

Inaccessible Entropy*

Iftach Haitner[†] and Omer Reingold[‡] and Salil Vadhan[§] and Hoeteck Wee[¶]

June 11, 2009

Abstract

We put forth a new computational notion of entropy, which measures the (in)feasibility of sampling high entropy strings that are consistent with a given protocol. Specifically, we say that the i th round of a protocol (A, B) has *accessible entropy* at most k , if no polynomial-time strategy A^* can generate messages for A such that the entropy of its message in the i th round has entropy greater than k when conditioned both on prior messages of the protocol and on prior coin tosses of A^* . We say that the protocol has *inaccessible entropy* if the total accessible entropy (summed over the rounds) is noticeably smaller than the real entropy of A 's messages, conditioned only on prior messages (but not the coin tosses of A).

As applications of this notion, we

- Give a much simpler and more efficient construction of statistically hiding commitment schemes from arbitrary one-way functions.
- Prove that constant-round statistically hiding commitments are necessary for constructing constant-round zero-knowledge proof systems for NP that remain secure under parallel composition (assuming the existence of one-way functions).

Keywords: computational complexity, cryptography, commitment schemes, interactive hashing, zero knowledge, one-way functions

*An extended abstract of this paper will appear in *STOC '09* [HRVW].

[†]Microsoft Research, New England Campus. E-mail: iftach@microsoft.com.

[‡]Department of Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel. E-mail: omer.reingold@weizmann.ac.il. Supported by US-Israel BSF grant 2006060.

[§]School of Engineering and Applied Sciences and Center for Research on Computation and Society, Harvard University. E-mail: salil@eecs.harvard.edu. Work done in part while visiting U.C. Berkeley, supported by the Miller Institute for Basic Research in Science and a Guggenheim Fellowship. Also supported by NSF grant CNS-0831289 and US-Israel BSF grant 2006060.

[¶]Queens College, CUNY. E-mail: hoeteck@cs.qc.cuny.edu. Part of this work was done while a post-doc at Columbia University, supported in part by NSF Grants CNS-0716245 and SBE-0245014.

1 Introduction

Computational analogues of information-theoretic notions have given rise to some of the most interesting phenomena in the theory of computation. For example, *computational indistinguishability* [GM], which is the computational analogue of statistical distance, enabled bypassing Shannon’s impossibility results on perfectly secure encryption [Sha], and provided the basis for the computational theory of pseudorandomness [BM, Yao]. A computational analogue of entropy, known as *pseudoentropy*, introduced by Håstad, Impagliazzo, Levin, and Luby [HILL], was the key to their fundamental result establishing the equivalence of pseudorandom generators and one-way functions, and has also now become a basic concept in complexity theory and cryptography.

In this work, we introduce another computational analogue of entropy, which we call *accessible entropy*, and present several applications of it to the foundations of cryptography. Before describing accessible entropy (and a complementary notion of *inaccessible entropy*), we recall the standard information-theoretic notion of entropy and the computational notion of pseudoentropy of Håstad et al.

1.1 Entropy and Pseudoentropy

Recall that the *entropy* of a random variable X is defined to be $H(X) := \mathbb{E}_{x \leftarrow X} [\log(1/\Pr[X = x])]$, which measures the number of “bits of randomness” in X (on average). We will refer to $H(X)$ as the *real entropy* of X to contrast with the computational analogues that we study. Håstad et al. [HILL] say that a random variable X has *pseudoentropy* (at least) k if there exists a random variable Y of entropy (at least) k such that X and Y are computationally indistinguishable.

The reason that pseudoentropy is interesting and useful is that there exist random variables X whose pseudoentropy is larger than their real entropy. For example, the output of a pseudorandom generator $G : \{0, 1\}^\ell \rightarrow \{0, 1\}^n$ on a uniformly random seed has entropy at most ℓ , but has pseudoentropy n (by definition). Håstad et al. proved that in fact, from *any* efficiently samplable distribution X whose pseudoentropy is noticeably larger than its real entropy, it is possible to construct a pseudorandom generator. By showing, in addition, how to construct such a distribution X from any one-way function, Håstad et al. prove their theorem that the existence of one-way functions implies the existence of pseudorandom generators.

The notion of pseudoentropy is only useful, however, as a lower bound on the “computational entropy” in a distribution. Indeed, it can be shown that every distribution on $\{0, 1\}^n$ is computationally indistinguishable from a distribution of entropy at most $\text{poly}(\log n)$. While several other computational analogues of entropy have been studied in the literature (cf., [BSW]), all of these are also meant to serve as ways of capturing the idea that a distribution “behaves like” one of higher entropy. In this paper, we explore a way in which a distribution can “behave like” one of much *lower* entropy.

1.2 Accessible Entropy

We motivate the idea of accessible entropy with an example. Consider the following 3-message protocol between parties (A, B):

1. B selects a random function $h : \{0, 1\}^n \rightarrow \{0, 1\}^m$ from a family of collision-resistant hash functions (where $m \ll n$) and sends h to A.

