

On the Composition of Two-Prover Commitments, and Applications to Multi-Round Relativistic Commitments

Serge Fehr and Max Fillinger*

Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
{`serge.fehr`, `max.fillinger`}@cwi.nl

Abstract. We consider the related notions of *two-prover* and of *relativistic* commitment schemes. In recent work, Lunghi *et al.* proposed a new relativistic commitment scheme with a *multi-round sustain phase* that keeps the binding property alive as long as the sustain phase is running. They prove security of their scheme against classical attacks; however, the proven bound on the error parameter is very weak: it blows up *double exponentially* in the number of rounds.

In this work, we give a new analysis of the multi-round scheme of Lunghi *et al.*, and we show a *linear* growth of the error parameter instead (also considering classical attacks only). Our analysis is based on a new *composition theorem* for two-prover commitment schemes. The proof of our composition theorem is based on a better understanding of the binding property of two-prover commitments that we provide in the form of new definitions and relations among them. As an additional consequence of these new insights, our analysis is actually with respect to a strictly *stronger* notion of security than considered by Lunghi *et al.*

1 Introduction

TWO-PROVER COMMITMENT SCHEMES. We consider the notion of *2-prover commitment schemes*, as originally introduced by Ben-Or, Goldwasser, Kilian and Wigderson in their seminal paper [2]. In a 2-prover commitment scheme, the prover (i.e., the entity that is responsible for preparing and opening the commitment) consists of two agents, P and Q , and it is assumed that these two agents cannot communicate with each other during the execution of the protocol. With this approach, the classical and quantum impossibility results [9] [11] for unconditionally secure commitment schemes can be circumvented.

* Supported by the *NWO Free Competition* grant 617.001.203.

©IACR 2016. This article is the final version submitted by the author(s) to the IACR and to Springer-Verlag on February 11th, 2016. The version published by Springer-Verlag is available at An extended version is available at <http://arxiv.org/abs/1507.00240>.

A simple 2-prover bit commitment scheme is the scheme proposed by Crépeau *et al.* [5], which works as follows. The verifier V chooses a uniformly random $a \in \{0, 1\}^n$ and sends it to P , who replies with $x := y + a \cdot b \in \{0, 1\}^n$, where b is the bit to commit to, and $y \in \{0, 1\}^n$ is a uniformly random string known (only) to P and Q . Furthermore, “+” is bit-wise XOR, and “ \cdot ” is scalar multiplication (of the scalar b with the vector a). To open the commitment (to b), Q sends y to V , and V checks if $x + y = a \cdot b$. This scheme is clearly hiding: the commitment $x = y + a \cdot b$ is uniformly random and independent of a no matter what b is. On the other hand, the binding property follows from the observation that in order to open the commitment to $b = 0$, Q needs to announce $y = x$, and in order to open to $b = 1$, he needs to announce $y = x + a$. Thus, in order to open to *both*, he must know x and $x + a$, and thus a , which is a contradiction to the no-communication assumption, because a was sent to P only.

RELATIVISTIC COMMITMENT SCHEMES. The idea of *relativistic commitment schemes*, as introduced by Kent [7], is to take a 2-prover commitment scheme as above and enforce the no-communication assumption by means of relativistic effects: place P and Q spatially far apart, and execute the scheme fast enough, so that there is not enough time for them to communicate. The obvious downside of such a relativistic commitment scheme is that the binding property stays alive only for a very short time: the opening has to take place almost immediately after the committing, before the provers have the chance to exchange information. This limitation can be circumvented by considering *multi-round* schemes, where after the actual commit phase there is a *sustain phase*, during which the provers and the verifier keep exchanging messages, and as long as this sustain phase is running, the commitment stays binding (and hiding), until the commitment is finally opened. Such schemes were proposed in [7] and [8], but they are rather inefficient, and the security analyses are informal (e.g., with no formal security definitions) and of asymptotic nature.

More recently, Lunghi *et al.* [10] proposed a new and simple multi-round relativistic commitment scheme, and provided a rigorous security analysis. Their scheme works as follows (see also Figure 1). The actual commit protocol is the commit protocol from the Crépeau *et al.* scheme: V sends a uniformly random string $a_0 \in \{0, 1\}^n$ to P , who returns $x_0 := y_0 + a_0 \cdot b$. Then, to sustain the commitment, before P has the chance to tell a_0 to Q , V sends a new uniformly random string $a_1 \in \{0, 1\}^n$ to Q who replies with $x_1 := y_1 + a_1 \cdot y_0$, where $y_1 \in \{0, 1\}^n$ is another random string shared between P and Q , and the multiplication $a_1 \cdot y_0$ is in a suitable finite field. Then, to further sustain the commitment, V sends a new uniformly random string $a_2 \in \{0, 1\}^n$ to P who replies with $x_2 := y_2 + a_2 \cdot y_1$, etc. Finally, after the last sustain round where $x_m := y_m + a_m \cdot y_{m-1}$ has been sent to V , in order to finally open the commitment, y_m is sent to V by the other prover. In order to verify the opening, V computes $y_{m-1}, y_{m-2}, \dots, y_0$ inductively in the obvious way, and checks if $x_0 + y_0 = a_0 \cdot b$.

What is crucial is that in round i (say for odd i), when preparing x_i , the prover Q must not know a_{i-1} , but he is allowed to know a_1, \dots, a_{i-2} . Thus, execution must be timed in such a way that between subsequent rounds there

is not enough time for the provers to communicate, but they may communicate over multiple rounds.




Fig. 1. The Lunghi *et al.* multi-round scheme (for $m = 3$).

As for the security of this scheme, it is obvious that the hiding property stays satisfied up to the open phase: every single message V receives is one-time-pad encrypted. As for the binding property, Lunghi *et al.* prove that the scheme with a m -round sustain phase is ε_m -binding against classical attacks, where ε_m satisfies $\varepsilon_0 = 2^{-n}$ (this is just the standard Crépeau *et al.* scheme) and $\varepsilon_m \leq 2^{-n-1} + \sqrt{\varepsilon_{m-1}}$ for $m \geq 1$. Thus, even when reading this recursive formula liberally by ignoring the 2^{-n-1} term, we obtain

$$\varepsilon_m \lesssim 2^m \sqrt{\varepsilon_0} = 2^{-\frac{n}{2^m}},$$

i.e., the error parameter blows up *double exponentially* in m .¹ In other words, in order to have a non-trivial ε_m we need that n , the size of the strings that are communicated, is *exponential* in m . This means that Lunghi *et al.* can only afford a very small number of rounds. For instance, in their implementation where they can manage $n = 512$ (beyond that, the local computation takes too long), asking for an error parameter ε_m of approximately 2^{-32} , they can do $m = 4$ rounds.² This allows them to keep a commitment alive for 2 ms.

¹ Lunghi *et al.* also provide a more complicated recursive formula for ε_m that is slightly better, but the resulting blow-up is still double exponential.

² Note that [10] mentions $\varepsilon_m \approx 10^{-5} \approx 2^{-16}$, but this is an error, as communicated to us by the authors, and as can easily be verified. Also, [10] mentions $m = 5$ rounds, but this is because they include the commit round in their counting, and we do not.

OUR RESULTS. Our main goal is to improve the bound on the binding parameter of the above multi-round scheme. Indeed, our results show that the binding parameter blows up only *linearly* in m , rather than double exponentially. Explicitly, our results show that (for classical attacks)

$$\varepsilon_m \leq (m + 1) \cdot 2^{-\frac{n}{2} + 2}.$$

Using the same n and error parameter as in the implementation of Lunghi *et al.*, we can now afford approximately $m = 2^{222}$ rounds. Scaling up the 2 ms from the Lunghi *et al.* experiment for 4 rounds gives us a time that is in the order of 10^{56} years. On top of having a hugely improved error parameter, our analysis is with respect to a *strictly stronger* definition of the binding property.

