Cryptology – EUROCRYPT '02*, Lecture Notes in Computer Science Vol. 2332, L. Knudsen ed., Springer-Verlag, 2002.
- [2] N. BARIĆ AND B. PFITZMANN. Collision-free accumulators and fail-stop signature schemes without trees. *Advances in Cryptology – EUROCRYPT '97*, Lecture Notes in Computer Science Vol. 1233, W. Fumy ed., Springer-Verlag, 1997.
- [3] M. BELLARE, M. FISCHLIN, S. GOLDWASSER AND S. MICALI. Identification protocols secure against reset attacks. *Advances in Cryptology – EUROCRYPT '01*, Lecture Notes in Computer Science Vol. 2045, B. Pfitzmann ed., Springer-Verlag, 2001.
- [4] M. BELLARE AND S. MINER. A forward-secure digital signature scheme. *Advances in Cryptology – CRYPTO '99*, Lecture Notes in Computer Science Vol. 1666, M. Wiener ed., Springer-Verlag, 1999.

- [5] M. BELLARE, C. NAMPREPRE, D. POINTCHEVAL AND M. SEMANKO. The one-more-RSA inversion problems and the security of Chaum’s blind signature scheme. Available as *IACR eprint archive Report 2001/002*, <http://eprint.iacr.org/2001/002/>. Preliminary version, entitled “The power of RSA inversion oracles and the security of Chaum’s RSA-based blind signature scheme,” in *Financial Cryptography ’01*, Lecture Notes in Computer Science Vol. 2339, P. Syverson ed., Springer-Verlag, 2001.
- [6] M. BELLARE AND G. NEVEN. Transitive signatures based on factoring and RSA. Manuscript, May 2002.
- [7] M. BELLARE AND A. PALACIO. GQ and Schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. Preliminary version of this paper, *Advances in Cryptology – CRYPTO ’02*, Lecture Notes in Computer Science Vol. ?? , M. Yung ed., Springer-Verlag, 2002.
- [8] M. BELLARE, D. POINTCHEVAL AND P. ROGAWAY. Authenticated key exchange secure against dictionary attacks. *Advances in Cryptology – EUROCRYPT ’00*, Lecture Notes in Computer Science Vol. 1807, B. Preneel ed., Springer-Verlag, 2000.
- [9] M. BELLARE AND P. ROGAWAY. Entity authentication and key distribution. *Advances in Cryptology – CRYPTO ’93*, Lecture Notes in Computer Science Vol. 773, D. Stinson ed., Springer-Verlag, 1993.
- [10] R. CANETTI, S. GOLDWASSER, O. GOLDBREICH AND S. MICALI. Resettable zero-knowledge. *Proceedings of the 32nd Annual Symposium on the Theory of Computing*, ACM, 2000.
- [11] R. CANETTI AND H. KRAWCZYK. Universally composable notions of key-exchange and secure channels. *Advances in Cryptology – EUROCRYPT ’02*, Lecture Notes in Computer Science Vol. 2332 , L. Knudsen ed., Springer-Verlag, 2002.
- [12] D. CHAUM. Blind signatures for untraceable payments. *Advances in Cryptology – CRYPTO ’82*, Lecture Notes in Computer Science, Plenum Press, New York and London, 1983, Aug. 1982.
- [13] R. CRAMER AND V. SHOUP. Signature schemes based on the strong RSA assumption. In *5th ACM Conference on Computer and Communications Security*, pages 46–51, Singapore, Nov. 1999. ACM Press.
- [14] U. FEIGE, A. FIAT AND A. SHAMIR. Zero knowledge proofs of identity. *Journal of Cryptology*, 1(2):77–94, 1988.
- [15] U. FEIGE AND A. SHAMIR. Witness indistinguishable and witness hiding protocols. *Proceedings of the 22nd Annual Symposium on the Theory of Computing*, ACM, 1990.
- [16] A. FIAT AND A. SHAMIR. How to prove yourself: Practical solutions to identification and signature problems. *Advances in Cryptology – CRYPTO ’86*, Lecture Notes in Computer Science Vol. 263, A. Odlyzko ed., Springer-Verlag, 1986.
- [17] E. FUJISAKI AND T. OKAMOTO. Statistical zero knowledge protocols to prove modular polynomial relations. *Advances in Cryptology – CRYPTO ’97*, Lecture Notes in Computer Science Vol. 1294, B. Kaliski ed., Springer-Verlag, 1997.
- [18] R. GENNARO, S. HALEVI, AND T. RABIN. Secure hash-and-sign signatures without the random oracle. *Advances in Cryptology – EUROCRYPT ’99*, Lecture Notes in Computer Science Vol. 1592, J. Stern ed., Springer-Verlag, 1999.
- [19] S. GOLDWASSER, S. MICALI AND C. RACKOFF. The knowledge complexity of interactive proof systems. *SIAM Journal of Computing*, 18(1):186–208, February 1989.
- [20] L. GUILLOU AND J. J. QUISQUATER. A “paradoxical” identity-based signature scheme resulting from zero-knowledge. *Advances in Cryptology – CRYPTO ’88*, Lecture Notes in Computer Science Vol. 403, S. Goldwasser ed., Springer-Verlag, 1988.
- [21] S. MICALI AND R. RIVEST. Transitive signature schemes. *Topics in Cryptology – CT-RSA ’02*, Lecture Notes in Computer Science Vol. 2271 , B. Preneel ed., Springer-Verlag, 2002.
- [22] T. OKAMOTO. Provably secure and practical identification schemes and corresponding signature schemes. *Advances in Cryptology – CRYPTO ’92*, Lecture Notes in Computer Science Vol. 740, E. Brickell ed., Springer-Verlag, 1992.

