Constructions and Bounds for Unconditionally Secure Commitment Schemes

C. Blundo and B. Masucci
Dipartimento di Informatica ed Applicazioni
Università di Salerno
Baronissi (SA), 84081
Italy
{carblu, masucci}@dia.unisa.it

D.R. Stinson
Department of Combinatorics and Optimization
University of Waterloo
Waterloo, Ontario N2L 3G1
Canada
dstinson@cacr.math.uwaterloo.ca

R. Wei
Department of Computer Science
Lakehead University
Thunder Bay, Ontario P7B 5E1
Canada
wei@ccc.math.lakeheadu.ca

August 15, 2000

Abstract

Commitment schemes have been extensively studied since they were introduced by Blum in 1982. Rivest recently showed how to construct unconditionally secure commitment schemes, assuming the existence of a trusted initializer. In this paper, we present a formal mathematical model for such schemes, and analyze their binding and concealing properties. In particular, we show that such schemes cannot be perfectly concealing: there is necessarily a small probability that Alice can cheat Bob by committing to one value but later revealing a different value. We prove several bounds on Alice's cheating probability, and present constructions of schemes that achieve optimal cheating probabilities. We also show a close link between commitment schemes and the classical "affine resolvable designs".

1 Introduction

Commitment schemes were first introduced by Blum [2] in 1982. In a commitment scheme, there is a sender, Alice, and a receiver, Bob. The scheme consists of a commit protocol and a reveal protocol. When Alice wants to commit a secret value x_0 to Bob, she uses the commit protocol to send some information, say y_0, to Bob. Although Bob should learn nothing about the value of x_0 at this stage, he can use the information y_0 for verification at a later time. Alice will use the reveal protocol to disclose the value of x_0. Bob should be able to reject the value x_0 disclosed by Alice if it is not consistent with the previously committed information y_0. A commitment scheme should satisfy the following (informal) properties:

- correctness: If both parties are honest and follow the protocols, then during the reveal protocol Bob will learn the (unique) value x_0 the Alice wished to commit to.
- concealing: Bob learns nothing about the value of x_0 during the commit protocol.
- binding: After the commit protocol has finished, there is only one value of x_0 that Bob will accept during the reveal protocol (i.e., Alice cannot "change her mind" regarding the value she committed to).

Commitment schemes have been studied by several researchers recently. It is indicated in [4] that there does not exist an unconditionally secure commitment scheme in the two-party scenario, assuming only noiseless communication. Therefore researchers have concentrated primarily on "partially" unconditionally secure schemes. Schemes in which one of Alice or Bob is computationally bounded have been proposed. Such schemes would be termed *computationally binding* or *computationally concealing*, respectively.

Rivest [5] suggested a three-party scenario for unconditionally secure commitment schemes. In this model, both Alice and Bob are not limited in computational power. As well, a "trusted initializer", Ted, is introduced. We assume that Ted is honest and there exist secure private channels connecting Ted to Alice and Bob. Ted's activity is limited to an initialize protocol, that takes place before the commit and reveal protocols, in which Ted provides some information to Alice and Bob. The information provided by Ted in the initialize protocol is assumed to be independent from the value x_0 that Alice will commit to. After initialize, Ted takes no further part in the scheme. The scheme is a "one-time" scheme in that a re-initialization is required for each value that Alice wishes to commit to.

Rivest's commitment scheme [5] allows Alice to commit to an element of $\mathbb{Z}_p$, for some prime p. The scheme works as follows:

Algorithm 1.1 (Rivest's Commitment Scheme)

initialize — Ted chooses random values $a \in \mathbb{Z}_p^*$, $b \in \mathbb{Z}_p$ and $x_1 \in \mathbb{Z}_p$.
Then he computes $y_1 = (ax_1 + b) \bmod p$.
Ted then privately sends (a, b) to Alice and (x_1, y_1) to Bob.

commit — Suppose Alice wants to commit to the value $x_0 \in \mathbb{Z}_p$.
Then she computes $y_0 = (ax_0 + b) \bmod p$ and sends y_0 to Bob.

reveal — Alice sends (a, b) and x_0 to Bob.
Bob verifies that $ax_1 + b \equiv y_1 \bmod p$ and $ax_0 + b \equiv y_0 \bmod p$.
If both these congruences hold, Bob accepts x_0, otherwise he rejects it.

