(Nearly) Round-Optimal Black-Box Constructions of Commitments Secure against Selective Opening Attacks

David Xiao[1,2,⋆]

[1] LIAFA, Université Paris 7, 75205 Paris Cedex 13, France
[2] Université Paris-Sud, 91405 Orsay Cedex, France
dxiao@liafa.fr

Abstract. Selective opening attacks against commitment schemes occur when the commitment scheme is repeated in parallel (or concurrently) and an adversary can choose depending on the commit-phase transcript to see the values and openings to some subset of the committed bits. Commitments are secure under such attacks if one can prove that the remaining, unopened commitments stay secret.

We prove the following black-box constructions and black-box lower bounds for commitments secure against selective opening attacks:

1. For parallel composition, 4 (resp. 5) rounds are necessary and sufficient to build computationally (resp. statistically) binding and computationally hiding commitments. Also, there are no perfectly binding commitments.
2. For parallel composition, $O(1)$-round statistically-hiding commitments are equivalent to $O(1)$-round statistically-binding commitments.
3. For concurrent composition, $\omega(\log n)$ rounds are sufficient to build statistically binding commitments and are necessary even to build computationally binding and computationally hiding commitments, up to $\log \log n$ factors.

Our lower bounds improve upon the parameters obtained by the impossibility results of Bellare *et al.* (EUROCRYPT '09), and are proved in a fundamentally different way, by observing that essentially all known impossibility results for black-box zero-knowledge can also be applied to the case of commitments secure against selective opening attacks.

Keywords: commitments, black-box lower bounds, zero knowledge, selective opening attacks, parallel composition, concurrent composition.

1 Introduction

Commitment schemes have a wide array of applications in cryptography, one of the most notable being the construction of zero knowledge protocols [14, 4].

⋆ Supported by the French ANR Defis program under contract ANR-08-EMER-012.

Y. Ishai (Ed.): TCC 2011, LNCS 6597, pp. 541–558, 2011.

A problem that arises in the use of commitment schemes is whether their hiding property holds when composed in parallel: if some subset of the committed messages are opened, do the remaining unopened messages remain secure? This question arose early in the study of zero knowledge protocols, and is also natural in other cryptographic contexts where commitments are used as building blocks for protocols that might be then used in parallel (*e.g.* secure multi-party computation, etc.).

Although naively one might think because commitments are hiding that no additional information should be leaked by composing them, nevertheless it is unknown how to prove that standard stand-alone commitments (*e.g.* [18]) remain hiding when composed.

More formally, a selective opening attack on a commitment scheme allows a cheating receiver to interact in k parallel (or concurrent) commitments, and then ask the sender to open some subset $I \subseteq [k]$ of the commitments. The question is whether the unopened messages remain hidden in the following sense: is there a simulator strategy for every cheating receiver strategy that outputs a commit-phase transcript, a set $I \subset [k]$, and decommitments to $(b_i)_{i \in I}$ that is indistinguishable from the output of the cheating receiver with an honest sender?

In this paper we show that techniques both for constructions and lower bounds from the study of zero knowledge protocols can be applied to the study of commitments secure against selective opening attacks. We study the minimal round complexity needed to construct such commitments, and give solutions for commitments secure against selective opening attacks that are optimal or nearly optimal up to small factors.

1.1 Our Results

We let PAR denote parallel composition and CC denote concurrent composition. We let CB (resp. SB, PB) denote computational (resp. statistical, perfect) binding and CH (resp. SH) denote computational (resp. statistical) hiding. We give the following constructions:

Theorem 1. *The following hold via fully black-box reductions:*

1. *One-way permutations imply* 4*-round* PAR-CBCH *commitments exist.*
2. t*-round stand-alone* SH *commitments imply* $(t + 3)$*-round* PAR-SB *commitments exist.*
3. t*-round stand-alone* SH *commitments imply* $\omega(t \log n)$*-round* CC-SB *commitments exist.*

In particular, Item 2 implies that collision-resistant hash functions (or even just 2-round statistically hiding commitments) suffice to construct 5-round PAR-SB commitments.

Assuming the proof of security for such a commitment scheme is given by a black-box simulator, we prove the following corresponding lower bounds:

Theorem 2 (Informal). *The following hold relative to any oracle:*

1. *There is no* 3*-round* PAR-CBCH *commitment.*
2. *There is no* 4*-round* PAR-SB *commitment.*
3. *There is a black-box reduction that uses a* $O(1)$*-round* PAR-SB *commitment to build a* $O(1)$*-round statistically hiding commitment.*
4. *There is no* $o(\log n/\log\log n)$*-round* CC-CBCH *commitment.*

We stress that besides the constraint that the simulator be black-box, these results are otherwise *unconditional*. Namely, Theorem 2 implies that no such commitments exist in the plain model (without oracles), but also implies that such commitments do not exist even in say the random oracle model (or stronger oracle models), where *a priori* one might have hoped to bypass impossibility results in the plain model.

Combining the second item of Theorem 2 with the main theorem of [15], which proves that there is no black-box reduction building a $o(n/\log n)$-round statistically hiding commitment from one-way permutations, we obtain the following corollary:

Corollary 1. *There is no black-box reduction that uses a one-way permutation to build a* $O(1)$*-round* PAR-SB *commitment.*

Wee [23] independently proved via different techniques a theorem similar to Corollary 1 for the very closely related case of trapdoor commitments.

In addition to the above impossibility results, we also prove:

Theorem 3 (Informal). *Relative to any oracle, there exists no* PAR-PB *commitments nor receiver public-coin* PAR-CBCH *commitments.*

1.2 Comparison to Previous Constructions

Notions related to security against selective opening attacks have previously been studied in the literature. Security against selective opening is closely related to chameleon blobs [5, 6], trapdoor commitments [11], and equivocable commitments [2, 9, 8]. Roughly speaking, these notions all allow a simulator that can generate commit-phase transcripts that can be opened in many ways. Indeed, our constructions will be based on the equivocable commitment of [8].

Security against selective opening may be weaker than the notions above, and was directly studied in [10, 3]. Bellare *et al.* [3] give a construction of a scheme that is CC-SB secure, but this construction is non-black-box and requires applying a concurrent zero knowledge proof on a statement regarding the code implementing a one-way permutation. In contrast, all constructions presented in this paper are fully black-box.

Equivalence of statistical hiding and statistical binding. In this work we only study commitments with computational hiding. [3] already noted that stand-alone SH commitments satisfy a notion of PAR-SH security based on indistinguishability (this notion is different from ours). Independent of our work, Zhang *et al.* [24] gave a black-box reduction that uses t-round stand-alone SH commitments and one-way permutations to construct $(t+3)$-round PAR-SH commitments (under

our definition of selective opening security). Their construction is an extension of a recent trapdoor commitment of Pass and Wee [19].