We use the following strategy to obtain our improved bound on ε_m . We observe that the first sustain round can be understood as committing on the opening information y_0 of the actual commitment, using an extended version of the Crépeau *et al.* scheme that commits to a *string* rather than to a bit. Similarly, the second sustain round can be understood as committing on the opening information y_1 of that commitment from the first sustain round, etc. Thus, thinking of the $m = 1$ version of the scheme, what we have to prove is that if we have two commitment schemes \mathcal{S} and \mathcal{S}' , and we modify the opening phase of \mathcal{S} in that we first commit to the opening information (using \mathcal{S}') and then open that commitment, then the resulting commitment scheme is still binding; note that, intuitively, this is what one would indeed expect. Given such a composition theorem, we can then apply it inductively and conclude security (i.e. the binding property) of the Lunghi *et al.* multi-round scheme.

Our main result is such a general composition theorem, which shows that if \mathcal{S} and \mathcal{S}' are respectively ε - and δ -binding (against classical attacks) then the composed scheme is $(\varepsilon + \delta)$ -binding (against classical attacks), under some mild assumptions on \mathcal{S} and \mathcal{S}' . Hence, the error parameters simply add up; this is what gives us the linear growth. The proof of our composition theorem crucially relies on a new definition of the binding property of 2-prover commitment schemes, which seems to be handier to work with than the $p_0 + p_1 \leq 1 + \varepsilon$ definition as for instance used by Lunghi *et al.* Our definition formalizes the intuitive requirement that after the commit phase, no matter how the provers behaved, there should exist a bit \hat{b} (or a *string* in case of a string commitment scheme) such that opening the commitment to $b \neq \hat{b}$ fails (with high probability). This new definition is *strictly stronger* than the $p_0 + p_1$ definition, and thus we improve the Lunghi *et al.* result also in that direction.

One subtle issue is that the extended version of the Crépeau *et al.* scheme to strings, as it is used in the sustain phase, is not a fully secure string commitment scheme. The reason is that for *any* y that may be announced in the opening phase, there exists a string s such that $x + y = a \cdot s$; as such, the provers can commit to some fixed string, and then can still decide to either open the commitment to that string (by running the opening phase honestly), or to open it to a random string that is out of their control (by announcing a random y). We deal with this by also introducing a *relaxed* version of the binding property (which

we call *fairly-binding*), which captures this limited freedom for the provers, and we show that it is satisfied by the (extended version of the) Crépeau *et al.* scheme and that our composition theorem holds for this relaxed version; finally, we observe that the composed fairly-binding string commitment scheme is a binding *bit* commitment scheme when restricting the domain to a bit.

As such, we feel that our techniques and insights not only give rise to an improved analysis of the Lunghi *et al.* multi-round scheme, but they significantly improve our understanding of the security of 2-prover commitment schemes, and as such are likely to find further applications.

OPEN PROBLEMS. Our work gives rise to a list of interesting and challenging open problems. For instance, our composition theorem only applies to pairs $\mathcal{S}, \mathcal{S}'$ of commitment schemes of a certain restricted form, e.g., only one prover should be involved in the commit phase (as it is the case in the Crépeau *et al.* scheme). Our proof crucially relies on this, but there seems to be no fundamental reason for such a restriction. Thus, we wonder if it is possible to generalize our composition theorem to a larger class of pairs of schemes, or, ultimately, to *all* pairs of schemes (that “fit together”).

Also, generalizing our composition theorem to the quantum setting is an interesting open problem. This seems particularly non-trivial because our definition for the binding property does not generalize (immediately) to the quantum setting. Furthermore, in order to obtain security of the Lunghi *et al.* multi-round scheme against quantum attacks, beyond a quantum version of the composition theorem, one also needs to prove security (of the string-commitment version) of the Crépeau *et al.* scheme with respect to a suitable definition of the binding property against quantum attacks.

CONCURRENT WORK. In independent and concurrent work, Chakraborty, Chailloux and Leverrier [3] showed (almost) the same linear bound for the Lunghi *et al.* scheme, but with respect to the original — and thus weaker — notion of security. Their approach is more direct and tailored to the specific scheme; our approach is more abstract and provides more insight, and our result applies much more generally.

2 Preliminaries

2.1 Basic Notation

PROBABILITY DISTRIBUTIONS. For the purpose of this work, a (*probability distribution*) is a function $p : \mathcal{X} \rightarrow [0, 1]$, $x \mapsto p(x)$, where \mathcal{X} is a finite non-empty set, with the property that $\sum_{x \in \mathcal{X}} p(x) = 1$. For specific choices $x_\circ \in \mathcal{X}$, we tend to write $p(x = x_\circ)$ instead of $p(x_\circ)$. For any subset $A \subset \mathcal{X}$, called an *event*, the probability $p(A)$ is naturally defined as $p(A) = \sum_{x \in A} p(x)$, and it holds that

$$p(A) + p(\Gamma) = p(A \cup \Gamma) + p(A \cap \Gamma) \leq 1 + p(A \cap \Gamma) \quad (1)$$

for all $A, \Gamma \subset \mathcal{X}$. For a distribution $p : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ on two (or more) variables, probabilities like $p(x=y)$, $p(x=f(y))$, $p(x \neq y)$ etc. are naturally understood as

$$p(x=y) = p(\{(x,y) \in \mathcal{X} \times \mathcal{Y} \mid x=y\}) = \sum_{\substack{x \in \mathcal{X}, y \in \mathcal{Y} \\ \text{s.t. } x=y}} p(x,y)$$

etc., and the *marginals* $p(x)$ and $p(y)$ are given by $p(x) = \sum_y p(x,y)$ and by $p(y) = \sum_x p(x,y)$, respectively. Finally, given that $p(y) > 0$, we write $p(x|y)$ for the *conditional distribution* $p(x|y) := p(x,y)/p(y)$.

PROTOCOLS. In this work, we will consider 3-party (interactive) *protocols*, where the parties are named P , Q and V (the two “provers” and the “verifier”). Such a protocol prot_{PQV} consists of a triple $(\text{prot}_P, \text{prot}_Q, \text{prot}_V)$ of L -round *interactive algorithms* for some $L \in \mathbb{N}$. Each interactive algorithm takes an input, and for every round $\ell \leq L$ computes the messages to be sent to the other algorithms/parties in that round as deterministic functions of its input, the messages received in the previous rounds, and the local randomness. In the same way, the algorithms produce their respective outputs after the last round. We write

$$(\text{out}_P \parallel \text{out}_Q \parallel \text{out}_V) \leftarrow (\text{prot}_P(\text{in}_P) \parallel \text{prot}_Q(\text{in}_Q) \parallel \text{prot}_V(\text{in}_V))$$

to denote the execution of the protocol prot_{PQV} on the respective inputs in_P, in_Q and in_V , and that the respective outputs $\text{out}_P, \text{out}_Q$ and out_V are produced. Clearly, for any protocol prot_{PQV} and any input $\text{in}_P, \text{in}_Q, \text{in}_V$, the probability distribution $p(\text{out}_P, \text{out}_Q, \text{out}_V)$ of the output is naturally well defined.

If we want to make the local randomness explicit, we write $\text{prot}_P[\xi_P](\text{in}_P)$ etc., and understand that ξ_P is correctly sampled. We write $\text{prot}_P[\xi_{PQ}](\text{in}_P)$ and $\text{prot}_Q[\xi_{PQ}](\text{in}_Q)$ to express that prot_P and prot_Q use *the same* randomness, in which case we speak of *joint randomness*.

We can *compose* two interactive algorithms prot_P and prot'_P in the obvious way, by applying prot'_P to the output of prot_P . The resulting interactive algorithm is denoted as $\text{prot}'_P \circ \text{prot}_P$. Composing the respective algorithms of two protocols $\text{prot}_{PQV} = (\text{prot}_P, \text{prot}_Q, \text{prot}_V)$ and $\text{prot}'_{PQV} = (\text{prot}'_P, \text{prot}'_Q, \text{prot}'_V)$ results in the composed protocol $\text{prot}'_{PQV} \circ \text{prot}_{PQV}$.

2.2 2-Prover Commitment Schemes

We formally introduce the notion of 2-prover commitment schemes and discuss the security properties. Defining the binding property is non-trivial; this will be further discussed in Section 3.