We observe that this scheme is not "perfectly" concealing. Suppose it happens that $y_0 = y_1$. Then Bob will know that $x_0 = x_1$ immediately after the commit protocol. On the other hand, if $y_0 \neq y_1$, then $x_0 \neq x_1$. In this situation, Bob knows that $x_0 \neq x_1$ immediately after the commit protocol. In any case, Bob obtains some partial information about x_0 after the commit protocol. It is not difficult to modify Rivest's scheme so that it is concealing; we will do this a bit later in the paper.

The remainder of this paper is organized as follows. In Section 2, we present our formal mathematical model. We also show in that section that there is no way to achieve binding with probability one in any unconditionally secure commitment scheme. (This fact was already observed in [5] in reference to the scheme presented there.) Several bounds on the binding probability are proven. In Section 3, we present a new scheme, which is a modification of Rivest's scheme that is concealing and which achieves optimal binding. In Sections 4 and 5, we study a special type of commitment scheme, based on a simplified protocol, and show that these schemes are closely related to "affine resolvable designs".

2 Properties of General Commitment Schemes

We begin with a formal definition of an unconditionally secure commitment scheme. Then we discuss the binding property, and show that there is always some positive probability that the sender (Alice) can deceive the receiver (Bob). The scheme will use the following notation:

- let E_A be the set of all possible *encoding keys for the sender*;
- let F_A be the set of all possible *authentication keys for the sender*;
- let E_B be the set of all possible *keys for the receiver*;
- let $E \subseteq E_A \times F_A \times E_B$ be the set of all possible *keys*;

- let X be the set of all possible *source states* that Alice may want to commit to Bob;
- let Y be the set of the *encoding states*;
- $e_A : X \to Y$ is an injective function for all $e_A \in E_A$;
- let R be the set of *authenticating states*;
- $f_A : X \to R$ for all $f_A \in F_A$;
- $\mathsf{test} : X \times Y \times R \times E_B \to \{\mathsf{true}, \mathsf{false}\}$.

In general, a commitment scheme can be described as follows:

Algorithm 2.1 (General Commitment Scheme)

initialize	The trusted initializer, Ted, randomly chooses a key $e = (e_A, f_A, e_B) \in E$. e_A and f_A are sent to Alice, and e_B is sent to Bob. (After the initialize protocol, T becomes inactive.)
commit	Suppose Alice wants to commit to the source state $x \in X$. She computes the value $y = e_A(x)$, and sends it to Bob.
reveal	Alice computes the value $r = f_A(x)$ and sends x and r to Bob. Bob accepts x if $\mathsf{test}(x, y, r, e_B) = \mathsf{true}$, and rejects x, otherwise.

Remark. We require that each function e_A be injective in order to avoid the following situation: Suppose $e_A(x) = e_A(x')$ for some $x' \neq x$. Then Alice can commit to the value $y = e_A(x)$, but later reveal either of x or x'. This situation is not defined as cheating, but we still do not want it to occur. This is why we stipulate that each e_A be injective.

2.1 Binding Probabilities

We are interested in the probability that Alice can commit to the value x by sending the value $y = e_A(x)$ to Bob in the commit phase, but later reveal values x' ($x \neq x$) and r such that Bob accepts y as being a commitment of x'. For any $x' \in X$ and $y \in Y$, and for any $r \in R$, define

$$\mathsf{accept}(x', y, r) = \{e_B \in E_B : \mathsf{test}(x', y, r, e_B) = \mathsf{true}\}.$$

The set $\mathsf{accept}(x', y, r)$ consists of all the keys e_B which will cause Bob to accept y as a commitment of x' when r is specified as the authenticating state.

Suppose that $x' \neq x$ and $y = e_A(x)$. The fact that e_A is injective implies that $y \neq e_A(x')$. Therefore, Alice will be able to cheat Bob by giving him the values x', y and r if and only if

$$e_B \in \mathsf{accept}(x', e_A(x), r).$$

The probability that the choices $x \neq x'$ and r will deceive Bob is computed to be

$$\mathsf{cheat}(e_A, f_A; x, x', r) = \sum_{e_B \in \mathsf{accept}(x', e_A(x), r)} pr(e_B | e_A, f_A). \tag{1}$$

For a fixed pair (e_A, f_A), Alice will maximize her probability of deceiving Bob by choosing $x' \neq x$ and r so that $\mathsf{cheat}(e_A, f_A; x, x', r)$ is maximized.

