



HAL
open science

Towards A Security Model for Computational Puzzle Schemes

Qiang Tang, Arjan Jeckmans

► **To cite this version:**

Qiang Tang, Arjan Jeckmans. Towards A Security Model for Computational Puzzle Schemes. International Journal of Computer Mathematics, 2011, pp.1. 10.1080/00207160.2010.543951 . hal-00691330

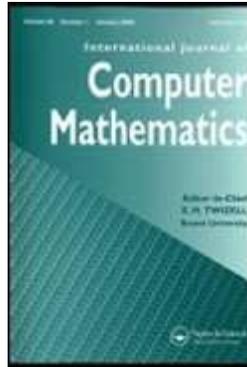
HAL Id: hal-00691330

<https://hal.science/hal-00691330>

Submitted on 26 Apr 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Towards A Security Model for Computational Puzzle Schemes

Journal:	<i>International Journal of Computer Mathematics</i>
Manuscript ID:	GCOM-2010-0642-A
Manuscript Type:	Original Article
Date Submitted by the Author:	24-Aug-2010
Complete List of Authors:	Tang, Qiang; University of Twente, Computer Science Jeckmans, Arjan; University of Twente, CS
Keywords:	computational puzzle, deterministic computation, parallelization, security model, repeated squaring

SCHOLARONE™
Manuscripts

Towards a Security Model for Computational Puzzle Schemes

Qiang Tang and Arjan Jeckmans

DIES, Faculty of EEMCS
University of Twente, the Netherlands
{q.tang, a.j.p.jeckmans}@utwente.nl

August 24, 2010

Abstract

In the literature, computational puzzle schemes have been considered as a useful tool for a number of applications, such as constructing timed cryptography, fighting junk emails, and protecting critical infrastructure from denial-of-service (DoS) attacks. However, there is a lack of a general security model for studying these schemes. In this paper, we propose such a security model and formally define two properties, namely the determinable difficulty property and the parallel computation resistance property. Furthermore, we prove that a variant of the RSW scheme, proposed by Rivest, Shamir, and Wagner, achieves both properties.

Keywords: computational puzzle, determinable difficulty, parallelization

1 Introduction

A computational puzzle scheme [13, 19, 20, 30] enables a prover to prove to a verifier that a certain amount of resources has been dedicated to solve a puzzle. To motivate our discussions, consider the following two toy examples.

1. Let H be a one-way hash function. The verifier selects a random number x and set the puzzle to be $H(x)$, while the prover searches for a number x' , where $H(x) = H(x')$, as the puzzle solution. The verifier

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

verifies a submitted puzzle solution x' through checking the equality $H(x) = H(x')$.

2. Let H be a one-way hash function. The verifier selects d random numbers r_i ($1 \leq i \leq d$) as the puzzle, while the prover computes $H(r_i)$ ($1 \leq i \leq d$) as the puzzle solution. The verifier verifies a submitted puzzle solution h_i ($1 \leq i \leq d$) through checking the equalities $H(r_i) = h_i$ ($1 \leq i \leq d$).

In the first scheme, the verifier only needs to compute one hash in order to verify the solution, regardless of the amount of computation that the prover needs to perform. In contrast, in the second scheme, the verifier needs to perform the same amount of computation as that of the prover in order to verify the solution. Intuitively, the first scheme seems to be the preferred one, but it still has some drawbacks.

- First, the amount of computation the prover needs to perform is probabilistic. Suppose x is chosen from the domain $[1, N]$. In the best case the prover only needs to compute one hash, while in the worst case it needs to compute N hashes until find the right answer. As a consequence, the verifier never knows what is the exact amount of computation required to solve a puzzle, although the average is $\frac{N}{2}$ hashes.
- Second, if a prover has access to more than one computer, then it can speed up the puzzle solving process by having all of them work in parallel. For example, client A, which has access to x computers, could be $x - 1$ times faster in finding a solution than client B, which has access to one computer. In practice, it is very difficult for a verifier to determine the amount of computing resources a prover can access, especially in the presence of malicious provers which control a large number of Zombie computers. In some cases such as timed cryptography [5, 17, 30], this will create an unfair situation for different provers. Arguably, in many other cases such as fighting against DoS attacks [20], the parallelism property of a computational puzzle scheme may also be undesirable.

1.1 Contribution

Most of existing computational puzzle schemes, as surveyed in Section 2, possess similar drawbacks. In fact, there is a lack of a general security

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

model for formally studying relevant properties. In this paper, we propose a security model for computational puzzle schemes in general. In particular, in correspondence to the above drawbacks, we formally define two properties, namely *determinable difficulty* and *parallel computation resistance*.

- The determinable difficulty property implies that the verifier can precisely determine the required resource required from the prover in solving a puzzle.
- The parallel computation resistance property implies that the prover cannot accelerate the puzzle solving process by letting more than one computer work in parallel.

These two properties describe the fundamental characteristics of computational puzzles. This is the first model of this kind, though there are security models for some specific categories of computational puzzle schemes (e.g. [9]). In addition, we also provide discussions on some other properties such as computation disparity between the verifier and a prover, puzzle hardness granularity, and puzzle statefulness.