2. A selects a random $x \xleftarrow{R} \{0, 1\}^n$, sets $y = h(x)$, and sends y to B.
3. A sends x to B.

Now, information-theoretically, A’s third message (namely x) has entropy at least $n - m$ conditioned on the previous messages h, y , because $y = h(x)$ reveals on m bits of information about x . However, the collision-resistance property says that given the *state* of A after the second message, there is at most one consistent value of x that A can reveal with nonnegligible probability. (Otherwise, A would be able find two distinct messages $x \neq x'$ such that $h(x) = h(x')$.) This holds even if A is replaced by any polynomial-time cheating strategy A^* . Thus, there is “real entropy” in x (conditioned on the history) but it is “computationally inaccessible” to A^* , to whom x effectively has entropy 0.

We generalize this basic idea to allow the upper bound on the “accessible entropy” to be a parameter k , and to consider both the real and accessible entropy accumulated over several rounds of a protocol. In more detail, consider an m -round protocol (A, B), and let $(B_1, A_1, \dots, B_m, A_m)$ be random variables denoting the messages sent by A and B in an interaction where their coin tosses are chosen uniformly at random. We define the *real entropy* of A when interacting with B to be

$$\sum_i H(A_i | B_1, A_1, \dots, B_i),$$

where $H(X|Y) = E_{y \leftarrow Y} [H(X|Y=y)]$ is the standard notion of conditional entropy.

To define *accessible entropy*, consider a probabilistic polynomial-time cheating strategy A^* that in each round, tosses some fresh random coins s_i , computes and sends a message a_i , and also locally outputs a string w_i that is supposed to be a “witness” to the fact that A^* is behaving consistently with the honest strategy A. Specifically, for A^* to “succeed”, each w_i should be a sequences of coin tosses for A that is consistent with all the messages a_i sent so far. For simplicity here in the introduction, we assume that A^* always outputs consistent witness strings w_i . Now, let $(B_1, S_1, A_1, W_1, \dots, B_m, S_m, A_m, W_m)$ be random variables corresponding to the view of A^* when interacting with B. Then we define the *accessible entropy* achieved by A^* to be

$$\sum_i H(A_i | B_1, S_1, A_1, W_1, \dots, B_i).$$

The key point is that now we compute the entropy conditioned not just on the previous messages exchanged, but also on everything in the local state/view of A^* prior to the i ’th round.

The collision resistance example given earlier shows there are protocols where the computationally accessible entropy is much smaller than the real Shannon entropy. Indeed, in that protocol, the real entropy of A’s messages is n (namely, the total entropy in x), but the computationally accessible entropy is at most $m + \text{neg}(n)$, where $m \ll n$ is the output length of the collision-resistant hash function. (Here we are counting the conditional entropy in all of A’s messages for simplicity, but the definitions generalize naturally if we only want to sum the conditional entropies over some subset of rounds.) Thus, in contrast to pseudoentropy, accessible entropy is useful for expressing the idea that the “computational entropy” in a distribution is *smaller* than its real entropy. We refer to the difference (real entropy) – (accessible entropy) as the *inaccessible entropy* of the protocol.

The above informal definitions are simplified or restricted compared to our actual definitions in several ways. First, we need to determine how to measure the entropy in case the adversary A^* fails to provide a consistent witness w_i . Second, in some of our results it is beneficial to work

with real *min*-entropy and/or accessible *max*-entropy rather than real and accessible Shannon entropy as defined above, and formulating conditional versions of these measures is a bit more delicate. Third, in cryptographic applications, one might also want a definition of real entropy that holds even if \mathbf{B} is replaced by a cheating strategy \mathbf{B}^* . The definitions generalize naturally to this case, but we do not consider them in this extended abstract for sake of simplicity. (In our applications below, we handle cheating strategies by applying a known compiler at the end of our constructions [HHK⁺].)

1.3 Applications

Our main applications of accessible entropy are to the construction of commitment schemes, so we begin by describing those.

Commitment Schemes. A *commitment scheme* is the cryptographic analogue of a safe. It is a 2-party protocol between a *sender* \mathbf{S} and a *receiver* \mathbf{R} that consists of two stages. The *commit stage* corresponds to putting an object in a safe and locking it. In it, the sender “commits” to a private message m . The *reveal stage* corresponds to unlocking and opening the safe. In it, the sender “reveals” the message m and “proves” that it was the value committed to in the commit stage (without loss of generality by revealing coin tosses consistent with m and the transcript of the commit stage).

Commitment schemes have two security properties. The *hiding* property informally says that at the end of the commit stage, an adversarial receiver has learned nothing about the message m , except with negligible probability. The *binding* property says that after the commit stage, an adversarial sender cannot produce valid openings for two distinct messages, except with negligible probability. Both of these security properties come in two flavors — *statistical*, where we require security even against a computationally unbounded adversary, and *computational*, where we only require security against feasible (e.g. polynomial-time) adversaries.