With Item 2 of Theorem 2, this implies that $O(1)$-round statistical hiding and $O(1)$-round statistical binding are *equivalent* via black-box reductions when security against selective opening attacks is required. This contrasts sharply with the stand-alone case, as 2-round statistically binding commitments are equivalent to one-way functions, but no black-box reduction can build $o(n/\log n)$-round statistically hiding commitment from one-way functions [15].

1.3 Comparison to Previous Lower Bounds

Bellare *et al.* [3] proved that non-interactive commitments and perfectly binding commitments secure against selective opening attacks cannot be based on *any* black-box cryptographic assumption. Our lower bounds are stronger than theirs in that we can rule out 3- or 4-round rather than non-interactive commitments, as well as ruling out certain types of commitment with non-zero statistical binding error. However, our proof *technique* is incomparable to theirs.

Ways in which our lower bounds are stronger: first, the lower bounds of [3] assume black-box access to a cryptographic primitive, and therefore do not apply to constructions based on *concrete assumptions* (*e.g.* factoring, discrete log, lattice problems) where one might hope to exploit the specific structure of those problems to achieve security. In contrast, our results immediately rule out all constructions in the plain model.

Second, the lower bounds of [3] prove that non-interactive and perfectly binding commitments secure against selective opening attacks are impossible with respect to a very specific message distribution *that is defined in terms of a random oracle*. One could argue that the message distribution they consider is artificial and would not arise in applications of these commitments. In particular, it may suffice for applications to build commitments that are secure only for particular natural message distributions, such as the uniform distribution or the distributions encountered when using commitments to build zero knowledge proofs for **NP**. [3] does not rule out the existence of commitments that are secure only for these message distributions, while our impossibility results do and in fact apply simultaneously to all message distributions satisfying what we argue are very natural constraints (see Definition 5). In particular, the results of [3] also use the assumptions in Definition 5.

Ways in which our lower bounds are weaker: our results are weaker because they only apply to constructions with black-box simulators, *i.e.* we require that there exists a single simulator that works given black-box access to any cheating receiver. The results of [3] hold even for slightly non-black-box simulation techniques: they only require that for every cheating receiver oracle algorithm $(\mathsf{Rec}')^{(\cdot)}$ that accesses the underlying crypto primitive as a black-box, there exists an efficient oracle algorithm $\mathsf{Sim}^{(\cdot)}$ that accesses the underling crypto primitive as a black box that generates an indistinguishable transcript. However, [3] do *not* rule out techniques such as Barak's simulator [1], because the simulator there includes

a PCP encoding of the code of the underlying cryptographic primitive, and thus treats the *crypto primitive itself* in a non-black-box way.

1.4 Our Techniques

Our constructions for parallel composition are essentially the equivocable commitment scheme of [8], while the case for concurrent composition follows in a straight-forward way by combining the commitment of [8] with the preamble from the concurrent zero knowledge proof of [21].

Our lower bounds are proven by observing that most known lower bounds for zero knowledge (*e.g.* [13, 17, 7, 16, 20]) extend naturally to the case of commitment schemes. Lower bounds for zero knowledge show that if a zero knowledge proof for L satisfies certain restrictions (*e.g.* 3 rounds, constant-round public coin [13], etc.), then $L \in \mathbf{BPP}$.

As was observed by [10, 3], plugging a t-round PAR-CBCH commitment into the GMW zero knowledge protocol for **NP** allows the zero knowledge property to be preserved under parallel repetition, thus allowing one to reduce soundness error while preserving zero knowledge and without increasing round complexity. Furthermore, the resulting protocol has $t + 2$ rounds, and has a black-box simulator if the commitment had a black-box simulator. This immediately implies the following:

Proposition 1 ([13], weak impossibility of PAR-CBCH, informal). *In the plain model, there exist no black-box simulator non-interactive or constant-round public-coin PAR-CBCH commitment schemes.*

To see why, suppose there were such a scheme, then by the above discussion one would obtain either a 3-round or constant-round public-coin zero knowledge argument for **NP** with a black-box simulator that remains zero knowledge under parallel repetition. By [13], this implies that $\mathbf{NP} = \mathbf{BPP}$. But this contradicts the existence of a PAR-CBCH commitment scheme, since by the Cook-Levin reduction we can use an algorithm solving **NP** to break any commitment.

Our results improve upon Proposition 1 as they apply to broader categories of commitments (*e.g.* 3-round vs. non-interactive). In addition, Proposition 1 uses the Cook-Levin reduction and therefore does not apply when considering schemes that might use random oracles. In contrast, Theorem 2 does hold relative to any oracle, and in the case of Item 3 of Theorem 2, is *black-box*. This is important for two reasons: first, Proposition 1 does not say whether such constructions are possible in the random oracle model, which is often used to prove the security of schemes for which we cannot prove security in the plain model. Second, if we want to compose our impossibility result with other black-box lower bounds, then our impossibility result had better also be black-box. For example, in order to obtain Corollary 1 we must combine Item 3 of Theorem 2 with the black-box lower bound of Haitner *et al.*. This is only possible if Item 3 of Theorem 2 is a black-box reduction, which would not be true using the approach of the weak impossibility result Proposition 1.

To prove Theorem 2, we construct what we call "equivocal senders": senders that run the commit phase without knowing the bits that must be revealed. We show that the existence of such equivocal senders implies that binding can be broken. We then construct equivocal senders for various kinds of protocols by applying the proof strategy for zero knowledge lower bounds originally outlined by Goldreich and Krawczyk [13]. By arguing directly, we avoid the Cook-Levin step in Proposition 1 and therefore our results hold relative to any oracle.

2 Preliminaries

For a random variable X, we let $x \leftarrow_{\mathrm{R}} X$ denote a sample drawn according to X. We let U_k denote the uniform distribution over $\{0,1\}^k$. For a set S, we let $x \leftarrow_{\mathrm{R}} S$ denote a uniform element of S. Let 2^S denote the set of all subsets of S. All security definitions in this paper are with respect to non-uniform circuits. We say that an event occurs with overwhelming probability if it occurs with probability $1 - n^{-\omega(1)}$, and that it occurs with negligible probability if it occurs with probability $n^{-\omega(1)}$. Two families of random variables $(X_n)_{n\in\mathbb{N}}, (Y_n)_{n\in\mathbb{N}}$ over $\{0,1\}^n$ are computationally indistinguishable, or equivalently $X \approx_c Y$, if for all circuits C of size $\mathrm{poly}(n)$ it holds that $|\Pr[C(X) = 1] - \Pr[C(Y) = 1]| \leq n^{-\omega(1)}$.

2.1 Commitment Schemes

We formally define commitments for single-bit messages; since we will be concerned with commitments that are composable, multi-bit messages can be handled by just repeating the single-bit protocol in parallel or concurrently.