Let $0 < \epsilon < 1$. We will say that a commitment scheme is $(1-\epsilon)$-*binding* if Alice's cheating probability is at most ϵ, regardless of the keys (e_A, f_A) that she holds. The formal properties that a $(1-\epsilon)$-binding commitment scheme should satisfy are the following:

correctness For all $e = (e_A, f_A, e_B)$ such that $pr(e_A, f_A, e_B) > 0$, and for all $x \in X$, $\mathsf{test}(x, e_A(x), f_A(x), e_B) = \mathsf{true}$. [If Alice and Bob follow the protocol, Bob will accept the value x that Alice commits to.]

concealing For all $e = (e_A, f_A, e_B)$ such that $pr(e_A, f_A, e_B) > 0$, and for all $x, x' \in X$, there exists a value $r \in R$ such that $\mathsf{test}(x', e_A(x), r, e_B) = \mathsf{true}$. [If Alice and Bob follow the protocol, Bob cannot rule out any possible values of x before the reveal step is performed.]

$(1-\epsilon)$-binding For all $e = (e_A, f_A, e_B)$ such that $pr(e_A, f_A, e_B) > 0$, for all $x, x' \in X$ with $x' \neq x$, and for all $r \in R$, it holds that $\mathsf{cheat}(e_A, f_A; x, x', r) \leq \epsilon$. [Alice cannot successfully cheat Bob with probability more than ϵ, regardless of the keys she and Bob hold.]

For all (e_A, f_A) with $pr(e_A, f_A) > 0$, define

$$\mathsf{possible}(e_A, f_A) = \{e_B : pr(e_A, f_A, e_B) > 0.\}$$

The following theorem establishes a lower bound on ϵ.

Theorem 2.1 *In any* $(1-\epsilon)$-binding *commitment scheme, it holds that*

$$\epsilon \geq \max\{1/|\mathsf{possible}(e_A, f_A)| : pr(e_A, f_A) > 0\}.$$

Proof. Suppose that Alice receives the keys e_A, f_A. We describe how Alice can cheat Bob with probability at least $1/|\mathsf{possible}(e_A, f_A)|$. First, Alice chooses $e_{B_0} \in \mathsf{possible}(e_A, f_A)$ such that $pr(e_{B_0}|e_A, f_A)$ is maximized. Note that $pr(e_{B_0}|e_A, f_A) \geq 1/|\mathsf{possible}(e_A, f_A)|$. Next, Alice picks arbitrary values $x, x' \in X$ such that $x \neq x'$. Then Alice determines a value r such that $\mathsf{test}(x', e_A(x), r, e_{B_0}) = \mathsf{true}$ (such an r exists by the concealing condition). Then it is clear that $e_{B_0} \in \mathsf{accept}(x', e_A(x), r)$. Since $pr(e_{B_0}|e_A, f_A) \geq 1/|\mathsf{possible}(e_A, f_A)|$, we are done. □

2.2 Average Binding Probabilities

We defined a scheme to be $(1-\epsilon)$-binding if Alice's cheating probability is at most ϵ, regardless of the keys she holds. Another way to measure the binding property of a scheme is to consider Alice's average probability of cheating, where the average is computed over all possible keys (e_A, f_A) that she might receive, assuming that Alice adopts an optimal cheating strategy for each possible choice of (e_A, f_A). For each (e_A, f_A) with $pr(e_A, f_A) > 0$, define

$$\mathsf{optimalcheat}(e_A, f_A) = \max_{x,x',r}\{\mathsf{cheat}(e_A, f_A; x, x', r)\}.$$

Then the average probability that Alice can cheat Bob, denoted $\overline{\epsilon}$, is given by the following formula:

$$\overline{\epsilon} = \sum_{\{(e_A,f_A):pr(e_A,f_A)>0\}} pr(e_A, f_A)\,\mathsf{optimalcheat}(e_A, f_A). \tag{2}$$