- [23] H. ONG AND C. P. SCHNORR. Fast signature generation with a Fiat Shamir–like scheme. *Advances in Cryptology – EUROCRYPT ’90*, Lecture Notes in Computer Science Vol. 473, I. Damgård ed., Springer-Verlag, 1990.
- [24] D. POINTCHEVAL. New public key cryptosystems based on the dependent-RSA problems. *Advances in Cryptology – EUROCRYPT ’99*, Lecture Notes in Computer Science Vol. 1592, J. Stern ed., Springer-Verlag, 1999.
- [25] D. POINTCHEVAL AND J. STERN. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000.
- [26] C. P. SCHNORR. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.
- [27] C. P. SCHNORR. Security of the 2^t -root identification and signatures. *Advances in Cryptology – CRYPTO ’96*, Lecture Notes in Computer Science Vol. 1109, N. Koblitz ed., Springer-Verlag, 1996.
- [28] V. SHOUP. Lower bounds for discrete logarithms and related problems. *Advances in Cryptology – EUROCRYPT ’97*, Lecture Notes in Computer Science Vol. 1233, W. Fumy ed., Springer-Verlag, 1997.
- [29] V. SHOUP. On the security of a practical identification scheme. *Journal of Cryptology*, 12:247–260, 1999.
- [30] V. SHOUP. On formal models for secure key exchange (version 4). *IACR eprint archive Report 1999/012*, <http://eprint.iacr.org/1999/012/>.

A Proof of Theorem 5.1

The proof is analogous to the proof of Theorem 4.1. We assume wlog that \widehat{V} never repeats a request. Fix $k \in \mathbb{N}$ and let (q, g) be an output of \mathcal{K}_{dl} running on input k . Adversary I has access to a discrete log oracle $\text{DLog}_{G_{q,g}}(\cdot)$ and a challenge oracle \mathcal{O}_N that takes no inputs and returns a random challenge point $W \in G_q$ each time it is invoked. The adversary attempts to invert all the challenges returned by \mathcal{O}_N , while making fewer queries to its discrete log oracle than the number of challenge points.

A detailed description of the adversary is in Figure 5. I simulates an interaction between \widehat{V} and the prover clones. To do so, it first queries its challenge oracle obtaining a random group element $W_0 \in G_q$ and uses it to create a public key pk for the imp-ca adversary A . It then runs \widehat{V} and answers its requests. In response to a request of the form (ε, i) , I queries its challenge oracle \mathcal{O}_N and returns the answer to \widehat{V} . By the definition of prover P , from \widehat{V} ’s perspective, this is equivalent to picking a random tape R_i for prover clone i , initializing clone i with state pk, R_i , computing clone i ’s commitment W_i , and returning the commitment to \widehat{V} . I is not in possession of the secret key $sk = (q, \text{DLog}_{G_{q,g}}(W_0))$ corresponding to pk , which the prover clones would use to respond to \widehat{V} ’s requests of the form (c, i) , where $c \in \{0, 1\}^{l(k)}$, but it compensates using its access to the discrete log oracle to answer these requests. Specifically, in response to request (c, i) , I makes the query $W_i W_0^c$ to its discrete log oracle and returns the answer z_i to \widehat{V} . This is exactly the response that clone i would return to \widehat{V} because $z_i = \text{DLog}_{G_{q,g}}(W_i W_0^c) = \text{DLog}_{G_{q,g}}(W_i) + c \text{DLog}_{G_{q,g}}(W_0) \bmod q$. Hence I simulates the behavior of the prover clones perfectly.