Statistical security is preferable to computational security, but it is impossible to have commitment schemes that are both statistically hiding and statistically binding. So instead we have to settle for one of the two properties being statistical and the other being computational. Statistically binding (and computationally hiding) commitments have been well-understood for a long time. Indeed, Naor [Nao] showed how to build a 2-message statistically binding commitment using any pseudorandom generator; and thus, in combination with the construction of pseudorandom generators from any one-way function [HILL], we obtain 2-message statistically binding commitments from the minimal assumption that one-way functions exist.

As we will describe below, our understanding of *statistically hiding* commitments has lagged behind. In this paper, we show that they are closely connected with the notion of inaccessible entropy, that is, with protocols having a gap between real entropy and accessible entropy. One direction is easy to see. Consider a statistically hiding commitment scheme in which the sender commits to a message of length k , and suppose we run the protocol with the message m chosen uniformly at random in $\{0, 1\}^k$. Then, by the statistical hiding property, the *real entropy* of the message m after the commit stage is $k - \text{neg}(n)$. On the other hand, computational binding property says that the *accessible entropy* of m after the commit stage is at most $\text{neg}(n)$.

Our main technical contribution is a converse to the above observation.

Theorem 1.1 (inaccessible entropy to commitment, informal). *If there is an efficient protocol (\mathbf{A}, \mathbf{B}) in which the real entropy of \mathbf{A} 's messages is noticeably larger than their accessible entropy,*

then statistically hiding commitment schemes exist.

Actually, since it gives us better parameters in the applications, we don't prove the above theorem for accessible Shannon entropy (as defined above), but prove it for accessible *max-entropy* (defined in the body of the paper). Indeed, for accessible max-entropy, we can preserve the property that the protocol has a constant number of rounds.

Theorem 1.2 (inaccessible entropy to commitment in constant rounds, informal). *If there is an efficient constant-round protocol (A, B) in which the real entropy of A 's messages is noticeably larger than their accessible max-entropy, then constant-round statistically hiding commitment schemes exist.*

Our proof of this theorem proceeds in a few modular steps:

1. (Entropy Equalization) First, using sequential repetition with a “random offset,” we convert the protocol into one where we *know* the real entropy in each round (rather than just knowing the total entropy), and there remains a noticeable gap between the real entropy and the accessible (max-)entropy. This step blows up the number of rounds, so for constant-round protocols, we use a different approach: we try “all possibilities” for how the real entropy is divided among the rounds, and combine the resulting commitment schemes in a standard way at the end.
2. (Gap Amplification) We repeat the protocol many times in parallel, which has the effect of (a) converting the real entropy to real *min-entropy*, and (b) amplifying the gap between the real entropy and accessible (max-)entropy.
3. (m -phase Commitment) By applying a constant-round hashing protocol in each round (based on the interactive hashing protocol of [DHRS] and universal one-way hash functions [NY, Rom]), we obtain an *m-phase commitment scheme*. This consists of m sequentially executed commitment protocols such that each commit stage is statistically hiding and no polynomial-time strategy can break the binding in all m phases. (This definition is inspired by related, but more complex, notions introduced in [NV, HNO⁺].)
4. (Standard Commitment) We convert the m -phase commitment to a standard statistically hiding commitment scheme by running it many times in parallel, and in each execution having the receiver randomly decide which phase will be used for the actual commitment. (This is similar to a construction in [HNO⁺], except that we show that this conversion can be combined with parallel repetition to obtain full computational binding in one shot, rather than first obtaining weak binding and then amplifying by sequential repetition.)

Statistically Hiding Commitments from One-Way Functions. Recently, it was shown that statistically hiding commitment schemes can in fact be constructed from any one-way function [HNO⁺]. However, the construction was very complicated and inefficient. Here we obtain a much simpler and more efficient construction, by combining Theorem 1.1 with the following:

Theorem 1.3 (one-way function to entropy gap, informal). *Given any one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$, we can construct an $O(n/\log n)$ -round protocol (A, B) in which the real entropy of A 's messages is noticeably larger than their accessible (max-)entropy.*

The proof of this theorem uses a simple variant of the interactive hashing protocol of [NOVY], which was designed to construct statistically hiding commitments from one-way *permutations*. A (different) variant of the [NOVY] protocol was also used as the first step in the previous construction of statistically hiding commitments from one-way functions in [HNO⁺]. Specifically, it was used to obtain a “weakly hiding 2-phase commitment scheme” (for a slightly different notion of 2-phase commitment scheme than the one we use). The main complications there came from the process of amplifying the “weak hiding” property of this 2-phase commitment, which was done through a complex recursive construction. The main source of our simplification is that the property of having a gap between real entropy and accessible entropy is much more well-suited to amplification, and indeed it can be achieved through just parallel repetition as described above.

In addition to being simpler, our protocol is also more efficient. Specifically, we obtain an $O((n/\log n)^2)$ -round protocol, whereas the previous construction gave a large unspecified polynomial number of rounds. Moreover, if we allow the protocol to use nonuniform advice, we obtain $O(n/\log n)$ rounds, which is optimal for “black-box constructions” [HHRS].

This construction also conceptually unifies the construction of statistically hiding commitments from one-way functions with the construction of statistically binding commitments from one-way functions (the latter being due to [HILL, Nao]): the first step of both constructions is obtain a gap between real entropy and “computational entropy” (pseudoentropy in the case of statistical binding and accessible entropy in the case of statistical hiding), which is then amplified by repetitions and finally combined with various forms of hashing.