Definition 1. *A t-round (stand-alone) commitment protocol is a pair of efficient algorithms* Send *and* Rec. *Given a sender input $b \in \{0,1\}$, we define:*

1. *The* commit phase *transcript is $\tau = \langle \mathsf{Send}(b; \omega_{\mathsf{Send}}), \mathsf{Rec}(\omega_{\mathsf{Rec}}) \rangle$ where ω_{Send}, ω_{Rec} are the random coins of the sender and receiver, respectively. Exactly t messages are exchanged in the commit phase t.*
2. *The* decommit phase *transcript consists of* Send *sending* (b, open) *to* Rec. $\mathsf{Rec}(\tau, b, \mathsf{open}) = 1$ *if* open *is a valid opening, and outputs* 0 *otherwise.*

Notation and variable definitions: We assume that a commitment scheme is put in a canonical form, where each party alternates speaking. We assume the number of rounds is even and the receiver speaks first. If the number of rounds is $2t$, then we label the sender's messages $\alpha_1, \ldots, \alpha_t$ and the receiver's messages $\beta_1, \ldots, \beta_t$, and we let $\alpha_{[i]} = (\alpha_1, \ldots, \alpha_i)$ and likewise for $\beta_{[i]}$. For a commitment protocol $(\mathsf{Send}, \mathsf{Rec})$, we write that the receiver's i'th response β_i is given by computing $\beta_{[i]} = \mathsf{Rec}(\alpha_{[i-1]}; \omega)$ where $\alpha_{[i-1]}$ are the first $i-1$ sender messages, and ω are the receiver's random coins. We let $\mathsf{Rec}(\perp; \omega) = \beta_1$ denote the first receiver message.

Let k denote the number of parallel or concurrent repetitions of a commitment protocol. Let n denote the security parameter of the protocol. Given a stand-alone commitment $(\mathsf{Send}, \mathsf{Rec})$, let Send^k denote the k-fold repeated sender (context will determine whether we mean parallel or concurrent composition). Let Rec^k denote the k-fold parallel receiver, and let Rec^k_Σ denote the k-fold concurrent receiver with schedule Σ. Underlined variables denote vectors of message bits (*e.g.* $\underline{b} \in \{0,1\}^k$) and plain letters with indices the bit at each coordinate (*e.g.* b_i is the i'th bit of $\underline{b}$).

Definition 2 (Binding). *A commitment scheme* $(\mathsf{Send}, \mathsf{Rec})$ *is computationally (resp. statistically) binding if for all polynomial-time (resp. unbounded) sender strategies* Send'*, only with negligible probability can* Send' *interact with an honest* Rec *to generate a commit-phase transcript* τ *and then produce* $\mathsf{open}, \mathsf{open}'$ *such* $\mathsf{Rec}(\tau, 0, \mathsf{open}) = 1$ *and* $\mathsf{Rec}(\tau, 1, \mathsf{open}') = 1$*. A scheme is* perfectly *binding if the above probability of cheating is* 0*.*

It is straight-forward to prove that all the variants of the binding property are preserved under parallel/concurrent composition.

Hiding under selective opening attacks. We only study the case of computational hiding. In the following, $\mathcal{I} \subseteq 2^{[k]}$ *is a family of subsets of* $[k]$*, which denotes the set of legal subsets of commitments that the receiver is allowed to ask to be opened.*

Definition 3 (Hiding under selective opening: k-fold parallel composition security game). *Sender input:* $\underline{b} \in \{0,1\}^k$*. Let* Rec' *be the (possibly cheating) sender.*

1. $\mathsf{Send}^k, \mathsf{Rec}'$ *run* k *executions of the commit phase in parallel using independent random coins, obtaining* k *commit-phase transcripts* $\tau^k = (\tau_1, \ldots, \tau_k)$*.*
2. Rec' *chooses a set* $I \leftarrow_R \mathcal{I}$ *and sends it to* Send^k*.*
3. Send^k *sends* (b_i, ω_i) *for all* $i \in I$*, where* ω_i *is an opening of the* i*'th commitment.*

In Item 2, the honest receiver is defined to pick $I \in \mathcal{I}$ uniformly, while a malicious receiver may pick I adversarially.

Definition 4 (Hiding under selective opening, parallel composition). *Let* $\mathcal{I} \subseteq 2^{[k]}$ *be a family of subsets and* $\underline{\mathcal{B}}$ *be a family of message distributions over* $\{0,1\}^k$ *for all* k*. Let* $(\mathsf{Send}, \mathsf{Rec})$ *be a commitment and* Sim_k *be a simulator. We say that* $(\mathsf{Send}, \mathsf{Rec})$ *is secure against selective opening attacks for* $(\mathcal{I}, \underline{\mathcal{B}})$ *if for all* $k \leq \mathrm{poly}(n)$*:*

- *Let* $\langle \mathsf{Send}^k(\underline{b}), \mathsf{Rec}' \rangle = (\tau^k, I, \{(b_i, \omega_i)\}_{i \in I})$ *be the complete interaction between* Rec' *and the honest sender, including the commit-phase transcript* τ^k*, the subset* I *of coordinates to be opened and the openings* $(b_i, \omega_i)_{i \in I}$*.*
- *Let* $(\mathsf{Sim}_k^{\mathsf{Rec}'} \mid \underline{b})$ *denote the following: first,* $\mathsf{Sim}_k^{\mathsf{Rec}'}$ *interacts with* Rec' *(without knowledge of* $\underline{b}$*) and outputs a subset* I *of bits to be opened. Then* Sim_k *is given* $\{b_i\}_{i \in I}$*. Using this,* Sim_k *interacts with* Rec' *some more and outputs a commit-phase transcript* τ^k*, the set* I*, and the openings* $\{(b_i, \omega_i)\}_{i \in I}$*.*
- *It holds that* $(\mathsf{Sim}_k^{\mathsf{Rec}'} \mid \underline{b}) \approx_c \langle \mathsf{Send}^k(\underline{b}), \mathsf{Rec}' \rangle$ *where* $\underline{b} \leftarrow_R \underline{\mathcal{B}}$*.*

Definition 5. *We say that* $(\mathcal{I}, \underline{\mathcal{B}})$ *is* non-trivial *if (the uniform distribution over)* $\mathcal{I}, \underline{\mathcal{B}}$ *are efficiently samplable, and (1)* $|\mathcal{I}| = k^{\omega(1)}$ *and (2)* $\Pr_{I \leftarrow_R \mathcal{I}}[H_\infty(\underline{\mathcal{B}}_I) \geq 1/\mathrm{poly}(n)] \geq 1/\mathrm{poly}(n)$.

Here $\underline{\mathcal{B}}_I$ is the joint distribution of bits $\underline{\mathcal{B}}_i$ for $i \in I$. Property 1 says that if the receiver asks for a random set in $\mathcal{I}$ to be opened, then the sender cannot guess the set with noticeable probability. This restriction is natural because in many contexts if the sender can guess the set to be opened then it can cheat in the larger protocol where the commitment is being used (*e.g.* in a zero knowledge proof). Property 2 says that with noticeable probability over the choice of I, there is non-negligible entropy in the bits revealed. This is very natural as otherwise any receiver is trivially simulable since it always sees the same constant bits.