We prove that a variant of the RSW computational puzzle scheme, proposed by Rivest, Shamir, and Wagner, achieves the determinable difficulty and parallel computation resistance properties. This not only shows that the defined security properties are achievable but also solves the open issue about analyzing the security of the well-known RSW scheme in a rigorously-defined security model. To our knowledge, most existing computational puzzles do not achieve the parallel computation resistance property.

1.2 Organization

The rest of the paper is organized as follows. In Section 2, we present a brief literature review on computational puzzle schemes and present a formal definition. In Section 3 we propose a general security model and define the properties. In Section 4, we present a variant of the RSW computational puzzle scheme and prove its security. In Section 5, we conclude the paper.

2 Preliminary of Computational Puzzles

In this section, we briefly review the literature of computational puzzle schemes, and then present a general and formal definition of such schemes to facilitate our following discussions.

2.1 Literature Review

Merkle [27] was the first to introduce the notion of *puzzle* which led to the invention of public key cryptography. In this context, the puzzle is required to be unsolvable by any polynomial-time entity. Dwork and Naor [13] proposed the concept of *pricing function* to combat junk emails. Rivest, Shamir, and Wagner [30] proposed the concept of *timed-lock puzzle*, which serves as a tool to realize the concept of *timed-release crypto*. Juels and Brainard [20] proposed the concept *client puzzle* and suggested to use it to prevent Denial of Service (DoS) attacks. Regardless of the different notations, *pricing function*, *timed-lock puzzle*, *proof of work*, and *client puzzle* share the same characteristic: they can be regarded as another type of *puzzle* (different from that of Merkle [27]) which is moderately hard in the sense that a polynomial-time entity can successfully find a solution by spending a certain amount of resources. Without loss of generality, we use the term *computational puzzle* to refer to the second type of puzzle.

Roughly speaking, there are two types of computational puzzle schemes. One type is CPU-bound, where the computation is measured by the amount of CPU cycles needed to solve a puzzle. Some examples are those in [4, 9, 13, 19, 20, 30, 32, 36], which form the majority of the existing computational puzzle schemes. Abadi *et al.* [1] first noticed the fact that CPU power varies a lot for different computers (such as PC, PDA, and Workstation), and introduced memory-bound computational puzzle schemes, where the computation is measured by the amount of memory look-ups needed to solve a puzzle. The schemes in [11, 12] fall into this category. Regardless of the nature of different forms of computations involved in solving puzzles, the essence is the same, namely the prover needs to spend a certain amount of resources (either CPU cycles or memory look-ups) in finding a solution.

Computational puzzle schemes have been explored to realize a number of secure functionalities, such as timed encryption [30], timed bit commitment [5], timed release of digital signatures [17], uncheatable benchmarks [8], trustworthy website usage metering [15], and generating delays in lottery applications [18]. The rationale is that the process of spending resources in solving a puzzle automatically results in a time delay, which enables the verifier to control when the prover is able to access the functionalities. In other words, the functionalities are masked by puzzles, which need to be solved first before the prover can access the functionalities. Besides realizing various security functionalities, researchers have also applied computational puzzle schemes to mitigate a wide range of DoS attacks, such as fighting junk emails [13, 12], protecting authentication protocols [2, 23, 29],

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

protecting IP networks [10, 14, 16, 20, 33, 34, 35], protecting wireless networks [24, 26], and preventing Sybil attacks [6]. With a computational puzzle scheme implemented, the server (playing the role of verifier) can mitigate an attack by asking every client (playing the role of prover) to solve a puzzle before allocating any required resource. The rationale is that, the number of “valid” requests from a malicious client will drop to some extent because the client has only limited resource to find puzzle solutions. In the literature, there has been some debate on whether computational puzzle schemes are really helpful to defend DoS attacks. Based on the collected results from ISPs in UK, Laurie and Clayton [21] claimed that computational puzzle schemes are hardly effective in combating junk emails in practice. While, Liu and Camp [25] argued that computational puzzle schemes could be helpful if such schemes are used in combination with reputation systems.

Chen *et al.* [9] proposed a security model for computational puzzle schemes tailored for defeating DoS attacks. In particular, their definition has focused on the unforgeability and difficulty of puzzles. However, their definition of puzzle difficulty implies neither the determinable difficulty property nor the parallel computation resistance property.

2.2 Definition of Computational Puzzle Schemes

In the design and analysis of computational puzzle schemes, we typically do not directly talk about how many CPU cycles or memory lookups are required for puzzle generation, puzzle solving, and puzzle solution verification. Instead, we often use the number of some generic operations (such as a hash or a multiplication in some group) as a puzzle complexity metric. For the notation purpose, we assume the generic operation to be denoted as *Func*.

A computational puzzle scheme consists of four (probabilistic) polynomial-time algorithms (*Setup*, *PuzzleGen*, *PuzzleSol*, *PuzzleVer*).

- *Setup*(ℓ, D): Run by the verifier, this algorithm takes a security parameter ℓ and an upper bound D on the puzzle hardness (see the definition below) as input, and outputs the public system parameter *params* and a private key *mk*. The public system parameter *params* should also include a specification about the metric function *Func*. This system parameter is implicitly part of the input to other algorithms, and we omit it in the descriptions for simplicity reasons.