Commitments and Constant-Round Zero Knowledge. One of the main applications of commitment schemes is to the construction of zero-knowledge proof systems. (Throughout this discussion, we refer to zero-knowledge proofs where the soundness property is statistical, as in the standard definition of interactive proof systems (as opposed to argument systems), but the zero-knowledge property is computational.) The basic zero-knowledge protocol for 3-Coloring and hence all of NP [GMW] utilizes statistically *binding* commitments, and hence the protocol can be implemented in a constant number of rounds assuming the existence of one-way functions (since one-way functions imply 2-message statistically binding commitments [HILL, Nao]). Unfortunately, this protocol has a large soundness error. It is natural to try to use parallel repetition to reduce the soundness error, but zero knowledge is not preserved under parallel repetition in general [FS, GK2]. However, we do know how to construct zero-knowledge proofs for NP that remain secure under parallel composition [GK1, Gol] utilizing statistically *hiding* commitments (used for the verifier to commit to its challenges in advance). Thus, assuming the existence of constant-round statistically hiding commitment schemes, we obtain constant-round zero-knowledge proofs for NP that remain zero knowledge under parallel composition.

It was unknown, however, whether constant-round statistically hiding commitments are *necessary* for constant-round zero-knowledge proofs that remain zero knowledge under parallel composition (or even just have negligible soundness error), or if such zero-knowledge proofs could be constructed from weaker assumptions (such as the existence of one-way functions). We show that that they are in fact necessary, if we restrict to zero knowledge proven via “black-box simulation”.

Theorem 1.4 (zero knowledge to commitments in constant rounds, informal). *Suppose that one-way functions exist and that NP has constant-round interactive proofs that are black-box zero knowledge under parallel composition. Then there exist constant-round statistically hiding commitment schemes.*

We leave as interesting open questions whether constant-round statistically hiding commitment schemes are necessary to just achieve negligible soundness error, and whether the requirement of black-box simulation can be eliminated.

There have been several other results deducing the existence of commitment schemes from zero-knowledge proofs. The first is the result of Ostrovsky and Wigderson [OW], which shows that if there is a zero-knowledge proof for a “hard-on-average” problem, then one-way functions (and hence commitment schemes) exist. In contrast, here we are willing to assume the existence of one-way functions, and are interested in understanding whether certain kinds of zero-knowledge proofs require stronger primitives (such as constant-round statistically hiding commitments). More closely related are the results of Ong and Vadhan [OV], which imply that if there is a *statistical* zero-knowledge proof for a hard-on-average problem, then constant-round statistically hiding commitment schemes exist. Our result is incomparable. On one hand, our result applies even to *computational* zero-knowledge proofs. On the other, we assume the existence of one-way functions and require a zero-knowledge proof for specific NP language (based on the one-way function) with many additional properties.¹

The proof of this theorem roughly proceeds by showing that the zero-knowledge protocol has gap between the real entropy of the verifier’s messages and the accessible entropy of the verifier’s messages, and then applying the construction of Theorem 1.2. However, it turns out that we are not quite able to establish an upper bound on the accessible max-entropy in general, but only if we restrict attention to adversaries A^* that “know” when they have achieved high entropy and for which the high entropy property holds in an arbitrary context (i.e. when interacting with an arbitrary strategy B^* , not just the honest B). We refer to this notion as “context-independent accessible max-entropy,” and it turns out to suffice for the constructions of Theorems 1.1 and 1.2.

The intuition for the accessible entropy of the verifier’s messages being small is that an adversary V^* achieving high accessible entropy should be hard to simulate. Indeed, the only advantage a black-box simulator has over a prover is its ability to “rewind” the verifier. But a verifier V^* achieving accessible high accessible entropy can “resample” new messages that are distributed similarly to the real verifier’s messages every time it is rewound. The infeasibility of simulating such “resampling” verifiers is shown following the approach of Goldreich and Krawczyk [GK2], who considered 3-round protocols and (constant-round) public-coin protocols, settings in which perfect resampling is trivial. Recent applications of this technique in more complex settings related to ours include [Pas, Kat].

Theorem 1.4 can be interpreted either as a negative result about constructing constant-round parallel zero-knowledge proofs from one-way functions (since constructing constant-round statistically hiding commitments from one-way functions has been elusive, and in fact cannot be done via a black-box construction [HHRS]), or as a positive result about constructing constant-round statistically hiding commitments from one-way functions (the use of zero knowledge for NP makes the construction non-black-box in the one-way function, and hence may allow bypassing the lower bound of [HHRS]).

¹The results of [OV] also imply that every statistical zero-knowledge proof can be converted into one with the additional properties we require (namely, constant rounds, parallel composition, black-box simulation, and an efficient prover).

