Collapse-binding commitments in the standard model*

Dominique Unruh
University of Tartu

Introduction. Commitment schemes are one of the most fundamental primitives in cryptography. A commitment scheme is a two-party protocol consisting of two phases, the commit and the open phase. The goal of the commitment is to allow the sender to transmit information related to a message m during the commit phase in such a way that the recipient learns nothing about the message (hiding property). But at the same time, the sender cannot change his mind later about the message (binding property). Later, in the open phase, the sender reveals the message m and proves that this was indeed the message that he had in mind earlier (by sending some "opening information" u). Unfortunately, it was shown by [3] that the binding and hiding property of a commitment cannot both hold with information-theoretical security. Thus, one typically requires one of them to hold only against computationally-limited adversaries. Since the privacy of data should usually extend far beyond the end of a protocol run, and since we cannot tell which technological advances may happen in that time, we may want the hiding property to hold information-theoretically, and thus are interested in *computationally binding* commitments. Unfortunately, computationally binding commitments turn out to be a subtle issue in the quantum setting. As shown in [7], if we use the natural analogue to the classical definition of computationally binding commitments,[1] we get a definition that is basically meaningless (the adversary can open the commitment to whatever message he wishes). They suggested a new definition, "collapse binding" commitments, that better captures the idea of computationally binding commitments. This definition was shown to perform well in security proofs that use rewinding (e.g., they allow the construction of statistical zero-knowledge arguments of knowledge). (They studied classical non-interactive commitments, i.e., all exchanged messages are classical, but the adversary is quantum.)

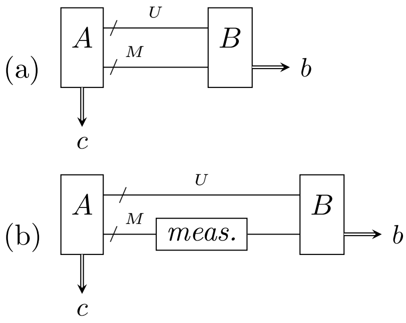

Figure 1: (a) and (b) should be indistinguishable, i.e., $\Pr[b = 1]$ negligibly close in both cases.

We describe basic idea of "collapse-binding" commitments: When committing to a message m using a commitment c, it should be impossible to produce a superposition of different messages m that the adversary can open to. Unfortunately, this requirement is too strong to achieve (at least for an information-theoretically hiding commitment).[2] Instead, we require something slightly weaker: Any superposition of different messages m that the adversary can open to should *look like* it is a superposition of only a single message m. Formally, if the adversary produces a classical commitment c, and a superposition of openings m, u in registers M, U, the

[*] An early draft of a full version is provided here: `http://www.ut.ee/~unruh/commit-std.pdf`. It provides proofs, but no motivation or explanations yet.

[1] This definition roughly states, that it is computationally hard to find a commitment c, two messages $m \neq m'$ and corresponding valid opening informations u, u'.

[2] The adversary can initialize a register M with the superposition of all messages, run the commit algorithm in superposition, and measure the resulting commitment c. Then M will still be in superposition between many messages m which the adversary can open c to.

adversary should not be able to distinguish whether M is measured in the computational basis or not measured. That is, for all quantum-polynomial-time A, B, the circuits (a) and (b) are indistinguishable (assuming A only outputs superpositions that contain only valid openings).

[7] further showed that in the quantum random oracle model, collapse-binding, information-theoretically hiding commitments can be constructed. However, they left open two big questions:

- Can collapse-binding commitments be constructed in the standard model? That is, without using the random oracle *heuristic*.
- One standard minimum requirement for commitments (called "sum-binding" in [7]) is that for quantum-polynomial-time A, $p_0 + p_1 \leq 1 + \textit{negligible}$ where p_b is the probability that A opens a commitment to b when he learns b only *after* the commit phase. Surprisingly, [7] left it open whether the collapse-binding property implies the sum-binding property.

Our contribution: collapse-binding in the standard model. We show that collapse-binding commitments exist in the standard model. More precisely, we construct a non-interactive, classical commitment in the public parameter model (i.e., we assume that some parameters are globally fixed), for arbitrarily long messages (the length of the public parameters and the commitment itself do not grow with the message length), information-theoretically hiding, and collapse-binding. The security assumption is the existence of lossy trapdoor functions [5], or alternatively that SIVP and GapSVP are hard for quantum algorithms to approximate within $\tilde{O}(d^c)$ factors for some constant $c > 5$.

The basic idea of our construction is the following: In [7], it was shown that information-theoretically hiding, collapse-binding commitments can be constructed from "collapsing" hash functions (using a classical construction from [2]). A function H is collapsing if an adversary that outputs h and a superposition M of H-preimages of h cannot distinguish whether M is measured or not. That is, the circuits (c) and (d) should be indistinguishable. So all we need to construct is a collapsing function in the standard model.

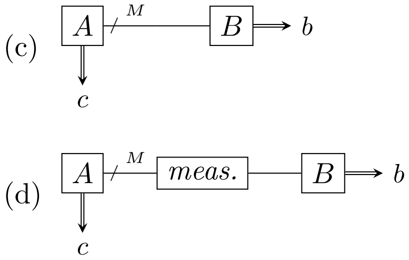

Figure 2: (c) and (d) should be indistinguishable.