Definition 2.1. A 2-prover (string) commitment scheme \mathcal{S} consists of two *interactive protocols*, the *commit protocol* $\text{com}_{PQV} = (\text{com}_P, \text{com}_Q, \text{com}_V)$ and the *opening protocol* $\text{open}_{PQV} = (\text{open}_P, \text{open}_Q, \text{open}_V)$ between the two provers P and Q and the verifier V , with the following *syntactics*. The *commit protocol* com_{PQV} uses *joint randomness* ξ_{PQ} for P and Q and takes a string $s \in \{0, 1\}^n$

as input for P and Q (and independent randomness and no input for V), and it outputs a commitment c to V and some state information to P and Q :

$$(state_P \| state_Q \| c) \leftarrow (\text{com}_P[\xi_{PQ}](s) \| \text{com}_Q[\xi_{PQ}](s) \| \text{com}_V).$$

The opening protocol open_{PQV} uses joint randomness η_{PQ} for P and Q , and outputs a string or a rejection symbol to V , and nothing to P and Q :

$$(\emptyset \| \emptyset \| s) \leftarrow (\text{open}_P[\eta_{PQ}](state_P) \| \text{open}_Q[\eta_{PQ}](state_Q) \| \text{open}_V(c))$$

with $s \in \{0, 1\}^n \cup \{\perp\}$. The set $\{0, 1\}^n$ is called the domain of \mathcal{S} ; if $n = 1$ then we refer to \mathcal{S} as a bit commitment scheme instead, and we tend to use b rather than s to denote the committed bit.

Remark 2.2. By convention, we assume throughout the paper that the commitment c output by V equals the *communication* that takes place between V and the provers during the commit phase. This is without loss of generality since, in general, c is computed as a (possibly randomized) function of the communication, which V just as well can apply in the opening phase.

Remark 2.3. Note that we specify that P and Q use *fresh* joint randomness η_{PQ} in the opening phase, and, if necessary, the randomness ξ_{PQ} from the commit phase can be “handed over” to the opening phase via $state_P$ and $state_Q$; this will be convenient later on. Alternatively, one could declare that P and Q *re-use* the joint randomness from the commit phase.

Whenever we refer to such a 2-prover commitment scheme, we take it as understood that the scheme is complete and hiding, as defined below, for “small” values of η and δ . Since our focus will be on the binding property, we typically do not make the parameters η and δ explicit.

Definition 2.4. A 2-prover commitment scheme is η -complete if in an honest execution V 's output s of open_{PQV} equals P and Q 's input s to com_{PQV} except with probability η , for any choice of P and Q 's input $s \in \{0, 1\}^n$.

The standard definition for the hiding property is as follows:

Definition 2.5. A 2-prover commitment scheme is δ -hiding if for any commit strategy $\overline{\text{com}}_V$ and any two strings s_0 and s_1 , the respective distributions of the commitments c_0 and c_1 , produced as

$$(state_P \| state_Q \| c_b) \leftarrow (\text{com}_P[\xi_{PQ}](s_b) \| \text{com}_Q[\xi_{PQ}c](s_b) \| \overline{\text{com}}_V)$$

for $b \in \{0, 1\}$, have statistical distance at most δ . A 0-hiding scheme is also called perfectly hiding.

Defining the binding property is more subtle. First, note that an *attack* against the binding property consists of an “allowed” commit strategy $\overline{\text{com}}_{PQ} = (\overline{\text{com}}_P, \overline{\text{com}}_Q)$ and an “allowed” opening strategy $\overline{\text{open}}_{PQ} = (\overline{\text{open}}_P, \overline{\text{open}}_Q)$ for

P and Q . Any such attack fixes $p(s)$, the distribution of $s \in \{0, 1\}^n \cup \{\perp\}$ that is output by V after the opening phase, in the obvious way.

What exactly “allowed” means may depend on the scheme and needs to be specified. Typically, in the 2-prover setting, we only allow strategies $\overline{\text{com}}_{PQ}$ and $\overline{\text{open}}_{PQ}$ with *no communication* at all between the two provers during the course of the scheme, but we may also be more liberal and allow some *well-controlled* communication, as in the Lunghi *et al.* multi-round scheme. Furthermore, in this work, we focus on *classical* attacks, where $\overline{\text{com}}_P, \overline{\text{com}}_Q, \overline{\text{open}}_P$ and $\overline{\text{open}}_Q$ are classical interactive algorithms as specified in the previous section, with access to joint randomness, but one could also consider *quantum* attacks, where the provers can perform measurements on an entangled quantum state.

A somewhat accepted definition for the binding property of a 2-prover *bit* commitment scheme, as it is for instance used in [5], [6] or [10] (up to the factor 2 in the error parameter), is as follows. Here, we assume it has been specified which attacks are *allowed*, e.g., those where P and Q do not communicate during the course of the scheme.

Definition 2.6. *A 2-prover bit commitment scheme is ε -binding in the sense of $p_0 + p_1 \leq 1 + 2\varepsilon$ if for every allowed commit strategy $\overline{\text{com}}_{PQ}$, and for every pair of allowed opening strategies $\overline{\text{open}}_{PQ}^0$ and $\overline{\text{open}}_{PQ}^1$, which fix distributions $p(b_0)$ and $p(b_1)$ for V 's respective outputs, it holds that*

$$p(b_0=0) + p(b_1=1) \leq 1 + 2\varepsilon.$$

In the literature (see e.g. [5] or [10]), the two probabilities $p(b_0=0)$ and $p(b_1=1)$ above are usually referred to as p_0 and p_1 , respectively.

2.3 The \mathcal{CHSH}^n Scheme

Our main example is the bit commitment scheme by Crépeau *et al.* [5] we mentioned in the introduction, and which works as follows. The commit phase com_{PQV} instructs V to sample and send to P a uniformly random $a \in \{0, 1\}^n$, and it instructs P to return $x := r + a \cdot b$ to V , where r is the joint randomness, uniformly distributed in $\{0, 1\}^n$, b is the bit to commit to, and the opening phase open_{PQV} instructs Q to send $y := r$ to V , and V outputs the (smaller) bit b that satisfies $x + y = a \cdot b$, or $b := \perp$ in case no such bit exists.

It is easy to see that this scheme is 2^{-n} -complete and perfectly hiding (completeness fails in case $a = 0$). For *classical* provers that do not communicate during the course of the scheme, the scheme is 2^{-n-1} -binding in the sense of $p_0 + p_1 \leq 1 + 2^{-n}$, i.e. according to Definition 2.6. As for *quantum* provers, Crépeau *et al.* showed that the scheme is $2^{-n/2}$ -binding; this was recently minorly improved to $2^{-(n+1)/2}$ by Sikora, Chailloux and Kerenidis [12].

We also want to consider an extended version of the scheme, where the bit b is replaced by a string $s \in \{0, 1\}^n$ in the obvious way (where the multiplication $a \cdot s$ is then understood in a suitable finite field), and we want to appreciate this version as a 2-prover *string* commitment scheme. However, it is a priori not clear

what is a suitable definition for the binding property, especially because for this particular scheme, the dishonest provers can always honestly commit to a string s , and can then decide to correctly open the commitment to s by announcing $y := r$, or open to a *random* string by announcing a randomly chosen y — any y satisfies $x + y = a \cdot s$ for *some* s (unless $a = 0$, which almost never happens).³

Due to its close relation to the CHSH game [4], in particular to the arbitrary-finite-field version considered in [1], we will refer to this *string* commitment scheme as \mathcal{CHSH}^n .

3 On the Binding Property of 2-Prover Commitments

We introduce a new definition for the binding property of 2-prover commitment schemes. In the case of *bit* commitment schemes, it implies Definition 2.6, as we will show. Our new definition is not only stronger, but we also feel that it is closer to the intuition of what is expected from a commitment scheme, and as such it is easier to work with. Indeed, the proof of our composition result is heavily based on our new definition. Also, our new notion is more flexible in terms of tweaking it; for instance, we modify it to obtain a *relaxed* notion for the binding property, which captures the binding property that is satisfied by the string commitment scheme \mathcal{CHSH}^n .