2 Preliminaries

2.1 Random Variables

Let X and Y be random variables taking values in a discrete universe \mathcal{U} . We adopt the convention that when the same random variable appears multiple times in an expression, all occurrences refer to the same instantiation. For example, $\Pr[X = X]$ is 1. For an event E , we write $X|_E$ to denote the random variable X conditioned on E . The *support* of a random variable X is $\text{Supp}(X) := \{x : \Pr[X = x] > 0\}$. X is *flat* if it is uniform on its support. For an event E , we write $I(E)$ for the corresponding indicatory random variable, i.e. $I(E)$ is 1 when E occurs and is 0 otherwise.

We write $\Delta(X, Y)$ to denote the *statistical difference* (a.k.a. variation distance) between X and Y , i.e.

$$\Delta(X, Y) = \max_{T \subseteq \mathcal{U}} |\Pr[X \in T] - \Pr[Y \in T]|.$$

If $\Delta(X, Y) \leq \varepsilon$ (respectively, $\Delta(X, Y) > \varepsilon$), we say that X and Y are ε -close (resp., ε -far).

2.2 Entropy Measures

We will refer to several measures of entropy in this work. The relation and motivation of these measures is best understood by considering a notion that we will refer to as the *sample-entropy*: For a random variable X and $x \in \text{Supp}(X)$, we define the sample-entropy of x with respect to X to be the quantity

$$H_X(x) := \log(1/\Pr[X = x]).$$

The sample-entropy measures the amount of “randomness” or “surprise” in the specific sample x , assuming that x has been generated according to X . Using this notion, we can define the *Shannon entropy* $H(X)$ and *min-entropy* $H_\infty(X)$ as follows:

$$\begin{aligned} H(X) &:= \mathbb{E}_{x \stackrel{R}{\leftarrow} X} [H_X(x)] \\ H_\infty(X) &:= \min_{x \in \text{Supp}(X)} H_X(x) \end{aligned}$$

We will also discuss the *max-entropy* $H_0(X) := \log(1/|\text{Supp}(X)|)$. The term “max-entropy” and its relation to the sample-entropy will be made apparent below.

It can be shown that $H_\infty(X) \leq H(X) \leq H_0(X)$ with equality if and only if X is flat. Thus, saying $H_\infty(X) \geq k$ is a strong way of saying that X has “high entropy” and $H_0(X) \leq k$ a strong way of saying that X as “low entropy”.

Smoothed Entropies. Shannon entropy is robust in that it is insensitive to small statistical differences. Specifically, if X and Y are ε -close then $|H(X) - H(Y)| \leq \varepsilon \cdot \log |\mathcal{U}|$. For example, if $\mathcal{U} = \{0, 1\}^n$ and $\varepsilon = \varepsilon(n)$ is a negligible function of n (i.e., $\varepsilon = n^{-\omega(1)}$), then the difference in Shannon entropies is vanishingly small (indeed, negligible). In contrast, min-entropy and max-entropy are brittle and can change dramatically with a small statistical difference. Thus it is common to work with “smoothed” versions of these measures, whereby we consider a random variable X to have high

entropy (respectively, low entropy) if X is ε -close to some X' with $H_\infty(X) \geq k$ (resp., $H_0(X) \leq k$) for some parameter k and a negligible ε .²

These smoothed versions of min-entropy and max-entropy can be captured quite closely (and more concretely) by requiring that the sample-entropy is large or small with high probability:

Lemma 2.1. 1. Suppose that with probability at least $1 - \varepsilon$ over $x \stackrel{R}{\leftarrow} X$, we have $H_X(x) \geq k$. Then X is ε -close to a random variable X' such that $H_\infty(X') \geq k$.

2. Suppose that X is ε -close to a random variable X' such that $H_\infty(X') \geq k$. Then with probability at least $1 - 2\varepsilon$ over $x \stackrel{R}{\leftarrow} X$, we have $H_X(x) \geq k - \log(1/\varepsilon)$.

Proof Sketch.

1. We can modify X on an ε fraction of the probability space (corresponding to when X takes on a value x such that $H_X(x) \geq k$) so as to bring all probabilities smaller than or equal to 2^{-k} .

2. Let $S = \{x : H_X(x) < k - \log(1/\varepsilon)\}$. Note that $|S| < 2^{k - \log(1/\varepsilon)}$. Then

$$\Pr[X \in S] \leq \Pr[X' \in S] + \varepsilon \leq 2^{-k} \cdot |S| + \varepsilon < 2\varepsilon.$$

□

Lemma 2.2. 1. Suppose that with probability at least $1 - \varepsilon$ over $x \stackrel{R}{\leftarrow} X$, we have $H_X(x) \leq k$. Then X is ε -close to a random variable X' such that $H_0(X') \leq k$.

2. Suppose that X is ε -close to a random variable X' such that $H_0(X') \leq k$. Then with probability at least $1 - 2\varepsilon$ over $x \stackrel{R}{\leftarrow} X$, we have $H_X(x) \leq k + \log(1/\varepsilon)$.

Proof Sketch.

1. Let $S = \{x : H_X(x) \leq k\}$. Note that $|S| \leq 2^k$. By modifying an ε fraction of X , we can obtain a random variable X' whose support is contained in S .