- **PuzzleGen**(mk, d, req): Run by the verifier, this algorithm takes the private key mk , a hardness parameter d , and some additional information req as input, and outputs a puzzle puz and some relevant information $info$. The hardness parameter d is an integer which indicates the total number of **Func** operations required. The verifier sends puz to the prover, and keeps $info$ for verifying the solution.
- **PuzzleSol**(puz): Run by a prover, this algorithm takes a puzzle puz as input and outputs a puzzle solution sol .
- **PuzzleVer**($mk, info, sol$): Run by the verifier, this algorithm takes the private key mk , the puzzle information $info$, and the solution sol as input, and outputs 1 if sol is correct or 0 otherwise.

It is worth noting that a prover may not follow the **PuzzleSol** algorithm to find a solution, and in fact it can use any means to find the solution. But the properties *determinable difficulty* and *parallel computation resistance*, which will be defined in next subsection, imply that a prover need to perform d operations of **Func** in order to find a solution whether or not it follows the **PuzzleSol** algorithm. In addition, in some schemes, a long-term private key mk may be unnecessary or the additional information req may not be required as part of the input. In some application scenarios such as timed cryptography [5, 17, 30], a puzzle verification algorithm **PuzzleVer** is not explicitly required. Nonetheless, the above definition is meant to be general enough to cover most existing computational puzzle schemes.

3 Property Formulations for Computational Puzzles

Similar to other types of cryptographic schemes, soundness is a basic requirement for a computational puzzle scheme. In our case, soundness is defined to be that, given a puzzle with the hardness parameter d , the prover can find a correct solution by performing a total number of d operations of **Func**.

In the following, we formalize the properties *determinable difficulty* and *parallel computation resistance*. It is worth noting that, as in the case of most cryptographic formulations, we only consider a polynomial time adversary in our formulation and define that a property is achieved if an adversary has only a negligible advantage in the attack. The notion of negligibility is defined as follows.

Definition 1 *The function $P(k) : \mathbb{Z} \rightarrow \mathbb{R}$ is said to be negligible with respect to k if, for every polynomial $f(k)$, there exists an integer N_f such that $P(k) < \frac{1}{f(k)}$ for all $k \geq N_f$.*

3.1 The Determinable Difficulty Property

Informally, the determinable difficulty property implies that the verifier can precisely determine the required resource from the prover in solving a puzzle. Basically, this has two implications.

- One is that puzzles, generated by the verifier, are independent from each other in the sense that solving one puzzle does not help solve another puzzle. In other words, any puzzle generated by the verifier will require the prover to perform d Func operations.
- The other is that puzzles, generated by the verifier, are unpredictable in the sense that any puzzle generated by the verifier looks fresh so that the prover is unable to pre-compute the solution.

An attack succeeds if, after receiving a puzzle of hardness d , the adversary finds a solution by performing less than d Func operations. As a standard practice, attacks against the determinable difficulty property is simulated through the following three-phase game between an adversary and a challenger, as depicted in Fig. 1. In the attack game, the adversary plays the role of a malicious prover and the challenger plays the role of the verifier.

In practice, the only information to a prover is the puzzles generated by the verifier and the public system parameter. It is straightforward to check that the adversary has all the same privileges as a malicious prover in practice. Formally, the determinable difficulty property is defined as follows.

Definition 2 *Suppose that the adversary has performed d^\dagger Func operations in the Response phase of the attack game (defined in Fig. 1). A computational puzzle scheme achieves the determinable difficulty property if the probability $d^\dagger < d^*$ is negligible with respect to the security parameter ℓ . The probability is computed based on the (random) coin tosses of the challenger and the adversary.*

1. Setup phase: The challenger takes the security parameter and an upper bound D on the puzzle hardness as input, and runs the Setup algorithm to generate the master key mk and the public system parameter $params$.
2. Challenge phase: The adversary issues queries to the PuzzleGen oracle with its chosen d and req , and obtains the reply $puz = \text{PuzzleGen}(mk, d, req)$ from the challenger. At some point, the adversary sends d^* and req^* to the challenger, which sends $puz^* = \text{PuzzleGen}(mk, d^*, req^*)$ back as a challenge.
3. Response phase: The adversary can issue queries to the PuzzleGen oracle as in the challenge phase. At some point, the adversary terminates by outputting a correct solution sol^* to the puzzle puz^* .

Figure 1: The Attack Game

3.2 The Parallel Computation Resistance Property

As illustrated by the example schemes in Section 1, if a prover has access to a number of computers, it is able to solve a puzzle much faster than others by dividing the workload and letting the computers work in parallel. Informally, the parallel computation resistance property implies that a prover cannot accelerate the puzzle solving process by exploiting the parallelism. More concretely, the parallel computation resistance property requires that, after receiving a puzzle of hardness d from the verifier, a prover needs to *sequentially* perform d Func operations to find a solution. It means that the prover is unable to speed up the process by letting more than one computer work in parallel. This also implies that the best strategy for the prover is to use its fastest computer to solve the puzzle. We first introduce the notion of *sequential computing time* as follows.