Stronger definitions of hiding. Our definitions are chosen to be as weak as possible in order to make our lower bounds stronger. Nevertheless, our positive results also satisfy a stronger definition of security, where security holds simultaneously for all $\mathcal{I}, \underline{\mathcal{B}}$. For such a notion, we prepend STR to the name of the security property (*e.g.* STR-PAR-SB).

Definition 6 (Security game for k-fold concurrent composition). *Identical to the parallel case, except that the receiver has the power to schedule messages as he wishes, rather than sending them in parallel. In addition, we allow the receiver to pick the set I incrementally subject to the constraint that at the end, $I \in \mathcal{I}$. For example, the receiver can generate one commit-phase transcript, ask the sender to decommit that instance, then use this information in its interaction to generate the second commit-phase transcript, and so forth.*

Definition 7 (Hiding under selective opening, concurrent composition) *Same as the parallel case, except that the simulator can incrementally ask for the values $(b_i)_{i \in I}$ before completing all commit-phase executions, subject to the constraint that at the end $I \in \mathcal{I}$.*

Discussion of definitional choices: One could weaken Definition 6 to require that although all the commit-phase transcripts may be generated concurrently, the openings happen simultaneously. Indeed, this was the definition used in [3]. We do not work with this weakening because it makes the definition not truly concurrent: forcing all the openings to occur simultaneously "synchronizes" the sessions.

3 Constructions

Di Crescenzo and Ostrovsky [8] (see also [9]) showed how to build an *equivocable* commitment scheme. Equivocable means that for every cheating receiver Rec', there exists a simulator that generates a commit-phase transcript that is computationally indistinguishable from a real transcript, but which the simulator can decommit to both 0 and 1. Equivocation seems even stronger than STR-PAR-CBCH

security, except that STR-PAR-CBCH explicitly requires security to hold in many parallel sessions. Although it is not clear how to generically convert any stand-alone equivocable commitment to an equivocable commitment that is composable in parallel/concurrently, the particular construction of Di Crescenzo and Ostrovsky can be composed by using a suitable preamble.

The DO construction consists of a preamble, which is a coin-flipping scheme that outputs a random string, followed by running Naor's commitment based on OWF [18] using the random string of the preamble as the receiver's first message. Depending on how the preamble is constructed, we get either a STR-PAR-CBCH, STR-PAR-SB, or STR-CC-SB commitment. Therefore, Theorem 1 follows by Theorem 6 and Theorem 8 about the following protocol.

Protocol 4 ([8, 9, 18]). *Sender's bit: b. Let $G : \{0,1\}^n \to \{0,1\}^{3n}$ be a PRG.*

> ***Preamble:*** *Use a coin-flipping protocol to obtain $\sigma \leftarrow_R \{0,1\}^{3n}$.*
> ***Commit phase:*** *The sender picks random $s \leftarrow_R \{0,1\}^n$ and sends $c = (\sigma \wedge b) \oplus G(s)$ (where we use the notation $(\sigma \wedge b)_i = \sigma_i \wedge b$).*
> ***Decommit phase:*** *The sender sends b, s. Receiver checks that $c = (\sigma \wedge b) \oplus G(s)$.*

We now present three different preambles that when used in the protocol above, provide STR-PAR-CBCH, STR-PAR-SB, and STR-CC-SB security, respectively.

Protocol 5 ([8]). *Preambles for STR-PAR-CBCH or STR-PAR-SB:*

1. *Using the non-interactive stand-alone CH commitment based on one-way permutations (to achieve STR-PAR-CBCH) or a t-round stand-alone SH commitment (to achieve STR-PAR-SB), the receiver sends a commitment to $\alpha \leftarrow_R \{0,1\}^{3n}$.*
2. *The sender replies with $\beta \leftarrow_R \{0,1\}^{3n}$.*
3. *The receiver opens α.*
4. *Output $\sigma = \alpha \oplus \beta$.*

Theorem 6 ([8]). *Protocol 4 with the STR-PAR-CBCH (resp. STR-PAR-SB) version of the preamble of Protocol 5 gives a STR-PAR-CBCH (resp. STR-PAR-SB) commitment.*

Proof (Proof sketch of Theorem 6). We include a proof sketch for the sake of completeness, and refer the reader to [18, 12, 8] for full proofs. The binding properties are easy to verify, given the fact that Naor's commitment scheme is statistically binding.

The following simulator proves security against selective opening attacks for both the computational and statistical binding variants. Consider the k-fold repetition $\mathsf{Send}^k, \mathsf{Rec}^k$ of the protocol. Following the proof of Goldreich and Kahan [12], one can construct a simulator such that, by rewinding the first step of the preamble (*i.e.* Step 1 of Protocol 5), can learn the value of the $\alpha_1, \ldots \alpha_k$ used in each of the k parallel sessions. Care must be taken to ensure this finishes in expected polynomial time, but the same technique as in [12] works in our setting and we refer the reader to that paper for details.

Now for each $i \in [k]$ in the i'th session the simulator can sample $s_0, s_1 \leftarrow_R \{0,1\}^n$ and reply with $\beta_i = G(s_0) \oplus G(s_1) \oplus \alpha_i$. This sets $\sigma_i = G(s_0) \oplus G(s_1)$. Then the simulator sets $c = G(s_0)$. Now the simulator can decommit to both 0 (by sending s_0) and to 1 (by sending s_1). ■

Protocol 7 ([21]). *Preamble for STR-CC-SB:*

1. *The receiver picks* $\alpha \leftarrow_R \{0,1\}^{3n}$ *and for* $\ell = \omega(\log n)$ *picks* $\alpha^0_{i,j} \leftarrow_R \{0,1\}^{3n}$ *for* $i, j \in [\ell]$ *and sets* $\alpha^1_{i,j} = \alpha \oplus \alpha^0_{i,j}$. *The receiver commits in parallel to* $\alpha, \alpha^0_{i,j}, \alpha^1_{i,j}$ *via a* t*-round statistically hiding commitment.*
2. *For each* $j = 1$ *to* ℓ *sequentially, do the following:*
   (a) *The sender sends* $q_1, \ldots, q_\ell \leftarrow_R \{0,1\}$.
   (b) *The receiver opens the commitment to* $\alpha^{q_i}_{i,j}$ *for all* $i \in [\ell]$.
3. *The sender sends* $\beta \leftarrow_R \{0,1\}^{3n}$.
4. *The receiver opens the commitment to* $\alpha, \alpha^0_{i,j}, \alpha^1_{i,j}$ *for all* $i, j \in [\ell]$.
5. *The sender checks that indeed* $\alpha = \alpha^0_{i,j} \oplus \alpha^1_{i,j}$ *for all* $i, j \in [\ell]$. *If so output* $\sigma = \alpha \oplus \beta$, *otherwise abort.*

Theorem 8 ([21, 22]). *Protocol 4 using the preamble of Protocol 7 gives a STR-CC-SB commitment.*

Proof. Binding is straightforward. For hiding, observe that this is the preamble of the concurrent zero knowledge proof of Prabhakaran *et al.* [21]. They prove the following:

Theorem 9 (Theorem 5.2 of [21], informal). *There is black-box simulator strategy that, given access to any efficient receiver for Protocol 7 with any concurrent scheduling, outputs with high probability in every session a string* α before *Step* 3 *such that the receiver opens to* α *in step* 5.