To do so, we use a lossy trapdoor function (we do not actually need the trapdoor part, though). A lossy function $f_s : A \to B$ is parametrized by a public parameter s. There are two kinds of parameters, which are assumed to be indistinguishable: We call s lossy if $|\operatorname{im} f(A)| \ll |A|$. We call s injective if f_s is injective.

If s is injective, then it is easy to see that f_s is collapsing: There can be only one preimage of h on register M, so measuring M will not disturb M. But since lossy and injective s are indistinguishable, it follows that f_s is also collapsing for lossy s. Note, however, that f_s is not yet useful on its own, because its range B is much bigger than A, while we want a compressing hash functions (output smaller than input).

However, for lossy s, $|f_s(A)| \ll |A|$. Let $h_r : A \to C$ be a universal hash function, indexed by r, with $|f_s(A)| \ll |C| \ll |A|$. We can show that with overwhelming probability h_r is injective on $f_s(A)$, for suitable choice of C. Hence h_r is collapsing (on $f_s(A)$). The composition of two collapsing functions is collapsing, thus $H_{(r,s)} := h_s \circ f_s$ is collapsing for lossy s.

Thus far, we have found a collapsing $H_{(r,s)} : A \to C$ that is compressing. But we need something stronger, namely a collapsing hash function $\{0,1\}^* \to C$, i.e., applicable to arbitrary long inputs. A well-known construction (in the classical setting) is the Merkle-Damgård construction, that transforms a compressing collision-resistant function H into a collision-resistant one with domain $\{0,1\}^*$. We prove that the Merkle-Damgård construction also preserves the collapsing property. Applying this result to $H_{(r,s)}$, we get a collapsing hash function $\mathrm{MD}_{(r,s)} : \{0,1\}^* \to C$. And from this, we get collapse-binding commitments.

We present all our proofs with concrete security bounds, and our reductions have only constant factors in the runtime, and the security level only has an $O(\textit{message length})$ factor.

We stress that the security proof for the Merkle-Damgård construction has an additional benefit: It shows that existing hash function like SHA-2 [4] are collapsing, assuming that the compression function is collapsing (which in turn is suggested by the random oracle results in [7]). Since we claim that collapsing is a desirable and natural analogue to collision-resistance in the post-quantum setting, this gives evidence for the post-quantum security of SHA-2.

Our contribution: Collapse-binding implies sum-binding. In the classical setting, it relatively straightforward to show that a computationally binding *bit* commitment satisfies the (classical) sum-binding condition. Namely, assume that the adversary breaks sum-binding, i.e., $p_0 + p_1 > 1 + \textit{negligible}$. Then one runs the adversary, lets him open the commitment as $m = 0$ (which succeeds with probability p_0), then rewinds the adversary, and lets him open the same commitment as $m = 1$ (which succeeds with probability p_1). So the probability that both runs success is at least $p_0 + p_1 - 1 > \textit{negligible}$, which is a contradiction to the computational binding property.

Since collapse-binding commitments work well with rewinding, one would assume that a similar proof works using the rewinding technique from [6]. Indeed, the proof goes through; but due to the worse bounds achieved by quantum rewinding, we can only conclude $p_0 + p_1 \leq \sqrt{2} + \textit{negigible}$. (This is better than the trivial bound $p_0 + p_1 \leq 2$, and possibly sufficient in many applications, but not enough to achieve sum-binding.)

To show that a collapse-binding commitment is sum-binding, another proof technique is needed. The basic idea is, instead of simulating two executions of the adversary (opening $m = 0$ and opening $m = 1$) after each other, we perform the two executions in superposition. For space reasons, we omit the details here, but in the end we can show that if $p_0 + p_1 - 1 > \textit{negligible}$, we can transform the this superposition of executions into an adversary against the collapse-binding property.

The same proof technique can be used to show that a collapse-binding *string* commitment satisfies the generalization of sum-binding presented in [1].

Possibly the technique of "rewinding in superposition" used here might be a special case of a more general new quantum rewinding technique (other than [8, 6]), we leave this as an open question.

Bibliography

[1] C. Crépeau, P. Dumais, D. Mayers, and L. Salvail. Computational collapse of quantum state with application to oblivious transfer. In *TCC 2004*, volume 2951 of *LNCS*, pages 374–393. Springer, 2004.

[2] S. Halevi and S. Micali. Practical and provably-secure commitment schemes from collision-free hashing. In N. Koblitz, editor, *Crypto '96*, volume 1109 of *LNCS*, pages 201–215. Springer, 1996.

[3] D. Mayers. Unconditionally Secure Quantum Bit Commitment is Impossible. *Physical Review Letters*, 78(17):3414–3417, 1997.

[4] National Institute of Standards and Technology (NIST). Secure hash standard (SHS). FIPS PUBS 180-4, 2015.

[5] C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196, New York, NY, USA, 2008. ACM.

[6] D. Unruh. Quantum proofs of knowledge. In *Eurocrypt 2012*, volume 7237 of *LNCS*, pages 135–152. Springer, April 2012.

[7] D. Unruh. Computationally binding quantum commitments. In *Eurocrypt 2016*, 2016. To appear. Preprint on IACR ePrint 2015/361.

[8] J. Watrous. Zero-knowledge against quantum attacks. *SIAM J. Comput.*, 39(1):25–58, 2009.