Definition 3 Suppose that a prover has access to a number of computers which can work in parallel, and the Func function is implemented as a program in every computer. Let the each evaluation of the Func function be denoted as an event, marked by the starting time and ending time of the evaluation, and the set of events be denoted by g_i ($1 \leq i \leq N$) in a computation task. Consider all sequences of the

events of the following form

$$Seq_x = \{q_1, q_2, \dots, q_x\},$$

where, $1 \leq x \leq N$ and for all $1 \leq i \leq x - 1$, the q_{i+1} -th Func evaluation is started by a computer after the q_i -th Func evaluation has ended (by any computer). The sequential computing time of the computation task is the maximum value of the sequence indexes, namely \max_x .

Formally, the parallel computation resistance property is defined as follows.

Definition 4 Suppose that the sequential computing time of the computation task of finding a correct solution sol^* is d^\dagger in the Response phase of the attack game (defined in Fig. 1). A computational puzzle scheme achieves the parallel computation resistance property if the probability $d^\dagger < d^*$ is negligible with respect to the security parameter ℓ . The probability is computed based on the (random) coin tosses of the challenger and the adversary.

According to our definition, the parallel computation resistance property implies the determinable difficulty property. However, the reverse is clearly not true.

Parallel computation resistance is a very desirable property when we want to limit the disparity of computing powers of potential provers of a computational puzzle scheme. Moreover, in some cases, this property is crucial to mitigate attacks from malicious provers which control a cluster of Zombie computers. However, the determinable difficulty property is of independent interest in evaluating the security of computational puzzle schemes when the property parallel computation resistance is unnecessary. There are straightforward ways to convert a scheme, which achieves the parallel computation resistance property, into a scheme which achieves determinable difficulty and yet supports parallelism. For example, to generate a puzzle of hardness d , the verifier generates x independent sub-puzzles $\{puz_1, \dots, puz_x\}$ such that each of the sub-puzzle has the hardness $\frac{d}{x}$. By doing so, the total hardness is still $d = \frac{d}{x} \cdot x$, but a prover can solve the sub-puzzles in parallel.

3.3 Further Remarks

When designing computational puzzle schemes, a naturally desirable property is the *computation disparity* between the verifier and a prover: it should

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

take very little resource for the verifier to generate a puzzle and verify a solution, while it should take the prover a certain amount of resources (determined by the verifier) to find a solution. Formally, this property can be evaluated by a ratio parameter $\frac{v}{d}$, where v is the total number of Func operations required in generating a puzzle and verifying a solution, and d is the puzzle hardness.

With respect to puzzle hardness d , it is desirable that a computational puzzle scheme provides finer granularity support. Take the hash-based client puzzle scheme as an example ¹: a puzzle is in the form of (h, x_2) , where $h = H(x_1 || x_2)$, H is a collision-resistant hash function, and x_1 has bit-length k ; a solution is a k bit x'_1 such that $h = H(x'_1 || x_2)$. Clearly, the puzzle hardness can only be set exponentially with respect to k in this case. However, in the RSW computational puzzle scheme [30] and the variant in Section 4, the puzzle hardness can be set linearly in the number of multiplications. Clearly, with the latter, the verifier can more flexibly set d according to the requirements of the underlying applications.

According to the definition of a computational puzzle scheme given in Section 2, the verifier may store some state information *info* for each puzzle. In this case, the scheme is said to be stateful. In practice, a stateful computational puzzle scheme may be considered to be inefficient because the verifier needs to store a piece of information for every unverified puzzle. Nevertheless, it is straightforward to turn a stateful scheme into a stateless one by sending both *puz* and *info* to the client, which should send *sol* and *info* back to the server for verification. In the case that *info* compromises the determinable difficulty property or the parallel computation resistance property, the verifier can protect *info* with a symmetric key encryption algorithm and a message authentication code algorithm. It is worth noting that a completely stateless computational puzzle scheme may be prone to replay attacks, in which an adversary replays a puzzle and its solution to the verifier. A more detailed discussion of this issue is beyond the scope of this paper, and should be addressed in the deployment of a computational puzzle scheme.

4 Analysis of a Variant of the RSW Scheme

In this section, we analyse a variant of the RSW computational puzzle scheme and show that it achieves the parallel computation resistance property. In the literature, Tritilanunt *et al.* [32] proposed a scheme towards

¹This is a simplified version of the scheme in [20].

achieving the parallel computation resistance property, but they did not give a rigorous security proof. It is unclear whether their scheme can be proven in our security model.

4.1 Description of the Scheme

The algorithms of the new scheme are defined as follows.

- **Setup**(ℓ, D): This algorithm selects two random large primes p, q and a cryptographic hash function $H : \{0, 1\}^* \rightarrow \mathbb{Z}_{pq}^*$. The public parameter is (ℓ, D, pq) , and the master key is $mk = (p, q)$. In addition, the metric function **Func** is multiplication in \mathbb{Z}_{pq}^* . Given ℓ , the prime numbers p, q should be chosen in such a way that the required assumptions hold, as stated in next subsection.
- **PuzzleGen**(mk, d, req): This algorithm chooses $r \in_R \mathbb{Z}_{pq}$ and computes $g = H(r||req)$, and outputs the puzzle $puz = (g, d)$. The related puzzle information is $info = (r, d, req)$.
- **PuzzleSol**(puz): This algorithm outputs $sol = g^{2^d} \pmod{pq}$.
- **PuzzleVer**($mk, info, sol$): This algorithm returns 1 if $sol \equiv (H(r||req))^{2^d} \pmod{\phi(pq)}$ (\pmod{pq}), and returns 0 otherwise.