We will prove an entropy bound for $\overline{\epsilon}$, but first we require a few definitions. For an arbitrary finite set X, we use $\mathbf{X}$ to denote a random variable taking on values from X. For $x \in X$ we denote by $pr(\mathbf{X} = x)$ the probability that the random variable $\mathbf{X}$ takes on the value x. We often write $pr(x)$ for $pr(\mathbf{X} = x)$, for short. The *entropy* of $\mathbf{X}$ is defined to be the quantity

$$H(\mathbf{X}) = -\sum_{x \in X} pr(x) \log_2 pr(x).$$

Suppose we have two finite sets X and Y with associated random variables $\mathbf{X}$ and $\mathbf{Y}$. For any $y \in Y$, define

$$H(\mathbf{X}|y) = -\sum_{x \in X} pr(x|y) \log_2 pr(x|y).$$

Then define

$$H(\mathbf{X}|\mathbf{Y}) = -\sum_{y \in Y} pr(y) H(\mathbf{X}|y).$$

$H(\mathbf{X}|\mathbf{Y})$ is the *conditional entropy* of the random variable $\mathbf{X}$, given $\mathbf{Y}$.

In the proof of Theorem 2.1, we used the fact that, if Alice can guess Bob's key, e_B, then she can deceive Bob with probability one. Therefore, for each (e_A, f_A), we have that

$$\mathsf{optimalcheat}(e_A, f_A) \geq \max_{e_B}\{pr(e_B|e_A, f_A)\}.$$

Now it is easy to see that

$$\max_{e_B}\{pr(e_B|e_A, f_A)\} \geq 2^{-H(\mathbf{E_B}|e_A, f_A)}.$$

Substituting the above inequalities into equation (2) and taking logarithms, we obtain

$$\begin{aligned}
\log_2 \overline{\epsilon} &\geq \log_2 \left(\sum_{\{(e_A,f_A): pr(e_A,f_A)>0\}} pr(e_A, f_A) 2^{-H(\mathbf{E_B}|e_A,f_A)} \right) \\
&\geq \sum_{\{(e_A,f_A): pr(e_A,f_A)>0\}} pr(e_A, f_A) \log_2 2^{-H(\mathbf{E_B}|e_A,f_A)} \\
&= - \sum_{\{(e_A,f_A): pr(e_A,f_A)>0\}} pr(e_A, f_A) H(\mathbf{E_B}|e_A, f_A) \\
&= -H(\mathbf{E_B}|\mathbf{E_A}, \mathbf{F_A}).
\end{aligned}$$

We have therefore proven the following result.

Theorem 2.2 *Alice's average cheating probability, $\overline{\epsilon}$, is at least $2^{-H(\mathbf{E_B}|\mathbf{E_A},\mathbf{F_A})}$, assuming she uses an optimal cheating strategy.*

In Sections 3 and 4, we will construct schemes for which the bound of Theorem 2.2 is met with equality.

3 The Affine Plane Scheme

We present a modification of Rivest's scheme which is perfectly concealing. The following scheme is based on an affine plane of order p (as is Rivest's scheme).

Algorithm 3.1 (Affine Plane Commitment Scheme)

initialize
: Ted chooses random values $a \in \mathbb{Z}_p$, $b \in \mathbb{Z}_p$ and $x_1 \in \mathbb{Z}_p$.
Then he computes $y_1 = (ax_1 + b) \bmod p$.
Ted then privately sends (a, b) to Alice and (x_1, y_1) to Bob.

commit
: Suppose Alice wants to commit to the value $x_0 \in \mathbb{Z}_p$.
Then she computes $y_0 = (x_0 + a) \bmod p$ and sends y_0 to Bob.

reveal
: Alice sends (a, b) and x_0 to Bob.
Bob verifies that $ax_1 + b \equiv y_1 \bmod p$ and $x_0 + a \equiv y_0 \bmod p$.
If both these congruences hold, Bob accepts x_0, otherwise he rejects it.

It is clear that the above scheme satisfies the correctness condition. The following theorem establishes the concealing property of the scheme.