Throughout this section, when quantifying over attacks against (the binding property of) a scheme, it is always understood that there is a notion of *allowed* attacks for that scheme (e.g., all attacks for which P and Q do not communicate), and that the quantification is over all such allowed attacks.

3.1 Defining The Binding Property

Intuitively, we say that a scheme is binding if after the commit phase there exists a string \hat{s} so that no matter what the provers do in the opening phase, the verifier will output either $s = \hat{s}$ or $s = \perp$ (except with small probability). Formally, we require that for every possible commit strategy, such a string \hat{s} is uniquely determined by the commitment c and the provers' joint randomness.

Definition 3.1 (Binding property). *A 2-prover commitment scheme \mathcal{S} is ϵ -binding if for every commit strategy $\overline{\text{com}}_{PQ}[\bar{\xi}_{PQ}]$ there exists a function $\hat{s}(\bar{\xi}_{PQ}, c)$ of the joint randomness $\bar{\xi}_{PQ}$ and the commitment c such that for every opening strategy $\overline{\text{open}}_{PQ}$ it holds that $p(s \neq \hat{s}(\bar{\xi}_{PQ}, c) \wedge s \neq \perp) \leq \epsilon$. In short:*

$$\forall \overline{\text{com}}_{PQ} \exists \hat{s}(\bar{\xi}_{PQ}, c) \forall \overline{\text{open}}_{PQ} : p(s \neq \hat{s} \wedge s \neq \perp) \leq \epsilon. \quad (2)$$

The string commitment scheme \mathcal{CHSH}^n does *not* satisfy this definition (the bit commitment version does, as we will show): after the commit phase, the provers

³ This could easily be prevented by asking Q to also announce s (rather than letting V compute it), but we want the information announced during the opening phase to fit into the domain of the commitment scheme.

can still decide to open the commitment to a *fixed* string, chosen before the commit phase, or to a *random* string that is out of their control. We capture this by the following relaxed version of the binding property. In this relaxed version, we allow V 's output s to be different from \hat{s} and \perp , but in this case the provers should have little control over s : for any *target string* s_o (computed as a function of the provers' randomness), it should be unlikely that $s = s_o$. Formally, this is captured as follows; we will show in Section 3.3 that \mathcal{CHSH}^n is fairly-binding in this sense.

Definition 3.2 (Fairly binding property). *A 2-prover commitment scheme \mathcal{S} is ε -fairly-binding if for every commit strategy $\overline{\text{com}}_{PQ}[\bar{\xi}_{PQ}]$ there exists a function $\hat{s}(\bar{\xi}_{PQ}, c)$ such that for every opening strategy $\overline{\text{open}}_{PQ}[\bar{\eta}_{PQ}]$ and all functions $s_o(\bar{\xi}_{PQ}, \bar{\eta}_{PQ})$ it holds that $p(s \neq \hat{s}(\bar{\xi}_{PQ}, c) \wedge s = s_o(\bar{\xi}_{PQ}, \bar{\eta}_{PQ})) \leq \varepsilon$. In short:*

$$\forall \overline{\text{com}}_{PQ} \exists \hat{s}(\bar{\xi}_{PQ}, c) \forall \overline{\text{open}}_{PQ} \forall s_o(\bar{\xi}_{PQ}, \bar{\eta}_{PQ}) : p(s \neq \hat{s} \wedge s = s_o) \leq \varepsilon. \quad (3)$$

Remark 3.3. By means of standard techniques, one can easily show that it is sufficient for the (fairly) binding property to consider *deterministic* provers. In this case, \hat{s} is a function of c only, and, in the case of fairly-binding, s_o runs over all *fixed* strings.

Remark 3.4. Clearly, the ordinary binding property (i.e., as in Definition 3.1) implies the fairly-binding property. Also, in the case of *bit* commitment schemes it obviously holds that $p(b \neq \hat{b} \wedge b \neq \perp) = p(b \neq \hat{b} \wedge b = 0) + p(b \neq \hat{b} \wedge b = 1)$, and thus the fairly-binding property implies the ordinary one, up to a factor-2 loss. Furthermore, every fairly-binding *string* commitment scheme gives rise to an ordinary-binding *bit* commitment scheme in a natural way, as shown by the following proposition.

Proposition 3.5. *Let \mathcal{S} be an ε -fairly-binding string commitment scheme. Fix any two distinct strings $s_0, s_1 \in \{0, 1\}^n$ and consider the bit-commitment scheme \mathcal{S}' obtained as follows. To commit to $b \in \{0, 1\}$, the provers commit to s_b using \mathcal{S} , and in the opening phase V checks if $s = s_b$ for some $b \in \{0, 1\}$ and outputs this bit if it exists and else outputs $b = \perp$. Then, \mathcal{S}' is 2ε -binding.*

Proof. Fix some commit strategy $\overline{\text{com}}_{PQ}$ for \mathcal{S}' and note that it can also be used to attack \mathcal{S} . Thus, there exists a function $\hat{s}(\bar{\xi}_{PQ}, c)$ as in Definition 3.2. We define

$$\hat{b}(\bar{\xi}_{PQ}, c) = \begin{cases} 0 & \text{if } \hat{s}(\bar{\xi}_{PQ}, c) = s_0 \\ 1 & \text{otherwise} \end{cases}$$

Now fix an opening strategy $\overline{\text{open}}_{PQ}$ for \mathcal{S}' , which again is also a strategy against \mathcal{S} . Thus, we have $p(\hat{s} \neq s = s_o) \leq \varepsilon$ for any s_o (and in particular $s_o = s_0$ or s_1). This gives us

$$\begin{aligned} p(\hat{b} \neq b \neq \perp) &= p(\hat{b} = 1 \wedge b = 0) + p(\hat{b} = 0 \wedge b = 1) \\ &= p(\hat{s} \neq s_0 \wedge s = s_0) + p(\hat{s} = s_0 \wedge s = s_1) \end{aligned}$$

$$\begin{aligned} &\leq p(\hat{s} \neq s_0 \wedge s = s_0) + p(\hat{s} \neq s_1 \wedge s = s_1), \\ &\leq 2\varepsilon \end{aligned}$$

and thus \mathcal{S}' is a 2ε -binding bit-commitment scheme. \square

Remark 3.6. The proof of Proposition 3.5 generalizes in a straightforward way to k -bit string commitment schemes: given an ε -fairly-binding n -bit string commitment scheme \mathcal{S} , for $k < n$, we define a k -bit string commitment scheme \mathcal{S}_k as follows: to commit to a k -bit string, the provers pad the string with $n - k$ zeros and then commit to the padded string using \mathcal{S} . In the opening phase, the verifier outputs the first k bits of s if the remaining bits in s are all zeros, and \perp otherwise. Then, \mathcal{S}' is $2^k\varepsilon$ -binding.

3.2 Relation To The Standard Definition

For bit commitment schemes, our binding property implies the $(p_0 + p_1)$ -definition.

Theorem 3.7. *A 2-prover bit-commitment scheme that is ε -binding (in the sense of Definition 3.1) is ε -binding in the sense of $p_0 + p_1 \leq 1 + 2\varepsilon$.*

Proof. Consider a scheme that is ε -binding. Fix $\overline{\text{com}}_{PQ}$ and let $\hat{b}(\bar{\xi}_{PQ}, c)$ be a function as promised by Definition 3.1, i.e., such that for every opening strategy $\overline{\text{open}}_{PQ}$ we have $p(b \neq \hat{b} \wedge b \neq \perp) \leq \varepsilon$. Now, fix two opening strategies $\overline{\text{open}}_{PQ}^0$ and $\overline{\text{open}}_{PQ}^1$, and consider the two respective output bits b_0 and b_1 . It holds that $p(\hat{b} \neq b_i \neq \perp) \leq \varepsilon$ for $i \in \{0, 1\}$, and thus

$$\begin{aligned} p(b_0 = 0) + p(b_1 = 1) &= p(b_0 = 0 \wedge \hat{b} = 0) + p(b_0 = 0 \wedge \hat{b} = 1) \\ &\quad + p(b_1 = 1 \wedge \hat{b} = 0) + p(b_1 = 1 \wedge \hat{b} = 1) \\ &\leq p(\hat{b} = 0) + p(\hat{b} \neq b_0 \neq \perp) + p(\hat{b} \neq b_1 \neq \perp) + p(\hat{b} = 1) \\ &\leq 1 + 2\varepsilon \end{aligned}$$

which proves our claim. \square

On the other hand, our Definition 3.1 is *strictly* stronger than the $p_0 + p_1$ based Definition 2.6. Consider the following (artificial and very non-complete) scheme: in the commit phase, V chooses a uniformly random bit and sends it to the provers, and then accepts everything or rejects everything during the opening phase, depending on that bit. Then, $p_0 + p_1 = 1$, yet a commitment can be opened to $1 - \hat{b}$ (no matter how \hat{b} is defined) with probability $\frac{1}{2}$.

