A new notion of soundness in bare public-key model *

ZHAO Yunlei[1,2**] and ZHU Hong[1]

(1. Department of Computer Science, Fudan University, Shanghai 200433, China; 2. Department of Computer Science, City University of Hong Kong, Hong Kong, China)

Received February 8, 2003; revised July 28, 2003

Abstract A new notion of soundness in bare public-key (BPK) model is presented. This new notion just lies in between one-time soundness and sequential soundness and its reasonableness is justified in the context of resettable zero-knowledge when resettable zero-knowledge prover is implemented by smart card.

Keywords: bare public-key model, resettable zero-knowledge, pseudorandom functions, complexity leveraging.

The bare public-key (BPK) model was introduced by Canetti et al. in the context of resettable zero-knowledge (rZK)[1]. The BPK model simply assumes that all users have deposited a public-key in a file that is accessible by all users at all times. The only assumption of this file is that all entries in this file are guaranteed to be deposited before any interaction taking place by the users. The BPK model is very simple, and it is in fact a weak version of the frequently used public-key infrastructure (PKI) model, which underlies any public-key cryptosystem or digital signature scheme. Despite its apparent simplicity, the BPK model is quite powerful. While rZK protocols exist both in the standard and in the BPK model, only in the latter case can they be constant-round, at least in a black-box sense[1~3]. For recent advances of resettable zero-knowledge readers are referred to Ref. [4, 5].

Micali and Reyzin first noted and clarified the soundness of protocols in BPK model[2]. In BPK model, the verifier has a secret key SK, corresponding to its public key PK. Thus, a malicious prover could potentially gain some knowledge about SK from an interaction with the verifier, and this gained knowledge might be helpful in convincing the verifier of a false theorem in a subsequent interaction. In Ref. [2] four notions of soundness in the public-key model are presented, each of which implies the previous one:

(1) One-time soundness: A potential malicious prover is allowed a single interaction with the verifier per theorem statement.

(2) Sequential soundness: A potential malicious prover is allowed multiple but sequential interactions with the verifier.

(3) Concurrent soundness: A potential malicious prover is allowed multiple interleaved interactions with the verifier.

(4) Resettable soundness: A potential malicious prover is allowed to reset the verifier with the same random tape and interact with it concurrently.

The four notions above are not only meaningful, but also distinct. That is, for each soundness notion, there exists a protocol that satisfies this soundness condition but not satisfies the next one[2].

In this paper, we note that besides the above four soundness notions, there exists another soundness notion in BPK model, which we name it weak sequential soundness. We prove that weak sequential soundness is distinct from the above four soundness notions and justify its reasonableness in the context of resettable zero-knowledge when resettable zero-knowledge prover is implemented by smart card.

1 Weak sequential soundness in BPK model

Roughly, weak sequential soundness in BPK model means that a potential malicious prover is infeasible to convince an honest verifier of a false statement

* Supported by the National Natural Science Foundation of China (Grant No. 60273045) and the Ministry of Science and Technology of China (Grant No. 2001CCA03000)

** To whom correspondence should be addressed. E-mail: CSYLZHAO@cityu.edu.hk

even it is allowed to interact with the verifier a priori bounded polynomial times per theorem statement. We present its formal definition in the following. For the formal definition of the original four soundness notions the reader is referred to Ref. [2].

1.1 Honest players in BPK model

Let F be a public file that consists of a polynomial-size collection of records (id, PK_{id}), where id is a string identifying a verifier, and PK_{id} is its (alleged) public key; let P be an honest prover (for a language L). It is an interactive deterministic polynomial-time turing machine (TM) that is given as inputs: (1) a security parameter 1^n, (2) an n-bit string $x \in L$, (3) an auxiliary input y, (4) a public file F, (5) a verifier identity id and, (6) a random tape w. Let V be an honest verifier. It is an interactive deterministic polynomial-time TM that works in two stages. In stage one (the key-generation stage), on input a security parameter 1^n and random tape r, V outputs a public key PK and the corresponding secret key SK. In stage two (the verification stage), on input SK, an n-bit string x and a random tape ρ, V performs an interactive protocol with a prover, and outputs either "accept x" or "reject x".

1.2 Weak sequential soundness

For an honest verifier V with public key PK and secret key SK, a (t, U)-weak sequential malicious prover P^* for a positive polynomial t and a priori bounded polynomial U is a probabilistic polynomial-time TM that on first input 1^n (security parameter) and PK runs sequentially in at most $t(n)$ rounds as follows.