Theorem 3.1 *The commitment scheme presented in Algorithm 3.1 is unconditionally concealing.*

Proof. After the initialize and commit protocols, Bob only knows x_1, y_1 and y_0. Now, for any $x_0' \in \mathbb{Z}_p$, there exist unique values $a', b' \in \mathbb{Z}_p$, such that $x_0' + a' \equiv y_0 \bmod p$ and $a'x_1 + b' \equiv y_1 \bmod p$ (namely, $a' = (y_0 - x_0') \bmod p$ and $b' = (y_1 - (y_0 - x_0')x_1) \bmod p$). Thus Bob obtains no information about x_0: his probability of guessing the value of x_0 after the commit protocol is the same as his probability of guessing x_0 before the commit protocol, namely, $1/p$. ∎

3.1 Binding Probability of the Affine Plane Scheme

The results of this section will show that the affine plane commitment scheme achieves optimal binding. We first observe that Alice can deceive Bob with probability $1/p$ (which proves that the probability of binding is at most $1-1/p$). This observation is an immediate application of Theorem 2.1, which says that the binding probability is at most $1-\epsilon$, where

$$\epsilon = \max\{1/|\mathsf{possible}(e_A, f_A)| : pr(e_A, f_A) > 0\}.$$

In Algorithm 3.1, we have that

$$\mathsf{possible}(e_A, f_A) = \{(x_1, y_1) \in \mathbb{Z}_p \times \mathbb{Z}_p : y_1 \equiv ax_1 + b \bmod p\},$$

so $|\mathsf{possible}(e_A, f_A)| = p$. Therefore, $1-\epsilon = 1-1/p$, as desired.

We now prove that there is no way for Alice to cheat Bob with higher probability.

Theorem 3.2 *In the affine plane commitment scheme, the binding probability is equal to* $1-1/p$.

Proof. Let $a, b \in \mathbb{Z}_p$ be fixed; let $x_0, x_0' \in \mathbb{Z}_p$ with $x_0 \neq x_0'$; and let $y_0 = (x_0 + a) \bmod p$. Define $a' = (y_0 - x_0') \bmod p$, and let $b' \in \mathbb{Z}_p$. We need to determine the set $\mathsf{accept}(x_0', y_0, (a', b'))$. This set is

$$\mathsf{accept}(x_0', y_0, (a', b')) = \{(x_1, y_1) \in \mathbb{Z}_p \times \mathbb{Z}_p : y_1 \equiv ax_1 + b \bmod p \text{ and } y_1 \equiv a'x_1 + b' \bmod p\}.$$

Let $\mathcal{L}$ be the line $y = (ax+b) \bmod p$ and let $\mathcal{L}'$ be the line $y = (a'x + b') \bmod p$. Since $a \not\equiv a' \bmod p$, it follows that the intersection of $\mathcal{L}$ and $\mathcal{L}'$ consists of a single point, say $\mathcal{L} \cap \mathcal{L}' = \{(x_1, y_1)\}$. Then $\mathsf{accept}(x_0', y_0, (a', b')) = \{(x_1, y_1)\}$, and we have that

$$\mathsf{cheat}(e_A, f_A; x_0, x_0', (a', b')) = pr((x_1, y_1)|(a, b)) = \frac{1}{p},$$

since x_1 is chosen randomly from $\mathbb{Z}_p$, and y_1 is determined uniquely, given x_1. □

4 A Simplified Protocol

In this section, we study a certain type of simplified commitment protocol, where $r = e_A$ (independent of the value x that Alice commits to). In this situation, we do not require an authenticating key, f_A. Therefore, the set of keys $E \subseteq E_A \times E_B$, and we define

$$\mathsf{possible}(e_A) = \{e_B : pr(e_A, e_B) > 0.\}$$

Suppose that $e_B \in \mathsf{possible}(e_A)$ and $y = e_A(x)$. Then $\mathsf{correctness}$ implies that $\mathsf{test}(x, y, e_A, e_B) = \mathsf{true}$. Conversely, it seems reasonable that Bob should not accept e_A as Alice's key if $e_B \notin \mathsf{possible}(e_A)$, and Bob should not accept y as a commitment of x if $y \neq e_A(x)$. Therefore, we will assume that the function test is defined as follows:

$$\mathsf{test}(x, y, e_A, e_B) = \mathsf{true} \Leftrightarrow e_B \in \mathsf{possible}(e_A) \text{ and } y = e_A(x).$$

The following is our simplified commitment scheme.

Algorithm 4.1 (Simplified Commitment Scheme)

initialize	The trusted initializer, Ted, randomly chooses a key $e = (e_A, e_B) \in E$. e_A is sent to Alice, and e_B is sent to Bob. (After the initialize protocol, T becomes inactive.)
commit	Suppose Alice wants to commit to the source state $x \in X$. She computes the value $y = e_A(x)$, and sends it to Bob.
reveal	Alice sends x and e_A to Bob. Bob accepts x if $y = e_A(x)$ and $e_B \in \mathsf{possible}(e_A)$, and rejects x, otherwise.