2. Let $S = \{x : H_X(x) > k + \log(1/\varepsilon)\}$. Then

$$\Pr[X \in S] \leq \Pr[X \in S \cap \text{Supp}(X')] + \Pr[X \notin \text{Supp}(X')] < 2^{-(k + \log(1/\varepsilon))} \cdot |\text{Supp}(X')| + \varepsilon = 2\varepsilon.$$

□

Think of ε as inverse polynomial or a slightly negligible function in $n = \log(|\mathcal{U}|)$. The above lemmas show that up to negligible statistical difference and a slightly superlogarithmic number of entropy bits, the min-entropy (resp. max-entropy) is captured by lower (resp. upper) bound on sample-entropy.

²The term “smoothed entropy” was coined by Renner and Wolf [RW] but the notion of smoothed min-entropy has commonly been used (without a name) in the literature on randomness extractors [NZ].

Conditional Entropies. We will also be interested in conditional versions of entropy. For jointly distributed random variables (X, Y) and $(x, y) \in \text{Supp}(X, Y)$, we define the *conditional sample-entropy* to be $H_{X|Y}(x|y) = \log(1/\Pr[X = x|Y = y])$. Then the standard *conditional Shannon entropy* can be written as:

$$H(X|Y) = \mathbb{E}_{(x,y) \stackrel{R}{\leftarrow} (X,Y)} [H_{X|Y}(x|y)] = \mathbb{E}_{y \stackrel{R}{\leftarrow} Y} [H(X|Y=y)] = H(X, Y) - H(Y).$$

There is no standard definition of conditional min-entropy and max-entropy, or even their smoothed versions. For us, it will be most convenient to generalize the sample-entropy characterizations of smoothed min-entropy and max-entropy given above. Specifically we will think of X as having “high min-entropy” (resp., “low max-entropy”) given Y if with probability at least $1 - \varepsilon$ over $(x, y) \stackrel{R}{\leftarrow} (X, Y)$, we have $H_{X|Y}(x|y) \geq k$ (resp., $H_{X|Y}(x|y) \leq k$).

Flattening Shannon Entropy. In [HILL] it was shown how to convert the Shannon entropy of a random variable to min-entropy (up to small statistical distance) by taking independent copies of this variable. This will be a useful tool our work as well.

Lemma 2.3 ([HILL], Prop 4.9). *1. For every $t \geq 1$, X^t is $2^{-t^{1/3}}$ -close to a random variable X' such that $H_\infty(X') \geq tH(X) - t^{2/3}|\log \mathcal{U}|$.*

2. For any $t \geq 1$, (X^t, Y^t) is $2^{-t^{1/3}}$ -close to a pair of jointly distributed random variables (X', Y') such that for all $y \in \text{Supp}(Y')$, $H_\infty(X' | Y' = y) \geq tH(X|Y) - t^{2/3}|\log \mathcal{U}|$.

2.3 Entropy with Failure

In defining accessible entropy, we will have an adversary A^* attempting to generate a string x with maximum possible entropy, and the adversary will also have to “justify” that the sample generated is consistent with a given “honest” algorithm A . In case the adversary fails to provide a proof, we would not want x to contribute to the entropy. To account for this, we consider the adversary to be generating a random variable X taking values in $\mathcal{U} \cup \{\perp\}$, where \perp is used whenever the adversary fails to provide a justification. Now, we do not simply want to measure the entropy of X itself, because then an adversary may be able to increase the entropy by sometimes refusing to provide a proof. For example, suppose that the string generated by the adversary is always 0^n , but the adversary refuses to provide a justification half of the time. Then $H(X) = 1$ but intuitively we should count the entropy as 0 (since the underlying string is always fixed).

To handle this, we consider modified variants of entropy that treat the “failure” value \perp in a special way. In first reading, the reader may choose to ignore the issue of entropy with failures altogether (and simply concentrate on A^* that always provides valid justification). Nevertheless, the following definitions may be useful even beyond our context.

For a random variable X taking values in $\mathcal{U} \cup \{\perp\}$ and $x \in \mathcal{U} \cup \{\perp\}$, we define the (*modified*) *sample-entropy* to be

$$H_X^*(x) := \begin{cases} \log \frac{1}{\Pr[X=x|X \neq \perp]} = \log \frac{\Pr[X \neq \perp]}{\Pr[X=x]} & \text{if } x \neq \perp \\ 0 & \text{if } x = \perp. \end{cases}$$

We define the (*modified*) *Shannon entropy* of X to be

$$\begin{aligned} \mathbf{H}^*(X) &= \mathbb{E}_{x \stackrel{R}{\leftarrow} X} [\mathbf{H}_X^*(x)] \\ &= \mathbf{H}(X|I(X = \perp)), \end{aligned}$$

where $I(X = \perp)$ is the indicator random variable for $X = \perp$. This way of measuring entropy with respect to failure behaves as we would expect, in that it agrees with Shannon entropy when there is no failure, and entropy cannot be increased by failing more often.

Lemma 2.4. 1. If $\Pr[X = \perp] = 0$, then $\mathbf{H}^*(X) = \mathbf{H}(X)$.