Namely, [21] show that by using an appropriate rewinding schedule, the simulator can obtain the value of α in *all* of the concurrent executions before the sender is supposed to send β, regardless of how the receiver schedules the messages. Once the simulator knows α, one can apply the simulator strategy of [12, 8], as in the proof sketch of Theorem 6. ■

4 Optimality of Constructions

We now define our main tool for proving lower bounds, *equivocal senders*. Intuitively, an equivocal sender must run its commit phase without knowing what it is committing to, so if it can cause the receiver to accept with non-negligible probability, then it must be able to open its commitments in many ways.

4.1 Equivocal Senders

For a pair of algorithms $T = (T_{com}, T_{decom})$, define the following game:

1. $\langle T_{com}, \mathsf{Rec}^k \rangle = (\tau^k, I, \mathsf{state}_{com})$. Here, state_{com} is the internal state of T_{com} to be transmitted to T_{decom}. I is the set Rec^k asks to be opened. Notice T_{com} runs without knowledge of $\underline{b}$, hence T is "equivocal" during the commit phase.
2. $T_{decom}(\underline{b}, \tau^k, I, \mathsf{state}_{com}) = \{(b_i, \mathsf{open}_i)\}_{i \in I}$.

The overall transcript is $(\langle T, \mathsf{Rec}^k \rangle \mid \underline{b}, \mathsf{NoAbort}_T) = (\tau^k, I, \{(b_i, \mathsf{open}_i)\}_{i \in I})$, where $\mathsf{NoAbort}_T$ denotes the event that T does not abort. Say that $(\tau^k, I, \mathsf{state}_{com})$ is δ-*openable* if with probability at least δ over the choice of $\underline{b}$, Rec^k accepts $(\tau^k, I, \{(b_i, \mathsf{open}_i)\}_{i \in I})$, where $\{(b_i, \mathsf{open}_i)\}_{i \in I} = T_{decom}(\underline{b}, \tau^k, I, \mathsf{state}_{com})$..

Definition 8 (Equivocal sender). *We say that* $T = (T_{com}, T_{decom})$ *is a* k*-equivocal sender for* $(\mathsf{Send}, \mathsf{Rec}, \mathsf{Sim}_k)$ *if it holds that*

$$\Pr[(\tau^k, I, \mathsf{state}_{com}) = \langle T_{com}, \mathsf{Rec}^k \rangle \textit{ is } (1 - n^{-\omega(1)})\textit{-openable} \wedge \mathsf{NoAbort}_T] \geq 1/\mathrm{poly}(n)$$

Using equivocal senders to break binding. To prove our impossibility results, we will apply the following theorem, which says that the existence of equivocal senders imply that a commitment is not secure.

Theorem 10. *Fix any non-trivial* $(\mathcal{I}, \underline{\mathcal{B}})$ *and* k*-fold repeated commitment scheme* $(\mathsf{Send}^k, \mathsf{Rec}^k)$ *with a simulator* Sim_k *that proves computational hiding. If this commitment has a* k*-equivocal sender* $T = (T_{com}, T_{decom})$ *for any* $k \leq \mathrm{poly}(n)$*, then this commitment cannot be statistically binding. If furthermore* T *is efficient, then this commitment cannot be computationally binding.*

Proof. The idea is to convert a k-equivocal T sender into a sender Send' that breaks binding in a single execution of the commitment, Send' emulates T internally and chooses one of the k parallel instances to insert its interaction with the real receiver Rec. By the non-triviality of $(\mathcal{I}, \underline{\mathcal{B}})$, with high probability over $I \leftarrow_{\mathrm{R}} \mathcal{I}$ the coordinates in I have significant min-entropy, and in particular some coordinate must have significant min-entropy. Therefore if Send' picks this coordinate, then since T is able to open its commitment with non-trivial probability for $I \leftarrow_{\mathrm{R}} \mathcal{I}$ and $\underline{b} \leftarrow_{\mathrm{R}} \underline{\mathcal{B}}$, it follows that Send' can open its commitment to both 0 and 1 with non-negligible probability.

We now proceed formally by constructing a malicious sender Send' and proving that this sender breaks binding.

Algorithm 11

Malicious sender Send'*, interacting with a single honest receiver* Rec*:*

1. *Pick a random* j*. For each* $j' \neq j$*, sample random coins* $\omega^{(j')}$ *to run an honest receiver.*
2. *Respond to the* i*'th message* β_i *from* Rec *as follows.*
   (a) *If* $i > 1$*, let* $(\alpha^{(1)}_{[i-1]}, \ldots, \alpha^{(k)}_{[i-1]})$ *be* T_{com}*'s response from previous queries.*
   (b) *For* $j' \neq j$*, compute* $\beta_i^{(j')} = \mathsf{Rec}(\alpha^{(j')}_{[i-1]}; \omega^{(j')})$*. Set* $\beta_i^{(j)} = \beta_i$*.*
   (c) *Feed* $(\beta_i^{(1)}, \ldots, \beta_i^{(k)})$ *to* T_{com} *to obtain response* $(\alpha^{(1)}_{[i]}, \ldots, \alpha^{(k)}_{[i]})$ *(assuming* T_{com} *does not abort).*

(d) Forward $\alpha_i^{(j)}$ *back to* Rec.

3. *If* T_{com} *does not abort,* Send' *successfully generates a commit-phase transcript distributed according to* $\langle T_{com}, \mathsf{Rec}^k \rangle$. Send' *picks a random* $I \leftarrow_R \mathcal{I}$ *to be opened.*
4. *If* $j \notin I$, Send' *aborts. Otherwise, it independently picks two* $\underline{b}, \underline{b}' \leftarrow_R \underline{\mathcal{B}}$, *and runs* $T_{decom}(\underline{b}, I)$ *to obtain a decommitment for* $(b_i)_{i \in I}$ *and runs* $T_{decom}(\underline{b}', I)$ *to obtain openings for* $(b'_i)_{i \in I}$. *In particular, the malicious sender obtains openings for* b_j *and* b'_j.

Analyzing Send'**:** By hypothesis, T is a $(k, \varepsilon, 1 - n^{-\omega(1)})$-equivocal server for some $\varepsilon = 1/\mathrm{poly}(n)$. This implies that with probability at least ε, $\langle T_{com}, \mathsf{Rec}^k \rangle$ produces an $(1 - n^{-\omega(1)})$-openable $(\tau^k, I, \mathsf{state}_{com})$. Therefore, since the probability of producing an accepting opening for a random $\underline{b}$ at least $(1 - n^{-\omega(1)})$, it holds with probability at least $\varepsilon(1 - n^{-\omega(1)})^2$ that Rec^k accepts both openings $T_{decom}(\underline{b}, \tau^k, I, \mathsf{state}_{com})$ and $T_{decom}(\underline{b}', \tau^k, I, \mathsf{state}_{com})$.