In each round i $(1 \leqslant i \leqslant t(n))$, P^* selects on the fly a common input x_i (which may be equal to x_j for $1 \leqslant j \leqslant i$) and interacts with the verification stage of $V(\mathrm{SK}, x_i, \rho_i)$ with the following restriction that the same x_i cannot be used by P^* in different rounds more than $U(n)$ times. We note that in different rounds V uses independent random tapes in its verification stage (that is, $\rho_1, \rho_2, \cdots, \rho_{t(n)}$ are independent random strings).

We then say a protocol $\langle P, V \rangle$ satisfies weak sequential soundness if for any honest verifier V, all positive polynomial t and all a priori bounded polynomial U, and all (t, U)-weak sequential malicious prover P^*, and the probability that there exists an i $(1 \leqslant i \leqslant t(n))$ such that $V(\mathrm{SK}, x_i, \rho_i)$ outputs "accept x_i" while $x_i \notin L$ is negligible in n.

1.3 Motivations, implementations and applications of weak sequential soundness

As an extension and generalization of one-time soundness, roughly speaking, almost all the ways to implement one-time soundness presented in Ref. [2] can also be used to implement weak sequential soundness. A simple way is to just let the honest verifier keep a counter for each common input on which it has been invoked. The upper-bound of each counter is set to be $U(n)$ and an honest verifier refuses to take part in other interactions with respect to the same common input once the reading of the corresponding counter reaches its upper-bound.

The BPK model was introduced in the literature of resettable zero-knowledge. Note that resettable zero-knowledge makes it possible to implement zero-knowledge prover by those devices that may be possibly maliciously resetted to its initial condition or cannot afford to generate fresh randomness for each invocation. A notable example of such devices is the widely used smart card. Also note that resettable zero-knowledge provides the basis for resettable identification protocols[6]. Then we consider the distributed client/server setting when the clients are smart cards. We remark that this setting is widely used in practice, especially in E-commerce over Internet. When a resettable identification scheme is executed in this setting we can view the identifier of each smart card as the common input. An adversary may hold many smart cards but according to the definition of weak sequential soundness we require that each smart card can be used by the adversary at most a priori bounded polynomial times. Note that in practice each smart card has an expiration date that can be viewed as the correspondence to the a priori bounded polynomial required in weak sequential soundness. In this smart-card/server setting there is a central server that may be located in a central bank or other organizations and plays the verifier's role. This central server keeps a record for each smart card and dynamically updates its information. It is easy for this central server to keep a counter in each record.

2 Separation of the various soundness notions in BPK model

In this section, we will show that the weak

sequential soundness is not only meaningful but also distinct from the original four soundness notions presented by Micali and Reyzin. We remark that as noted by Micali and Reyzin, separation of various soundness notions in public-key model is of conceptually important[2].

2.1 Previous results

In Ref. [2] Micali and Reyzin have proved the following theorems.

Theorem 1.[2] If one-way functions exist, there is a compiler-type algorithm that, for any language L, and any interactive argument system for L satisfying one-time soundness, produces another interactive argument system for the same language L that satisfies one-time soundness but not weak sequential soundness.

Here by interactive argument system we mean an interactive protocol in which the soundness of this protocol is held against probabilistic polynomial-time provers rather than computational power unbounded provers. We note that although weak sequential soundness is not considered in Ref. [2] but in their proof the malicious prover in the compiled protocol can convince the honest verifier of a false statement (common input) with probability 1 if it can invoke the honest verifier using the same common input twice. It means the compiled arguments system does not satisfy weak sequential soundness.

Theorem 2.[2] If one-way functions exist, there is a compiler-type algorithm that, for any language L, and any interactive argument system for L satisfying sequential soundness, produces another interactive argument system for the same language L that satisfies sequential soundness but not concurrent soundness.

Theorem 3.[2] There exists a compiler-type algorithm that, for any language L, and any interactive proof (or argument) system for L satisfying concurrent soundness, produces another interactive proof (respectively, argument) system for the same language L that satisfies concurrent soundness but not resettable soundness.

2.2 Our result

We first introduce the main tool used in our construction: pseudorandom function (PRF). In this paper we will use a stronger version of PRF in which the pseudorandomness is guaranteed against sub-exponential size adversaries rather than polynomial size ones.