In this scheme, g is computed as $g = H(r||req)$, which can be regarded as a randomly chosen element of \mathbb{Z}_{pq}^* if H is modeled as a random oracle (as shown in our security proof). By doing so, if needed, this element can be bound to situational information (such as the identity information of the prover) contained in req . Furthermore, the modification also facilitates the proof in our security model. Note that, in [30], g is directly assumed to be a randomly chosen element of \mathbb{Z}_{pq}^* . We further note that the hash function H can be replaced with any function that maps $r||req$ to a random element in \mathbb{Z}_{pq}^* .

The puzzle hardness d can be set linearly in the number of multiplications in the group \mathbb{Z}_{pq}^* . It is clear that the workload of the verifier is constant regardless of the puzzle hardness. With respect to the verification complexity of the verifier, we omit the computation $2^d \pmod{\phi(pq)}$ for two reasons. One is that it could be pre-computed and stored by the verifier. The other is that, in many cases, multiple puzzles might share the same hardness so that the computation only needs to be done once. As a consequence, it is

straightforward to calculate that the average verification complexity for the verifier is approximately $\frac{3L}{2}$ multiplications in \mathbb{Z}_{pq}^* , where L is the bit-length of $\phi(pq)$. The computation disparity ratio parameter is $\frac{3L}{2d}$.

4.2 Security Analysis

In [22], Leander and Rupp model \mathbb{Z}_{pq} as a generic ring structure in order to prove the equivalence of RSA and factoring. In this paper, we assume the same structure. For our purpose, we only describe a generic algorithm according to the multiplication operation in \mathbb{Z}_{pq}^* in the following. It is worth noting that there are different ways of giving the definition (e.g. [28, 31]), and we follow that of Shoup [31].

Let $g \in_R \mathbb{Z}_{pq}^*$ and σ be an encoding function of $\mathbb{G} = \{g^i \mid i \in \mathbb{N}\}$ on $\{0, 1\}^{|pq|}$, where \mathbb{N} is the set of integers and $|pq|$ means the bit-length of pq . Suppose \mathcal{O} is a multiplication oracle, which, for any $r \in \mathbb{G}$, computes $\sigma(r)$ as follows: if r has been calculated before, then the same value of $\sigma(r)$ is returned; otherwise, it sets $\sigma(r)$ to be a random value from $\{0, 1\}^{|pq|} \setminus \mathcal{S}$, where \mathcal{S} is a set initialized to be $\{\sigma(g)\}$. A generic algorithm \mathcal{A} is a probabilistic algorithm, which takes \mathcal{S} and pq as input, and behaves as follows. At any time, \mathcal{A} can send a query with the input (x, y, b) , where $x, y \in \mathcal{S}$ and $b \in \{1, -1\}$, to \mathcal{O} , and will receive $\sigma(x \cdot y^b)$ as the reply. After every query, the result is added to the set \mathcal{S} .

Remark 1 *The concept of generic algorithms, which traces back to Babai and Szemeredi [3], has been extensively used to model algebraic objects, such as groups and rings. It has made it possible to prove the security of many cryptographic protocols. Note that this methodology does possess an inherent limitation, namely a lot of attacks against cryptographic protocols have actually been subject to the successful exploit of the structures of the underlying algebraic objects. Nevertheless, a security proof can still provide a certain level of confidence in the security of the studied protocols.*

To complete our proof, we further need the following computational assumption associated with the multiplication group \mathbb{Z}_{pq}^* . This assumption is similar to the representation assumption in groups of prime order made by Brands [7].

Definition 5 *Let p, q be two prime numbers. The extended discrete logarithm assumption holds for \mathbb{Z}_{pq}^* if the following event only occurs with a negligible probability with respect to a security parameter ℓ : Given $(pq, g, g_i (1 \leq i \leq V))$,*

where g, g_i ($1 \leq i \leq V$) are chosen uniformly at random from \mathbb{Z}_{pq}^* and V is any polynomial in the security parameter, a polynomial-time adversary finds $x \neq 0, x_i$ ($1 \leq i \leq V$) $\in \mathbb{N}$ such that

$$g^x \equiv \prod_{\substack{1 \leq i \leq V \\ x_i \in \mathbb{N}}} g_i^{x_i} \pmod{pq}.$$

In the following theorem, we prove that the variant scheme achieves the parallel computation resistance property (and certainly also the determinable difficulty property).

Theorem 1 *If the adversary is modeled as a generic algorithm as defined above, then the variant scheme achieves the parallel computation resistance property based on the extended discrete logarithm assumption (given in Definition 5) in the random oracle model.*

Proof sketch. In order to prove the theorem, we need to show that the adversary's advantage in the attack game (defined in Fig. 1) is negligible according to Definition 4. Since the adversary is modeled as a generic algorithm, there are two types of oracle queries it may issue. One is the PuzzleGen query which can be issued to the challenger, and the other is multiplication oracle query which can be issued to the generic group oracle \mathcal{O} .

We first consider a simple situation where the adversary does not issue any PuzzleGen query in the game. In this situation, the best strategy for the adversary to issue oracle queries to \mathcal{O} in the Response phase is the following.