Since $(\mathcal{I}, \underline{\mathcal{B}})$ is non-trivial, it follows that $\Pr_{\underline{b}, \underline{b}', I}[\forall\ i\ \in\ I,\ b_i = b'_i] \leq n^{-\omega(1)}$. Therefore with probability $\varepsilon(1 - n^{-\omega(1)})^2 - n^{-\omega(1)}$, T produces accepting openings for $\underline{b}$ and $\underline{b}'$ and furthermore there exists i such that $\underline{b}_i \neq \underline{b}'_i$. Since the sender picked at random the coordinate j that contains the real interaction, with probability $1/k$ it chooses $j = i$ and therefore with non-negligible probability produces decommitments for both 0 and 1 in an interaction with the real receiver, breaking binding. ∎

4.2 Impossibility Results for Parallel Composition

We present the proofs for the case of 3-round PAR-CBCH and 4-round PAR-SB commitments, while the cases in Theorem 3 are deferred to the full version.

We construct equivocal senders using the strategy of Goldreich and Krawczyk [13]. Intuitively, the idea is to construct a sender T whose output distribution is the same as $\mathsf{Sim}_k^{\mathsf{Rec}_h}$. Here, Rec_h is intuitively a cheating receiver that, for each sender message, uses its hash function h to generate a response that looks completely random, and therefore Sim_k gains no advantage by rewinding Rec_h. From this cheating property, we will be able to conclude that T satisfies Definition 8.

Goldreich and Krawczyk [13] observe that we can make the following simplifying assumptions *w.l.o.g.*: **(1)** Sim_k makes exactly $p(n) = \mathrm{poly}(n)$ queries to its receiver black box, **(2)** all queries made by Sim_k are distinct, and **(3)** Sim_k always outputs a transcript τ^k that consists of queries it made to the receiver and the corresponding receiver responses.

The following lemma from [13] says that simply by guessing uniformly at random, one can pick with some noticeable probability the queries/responses that the simulator outputs as its final transcript.

Lemma 1 ([13]). *Fix a black-box simulator* Sim_k *for a protocol with* t *sender messages, and suppose* Sim_k *makes* $p(n)$ *queries. Draw* $u_1, \ldots, u_t \leftarrow_R [p(n)]$, *then with probability* $\geq 1/p(n)^t$, *the final transcript output by* Sim_k *consists of the* $u_1, \ldots, u_t$*'th queries (along with the corresponding receiver responses).*

3-Round Commitments

Theorem 12. *For all non-trivial $(\mathcal{I}, \underline{\mathcal{B}})$ and relative to any oracle, there exists no 3-round PAR-CBCH commitment protocol secure for $(\mathcal{I}, \underline{\mathcal{B}})$.*

Proof. We construct a polynomial-time k-equivocal sender for $(\mathsf{Send}, \mathsf{Rec})$ for $k = n$. By Theorem 10, this contradicts the binding property of the commitment.

Algorithm 13
Equivocal sender $T = (T_{com}, T_{decom})$ for 3-round commitments:

1. *T_{com} picks $u_1, u_2 \leftarrow_R [p(n)]$.*
2. *T_{com} internally runs Sim_k, answering its queries as follows:*
   - *For the u_1, u_2'th queries, if the u_1'th query is a first sender message α_1 and the u_2'th query is a second sender message $\alpha_{[2]}$ that extends α_1, then T_{com} forwards them to the real receiver and forwards the receiver's responses to the simulator. Otherwise, T_{com} aborts.*
   - *For all other queries: if the query is α_1, then T_{com} returns $\mathsf{Rec}^k(\alpha_1; \omega)$ for uniform ω. If the query is $\alpha_{[2]}$ then T returns a random $I \leftarrow_R \mathcal{I}$.*
3. *When Sim_k requests that a subset I of bits be revealed, T_{com} checks to see if I equals the set that the real receiver asked to be opened. If not, T_{com} aborts.*
4. *In the opening phase, T_{decom} receives $\underline{b}$ and feeds $(b_i)_{i \in I}$ to the simulator and obtains $(\tau^k, I, (b_i, \mathsf{open}_i)_{i \in I})$. T_{decom} checks that τ^k and I consists of queries to/from the real receiver, and if not aborts. Otherwise it outputs these openings.*

Analyzing equivocal sender T. It is clear that T runs in polynomial time. Lemma 1 implies that with probability $1/p(n)^2$, Sim_k picks the set to be revealed I using the receiver's responses to the guessed queries u_1, u_2.

Lemma 2. *The probability that Sim_k makes two queries $\alpha_{[2]}, \alpha'_{[2]}$ that are both answered with the same I is negligible*

This claim holds because $|\mathcal{I}| = n^{\omega(1)}$ and Sim_k makes at most $p(n) = \mathrm{poly}(n)$ queries. Lemma 2 implies that when T emulates Sim_k, Sim_k cannot pick I using the real receiver's messages but then find a different commit-phase transcript that leads to the same set I. Therefore the probability that T does not abort and outputs the queries to and responses from the real receiver is at least $1/p(n)^2 - n^{-\omega(1)} \geq 1/\mathrm{poly}(n)$.

Lemma 3. Rec^k *accepts* $(\langle T, \mathsf{Rec}^k \rangle \mid \underline{b}, \mathsf{NoAbort}_T)$ *with overwhelming probability.*

This claim combined with the above assertion that T does not abort with non-negligible probability implies that T satisfies Definition 8.

We now prove Lemma 3 by comparing the output of T to $(\mathsf{Sim}_k^{\mathsf{Rec}_h} \mid \underline{b})$ where Rec_h is defined as follows: h is a $p(n)$-wise independent hash function, it responds to first sender queries α_1 by computing $\beta_1 = \mathsf{Rec}(\alpha_1; h(\alpha_1))$ and to second sender queries $\alpha_{[2]}$ by sampling uniform $I \leftarrow_R \mathcal{I}$ using $h(\alpha_{[2]})$ as random coins.

As observed by [13], $(\langle T, \mathsf{Rec} \rangle \mid \underline{b}, \mathsf{NoAbort}_T) = (\mathsf{Sim}_k^{\mathsf{Rec}_h} \mid \underline{b})$ for a uniform choice of h. Since Rec_h is efficient, by the hiding property this is indistinguishable

from $\langle \mathsf{Send}^k(\underline{b}), \mathsf{Rec}_h \rangle$. This in turn is equal to a true interaction $\langle \mathsf{Send}^k(\underline{b}), \mathsf{Rec}^k \rangle$, since by the definition of Rec_h the two receivers Rec_h and Rec^k behave identically when there is no rewinding. Since Rec^k always accepts a real interaction, therefore Rec^k accepts $(\langle T, \mathsf{Rec} \rangle \mid \underline{b}, \mathsf{NoAbort}_T)$ with overwhelming probability. ■

4-Round Commitments

Theorem 14. *For all non-trivial $(\mathcal{I}, \underline{\mathcal{B}})$ and relative to any oracle, there exists no 4-round PAR-SB commitment protocol secure for $(\mathcal{I}, \underline{\mathcal{B}})$.*

Proof (Proof). As before, it suffices to construct a k-equivocal sender for $k = n$.