2. If X, X' are jointly distributed random variables taking values in $\mathcal{U} \cup \{\perp\}$ such that $\Pr[X' = X \vee X' = \perp] = 1$, then $\mathbf{H}^*(X') \leq \mathbf{H}^*(X)$.

Proof Sketch.

1. Immediate from definitions.

2.

$$\begin{aligned} \mathbf{H}^*(X) &= \mathbf{H}(X|I(X = \perp)) \\ &\geq \mathbf{H}(X|I(X = \perp), I(X' = \perp)) \\ &\geq \Pr[X' \neq \perp] \cdot \mathbf{H}(X|_{X' \neq \perp}) \quad (\text{because } X' \neq \perp \Rightarrow X \neq \perp) \\ &= \Pr[X' \neq \perp] \cdot \mathbf{H}(X'|_{X' \neq \perp}) \\ &= \mathbf{H}^*(X') \end{aligned}$$

□

Max-Entropy with Failures. We will also consider a version of max-entropy that handles failure. Here we will simply require that with probability at least $1 - \varepsilon$ over $x \stackrel{R}{\leftarrow} X$, we have $\mathbf{H}_X^*(x) \leq k$. For this notion, it can be shown that failing more often cannot increase entropy by much:

Lemma 2.5. Let X, X' be jointly distributed random variables taking values in $\mathcal{U} \cup \{\perp\}$ such that $\Pr[X' = X \vee X' = \perp] = 1$,

1. For every $\varepsilon > 0$, with probability at least $1 - \varepsilon$ over $x \stackrel{R}{\leftarrow} X'$, $\mathbf{H}_{X'}^*(x) \leq \mathbf{H}_X^*(x) + \log(1/\varepsilon)$.

2. Suppose that with probability at least $1 - \varepsilon$ over $x \stackrel{R}{\leftarrow} X$, we have $\mathbf{H}_X^*(x) \leq k$. Then with probability at least $1 - 2\varepsilon$ over $x \stackrel{R}{\leftarrow} X'$, we have $\mathbf{H}_{X'}^*(x) \leq k + \log(1/\varepsilon)$.

Proof. 1. Let $S = \{x : \mathbf{H}_{X'}^*(x) > \mathbf{H}_X^*(x) + \log(1/\varepsilon)\}$.

$$\begin{aligned} \Pr[X' \in S] &\leq \sum_{x \in S} 2^{-\mathbf{H}_{X'}^*(x)} \cdot \Pr[X' \neq \perp] \\ &\leq \sum_{x \in S} \varepsilon \cdot 2^{-\mathbf{H}_X^*(x)} \cdot \Pr[X' \neq \perp] \\ &= \varepsilon \cdot \frac{\Pr[X \in S]}{\Pr[X \neq \perp]} \cdot \Pr[X' \neq \perp] \\ &\leq \varepsilon. \end{aligned}$$

2. Follows from Part 1. □

Useful Properties. Both the modified Shannon entropy and the modified max-entropy satisfy similar subadditivity properties as (standard) Shannon entropy:

Lemma 2.6. *Let (X_1, \dots, X_t) be jointly distributed random variables, each taking values in $\mathcal{U} \cup \{\perp\}$, and define \tilde{X} to equal \perp if at least one of the X_j 's equals \perp and to equal (X_1, \dots, X_t) otherwise. Then:*

1. $H^*(\tilde{X}) \leq \sum_j H^*(\tilde{X}_j)$.
2. With probability at least $1 - \varepsilon$ over $(x_1, \dots, x_t, \tilde{x}) \stackrel{R}{\leftarrow} (X_1, \dots, X_t, \tilde{X})$,

$$H_{\tilde{X}}^*(\tilde{x}) \leq \sum_j H_{X_j}^*(x_j) + \log(1/\varepsilon).$$

Proof.

$$\begin{aligned} H^*(\tilde{X}) &= \Pr[\tilde{X} \neq \perp] \cdot H(\tilde{X}|_{\tilde{X} \neq \perp}) \\ &\leq \Pr[\tilde{X} \neq \perp] \cdot \sum_j H(X_j|_{\tilde{X} \neq \perp}) \\ &= \sum_j \Pr[\tilde{X} \neq \perp] \cdot H(X_j|_{\tilde{X} \neq \perp}) \\ &\leq \sum_j H(X_j|I(\tilde{X} = \perp), I(X_j = \perp)) \\ &\leq \sum_j H(X_j|I(X_j = \perp)) \\ &= \sum_j H^*(X_j) \end{aligned}$$

Let $S = \{(x_1, \dots, x_t) : H_{\tilde{X}}^*(\tilde{x}) > \sum_j H_{X_j}^*(x_j) + \log(1/\varepsilon)\}$. For each $(x_1, \dots, x_t) \in S$, we have

$$\Pr[\tilde{X} = (x_1, \dots, x_t) | \forall i X_j \neq \perp] < \varepsilon \cdot \prod_j \Pr[X_j = x_j | X_j \neq \perp].$$

Summing both sides over all $(x_1, \dots, x_t) \in S$, we get

$$\Pr[\tilde{X} \in S | \tilde{X} \neq \perp] < \varepsilon.$$

□

The next lemma is useful for bounding the entropy with failure of variables of small support.