The affine plane scheme described in Algorithm 3.1 can be viewed as an example of the simplified commitment scheme. This scheme can in fact be generalized using any resolvable design. A *t-design*, $S_\lambda(t,k,v)$, is a pair $(V,\mathcal{W})$, where V is a set of v elements called *points* and $\mathcal{W}$ is a collection of k-subsets of V called *blocks*, such that every t-subset of points occurs in exactly λ blocks. A *parallel class* in an $S_\lambda(t,k,v)$, $(V,\mathcal{W})$, is a set of v/k blocks that form a partition of V. $(V,\mathcal{W})$ is said to be *resolvable* if the block set $\mathcal{W}$ can be partitioned into parallel classes.

Suppose that $(V,\mathcal{W})$ is a resolvable $S_r(1,k,v)$, and let the parallel classes be denoted $\mathcal{P}_i$, $i=0,\ldots,r-1$. Define $E_A=\{e_W : W\in\mathcal{W}\}$ and $E_B=V$. Define $E=\{(W,v) : v\in W\in\mathcal{W}\}$, so $\mathsf{possible}(e_W)=W$. Finally, for each $W\in\mathcal{W}$ define the encoding rule $e_W:\mathbb{Z}_r\to\mathbb{Z}_r$ as follows:

$$e_W(x)=(x+i) \bmod r,$$

where $W\in\mathcal{P}_i$. The resulting scheme is as follows.

Algorithm 4.2 (Resolvable Design Commitment Scheme)

initialize Ted chooses a random pair $(W,v)\in E$.
Ted then sends W to Alice and v to Bob.

commit Suppose Alice wants to commit to the value $x_0\in\mathbb{Z}_r$.
Then she computes $y_0=(x_0+i) \bmod r$, where $W\in\mathcal{P}_i$, and sends y_0 to Bob.

reveal Alice sends W and x_0 to Bob.
Bob verifies that $W\in\mathcal{P}_i$, $y_0\equiv x_0+i \bmod r$ and $v\in W$.
If these conditions hold, Bob accepts x_0, otherwise he rejects it.

We present a tiny example, using a resolvable $(6,2,1)$-BIBD (i.e., a one-factorization of the complete graph K_6).

W	$e_W(x)$				
	$x=0$	1	2	3	4
$\{0,5\}$	0	1	2	3	4
$\{1,4\}$	0	1	2	3	4
$\{2,3\}$	0	1	2	3	4
$\{1,5\}$	1	2	3	4	0
$\{2,0\}$	1	2	3	4	0
$\{3,4\}$	1	2	3	4	0
$\{2,5\}$	2	3	4	0	1
$\{3,1\}$	2	3	4	0	1
$\{4,0\}$	2	3	4	0	1
$\{3,5\}$	3	4	0	1	2
$\{4,2\}$	3	4	0	1	2
$\{0,1\}$	3	4	0	1	2
$\{4,5\}$	4	0	1	2	3
$\{0,3\}$	4	0	1	2	3
$\{1,2\}$	4	0	1	2	3

Theorem 4.1 *The resolvable design commitment scheme is unconditionally concealing.*

Proof. (Sketch) For each $i\in\mathbb{Z}_r$ and each $v\in V$, there exists a block $W\in\mathcal{W}$ such that $v\in W\in\mathcal{P}_i$. Thus, given the point v held by Bob, the value of i is completely undetermined. Since $y_0\equiv x_0+i \bmod r$, it follows that the value of x_0, given y_0, is completely undetermined. ∎

For any (resolvable) $S_\lambda(1,k,v)$ $(V,\mathcal{W})$, define

$$\mu(V,\mathcal{W})=\max\{|W_1\cap W_2| : W_1,W_2\in\mathcal{W} : W_1\neq W_2\}.$$

Theorem 4.2 *The resolvable design commitment scheme has binding probability equal to* $1 - \mu/k$.