Algorithm 15
Equivocal sender $T = (T_{com}, T_{decom})$ for 4-round PAR-SB commitments

1. *T_{com} picks $u_1, u_2 \leftarrow_R [p(n)]$.*
2. *T_{com} receives the first message β_1 from the receiver.*
3. *T_{com} internally runs Sim_k, answering its queries as follows:*
   - *For the simulator's u_1, u_2'th queries, if the u_1'th query is a first sender message α_1 and the u_2'th query is a second sender message $\alpha_{[2]}$ that extends α_1, then T_{com} forwards them to the real receiver and forwards the receiver's responses to the simulator. Otherwise, T_{com} aborts.*
   - *For all other queries: if the query is α_1 then T_{com} samples a random $\omega' \leftarrow_R \{\omega \mid \mathsf{Rec}(\perp; \omega) = \beta_1\}$ and returns $\beta_2 = \mathsf{Rec}(\beta_1, \alpha_1; \omega')$ to the simulator. If the query is $\alpha_{[2]}$ then the simulator picks a random $I \leftarrow_R \mathcal{I}$ and returns it to the simulator.*
4. *When Sim_k requests that a subset I of bits be revealed, T_{com} checks to see if I equals the set that the real receiver asked to be opened. If not, T_{com} aborts.*
5. *In the opening phase, T_{decom} receives $\underline{b}$ and feeds $(b_i)_{i \in I}$ to the simulator and obtains $(\tau^k, I, (b_i, \mathsf{open}_i)_{i \in I})$. T_{decom} checks that τ^k and I consists of queries to/from the real receiver, and if not aborts. Otherwise it outputs the openings.*

Analyzing equivocal sender T. T may not run in polynomial time because sampling $\omega' \leftarrow_R \{\omega \mid \beta_1 = \mathsf{Rec}(\perp; \omega)\}$ may be inefficient. This implies the sender breaking binding given by Theorem 10 may be inefficient, which is why we can only handle PAR-SB commitments.

Applying Lemma 1, T does not abort with probability $\geq 1/p(n)^2$. Lemma 2 applies here for the same reason as in the proof of Theorem 12, therefore it holds with probability $1/p(n)^2 - n^{-\omega(1)} \geq 1/\mathrm{poly}(n)$ that T's messages to/from the receiver are exactly those in the output of its emulation of Sim_k.

We claim that Lemma 3 holds in this case as well, which would imply that T satisfies Definition 8. We prove Lemma 3 in this setting by comparing the output of T to $(\mathsf{Sim}_k^{\mathsf{Rec}_h^{\omega_1,\ldots,\omega_s}} \mid \underline{b})$, where we use the cheating receiver strategy $\mathsf{Rec}_h^{\omega_1,\ldots,\omega_s}$ defined by Katz [17]: s will be set below, and the ω_i are random coins for the honest receiver algorithm such that $\mathsf{Rec}(\perp; \omega_i) = \mathsf{Rec}(\perp; \omega_j)$ for all $i, j \in [s]$, and h is a $p(n)$-wise independent hash function with output range $[s]$. The first message of

$\mathsf{Rec}_h^{\omega_1,\ldots,\omega_s}$ is $\beta_1 = \mathsf{Rec}(\perp;\omega_1)$ and given sender message α_1, the second message is $\beta_2 = \mathsf{Rec}(\beta_1,\alpha_1;\omega_{h(\beta_1,\alpha_1)})$. Given sender messages $\alpha_{[2]}$, the set I to be opened is sampled using $\omega_{h(\beta_{[2]},\alpha_{[2]})}$ as random coins.

As observed in [17], for $s = 50p(n)^2/\delta$ it holds that the statistical distance between $(\langle T, \mathsf{Rec}^k\rangle \mid \underline{b}, \mathsf{NoAbort}_T)$ and $(\mathsf{Sim}_k^{\mathsf{Rec}_h^{\omega_1,\ldots,\omega_s}} \mid \underline{b})$ is at most δ, where the randomness is over uniform $p(n)$-wise independent h, uniform ω_1 and uniform $\omega_2,\ldots,\omega_s$ conditioned on $\mathsf{Rec}(\perp;\omega_j) = \mathsf{Rec}(\perp;\omega_1)$ for all $j \in [s]$. By the commitment's hiding property this is indistinguishable from $\langle \mathsf{Send}^k(\underline{b}), \mathsf{Rec}_h^{\omega_1,\ldots,\omega_s}\rangle$, which in turn is equal to $\langle \mathsf{Send}^k(\underline{b}), \mathsf{Rec}^k\rangle$ by the definition of $\mathsf{Rec}_h^{\omega_1,\ldots,\omega_s}$. Finally, since Rec^k always accepts a real interaction, therefore it accepts $(\langle T, \mathsf{Rec}^k\rangle \mid \underline{b}, \mathsf{NoAbort}_T)$ with probability $1 - \delta - n^{-\omega(1)}$.

We can apply the above argument for any $\delta \geq 1/\mathrm{poly}(n)$ to conclude that Rec^k accepts $(\langle T, \mathsf{Rec}^k\rangle \mid \underline{b}, \mathsf{NoAbort}_T)$ with probability $1 - \delta - n^{-\omega(1)}$ for all $\delta \geq 1/\mathrm{poly}(n)$.

Therefore Rec^k must accept $(\langle T, \mathsf{Rec}^k\rangle \mid \underline{b}, \mathsf{NoAbort}_T)$ with probability $1 - n^{-\omega(1)}$ and so T satisfies Definition 8. ∎

4.3 PAR-SB Commitments Imply (Stand-Alone) SH Commitments

To prove Item 2 of Theorem 2, we show that PAR-SB commitments can be used to generate a gap between real and accessible entropy [16]. Then we apply the transformation of [16] that converts an entropy gap into a statistically hiding commitment.

To simplify the statement of our result, we assume that $\mathcal{I} = 2^{[k]}$ and $\underline{\mathcal{B}} = U_k$; see the full version for the general treatment. We also defer the definitions of real min-entropy and accessible max-entropy and the formal proof of the main technical lemma (Lemma 4) to the full version.

Theorem 16. *For $(\mathcal{I} = 2^{[k]}, \underline{\mathcal{B}} = U_k)$, if there exists $O(1)$-round $(\mathsf{Send}, \mathsf{Rec})$ that is PAR-SB secure for $(\mathcal{I}, \underline{\mathcal{B}})$, then there exists $O(1)$-round statistically hiding commitments.*

Proof. Assume without loss of generality that Rec^k sends all his random coins at the end of the opening phase, and that Rec uses m random coins in a single stand-alone instance.