Since $n(k)$ is the number of prover clones \widehat{V} interacts with, when \widehat{V} stops, I has made $n(k)$ queries to its discrete log oracle and it needs to find the discrete log of $n(k) + 1$ challenge points. I attempts to extract from \widehat{P} , initialized with the output of \widehat{V} , the discrete log of challenge W_0 . It can then use this value to compute the discrete log of each of the other challenge points. To do so, I runs \widehat{P} obtaining its commitment, selects a challenge uniformly at random from $\{0, 1\}^{l(k)}$, runs \widehat{P} to obtain its response to this challenge, and evaluates the verifier’s decision predicate. It then selects another random challenge, re-runs \widehat{P} (with the same state as before) to obtain its response to the new challenge, and evaluates the verifier’s decision predicate. If the decision predicate evaluates to 1, meaning \widehat{P} makes the verifier accept, on both accounts and the challenges are different, then I extracts the discrete log of W_0

Adversary $I^{\text{DLog}_{G_q, g}(\cdot), \mathcal{O}_N}(q, g)$

Make a query to \mathcal{O}_N and let W_0 be the response ; $pk \leftarrow (q, g, W_0)$

Choose a random tape R for \widehat{V} ; Initialize \widehat{V} with (pk, R) ; $n \leftarrow 0$

Run \widehat{V} answering its requests as follows:

When \widehat{V} issues a request of the form (ε, i) do

$n \leftarrow n + 1$; Make a query to \mathcal{O}_N , let W_i be the response and return W_i to \widehat{V}

When \widehat{V} issues a request of the form (c, i) , where $c \in \{0, 1\}^{l(k)}$, do

$c_i \leftarrow c$; Make query $W_i W_0^{c_i}$ to $\text{DLog}_{G_q, g}(\cdot)$, let z_i be the response and return z_i to \widehat{V}

Until \widehat{V} outputs state information St and stops

$St \leftarrow (St, \varepsilon)$; $(Y, St) \leftarrow \widehat{P}(\varepsilon; St)$

$\text{CH}_1 \xleftarrow{R} \{0, 1\}^{l(k)}$; $(\text{RSP}_1, \overline{St}) \leftarrow \widehat{P}(\text{CH}_1; St)$; If $g^{\text{RSP}_1} = YW_0^{\text{CH}_1}$ then $d_1 \leftarrow 1$ else $d_1 \leftarrow 0$

$\text{CH}_2 \xleftarrow{R} \{0, 1\}^{l(k)}$; $(\text{RSP}_2, \overline{St}) \leftarrow \widehat{P}(\text{CH}_2; St)$; If $g^{\text{RSP}_2} = YW_0^{\text{CH}_2}$ then $d_2 \leftarrow 1$ else $d_2 \leftarrow 0$

If $(d_1 = 1 \text{ AND } d_2 = 1 \text{ AND } \text{CH}_1 \neq \text{CH}_2)$ then

$w_0 \leftarrow (\text{RSP}_1 - \text{RSP}_2)(\text{CH}_1 - \text{CH}_2)^{-1} \bmod q$; For $i = 1$ to n do $w_i \leftarrow z_i - c_i w_0 \bmod q$

Return w_0, w_1, \dots, w_n

else

Return \perp

Figure 5: omdl adversary I for the proof of Theorem 5.1.

as $w_0 = (\text{RSP}_1 - \text{RSP}_2)(\text{CH}_1 - \text{CH}_2)^{-1} \bmod q$. We observe that since $\text{CH}_1 \neq \text{CH}_2$ and q is prime, $\text{CH}_1 - \text{CH}_2$ has a multiplicative inverse in \mathbb{Z}_q . To prove that the computation yields the desired value, we show that $g^{w_0} = W_0$. Since RSP_1 is the discrete log of $YW_0^{\text{CH}_1}$ and RSP_2 is the discrete log of $YW_0^{\text{CH}_2}$, we have

$$\begin{aligned}
g^{w_0} &= g^{(\text{RSP}_1 - \text{RSP}_2)(\text{CH}_1 - \text{CH}_2)^{-1} \bmod q} \\
&= \left(g^{\text{RSP}_1} (g^{\text{RSP}_2})^{-1} \right)^{(\text{CH}_1 - \text{CH}_2)^{-1} \bmod q} \\
&= \left(YW_0^{\text{CH}_1} (YW_0^{\text{CH}_2})^{-1} \right)^{(\text{CH}_1 - \text{CH}_2)^{-1} \bmod q} \\
&= \left(W_0^{\text{CH}_1 - \text{CH}_2} \right)^{(\text{CH}_1 - \text{CH}_2)^{-1} \bmod q} \\
&= W_0.
\end{aligned}$$