1. At the beginning, the adversary issues two queries with the inputs $(\sigma(g^*), \sigma(g^*), 1)$ and $(\sigma(g^*), \sigma(g^*), -1)$. Clearly, until it receives the replies, namely $\sigma((g^*)^2)$ and $\sigma(1)$, from the oracle \mathcal{O} , it does not make sense for the adversary to send any other query.
2. After obtaining the replies, the set \mathcal{S} becomes $\{\sigma(1), \sigma(g^*), \sigma((g^*)^2)\}$. Then the adversary issues queries with the following inputs

$$(\sigma(g^*), \sigma((g^*)^2), 1), (\sigma((g^*)^2), \sigma((g^*)^2), 1),$$

$$(\sigma(1), \sigma(g^*), -1), (\sigma(1), \sigma((g^*)^2), -1).$$

Clearly, until it receives the replies, namely $\sigma((g^*)^3)$, $\sigma((g^*)^4)$, $\sigma((g^*)^{-1})$, and $\sigma((g^*)^{-2})$ from the oracle \mathcal{O} , it does not make sense for the adversary to send any other query.

3. After obtaining the replies, the set \mathcal{S} becomes

$$\{\sigma(1), \sigma(g^*), \sigma((g^*)^2), \sigma((g^*)^3), \sigma((g^*)^4), \sigma((g^*)^{-1}), \sigma((g^*)^{-2})\}.$$

Then the adversary issues queries with (x, y, b) , where $x, y \in \mathcal{S}$, $b \in \{1, -1\}$, and $\sigma(x \cdot y^b) \notin \mathcal{S}$. Clearly, until it receives the replies, it does not make sense for the adversary to send any other query.

4. The adversary continues the above process until it sends the response to the challenger.

Note that, we have enlarged the adversary's ability here, by allowing it to issue potentially an exponential number of queries to the oracle \mathcal{O} . In fact, a polynomial-time adversary can only issue a polynomial number of queries. Nonetheless, if such an extensively empowered adversary cannot win the game, then a polynomial-time adversary can neither.

Suppose, at the end of the game, the adversary performs $d^* - 1$ steps as above (which also means the *sequential computing time* is $d^* - 1$), then the set \mathcal{S} is a subset of

$$\{\sigma((g^*)^{-2^{d^*-1}}), \sigma((g^*)^{-2^{d^*-1}+1}), \dots, \sigma(1), \dots, \sigma((g^*)^{2^{d^*-1}-1}), \sigma((g^*)^{2^{d^*-1}})\}$$

In the game, the adversary could choose to submit a value $\sigma((g^*)^z)$ from \mathcal{S} or a value r from $\{0, 1\}^{|\mathcal{PQ}|} \setminus \mathcal{S}$ as the response to the challenger.

- In the first case, if $\sigma((g^*)^z) = \sigma((g^*)^{2^{d^*}})$, then the adversary has actually found $2^{d^*} - z$ such that $(g^*)^{2^{d^*}-z} \equiv 1 \pmod{pq}$. Based on the extended discrete logarithm assumption, the probability is negligible.
- In the second case, if $(g^*)^{2^{d^*}}$ has not been queried to the oracle, the probability $r = \sigma((g^*)^{2^{d^*}})$ is $\frac{Q}{pq}$ where Q is the total number of queries issued to the oracle \mathcal{O} . Clearly, this probability is negligible. Otherwise, if $(g^*)^{2^{d^*}}$ has been queried to the oracle \mathcal{O} , the probability $r = \sigma((g^*)^{2^{d^*}})$ is 0.

To summarize, in this simple situation, the adversary's advantage is negligible.

Next, we evaluate a more complex situation where the adversary is free to issue PuzzleGen oracle queries. This reflects the practical situation that a prover can ask the verifier for puzzles in a computational puzzle scheme. Our analysis will show that such a privilege does not help the adversary gain any benefit. Let $\sigma(g_1), \sigma(g_2), \dots, \sigma(g_V)$ be the elements resulted from

the replies of PuzzleGen oracle queries, where g_1, g_2, \dots, g_V are random elements from \mathbb{Z}_{pq}^* . Certainly, V is a polynomial in the security parameter. Suppose, at the end of the game, the adversary performs $d^* - 1$ steps in the Response phase, then the set \mathcal{S} is a subset of $\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_3$. The set \mathcal{S}_1 is the following.

$$\mathcal{S}_1 = \{\sigma((g^*)^{-2^{d^*-1}}), \sigma((g^*)^{-2^{d^*-1}+1}), \dots, \sigma(1), \dots, \sigma((g^*)^{2^{d^*-1}-1}), \sigma((g^*)^{2^{d^*-1}})\}$$

The set \mathcal{S}_2 contains a polynomial number of encodings of the form $\sigma(\prod_{x_i \in \mathbb{N}}^{1 \leq i \leq V} g_i^{x_i})$, and the set \mathcal{S}_3 contains a polynomial number of encodings of the form $\sigma(A \cdot B)$, where $\sigma(A) \in \mathcal{S}_1$ and $\sigma(B) \in \mathcal{S}_2$.

In the game, the adversary could choose to submit a value $\sigma((g^*)^z)$ from $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$, or a value r from $\{0, 1\}^{|pq|} \setminus \mathcal{S}$ as the response to the challenger.