Definition 1. (Pseudorandom function PRF[7]) A function $\mathrm{PRF}: \{0,1\}^n \times \{0,1\}^* \to \{0,1\}^n$ is a pseudorandom function if $\exists \alpha > 0$ such that for all sufficiently large k (security parameter, which is a polynomial of n) and all 2^{k^α}-gate adversaries ADV, the following difference is negligible in k:

$$\Pr[\mathrm{PRFKey} \xrightarrow{R} \{0,1\}^n : \mathrm{ADV}^{\mathrm{PRF}(\mathrm{PRFKey},\cdot)} = 1] - \Pr[F \xrightarrow{R} (\{0,1\}^n)^{\{0,1\}^*} : \mathrm{ADV}^{F(\cdot)} = 1].$$

The value α is called the pseudorandomness constant.

Such a PRF can be constructed assuming the security of RSA with large prime exponents against subexponentially-strong adversaries.

The main result of this section is the following theorem.

Theorem 4. Assuming the security of RSA with large exponents against sub-exponentially-strong adversaries, there is a compiler-type algorithm that, for any language L, and any interactive argument system for L satisfying weak sequential soundness, produces another interactive argument system for the same language L that satisfies weak sequential soundness but not sequential soundness.

Proof. This proof uses "complexity leveraging" technique introduced in Ref. [1].

Let γ be the following constant: for all sufficiently large n, the length of the NP-witness y for $x \in L$ of length n is upper-bounded by n^γ. We then set $\varepsilon > \frac{\gamma}{\alpha}$. For an NP-language L we construct an argument system for L on security parameter n while using a PRF with a (larger) security parameter $k = n^\varepsilon$. This ensures that one can enumerate all potential NP-witnesses y for $x \in L$ in time 2^{n^γ}, which is still less than the time it would take to break the pseudorandomness of PRF because $2^{n^\gamma} < 2^{k^\alpha}$.

Let F be a pseudorandom function as defined in Definition 1 and U be the a priori bounded polynomial guaranteed in the definition of weak sequential soundness. Denote by $F_s(X) = R_1 R_2 \cdots R_{U(|x|)+1}$

the output of F with seed s on input x, where for each i, $1 \leqslant i \leqslant U(|x|)+1$, $|R_i| = |x|$.

Let x of length n be the theorem that the prover is trying to prove to the verifier. Given any interactive argument system (P, V) for a language L in NP satisfying weak sequential soundness, we produce another interactive argument system (P', V') for the same language L that satisfies weak sequential soundness but not sequential soundness.

Add to key generation: Generate randomly an n-bit seed s and add s to the key SK.

Add *P* step: Set $\beta_1 = \beta_2 = \cdots = \beta_{U(n)+1} = 0$, and send $(\beta_1 \cdots \beta_{U(n)+1})$ to the verifier V'.

Add *V* step: If $F_s(x) = \beta_1 \cdots \beta_{U(n)+1}$, then accept and stop. Otherwise, randomly select i in $\{1, 2, \cdots, U(n)+1\}$ and send (i, R_i) to the prover P'.

Note that a malicious prover can get V' to accept a false statement x with overwhelming probability after $(U(n)+1)^2$ sequential interactions with the honest verifier V' on the same common input x.

However, if the malicious prover is restricted to use the same common input x at most a priori bounded polynomial, specifically $U(n)$ times, in its sequential interactions with the honest verifier we will prove in below that it is infeasible for the malicious prover to get the honest verifier to accept an x while $x \notin L$.

First, we assume the F used by V' is a truly random function. Then, it can be easily seen that the malicious prover is infeasible to convince the honest verifier of a false statement if for each common input x, x is used by the malicious prover at most $U(n)$ times in its interactions with the honest verifier.

Now, we deal with the real case in which the honest verifier uses a pseudorandom function rather than a truly random function. The subtle problem here is that in our context during the interactions between the malicious prover and the honest verifier the malicious prover selects many common inputs on the fly. The problem is how to distinguish $x \in L$ or $x \notin L$ for each common input selected by the malicious prover on the fly. In Ref. [8], to overcome this problem the authors of Ref. [8] required their protocol to be also an argument of knowledge. What saves us here is the "complexity leveraging". That is, for each common input selected by the malicious prover on the fly, we just enumerate all its NP-witnesses in time 2^{n^γ} and decide whether $x \in L$ or not.

Now we claim that the malicious prover is also infeasible to convince the honest verifier of a false statement in its weak sequential interactions with the honest verifier even the function F used by the honest verifier is a pseudorandom function. Otherwise, suppose the malicious prover can convince the honest verifier of a false statement with non-negligible probability in its weak sequential interactions when F is a pseudorandom function. Then we can construct a distinguisher with size less than 2^{k^α} that distinguishes F from a truly random function with non-negligible probability as follows. The distinguisher runs the malicious prover while oracle accessing to the honest verifier to simulate the real interactions between the malicious prover and the honest verifier. For each common input x selected by the malicious prover on the fly the simulator decides $x \in L$ or not in time 2^{n^γ} and for each $x \in L$ the distinguisher also feeds the corresponding NP witness to the malicious prover. If during the simulation the distinguisher finds that the honest verifier accepts a $x \notin L$ then it concludes that the oracle is using a pseudorandom function. Note that the size of this distinguisher is at most $\text{poly}(n) \cdot 2^{n^\gamma} < 2^{k^\alpha}$ which violates the pseudorandomness of the PRF F.