Since a non-complete separation example may not be fully satisfying, we note that it can be converted into a complete (but even more artificial) scheme. Fix a “good” (i.e., complete, hiding and binding with low parameters) scheme and call our example scheme above the “bad” scheme. We define a *combined* scheme as follows: at the start, the first prover can request either the “good” or “bad” scheme to be used. The honest prover is instructed to choose the former, guaranteeing completeness. The dishonest prover may choose the latter, so the combined scheme inherits the binding properties of the “bad” scheme: it is binding according to the $(p_0 + p_1)$ -definition, but not according to Definition 3.1.

3.3 Security of \mathcal{CHSH}^n

In this section, we show that \mathcal{CHSH}^n is a fairly-binding string commitment scheme.⁴ To this end, we introduce yet another version of the binding property and show that \mathcal{CHSH}^n satisfies this property. Then we show that this version of the binding property implies the fairly-binding property (up to some loss in the parameter, and under some mild restriction on the scheme).

This new binding property is based on the intuition that it should not be possible to open a commitment to two different values *simultaneously* (except with small probability). For this, we observe that, when considering a commit strategy $\overline{\text{com}}_{PQ}$, as well as *two* opening strategies $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$, we can run both opening strategies *simultaneously* on the produced commitment with two (independent) copies of open_V , by applying $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$ to two copies of the respective internal states of P and Q . This gives rise to a *joint* distribution $p(s, s')$ of the respective outputs s and s' of the two copies of open_V .

Definition 3.8 (Simultaneous opening). *A 2-prover commitment scheme \mathcal{S} is ε -fairly-binding in the sense of simultaneous opening⁵ if for all $\overline{\text{com}}_{PQ}$, all pairs of opening strategies $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$, and all pairs s_o, s'_o of distinct strings, we have $p(s = s_o \wedge s' = s'_o) \leq \varepsilon$.*

Remark 3.9. Also for this notion of fairly-binding, it is sufficient to consider *deterministic* strategies, as can easily be seen.

Proposition 3.10. *The commitment scheme \mathcal{CHSH}^n is 2^{-n} -fairly-binding in the sense of simultaneous opening.*

Proof. By Remark 3.9, it suffices to consider deterministic attack strategies. Fix a deterministic strategy $\overline{\text{com}}_{PQ}$ and two deterministic opening strategies $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$. The strategy $\overline{\text{com}}_{PQ}$ specifies P 's output x as a function $f(a)$ of the verifier's message a . The opening strategies are described by constants y and y' . By definition of \mathcal{CHSH}^n , $s = s_o$ implies $f(a) + y = a \cdot s_o$ and likewise, $s' = s'_o$ implies $f(a) + y' = a \cdot s'_o$. Therefore, $s = s_o \wedge s' = s'_o$ implies $a = (y - y') / (s_o - s'_o)$. It thus holds that $p(s = s_o \wedge s' = s'_o) \leq p(a = (y - y') / (s_o - s'_o)) \leq \frac{1}{2^n}$, which proves our claim. \square

Remark 3.11. It follows directly from (1) that every *bit* commitment scheme that is ε -fairly-binding in the sense of simultaneous opening is ε -binding in the sense of $p_0 + p_1 \leq 1 + 2\varepsilon$. The converse is not true though: the schemes described at the end of Section 3.2 again serve as counterexamples.

Theorem 3.12. *Let $\mathcal{S} = (\text{com}_{PQV}, \text{open}_{PQV})$ be a 2-prover commitment scheme. If \mathcal{S} is ε -fairly-binding in the sense of simultaneous opening and open_V is deterministic, then \mathcal{S} is $2\sqrt{\varepsilon}$ -fairly-binding.*

⁴ It is understood that the allowed attacks against \mathcal{CHSH}^n are those where the provers do not communicate during the course of the scheme.

⁵ We use “fairly” here to distinguish the notion from a possible “non-fairly” version with $p(\perp \neq s \neq s' \neq \perp) \leq \varepsilon$; however, we do not consider this latter version any further here.

Proof. By Remark 3.3, it suffices to consider deterministic strategies for the provers. We fix some deterministic commit strategy $\overline{\text{com}}_{PQ}$ and an enumeration $\{\overline{\text{open}}_{PQ}^i\}_{i=1}^N$ of all deterministic opening strategies. Since we assume that open_V is deterministic, for any fixed opening strategy for the provers, the verifier's output s is a *function* of the commitment c . Thus, for each opening strategy $\overline{\text{open}}_{PQ}^i$ there is a function f_i such that the verifier's output is $s = f_i(c)$. We will now define the function $\hat{s}(c)$ that satisfies the properties required by the fairly-binding property. Our definition depends on a parameter $\alpha > 0$ which we fix later. In order to define \hat{s} , we partition the set C of all possible commitments into *disjoint* sets $C = R \cup \bigcup_{s,i} C_{s,i}$ that satisfy the following three properties for every i and every s :

$$C_{s,i} \subseteq f_i^{-1}(\{s\}), \quad p(c \in C_{s,i}) \geq \alpha \text{ or } C_{s,i} = \emptyset, \quad \text{and} \quad p(c \in R \wedge f_i(c) = s) < \alpha.$$

The second property implies that there are at most α^{-1} non-empty sets $C_{s,i}$. It is easy to see that such a partitioning exists: start with $R = C$ and while there exist s and i with $p(c \in R \wedge f_i(c) = s) \geq \alpha$, let $C_{s,i} = \{c \in R \mid f_i(c) = s\}$ and remove the elements of $C_{s,i}$ from R . For any $c \in C$, we now define $\hat{s}(c)$ as follows. We set $\hat{s}(c) = s$ for $c \in C_{s,i}$ and $\hat{s}(c) = 0$ for $c \in R$.

Now fix some opening strategy $\overline{\text{open}}_{PQ}^i$ and a string s_o , and write s_i for the verifier's output. Using $C_{\neq s_o}$ as a shorthand for $\bigcup_{s \neq s_o} \bigcup_j C_{s,j}$, we note that if $\hat{s}(c) \neq s_o$ then $c \in R \cup C_{\neq s_o}$. Thus, it follows that

$$\begin{aligned} p(s_i \neq \hat{s}(c) \wedge s_i = s_o) &= p(\hat{s}(c) \neq s_o \wedge s_i = s_o) \\ &\leq p(c \in (R \cup C_{\neq s_o}) \wedge f_i(c) = s_o) \\ &= p(c \in R \wedge f_i(c) = s_o) + \sum_{s \neq s_o, j} p(c \in C_{s,j} \wedge f_i(c) = s_o) \\ &\leq p(c \in R \wedge f_i(c) = s_o) + \sum_{\substack{s \neq s_o, j \\ \text{s.t. } C_{s,j} \neq \emptyset}} p(f_j(c) = s \wedge f_i(c) = s_o) \\ &< \alpha + \alpha^{-1} \cdot \varepsilon \end{aligned}$$

where the final inequality holds because $p(c \in R \wedge f_i(c) = s_o) < \alpha$ by the choice of R , because $p(f_j(c) = s \wedge f_i(c) = s_o) \leq \varepsilon$ by the assumed binding property, and because the number of non-empty $C_{s,j}$'s is at most $1/\alpha$. It is easy to see that the upper bound $\alpha + \alpha^{-1} \cdot \varepsilon$ is minimized by setting $\alpha = \sqrt{\varepsilon}$. We conclude that $p(s_i \neq \hat{s}(c) \wedge s_i = s_o) < 2\sqrt{\varepsilon}$. \square

By combining Theorem 3.7 with Theorem 3.12, we obtain the following statement for the (fairly-)binding property of \mathcal{CHSH}^n .