Proof. (Sketch) Suppose that (W) is Alice's key, where $W \in \mathcal{P}_i$. Let $x \in \mathbb{Z}_r$, define $y = (x+i) \bmod r$, and let $x' \in \mathbb{Z}_r$, $x' \neq x$. Define $j = y - x' \bmod r$, and let $W' \in \mathcal{P}_j$. Then we have that $\mathsf{accept}(x', y, W') = W \cap W'$. It then follows that

$$\mathsf{cheat}(W; x, x', W') = \sum_{v \in W \cap W'} pr(v|W) = \sum_{v \in W \cap W'} \frac{1}{k} = \frac{|W \cap W'|}{k}.$$

Now $|W \cap W'| \leq \mu$, so the scheme is $(1-\epsilon)$-binding, where $\epsilon = \mu/k$. □

5 Optimal Schemes and Resolvable Designs

In this section, we address the problem of constructing simplified schemes that are optimal with respect to various parameters. First, we show that the general class of resolvable design schemes are, given some reasonable simplifying assumptions, optimal within the class of simplified schemes.

Associated with any simplified scheme is an incidence structure $(E_B, \{\mathsf{possible}(e_A) : e_A \in E_A\})$. Suppose we denote this incidence structure as $(V, \mathcal{W})$, as was done in describing the resolvable design scheme. For each $W \in \mathcal{W}$, let e_W be the encoding function associated with the block W. Then we have that $e_W : X \to Y$ for each $W \in \mathcal{W}$. Since each e_W is injective, it must be the case that $|Y| \geq |X|$.

Let $v \in V$, and let $\mathcal{W}(v) \subseteq \mathcal{W}$ denote the blocks that contain the point v. Define

$$\mathcal{Y}(v) = \{e_W(x) : W \in \mathcal{W}(v), x \in X\}.$$

Let $x_0 \in X$ be fixed, and consider the set $\{e_W(x_0) : W \in \mathcal{W}(v)\}$. Suppose that

$$y \in \mathcal{Y}(v) \backslash \{e_W(x_0) : W \in \mathcal{W}(v)\}.$$

Then it follows that the scheme is not concealing. Contrapositively, if the scheme is concealing, then

$$\mathcal{Y}(v) = \{e_W(x_0) : W \in \mathcal{W}(v)\}$$

for all $x_0 \in X$. It is also clear that $\mathcal{Y}(v) \subseteq Y$ for all $v \in V$.

Now, suppose that $\mathcal{Y}(v) = Y$ for all $v \in V$. Then it follows that

$$\{e_W(x_0) : W \in \mathcal{W}(v)\} = Y$$

for all $x_0 \in X$. Consequently, it must be the case that $|\mathcal{W}(v)| \geq |Y|$. Suppose we additionally stipulate that $|\mathcal{W}(v)| = |Y|$ for all v. Fix $x_0 \in X$, and, for each $y \in Y$, define

$$\mathcal{P}_y = \{W \in \mathcal{W} : e_W(x_0) = y\}.$$

Then each $\mathcal{P}_y$, $y \in Y$, is a parallel class, and the incidence structure $(V, \mathcal{W})$ is resolvable.

At this point, after making certain assumptions about the commitment scheme, we have obtained a resolvable incidence structure. This incidence structure is not necessarily an $S_r(1, k, v)$ because the blocks do not necessarily all have the same size. However, if we consider the binding probabilities of the scheme, we will see that is advantageous to consider incidence structures with constant block size.

We first prove a useful lemma.

Lemma 5.1 *Suppose that* $(V, \mathcal{W})$ *is a resolvable incidence structure with* $|V| = v$, $\mathcal{W} = b$, *and* r *parallel classes, denoted* $\mathcal{P}_i$, $0 \leq i \leq r-1$. *Then there exist two distinct blocks* W, W' *with* $|W \cap W'| \geq r^2 v/b^2$.

Proof. Let $b_i = |\mathcal{P}_i|$, $0 \leq i \leq r-1$. Suppose without loss of generality that $b_0 \leq \ldots \leq b_{r-1}$. The average size of a block in $\mathcal{P}_0$ is v/b_0, so there exists a block $W \in \mathcal{P}_0$ having size at least v/b_0. Now, the average size of an intersection of W with a block from $\mathcal{P}_1$ is $|W|/b_1$, so there exists a block $W' \in \mathcal{P}_1$ such that

$$|W \cap W'| \geq \frac{|W|}{b_1} \geq \frac{v}{b_0 b_1}.$$

We complete the proof by showing that

$$\frac{v}{b_0 b_1} \geq \frac{r^2 v}{b^2},$$

or equivalently, $b_0 b_1 \leq (b/r)^2$.

