Two-Prover Bit-Commitments: Classical, Quantum and Non-Signaling

Fillinger, M.J.

Summary

This dissertation makes different contributions to the theory of multi-prover commitment schemes; in particular relativistic commitment schemes. A *commitment scheme* is an important cryptographic primitive crucial for cryptographic protocols that allow two or more parties that do not fully trust each other to cooperate in a secure way. More concretely, we consider multi-prover commitment schemes whose security relies on an assumed restriction on the communication between the provers, and not on computational hardness assumptions.

A commitment scheme is a means to solve the following problem: Alice has selected a message which she wants to keep secret at the moment, but which she may want to reveal to Bob at a later time. However, if she simply sends the message to Bob later, he has no guarantee that the message he received is the same one that Alice selected earlier. A commitment scheme ensures the secrecy of the message until Alice chooses to reveal it, but also prevents Alice from revealing a different message than the one she selected originally.

Formally, a commitment scheme consists of a pair of interactive protocols between (usually) two parties called the *prover* and the *verifier*. The first protocol is called the *commit phase*; it takes a message from the prover as input and no input from the verifier. The second protocol, called the *opening phase*, then outputs either a message or the failure symbol $\perp$ to the verifier. If the output is a message m, we say that the prover opened the commitment to m; if the output is $\perp$, we say that the prover failed to open the commitment. Often, the opening phase just consists in the prover sending some *opening information* to the verifier who then computes the output locally.

To be secure, a commitment scheme needs to have the following three properties: It should be *complete*, meaning that the input to the commit phase and the output of the opening phase are equal if both parties follow the protocols. It should be *hiding*, meaning that the verifier cannot learn the prover's input message before the opening phase is executed, even if the verifier is dishonest and deviates from the protocols. Finally, it should be *binding*, meaning that after the commit phase, there is at most one message that the prover can successfully open to, even if the prover is dishonest and deviates from the protocols. The set of possible messages that a commitment

scheme can take as input is called its *domain*. If the domain is the set $\{0, 1\}$, we speak of a *bit*-commitment scheme.

The standard notion of commitment schemes can offer security only if the dishonest prover or the dishonest verifier is *computationally bounded*, i.e. limited in the amount of computation he can perform. However, this is only true for commitment schemes with a single prover. Ben-Or, Goldwasser, Kilian and Wigderson showed in 1988 how to overcome this limitation by considering a variant of the notion of commitment schemes where the prover is split into two (or more) separate entities and it is assumed that they cannot communicate during the execution of the commitment scheme. Related to this approach is the notion of *relativistic commitment schemes*, introduced by Kent in 1999, where this non-communication assumption is temporarily enforced through spatial separation of the provers.

The first main contribution of this dissertation is a set of new definitions of the binding property for multi-prover commitment schemes. These new definitions have several advantages over the *sum-binding* definition which has been used so far: They are not restricted to *bit*-commitment schemes but are applicable to commitment schemes with arbitrary finite domains. When restricted to bits, some of our definitions are *strictly stronger* than the sum-binding definition. Finally, our definitions are closer to the intuitive notion of a commitment scheme being binding and are more convenient to work with. We introduce these new definitions and study how they relate to each other.

As a testing ground for our new definitions, we consider the bit-commitment scheme $\mathcal{CHSH}^q$, introduced by Crépeau, Salvail, Simard and Tapp, which can be extended in a natural way to a commitment scheme with domain $\mathbb{F}_q$ (where q is a prime power). We analyze it with respect to this larger domain for the first time and show that different variations of the scheme satisfy our different definitions of the binding property.

Our new definitions enable us to prove a rather general *composition theorem* for two-prover commitment schemes, which is the second main contribution of this dissertation. We compose two commitment schemes by having the provers commit to the opening information of the first scheme instead of sending it to the verifier, and then they open this second commitment, revealing the opening information of the first scheme so that the original commitment is opened. Under some mild assumptions about the two original schemes, we prove that the composed scheme is binding if the two original schemes are binding (with the cheating probabilities adding up).

The purpose of this composition is to *delay* the opening of the commitment. This is important in the context of relativistic commitment schemes where the no-communication assumption is only enforced temporarily and so the binding property holds only for a limited time.

Very concretely, our composition theorem allows us to give a tight analysis of the relativistic commitment scheme introduced by Lunghi, Kaniewski, Brussières, Houlman, Tomamichel and Wehner in 2015. Their original anal-

ysis showed an upper bound on the cheating probability of dishonest provers that was doubly-exponential in the number of communication rounds where the latter determines for how long the scheme remains binding. The scheme can be understood as an iterated composition of $\mathcal{CHSH}^q$ with itself and so we can analyze its security using our composition theorem, achieving a significant improvement over the original analysis by Lunghi *et al.*: it follows that the cheating probability for dishonest provers can be bounded by a term that is only *linear* in the number of rounds, rather than double-exponential. We also show the optimality of our bound up to a small constant factor.

To put this difference into more concrete terms: Lunghi *et al.* implemented their scheme with provers in Bern and Geneva (distance: 129.2 km). Their analysis guaranteed that the commitment would stay binding (with a reasonably low cheating probability) for about 2 ms. Based on our analysis, this time scales up to 10^{56} years, or, speaking more practically, until the devices run out of memory.

The third main contribution is an impossibility result about two-prover commitment schemes with *general non-signaling adversaries*. As Crépeau, Salvail, Simard and Tapp pointed out, the assumption that the provers cannot communicate needs further specification. Some two-prover commitment schemes are secure against classical non-communicating provers but insecure if they have quantum capabilities and can use entangled quantum states shared among them. Furthermore, if we want to truly base security on the sole assumption that the provers cannot communicate, we need to consider general non-signaling provers, i.e., provers whose behavior may be correlated in arbitrary ways as long as no communication between them is implied. The $\mathcal{CHSH}^q$ scheme is secure against provers with quantum entanglement, but insecure against general non-signaling provers. This raises the question whether any other commitment schemes are secure against such general non-signaling provers.

We show that for *two-prover* commitment schemes the answer is no: any commitment scheme that is complete and hiding is by necessity not binding against general non-signaling provers. On the other hand, we show a positive answer for *three-prover* commitment schemes: we prove that a simple extension of $\mathcal{CHSH}^q$ to three provers is complete, hiding against an arbitrary dishonest verifier and binding against general non-signaling provers.