Corollary 3.13. *\mathcal{CHSH}^n is $2^{-\frac{n}{2}+1}$ -fairly-binding.*

4 Composing Commitment Schemes

4.1 The Composition Operation

We consider two 2-prover commitment schemes \mathcal{S} and \mathcal{S}' of a restricted form, and we compose them to a new 2-prover commitment scheme $\mathcal{S}'' = \mathcal{S} \star \mathcal{S}'$ in

a well-defined way; our composition theorem then shows that \mathcal{S}'' is secure if \mathcal{S} and \mathcal{S}' are. We start by specifying the restriction to \mathcal{S} and \mathcal{S}' that we impose.

Definition 4.1. *Let \mathcal{S} and \mathcal{S}' be two 2-prover string commitment schemes. We call the pair $(\mathcal{S}, \mathcal{S}')$ eligible if the following three properties hold, or they hold with the roles of P and Q exchanged.*

1. *The commit phase of \mathcal{S} is a protocol $\text{com}_{PV} = (\text{com}_P, \text{com}_V)$ between P and V only, and the opening phase of \mathcal{S} is a protocol $\text{open}_{QV} = (\text{open}_Q, \text{open}_V)$ between Q and V only. In other words, com_Q and open_P are both trivial and do nothing.⁶ Similarly, the commit phase of \mathcal{S}' is a protocol com'_{QV} between Q and V only (but both provers may be active in the opening phase).*
2. *The opening phase open_{QV} of \mathcal{S} is of the following simple form: Q sends a bit string $y \in \{0, 1\}^m$ to V , and V computes s deterministically as $s = \text{Extr}(y, c)$, where c is the commitment.⁷*
3. *The domain of \mathcal{S}' contains (or equals) $\{0, 1\}^m$.*

Furthermore, we specify that the allowed attacks on \mathcal{S} are so that P and Q do not communicate during the course of the entire scheme, and the allowed attacks on \mathcal{S}' are so that P and Q do not communicate during the course of the commit phase but there may be limited communication during the opening phase.

An example of an eligible pair of 2-prover commitments is $(\text{CHSH}^n, \mathcal{X}\text{CHSH}^n)$, where $\mathcal{X}\text{CHSH}^n$ coincides with scheme CHSH^n except that the roles of P and Q are exchanged.

Remark 4.2. For an eligible pair $(\mathcal{S}, \mathcal{S}')$, it will be convenient to understand open_Q and open_V as *non-interactive* algorithms, where open_Q produces y as its *output*, and open_V takes y as additional *input* (rather than viewing the pair as a protocol with a single one-way communication round).

We now define the composition operation. Informally, committing is done by means of committing using \mathcal{S} , and to open the commitment, Q uses open_Q to locally compute the opening information y and he commits to y with respect to the scheme \mathcal{S}' , and then this commitment is opened (to y), and V computes and outputs $s = \text{Extr}(y, c)$. Formally, this is captured as follows (see also Figure 2).

Definition 4.3. *Let $\mathcal{S} = (\text{com}_{PV}, \text{open}_{QV})$ and $\mathcal{S}' = (\text{com}'_{QV}, \text{open}'_{PQV})$ be an eligible pair of 2-prover commitment schemes. Then, their composition $\mathcal{S} \star \mathcal{S}'$ is defined as the scheme consisting of $\text{com}_{PV} = (\text{com}_P, \text{com}_V)$ and*

$$\text{open}''_{PQV} = (\text{open}'_P, \text{open}'_Q \circ \text{com}'_Q \circ \text{open}_Q, \text{open}_V \circ \text{open}'_V \circ \text{com}'_V).$$

⁶ Except that com_Q may output the shared randomness in order to hand it over to the opening protocol open_Q .

⁷ Our composition theorem also works for a randomized Extr , but for simplicity, we restrict to the deterministic case.

If in this composition the output in open'_V is $y = \perp$, we define the output of open_V to be $s = \perp$ as well.

When considering attacks against the binding property of the composed scheme $\mathcal{S} \star \mathcal{S}'$, we declare that the allowed deterministic attacks⁸ are those of the form $(\overline{\text{com}}_P, \overline{\text{open}}'_{PQ} \circ \text{ptoq}_{PQ} \circ \overline{\text{com}}'_Q)$, where $\overline{\text{com}}_P$ is an allowed deterministic commit strategy for \mathcal{S} , $\overline{\text{com}}'_Q$ and $\overline{\text{open}}'_{PQ}$ are allowed deterministic commit and opening strategies for \mathcal{S}' , and ptoq_{PQ} is the one-way communication protocol that communicates P 's input to Q (see also Figure 3).⁹

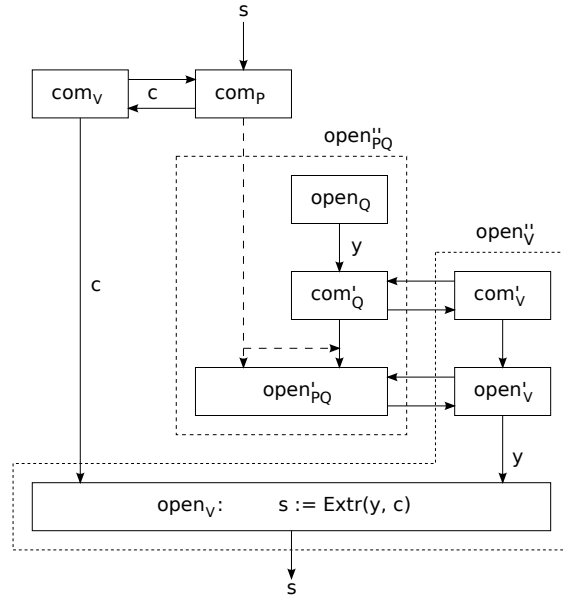


Fig. 2. The composition of $\mathcal{S} = (\text{com}_{PV}, \text{open}_{QV})$ and $\mathcal{S}' = (\text{com}'_{QV}, \text{open}'_{PQV})$. The dotted arrows indicate communication allowed to the dishonest provers.

Remark 4.4. It is immediate that $\mathcal{S} \star \mathcal{S}'$ is a commitment scheme in the sense of Definition 2.1, and that it is complete if \mathcal{S} and \mathcal{S}' are, with the error parameters adding up. Also, the hiding property is obviously inherited from \mathcal{S} ; however, the point of the composition is to keep the hiding property alive for longer, namely up to before the last round of the opening phase—recall that, using the terminology used in context of relativistic commitments, these rounds of the

⁸ The allowed *randomized* attacks are then naturally given as those that pick one of the deterministic attacks according to some distribution.

⁹ This one-way communication models that in the relativistic setting, sufficient time has passed at this point for P to inform Q about what happened during com_P .

opening phase up to before the last would then be referred to as the *sustain phase*. We show in Appendix A that $\mathcal{S} \star \mathcal{S}'$ is hiding up to before the last round, with the error parameters adding up.

It is intuitively clear that $\mathcal{S} \star \mathcal{S}'$ *should be* binding if \mathcal{S} and \mathcal{S}' are: committing to the opening information y and then opening the commitment allows the provers to *delay* the announcement of y (which is the whole point of the exercise), but it does not allow them to *change* y , by the binding property of \mathcal{S}' ; thus, $\mathcal{S} \star \mathcal{S}'$ should be (almost) as binding as \mathcal{S} . This intuition is confirmed by our composition theorem below.

Remark 4.5. We point out that the composition $\mathcal{S} \star \mathcal{S}'$ can be naturally defined for a *larger* class of pairs of schemes (e.g. where *both* provers are active in the commit phase of both schemes), and the above intuition still holds. However, our proof only works for this restricted class of (pairs of) schemes. Extending the composition result in that direction is an open problem.