Since

$$b = b_0 + \ldots + b_{r-1} \geq b_0 + (r-1)b_1$$

and $r \geq 2$, it follows that

$$\frac{b}{r} \geq \frac{b_0 + (r-1)b_1}{r} \geq \frac{b_0 + b_1}{2}.$$

Finally, we have that

$$\left(\frac{b}{r}\right)^2 \geq \left(\frac{b_0 + b_1}{2}\right)^2 \geq b_0 b_1,$$

applying the arithmetic mean – geometric mean inequality. □

The proof of Lemma 5.1 also tells us when equality can occur in the above bound.

Corollary 5.2 *Suppose that $(V, \mathcal{W})$ is a resolvable incidence structure with $|V| = v$, $\mathcal{W} = b$, and r parallel classes. Suppose further that $|W \cap W'| = r^2 v/b^2$ for all blocks $W \neq W'$. Then $|W| = rv/b$ for all $W \in \mathcal{W}$.*

Proof. In order to have $|W \cap W'| = r^2 v/b^2$ for all $W \neq W'$, it must be the case that $b_0 = \ldots = b_{r-1}$ and all blocks have the same size, say k. Then the equation $bk = vr$ yields the solution $k = rv/b$. □

A resolvable $S_r(1, k, v)$ is said to be an *affine* design if $|W_1 \cap W_2| = \mu$ for all $W_1 \in \mathcal{P}_i$ and $W_2 \in \mathcal{P}_j$ with $i \neq j$, where μ is a constant. From the above discussion, it must be the case that $\mu = r^2 v/b^2$. Since $b = rv/k$, we see that $\mu = k^2/v$. Writing $s = k\mu$, the parameters v and k of an affine resolvable $S_r(1, k, v)$ can be written in the form $v = s^2\mu$ and $k = s\mu$.

Remark. An affine resolvable $S_r(1, k, v)$ is equivalent to an $(s, r; \mu)$-net ([1, p. 124]), which is the same thing as the dual of a $\mathrm{TD}_\mu[r; s]$ ([1, p. 38]).

Corollary 5.3 *The resolvable design commitment scheme obtained from an affine resolvable $S_r(1, k, v)$ has binding probability equal to $1 - k/v$.*

Proof. We showed above that $\mu = k^2/v$ in an affine resolvable $S_r(1, k, v)$. □

An affine resolvable design minimizes the value of μ, which maximizes the binding probability in the corresponding commitment scheme. It is also of interest to maximize the value of r, which allows the set of source states to be as large as possible. It is well-known that, if an affine resolvable $S_r(1, s\mu, s^2\mu)$ exists, then $r \leq (s^2\mu - 1)/(s - 1)$. Further, in the case of equality, the design is a 2-design (see [1, p. 127] for proofs of these results). Thus the affine resolvable 2-designs provide commitment schemes which are optimal in terms of binding probabilities and the number of source states.

An affine resolvable 2-design is often denoted as $A_\mu(s)$. The only parameters for which $A_\mu(s)$ are known to exist are the following:

- an $A_{q^n}(q)$ exists for all prime powers q and all positive integers n;
- an $A_n(2)$ exists for all positive integers n such that a Hadamard matrix of order $4n$ exists.

Remark. An $A_1(s)$ is in fact an affine plane of order p.

Acknowledgements

D.R. Stinson's research is supported by NSERC grants IRC #216431-96 and RGPIN #203114-98. D.R. Stinson and R. Wei are supported by the MITACS project "Applied Cryptography".

References

[1] T. Beth, D. Jungnickel and H. Lenz. *Design Theory, Volume 1*, Second Edition, Cambridge University Press, 1999.

[2] M. Blum. Coin flipping by telephone: a protocol for solving impossible problems. In *24th IEEE Spring Computer Conference*, pp. 133–137, IEEE Press, 1982.

[3] C.J. Colbourn and J.H. Dinitz. *CRC Handbook of Combinatorial Designs*, CRC Press, 1996.

[4] I. Damgard, J. Kilian and L. Salvail. On the (im)possibility of basing oblivious transfer and bit commitment on weakened security assumptions. *Lecture Notes in Computer Science* **1592** (1999), 56–73 (CRYPTO '99 Proceedings).

[5] R.L. Rivest. Unconditionally secure commitment and oblivious transfer schemes using concealing channels and a trusted initializer, preprint.