- In the first case, from the analysis in the simple situation, the adversary's advantage is negligible.
- In the second and the third cases, if $\sigma((g^*)^z \cdot \prod_{x_i \in \mathbb{N}}^{1 \leq i \leq V} g_i^{x_i}) = \sigma((g^*)^{2^{d^*}})$, then the adversary has actually found $2^{d^*} - z$ such that $(g^*)^{2^{d^*} - z} \equiv \prod_{x_i \in \mathbb{N}}^{1 \leq i \leq V} g_i^{x_i} \pmod{pq}$. Based on the extended discrete logarithm assumption, the probability is negligible.
- In the fourth case, from the analysis in the simple situation, the adversary's advantage is negligible.

As a result, in this situation, the adversary's advantage is negligible. The theorem follows. \square

5 Conclusion

In this paper, we have revisited the concept of computational puzzle schemes, proposed a security model and presented formal definitions for two important properties, namely the determinable difficulty property and the parallel computation resistance property. We have proven that a variant of the RSW client puzzle scheme achieves both properties. To our knowledge, this is the first scheme which has been proven possessing the parallel computation resistance property. [Compared with many other schemes such as the hash-based ones \[4, 20\], the RSW scheme is computationally more expensive for the verifier, which needs to perform one exponentiation in verifying](#)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

a puzzle solution. Therefore, an interesting direction is to investigate how to improve the verifier's efficiency in practice. Another interesting direction is to analyse the memory-bound computational puzzle schemes in our security model.

References

- [1] M. Abadi, M. Burrows, M. Manasse, and T. Wobber. Moderately hard, memory-bound functions. *ACM Transactions on Internet Technology*, 5(2):299–327, 2005.
- [2] T. Aura, P. Nikander, and J. Leiwo. DoS-resistant authentication with client puzzles. In *Security Protocols, 8th International Workshop*, pages 170–177, 2000.
- [3] L. Babai and E. Szemerédi. On the complexity of matrix group problems i. *Foundations of Computer Science, Annual IEEE Symposium on*, pages 229–240, 1984.
- [4] A. Back. Hashcash — amortizable publicly auditable cost functions, 2002.
- [5] D. Boneh and M. Naor. Timed commitments. In *CRYPTO '00: Proceedings of the 20th Annual International Cryptology Conference on Advances in Cryptology*, pages 236–254. Springer, 2000.
- [6] N. Borisov. Computational puzzles as sybil defenses. In *P2P '06: Proceedings of the Sixth IEEE International Conference on Peer-to-Peer Computing*, pages 171–176. IEEE Computer Society, 2006.
- [7] S. Brands. Untraceable off-line cash in wallets with observers (extended abstract). In D. R. Stinson, editor, *Advances in Cryptology — CRYPTO 1993*, volume 773 of *Lecture Notes in Computer Science*, pages 302–318. Springer, 1993.
- [8] J. Cai, R. J. Lipton, R. Sedgewick, and A. C. Yao. Towards uncheatable benchmarks. In *Structure in Complexity Theory Conference*, pages 2–11, 1993.
- [9] L. Chen, P. Morrissey, N. Smart, and B. Warinschi. Security notions and generic constructions for client puzzles. In *Advances in Cryptology — Asiacrypt 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 505–523. Springer, 2009.

- 1
2
3
4
5
6
7
8
9
10
11
12
13 [10] D. Dean and A. Stubblefield. Using client puzzles to protect TLS. In
14 *SSYM'01: Proceedings of the 10th conference on USENIX Security Symposi-*
15 *um*, pages 1–1. USENIX Association, 2001.
- 16
17 [11] S. Doshi, F. Monrose, and A. D. Rubin. Efficient memory bound puzzles
18 using pattern databases. In *Applied Cryptography and Network Security,*
19 *4th International Conference, ACNS 2006*, pages 98–113, 2006.
- 20
21 [12] C. Dwork, A. Goldberg, and M. Naor. On memory-bound functions
22 for fighting spam. In Dan Boneh, editor, *Advances in Cryptology —*
23 *CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages
24 426–444. Springer, 2003.
- 25
26 [13] C. Dwork and M. Naor. Pricing via processing or combatting junk mail.
27 In E. F. Brickell, editor, *Advances in Cryptology — CRYPTO '92*, volume
28 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer, 1992.
- 29
30 [14] W. Feng, E. Kaiser, W. Feng, and A. Luu. The design and implemen-
31 tation of network puzzles. In *Proceedings of the 24th Annual Joint Con-*
32 *ference of the IEEE Computer and Communications Societies (INFOCOM*
33 *2005)*, pages 2372–2382, 2005.
- 34
35 [15] M. K. Franklin and D. Malkhi. Auditable metering with lightweight
36 security. In R. Hirschfeld, editor, *Financial Cryptography, First Interna-*
37 *tional Conference, FC '97, Proceedings*, volume 1318 of *Lecture Notes in*
38 *Computer Science*, pages 151–160. Springer, 1997.
- 39
40 [16] N.A. Fraser, D.J. Kelly, R.A. Raines, R.O. Baldwin, and B.E. Mullins.
41 Using client puzzles to mitigate distributed denial of service attacks in
42 the tor anonymous routing environment. *Communications, 2007. ICC*
43 *'07. IEEE International Conference on*, pages 1197–1202, 2007.
- 44
45 [17] J. A. Garay and M. Jakobsson. Timed release of standard digital sig-
46 natures. In M. Blaze, editor, *Financial Cryptography, 6th International*
47 *Conference, FC 2002*, volume 2357 of *Lecture Notes in Computer Science*,
48 pages 168–182. Springer, 2002.
- 49
50 [18] D. Goldschlag and S. Stubblebine. Publicly verifiable lotteries: Ap-
51 plications of delaying functions. In *FC '98: Proceedings of the Sec-*
52 *ond International Conference on Financial Cryptography*, pages 214–226.
53 Springer-Verlag, 1998.
- 54
55
56
57
58
59
60