Remark 4.6. We observe that if $(\mathcal{S}, \mathcal{X}\mathcal{S})$ is an eligible pair, where $\mathcal{X}\mathcal{S}$ coincides with \mathcal{S} except that the roles of P and Q are exchanged, then so is $(\mathcal{X}\mathcal{S}, \mathcal{S} \star \mathcal{X}\mathcal{S})$. As such, we can then compose $\mathcal{X}\mathcal{S}$ with $\mathcal{S} \star \mathcal{X}\mathcal{S}$, and obtain yet another eligible pair $(\mathcal{S}, \mathcal{X}\mathcal{S} \star \mathcal{S} \star \mathcal{X}\mathcal{S})$, etc. Applying this to the schemes $\mathcal{S} = \mathcal{CHSH}^n$, we obtain the multi-round scheme from Lunghi *et al.* [10]. As such, our composition theorem below implies security of their scheme — with a *linear* blow-up of the error term (instead of double exponential).

We point out that formally we obtain security of the Lunghi *et al.* scheme as a *2-prover commitment scheme* under an *abstract restriction* on the provers' communication: in every round, the active prover cannot access the message that the other prover received in the previous round. As such, when the rounds of the protocol are executed fast enough so that it is ensured that there is no time for the provers to communicate between subsequent rounds, then security as a *relativistic commitment scheme* follows immediately.

Before stating and proving the composition theorem, we need to single out one more relevant parameter.

Definition 4.7. *Let $(\mathcal{S}, \mathcal{S}')$ be an eligible pair, which in particular means that V 's action in the opening phase of \mathcal{S} is determined by a function Extr . We define $k(\mathcal{S}) := \max_{c,s} |\{y \mid \text{Extr}(y, c) = s\}|$.*

I.e., $k(\mathcal{S})$ counts the number of ys that are consistent with a given string s (in the worst case). Note that $k(\mathcal{CHSH}^n) = 1$: for every $a, x, s \in \{0, 1\}^n$ there is exactly one $y \in \{0, 1\}^n$ such that $x + y = a \cdot s$.

4.2 The Composition Theorem

In the following composition theorem, we take it as understood that the assumed respective binding properties of \mathcal{S} and \mathcal{S}' hold with respect to a well-defined respective classes of allowed attacks.

Theorem 4.8. *Let $(\mathcal{S}, \mathcal{S}')$ be an eligible pair of 2-prover commitment schemes, and assume that \mathcal{S} and \mathcal{S}' are respectively ε -fairly-binding and δ -fairly-binding. Then, their composition $\mathcal{S}'' = \mathcal{S} \star \mathcal{S}'$ is $(\varepsilon + k(\mathcal{S}) \cdot \delta)$ -fairly-binding.*

Proof. We first consider the case $k(\mathcal{S}) = 1$. We fix an attack $(\overline{\text{com}}_P, \overline{\text{open}}''_{PQ})$ against \mathcal{S}'' . Without loss of generality, the attack is deterministic, so $\overline{\text{open}}''_{PQ}$ is of the form $\overline{\text{open}}''_{PQ} = \overline{\text{open}}'_{PQ} \circ \text{ptoq}_{PQ} \circ \overline{\text{com}}'_Q$.

Note that $\overline{\text{com}}_P$ is also a commit strategy for \mathcal{S} . As such, by the fairly-binding property of \mathcal{S} , there exists a function $\hat{s}(c)$, only depending on $\overline{\text{com}}_P$, so that the property specified in Definition 3.2 is satisfied for every opening strategy $\overline{\text{open}}_Q$ for \mathcal{S} . We will show that it is also satisfied for the (arbitrary) opening strategy $\overline{\text{open}}''_{PQ}$ for \mathcal{S}'' , except for a small increase in ε : we will show that $p(\hat{s}(c) \neq s \wedge s = s_o) \leq \varepsilon + \delta$ for every fixed target string s_o . This then proves the claim.

In order to show this property on $\hat{s}(c)$, we “decompose and reassemble” the attack strategy $(\overline{\text{com}}_P, \overline{\text{open}}'_{PQ} \circ \text{ptoq}_{PQ} \circ \overline{\text{com}}'_Q)$ for \mathcal{S}'' into an attack strategy $(\overline{\text{com}}'_Q, \overline{\text{newopen}}'_{PQ})$ for \mathcal{S}' with $\overline{\text{newopen}}'_{PQ}$ formally defined as

$$\overline{\text{newopen}}'_{PQ}[c](\overline{\text{state}}'_Q) := \overline{\text{open}}'_{PQ}(\overline{\text{state}}_P(c) \| (\overline{\text{state}}_P(c), \overline{\text{state}}'_Q))$$

where

$$(\overline{\text{state}}_P(c) \| c) \leftarrow (\overline{\text{com}}_P \| \text{com}_V).$$

Informally, this means that ahead of time, P and Q *simulate* an execution of $(\overline{\text{com}}_P \| \text{com}_V)$ and take the resulting communication/commitment¹⁰ c as shared randomness, and then $\overline{\text{newopen}}'_{PQ}$ computes $\overline{\text{state}}_P$ from c as in $\overline{\text{com}}_P$, and runs $\overline{\text{open}}'_{PQ}$ (see Figure 3).¹¹ It follows from the fairly-binding property that there is a function $\hat{y}(c')$ of the commitment c' so that $p(\hat{y}(c') \neq y \wedge y = y_o(c)) \leq \delta$ for every function $y_o(c)$.

The existence of \hat{y} now gives rise to an opening strategy $\overline{\text{open}}_Q$ for \mathcal{S} ; namely, simulate the commit phase of \mathcal{S}' to obtain the commitment c' , and output $\hat{y}(c')$. By Definition 3.2, for $\tilde{s} := \text{Extr}(\hat{y}(c'), c)$ and every s_o , $p(\hat{s}(c) \neq \tilde{s} \wedge \tilde{s} = s_o) \leq \varepsilon$.

We are now ready to put things together. Fix an arbitrary target string s_o . For any c we let $y_o(c)$ be the unique string such that $\text{Extr}(y_o(c), c) = s_o$ (and some default string if no such string exists); recall, we assume for the moment that $k(\mathcal{S}) = 1$. Omitting the arguments in $\hat{s}(c)$, $\hat{y}(c')$ and $y_o(c)$, it follows that

$$\begin{aligned} p(\hat{s} \neq s \wedge s = s_o) &\leq p(\hat{s} \neq s \wedge s = s_o \wedge s = \tilde{s}) + p(s = s_o \wedge s \neq \tilde{s}) \\ &\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} = s_o) + p(\text{Extr}(y, c) \neq \text{Extr}(\hat{y}, c) \wedge \text{Extr}(y, c) = s_o) \\ &\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} = s_o) + p(y \neq \hat{y} \wedge y = y_o) \\ &\leq \varepsilon + \delta. \end{aligned}$$

¹⁰ Recall that by convention (Remark 2.2), the commitment c equals the communication between V and, here, P .

¹¹ We are using here that Q is inactive during $\overline{\text{com}}_P$ and P during $\overline{\text{com}}'_Q$, and thus the two “commute”.

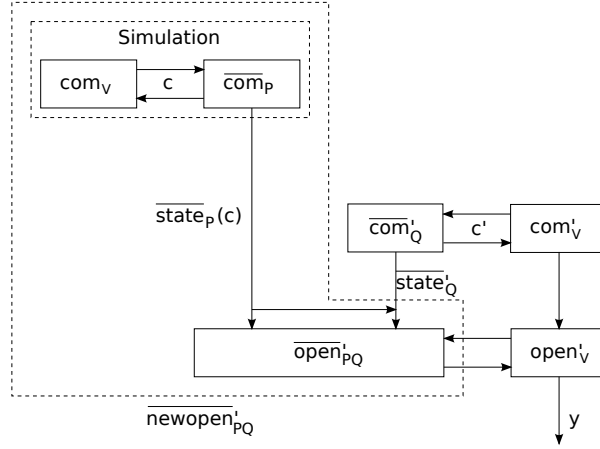


Fig. 3. Constructing the opening strategy $\overline{\text{newopen}}'_{PQ}$ against S' .