For $i = 1, \dots, n(k)$, I computes the discrete log of the i -th challenge point as $w_i = z_i - c_i w_0 \bmod q$. To prove that this computation is correct, we show that $g^{w_i} = W_i$. Since z_i is the discrete log of $W_i W_0^{c_i}$ and w_0 is the discrete log of W_0 , we have

$$g^{w_i} = g^{z_i - c_i w_0 \bmod q} = g^{z_i} (g^{w_0})^{-c_i} = W_i W_0^{c_i} W_0^{-c_i} = W_i.$$

If the decision predicate does not evaluate to 1 on both occasions or the challenges coincide, then I fails. Therefore, I wins if and only if $d_1 = 1$, $d_2 = 1$ and $\text{CH}_1 \neq \text{CH}_2$. We proceed to relate the probability of this event with the imp-ca advantage of adversary A .

We observe that pk has the same distribution as in the two-phase game that defines a concurrent attack. Since I simulates the environment provided to \widehat{V} in that game perfectly, \widehat{V} behaves as it does when performing a concurrent attack against \mathcal{ID} , and \widehat{P} is given state information with the same distribution as in that case. Therefore, the probability that $d_1 = 1$ is exactly $\mathbf{Adv}_{\mathcal{ID},A}^{\text{imp-ca}}(k)$.

Let $\text{acc}(St, pk)$ denote the probability that $d_1 = 1$ when the public key created by I is pk and the output of \widehat{V} is St . (This probability is over the choice of challenge CH_1 .) Let $\text{res}(St, pk)$ denote the probability that $d_1 = 1$, $d_2 = 1$ and $\text{CH}_1 \neq \text{CH}_2$ when the public key created by I is pk and the output of \widehat{V} is St . (The probability here is over the choice of challenges CH_1 and CH_2 .) Then, if $\mathbf{E}[\cdot]$ denotes the expectation of random variable \cdot over the choice of pk and St , the probability that $d_1 = 1$ is $\mathbf{E}[\text{acc}(St, pk)]$, and the probability that I wins is $\mathbf{E}[\text{res}(St, pk)]$. The Reset Lemma (Lemma 3.1) applied to cheating prover \widehat{P} with input St and verifier V with input pk , the latter being implemented by I , implies

$$\text{acc}(St, pk) \leq 2^{-l(k)} + \sqrt{\text{res}(St, pk)}.$$

We obtain Equation (2) as follows.

$$\begin{aligned} \mathbf{Adv}_{\mathcal{ID},A}^{\text{imp-ca}}(k) &= \mathbf{E}[\text{acc}(St, pk)] \\ &\leq \mathbf{E}\left[2^{-l(k)} + \sqrt{\text{res}(St, pk)}\right] \\ &= 2^{-l(k)} + \mathbf{E}\left[\sqrt{\text{res}(St, pk)}\right] \\ &\leq 2^{-l(k)} + \sqrt{\mathbf{E}[\text{res}(St, pk)]} \\ &= 2^{-l(k)} + \sqrt{\mathbf{Adv}_{\mathcal{K}_{\text{dl}},I}^{\text{omdl}}(k)}, \end{aligned}$$

where the last inequality follows from Jensen's inequality¹ applied to the concave function $f(x) = \sqrt{x}$.

To complete the proof of Theorem 4.1, it remains to justify the claim about the time complexity of adversary I . Consider the game that defines the omdl advantage of I . Our conventions for measuring time complexity imply that the cost of all the steps of this game before the execution of the final “If” in the algorithm of adversary I is at most $2t(k)$ plus the cost of evaluating the verifier's decision predicate twice. The latter involves computing two exponentiations of $|q|$ -bit exponents and two exponentiations of $l(k)$ -bit exponents. Since p is k bits long and q is at most k bits long, this is $O(k^3 + l(k) \cdot k^2)$. We now calculate the cost of the remaining operations performed by I . The computation of w_0 has cost $O(|q|^2) = O(k^2)$. The “For” loop has cost $n(k) \cdot O(k^2)$. The time complexity of I is then $2t(k) + O(k^3 + (l(k) + n(k)) \cdot k^2)$. ■