- 1
2
3
4
5
6
7
8
9
10
11
12
13 [19] M. Jakobsson and A. Juels. Proofs of work and bread pudding pro-
14 tocols. In *CMS '99: Proceedings of the IFIP TC6/TC11 Joint Working*
15 *Conference on Secure Information Networks*, pages 258–272. Kluwer, B.V.,
16 1999.
- 17
18 [20] A. Juels and J. G. Brainard. Client puzzles: A cryptographic coun-
19 termeasure against connection depletion attacks. In *Proceedings of the*
20 *Network and Distributed System Security Symposium, NDSS 1999*, pages
21 151–165, 1999.
- 22
23 [21] B. Laurie and R. Clayton. Proof-of-Work Proves not to Work. In *Third*
24 *Annual Workshop on Economics of Information Security (WEIS04)*, 2004.
- 25
26 [22] G. Leander and A. Rupp. On the equivalence of rsa and factoring
27 regarding generic ring algorithms. In X. Lai and K. Chen, editors,
28 *Advances in Cryptology — ASIACRYPT 2006*, volume 4284 of *Lecture*
29 *Notes in Computer Science*, pages 241–251. Springer, 2006.
- 30
31 [23] M. Lee and C. Fung. A public-key based authentication and key es-
32 tablishment protocol coupled with a client puzzle. *J. Am. Soc. Inf. Sci.*
33 *Technol.*, 54(9):810–823, 2003.
- 34
35 [24] Y. Lei, S. Pierre, and A. Quintero. Client puzzles based on quasi partial
36 collisions against DoS attacks in UMTS. *Proceedings of the 64th IEEE*
37 *Vehicular Technology Conference*, pages 1–5, 2006.
- 38
39 [25] D. liu and L. Jean Camp. Proof of Work can Work. In *Fifth Workshop on*
40 *the Economics of Information Security (WEIS06)*, 2006.
- 41
42 [26] I. Martinovic, F. A. Zdarsky, M. Wilhelm, C. Wegmann, and J. B.
43 Schmitt. Wireless client puzzles in IEEE 802.11 networks: Security
44 by wireless. In *WiSec '08: Proceedings of the first ACM conference on*
45 *Wireless network security*, pages 36–45. ACM, 2008.
- 46
47 [27] R. C. Merkle. Secure communications over insecure channels. *Commu-*
48 *nications of the ACM*, pages 294–299, 1978.
- 49
50 [28] V. I. Nechaev. Complexity of a determinate algorithm for the discrete
51 logarithm. *Math. Notes*, 55(2):91–101, 1994.
- 52
53 [29] P. Ning, A. Liu, and W. Du. Mitigating dos attacks against broadcast
54 authentication in wireless sensor networks. *ACM Transactions on Sensor*
55 *Networks*, 4(1), 2008.
- 56
57
58
59
60

- 1
2
3
4
5
6
7
8
9
10
11
12
13 [30] R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles
14 and timed-release crypto. Technical Report MIT/LCS/TR-684, Mas-
15 sachusetts Institute of Technology, 1996.
- 16
17 [31] V. Shoup. Lower bounds for discrete logarithms and related problems.
18 In *Advances in Cryptology — EUROCRYPT '97*, volume 1233 of *Lecture*
19 *Notes in Computer Science*, pages 256–266. Springer, 1997.
- 20
21 [32] S. Tritilanunt, C. Boyd, E. Foo, and J. M. González Nieto. Toward
22 non-parallelizable client puzzles. In *Cryptology and Network Security,*
23 *6th International Conference, CANS 2007*, pages 247–264, 2007.
- 24
25 [33] X. Wang and M. K. Reiter. Defending against denial-of-service attacks
26 with puzzle auctions. In *SP '03: Proceedings of the 2003 IEEE Symposium*
27 *on Security and Privacy*, page 78, Washington, DC, USA, 2003. IEEE
28 Computer Society.
- 29
30 [34] X. Wang and M. K. Reiter. Mitigating bandwidth-exhaustion attacks
31 using congestion puzzles. In *CCS '04: Proceedings of the 11th ACM*
32 *conference on Computer and communications security*, pages 257–267, 2004.
- 33
34 [35] X. Wang and M. K. Reiter. A multi-layer framework for puzzle-based
35 denial-of-service defense. *International Journal of Information Security,*
36 pages 243–263, 2008.
- 37
38 [36] B. Waters, A. Juels, J. A. Halderman, and E. W. Felten. New client
39 puzzle outsourcing techniques for DoS resistance. In *Proceedings of the*
40 *11th ACM Conference on Computer and Communications Security, CCS*
41 *2004*, pages 246–256, 2004.
- 42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60