Denote by $A < B$ if there exists a compiler-type algorithm that, for any language L, and any interactive argument system for L satisfying soundness notion A, produces another interactive argument system for the same language L that satisfies soundness A but not soundness notion B. Then all the five soundness notions in BPK model can be presented as follows: One-time soundness $<$ weak sequential soundness $<$ sequential soundness $<$ concurrent soundness $<$ resettable soundness.

3 Future investigations

Canetti et al. first gave a 5-round resettable black-box zero-knowledge argument with sequential soundness for NP in BPK model[1] and the round-complexity is further reduced to four by Micali and Reyzin[2]. For efficient four round zero-knowledge proof systems for NP in the standard model (without public-keys) readers are referred to Ref. [9]. Micali

and Reyzin also showed that any (resettable or not) auxiliary-input zero-knowledge protocol with sequential soundness for a language outside BPP needs at least three rounds. Note that black-box zero-knowledge implies auxiliary-input zero-knowledge[10]. In a stronger version, upper-bounded public-key (UPK) model in which the public-key of an honest verifier is restricted to be used at most a priori bounded polynomial times, Micali and Reyzin presented a 3-round black-box resettable zero-knowledge argument with sequential soundness for *NP*[11]. But we note that their protocol does not satisfy our weak sequential soundness. The reason is that in the definition of weak sequential soundness we do not restrict the number of times that the public-key of an honest verifier can be used. That is, the public-key of an honest verifier may be used by a malicious prover for any polynomial times. What is restricted in our weak sequential soundness is that for each common input selected by a malicious prover on the fly the same common input cannot be used more than a priori bounded polynomial times. If the public-key of an honest verifier can be used for any polynomial times then a malicious prover of Micali and Reyzin's protocol[11] can easily cheat the honest verifier with non-negligible probability. Then, it comes a question of whether 3-round resettable zero-knowledge argument with weak sequential soundness for *NP* exists in BPK model. We suggest that it should be an interesting open problem for future investigations.

Acknowledgment The first author is grateful to Leonid Reyzin for his kindly clarifications and valuable discussions.

References

1 Canetti, R. et al. Resettable zero-knowledge. In: Proceedings of the Annual ACM Symposium on Theory of Computing, New York: ACM Press, 2000, 235.
2 Micali, S. et al. Soundness in the public-key model. In: Proceedings of Advances in Cryptology: CRYPTO 2001, Berlin: Springer-Verlag, 2001, 542.
3 Canetti, R. et al. Black-box concurrent zero-knowledge requires $\Omega(\log n)$ rounds. In: Proceedings of the Annual ACM Symposium on Theory of Computing, New York: ACM Press, 2001, 570.
4 Zhao, Y. L. et al. Resettable zero-knowledge in the weak public-key model. In: Proceedings of Advances in Cryptology: EUROCRYPT 2003, Berlin: Springer-Verlag, 2003, 145.
5 Zhao, Y. L. et al. Reduction zero-knowledge. In: Proceedings of the Third International Conference on Security in Communication Networks, Berlin: Springer-Verlag, 2002, 314.
6 Bellare, M. et al. Identification protocols secure against reset attacks. In: Proceedings of Advances in Cryptology: EUROCRYPT 2001, Berlin: Springer-Verlag, 2001, 495.
7 Goldreich, O. How to construct random functions. Journal of the Association for Computing Machinery, 1986, 33(4): 792.
8 Barak, B. et al. Resettably-sound zero-knowledge and its applications. In: Proceedings of IEEE Symposium on Foundations of Computer Science, Las Vegas: IEEE Press, 2001, 116.
9 Zhao, Y. L. et al. Efficient 4-round zero-knowledge proof system for *NP*. Progress in Natural Science, 2002, 12(12): 948.
10 Goldreich, O. et al. Definitions and properties of zero-knowledge proof systems. Journal of Cryptology, 1994, 7(1): 1.
11 Micali, S. et al. Min-round resettable zero-knowledge in the public-key model. In: Proceedings of Advances in Cryptology: EUROCRYPT 2001, Berlin: Springer-Verlag, 2001, 373.