Lemma 4. *Rec^k has real min-entropy at least $km(1 - 1/k^{1/3})$ and has context-independent accessible max-entropy $\leq km - k/4$.*

Lemma 4 implies there is an entropy gap, and so the theorem follows by combining it with the black-box construction of statistically hiding commitments from entropy gaps given by Lemmas 6.7, 4.7, and 4.18 of [16]. ∎

Proof (Proof of Lemma 4.). The real min-entropy part of the claim follows from the definitions and amplification by parallel repetition (Proposition 3.8 in [16]). For the accessible entropy part, we use the following:

Claim. If there exists efficient A^* sampling context-independent max-entropy $> km - k/4$ for Rec^k, then there exists a k-equivocal sender.

By Theorem 10 this contradicts the binding property of the commitment, and so A^* cannot exist. The proof follows from ideas of [16] and we defer a formal proof to the full version. ■

4.4 Impossibility Results for Concurrent Composition

Again, we state our theorem for the natural case $\mathcal{I} = 2^{[k]}$ and $\underline{\mathcal{B}} = U_k$, and defer the general statement to the full version.

Theorem 17. *For $(\mathcal{I} = 2^{[k]}, \underline{\mathcal{B}} = U_k)$, and relative to any oracle, no $o(\log n / \log \log n)$-round commitment is CC-CBCH secure for $\mathcal{I}, \underline{\mathcal{B}}$.*

Proof. **Building the equivocal sender:** The message schedule we use is exactly that of [7], which we call Σ, and is defined in the full version. The high-level idea of T, also adapted from [7], is to execute Sim_k and to insert the real receiver's messages into one session j chosen at random, and where T aborts if the simulator tries to rewind queries in session j. Messages for other sessions are computed using the partial transcripts generated by the simulator so far. We defer a more explicit description to the full version.

Analyzing equivocal sender T: [7] prove the following lemma:

Lemma 5 ([7], informal). *It holds with non-negligible probability that there exists a "good session" in the execution of $\mathsf{Sim}_k^{\mathsf{Rec}_\Sigma^k}$,* i.e. *a session where Sim_k does not rewind Rec_Σ^k.*

The only place where T may abort is if in its emulation of Sim_k, the simulator tries to rewind the receiver in session j. Therefore, with probability $1/k$, T inserts the real receiver into the good session that is guaranteed to exist by Lemma 5 with non-negligible probability. Furthermore, since the k concurrent simulation is indistinguishable from a real interaction, it follows that Rec_Σ^k accepts $(\langle T, \mathsf{Rec}_\Sigma^k \rangle \mid b, \mathsf{NoAbort}_T)$ with overwhelming probability. ■

Acknowledgements

The author would like to thank Dennis Hofheinz and Salil Vadhan for helpful conversations.

References

[1] Barak, B.: How to go beyond the black-box simulation barrier. In: Proc. 42nd FOCS, pp. 106–115. IEEE, Los Alamitos (2001)
[2] Beaver, D.: Adaptive zero knowledge and computational equivocation (extended abstract). In: Proc. STOC 1996, pp. 629–638. ACM, New York (1996)

[3] Bellare, M., Hofheinz, D., Yilek, S.: Possibility and impossibility results for encryption and commitment secure under selective opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009)
[4] Brassard, G., Crépeau, C.: Zero-knowledge simulation of boolean circuits. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 223–233. Springer, Heidelberg (1987)
[5] Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. J. Comput. Syst. Sci. 37(2), 156–189 (1988)
[6] Brassard, G., Crépeau, C., Yung, M.: Everything in NP can be argued in *perfect* zero-knowledge in a *bounded* number of rounds. In: Quisquater, J.-J., Vandewalle, J. (eds.) EUROCRYPT 1989. LNCS, vol. 434, pp. 192–195. Springer, Heidelberg (1990)
[7] Canetti, R., Kilian, J., Petrank, E., Rosen, A.: Black-box concurrent zero-knowledge requires (almost) logarithmically many rounds. SIAM J. Comput. 32(1), 1–47 (2003)
[8] Di Crescenzo, G., Ostrovsky, R.: On concurrent zero-knowledge with pre-processing (Extended abstract). In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 485–502. Springer, Heidelberg (1999)
[9] Di Crescenzo, G., Ishai, Y., Ostrovsky, R.: Non-interactive and non-malleable commitment. In: Proc. STOC 1998, pp. 141–150. ACM, New York (1998)
[10] Dwork, C., Naor, M., Reingold, O., Stockmeyer, L.: Magic functions. J. ACM 50(6), 852–921 (2003)
[11] Fischlin, M.: Trapdoor Commitment Schemes and Their Applications. Ph.D. Thesis (Doktorarbeit), Department of Mathematics, Goethe-University, Frankfurt, Germany (2001)
[12] Goldreich, O., Kahan, A.: How to construct constant-round zero-knowledge proof systems for NP. Journal of Cryptology 9(3), 167–189 (1996)
[13] Goldreich, O., Krawczyk, H.: On the composition of zero-knowledge proof systems. SIAM J. of Com. 25(1), 169–192 (1996); Preliminary version appeared In: Paterson, M. (ed.) ICALP 1990. LNCS, vol. 443. Springer, Heidelberg (1990)
[14] Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. Journal of the ACM 38(3), 691–729 (1991); Preliminary version in FOCS 1986
[15] Haitner, I., Hoch, J.J., Reingold, O., Segev, G.: Finding collisions in interactive protocols - a tight lower bound on the round complexity of statistically-hiding commitments. In: Proc. FOCS 2007, pp. 669–679 (2007)
[16] Haitner, I., Reingold, O., Vadhan, S., Wee, H.: Inaccessible entropy. In: Proc. STOC 2009, pp. 611–620. ACM, New York (2009)
[17] Katz, J.: Which languages have 4-round zero-knowledge proofs? In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 73–88. Springer, Heidelberg (2008)
[18] Naor, M.: Bit commitment using pseudorandomness. Journal of Cryptology 4(2), 151–158 (1991); Preliminary version In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 128–136. Springer, Heidelberg (1990)
[19] Pass, R., Wee, H.: Black-box constructions of two-party protocols from one-way functions. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 403–418. Springer, Heidelberg (2009)
[20] Pass, R., Tseng, W.-L.D., Wikström, D.: On the composition of public-coin zero-knowledge protocols. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 160–176. Springer, Heidelberg (2009)
[21] Prabhakaran, M., Rosen, A., Sahai, A.: Concurrent zero knowledge with logarithmic round-complexity. In: Proc. 43rd FOCS, pp. 366–375. IEEE, Los Alamitos (2002)

[22] Rosen, A.: Concurrent Zero-Knowledge - With Additional Background by Oded Goldreich. Information Security and Cryptography. Springer, Heidelberg (2006)
[23] Wee, H.: On statistically binding trapdoor commitments from one-way functions (2008) (manuscript)
[24] Zhang, Z., Cao, Z., Zhu, H.: Constant-round adaptive zero knowledge proofs for NP (2009) (manuscript)