Thus, \hat{s} is as required.

For the general case where $k(S) > 1$, we can reason similarly, except that we then list the $k \leq k(S)$ possibilities $y_o^1(c), \dots, y_o^k(c)$ for $y_o(c)$, and conclude that $p(s \neq \tilde{s} \wedge s = s_o) \leq \sum_i p(y \neq \hat{y} \wedge y = y_o^i) \leq k(S) \cdot \delta$, which then results in the claimed bound. \square

Remark 4.9. Putting things together, we can now conclude the security (i.e., the binding property) of the Lunghi *et al.* multi-round commitment scheme. Corollary 3.13 ensures the fairly-binding property of \mathcal{CHSH}^n , i.e., the Crépeau *et al.* scheme as a string commitment scheme, with parameter $2^{-n/2+1}$. The composition theorem (Theorem 4.8) then guarantees the fairly-binding property of the m -fold composition as a string commitment scheme, with parameter $(m+1) \cdot 2^{-n/2+1}$. Finally, Proposition 3.5 implies that the m -fold composition of \mathcal{CHSH}^n with itself is a ε_m -binding bit commitment scheme with error parameter $\varepsilon_m = (m+1) \cdot 2^{-n/2+2}$ as claimed in the introduction, or, more generally, and by taking Remark 3.6 into account, a $(m+1) \cdot 2^{-n/2+k+1}$ -binding k -bit-string commitment scheme.

Finally, for completeness, we point out that the composition theorem also holds for regularly (i.e., “non-fairly”) binding schemes.

Theorem 4.10. *Let (S, S') be an eligible pair of 2-prover commitment schemes, and assume that S and S' are respectively ε -binding and δ -binding against classical attacks. Then, their composition $S'' = S \star S'$ is a $(\varepsilon + \delta)$ -binding 2-prover commitment scheme against classical attacks.*

Proof. The proof is almost the same as that of Theorem 4.8, except that now there are no s_o and y_o , and in the end we simply conclude that

$$p(s \neq \hat{s} \wedge s \neq \perp) \leq p(s \neq \hat{s} \wedge s \neq \perp \wedge s = \tilde{s}) + p(s \neq \tilde{s} \wedge s \neq \perp)$$

$$\begin{aligned} &\leq p(\tilde{s} \neq \hat{s} \wedge \tilde{s} \neq \perp) + p(y \neq \hat{y} \wedge y \neq \perp) \\ &\leq \varepsilon + \delta, \end{aligned}$$

where the second inequality holds since $y = \perp$ implies $s = \perp$. \square

Acknowledgments

We would like to thank Jędrzej Kaniewski for helpful discussions regarding [10], and for commenting on an earlier version of our work.

References

1. Mohammad Bavarian and Peter W. Shor. Information Causality, Szemerédi-Trotter and Algebraic Variants of CHSH. In Tim Roughgarden, editor, *ITCS 2015*, pages 123–132. ACM, 2015.
2. Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-Prover Interactive Proofs: How to Remove Intractability Assumptions. In Janos Simon, editor, *STOC 1988*, pages 113–131. ACM, 1988.
3. Kaushik Chakraborty, André Chailloux, and Anthony Leverrier. Arbitrarily Long Relativistic Bit Commitment. *ArXiv e-prints*, 2015. <http://arxiv.org/abs/1507.00239>.
4. John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed Experiment to Test Local Hidden-Variable Theories. *Physical Review Letters*, 23:880–884, 1969.
5. Claude Crépeau, Louis Salvail, Jean-Raymond Simard, and Alain Tapp. Two Provers in Isolation. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 407–430. Springer, 2011.
6. Serge Fehr and Max Fillinger. Multi-Prover Commitments Against Non-Signaling Attacks. In Rosario Genaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015, part II*, volume 9216 of *Lecture Notes in Computer Science*, pages 403–421, 2015.
7. Adrian Kent. Unconditionally Secure Bit Commitment. *Physical Review Letters*, 83(7):1447–1450, 1999.
8. Adrian Kent. Secure Classical Bit Commitment Using Fixed Capacity Communication Channels. *Journal of Cryptology*, 18(4):313–335, 2005.
9. Hoi-Kwong Lo and Hoi Fung Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78:3410–3413, 1997.
10. Tommaso Lunghi, Jędrzej Kaniewski, Felix Bussières, Raphael Houlmann, Marco Tomamichel, Stephanie Wehner, and Hugo Zbinden. Practical Relativistic Bit Commitment. *Physical Review Letters*, 115, 2015.
11. Dominic Mayers. Unconditionally Secure Quantum Bit Commitment is Impossible. *Physical Review Letters*, 18:3414–3417, 1997.
12. Jamie Sikora, André Chailloux, and Iordanis Kerenidis. Strong Connections Between Quantum Encodings, Non-Locality and Quantum Cryptography. *Physical Review A*, page 9, 2014.

A The Hiding Property of Composed Schemes

We already mentioned that the standard hiding property is not good enough for multi-round relativistic bit commitment schemes, where we want the hiding property to hold until the last round of communication. In this appendix, we define a variation of the hiding property that captures this requirement, and we prove that a composed scheme $\mathcal{S}'' = \mathcal{S} \star \mathcal{S}'$ is hiding “up to the last round” if both \mathcal{S} and \mathcal{S}' are (with the error parameters adding up).

Definition A.1. *Let $\mathcal{S} = (\text{com}_{PQV}, \text{open}_{PQV})$ be a 2-prover commitment scheme. We say that \mathcal{S} is ε -hiding until the last round if for any dishonest verifier V and any two inputs s_0 and s_1 to the honest provers, we have $d(p(v|s_0), p(v|s_1)) \leq \varepsilon$, where v is the verifier’s view immediately before the last round of communication in $(\text{open}_{PQ} \parallel \overline{\text{open}}_V) \circ (\text{com}_{PQ} \parallel \overline{\text{com}}_V)(s_b \parallel s_b \parallel \emptyset)$.*

Theorem A.2. *Let \mathcal{S} be an ε -hiding commitment scheme and \mathcal{S}' a scheme that is δ -hiding until the last round. If $(\mathcal{S}, \mathcal{S}')$ is eligible, then the composed scheme $\mathcal{S}'' = \mathcal{S} \star \mathcal{S}'$ is $(\varepsilon + \delta)$ -hiding until the last round.*

Proof. Fix some strategy against the hiding-until-the-last-round property of \mathcal{S}'' . We consider the distribution $p(v, y, v'|s)$ where s is the string that the provers commit to, v the verifier’s view after com_{PQV} has been executed, y the opening information to which Q commits using the scheme \mathcal{S}' , and v' the verifier’s view immediately before the last round of communication. We need to show that $d(p(v'|s_0), p(v'|s_1)) \leq \varepsilon + \delta$ for any s_0 and s_1 .

First, note that $p(v'|v, y, s_b) = p(v'|v, y)$ since v' is produced by P, Q and V acting on y and v only. From any strategy against \mathcal{S}'' , we can obtain a strategy against \mathcal{S}' by fixing v . Thus, by the hiding property of \mathcal{S}' , for any y_0 and y_1 , we have $d(p(v'|v, y = y_0), p(v'|v, y = y_1)) \leq \delta$ and it follows by the convexity of the statistical distance in both arguments that

$$p(v'|v, s_0) = \sum_y p(y|v, s_0) p(v'|v, y) \approx_\delta \sum_y p(y|v, s_1) p(v'|v, y) = p(v'|v, s_1)$$

where we use \approx_δ to indicate that the two distributions have statistical distance at most δ . Since we have $d(p(v|s_0), p(v|s_1)) \leq \varepsilon$ by the hiding property of \mathcal{S} , it follows that

$$\begin{aligned} p(v'|s_0) &= p(v, v'|s_0) = p(v|s_0) p(v'|v, s_0) \approx_\delta p(v|s_0) p(v'|v, s_1) \\ &\approx_\varepsilon p(v|s_1) p(v'|v, s_1) = p(v, v'|s_1) = p(v'|s_1) \end{aligned}$$

where the first and last equalities hold because v' contains v since v' is the view of V at a later point in time. \square