Lemma 2.7. *For every random variable X taking values in a universe $\mathcal{U} \cup \{\perp\}$ it holds that $\mathbb{E}_{x \stackrel{R}{\leftarrow} X} [2^{H_X^*(x)}] \leq |\mathcal{U}|$.*

1. *Proof.*

$$\begin{aligned}
\mathbb{E}_{x \leftarrow X} \left[2^{\mathbb{H}_X^*(x)} \right] &\leq \Pr[X = \perp] + \sum_{x \neq \perp} \Pr[X = x] \cdot (1/\Pr[X = x \mid X \neq \perp]) \\
&\leq \Pr[X = \perp] + |\mathcal{U}| \cdot \Pr[X \neq \perp] \\
&\leq |\mathcal{U}| .
\end{aligned}$$

□

In cases where X is with high probability either a constant or \perp , we use the following lemma.

Lemma 2.8. *Let X be a random variable taking values in a universe $\mathcal{U} \cup \{\perp\}$ and let $x_0 \in \mathcal{U}$. Assume that $\Pr[X \in \{x_0 \cup \perp\}] > 1 - \varepsilon$, then $\Pr_{x \leftarrow X}[\mathbb{H}_X^*(x) \leq \sqrt{\varepsilon}] \geq 1 - O(\sqrt{\varepsilon})$.*

Proof. We treat separately the case that $\mathbb{H}_X^*(x_0) \leq \sqrt{\varepsilon}$ and the complementary case.

$\mathbb{H}_X^*(x_0) \leq \sqrt{\varepsilon}$. Hence, $\Pr_{x \leftarrow X}[\mathbb{H}_X^*(x) \leq \sqrt{\varepsilon}] \geq \Pr[X \in \{x_0 \cup \perp\}] > 1 - \varepsilon$.

$\mathbb{H}_X^*(x_0) > \sqrt{\varepsilon}$. Hence $\Pr[X = x_0 \mid X \neq \perp] < 2^{-\sqrt{\varepsilon}} = 1 - \Omega(\sqrt{\varepsilon})$. Therefore,

$$\begin{aligned}
\Pr[X \neq \perp] &= \Pr[X \notin \{x_0 \cup \perp\}] + \Pr[X = x_0] \\
&= \varepsilon + \Pr[X \neq \perp] \cdot \Pr[X = x_0 \mid X \neq \perp] \\
&\leq \varepsilon + \Pr[X \neq \perp] \cdot (1 - \Omega(\sqrt{\varepsilon})) .
\end{aligned}$$

Hence, $\Pr[X \neq \perp] \leq \frac{\varepsilon}{\Omega(\sqrt{\varepsilon})} \in O(\sqrt{\varepsilon})$ and thus $\Pr_{x \leftarrow X}[\mathbb{H}_X^*(x) = 0] \geq \Pr[X = \perp] \geq 1 - O(\sqrt{\varepsilon})$.

□

Conditional Entropy With Failures. Finally, we can define conditional versions of these notions. Suppose (X, Y) are jointly distributed random variables taking values in $(\mathcal{U} \times \mathcal{V}) \cup \{(\perp, \perp)\}$. (For convenience we require that failure always occurs simultaneously for both random variables.) Then we define

$$\mathbb{H}_{X|Y}^*(x|y) = \begin{cases} \log \frac{1}{\Pr[X=x|Y=y, X \neq \perp]} = \log \frac{\Pr[X \neq \perp]}{\Pr[X=x|Y=y]} & \text{if } x \neq \perp \\ 0 & \text{if } x = \perp. \end{cases}$$

Again, we can obtain a corresponding form of Shannon entropy $\mathbb{H}^*(X|Y)$ by taking expectations, and a corresponding form of max-entropy by requiring that $\mathbb{H}_{X|Y}^*(x|y) \leq k$ with probability at least $1 - \varepsilon$.

2.4 Views and indistinguishability

View of an interaction. For an interactive protocol (A, B) , the random variable $\text{view}_A(A, B)$ denotes the collection of all messages exchanged and the coin tosses of A .

Statistical and computational indistinguishability. Two ensembles of distributions $X = \{X_n\}$ and $Y = \{Y_n\}$ are statistically (resp. computationally) indistinguishable if for all possibly unbounded algorithms (resp. nonuniform PPT) D , the quantity $|\Pr[D(1^n, X_n) = 1] - \Pr[D(1^n, Y_n) = 1]|$ is negligible in n .

3 Real vs. Accessible Entropy of Protocols

In this section we formalize the notions of real and accessible entropies of a protocol. As discussed in the introduction, these entropies and the gap between them (i.e., the inaccessible entropy of a protocol) play a pivotal role in our work. In addition, we will give tools for manipulating accessible and real entropies.

Let us briefly recall the intuition behind these notions of entropy. Let $(A, B)(1^n)$ be an m -round interactive protocol in which B sends the first message. The common input 1^n is the security parameter, which we will often omit from the notation. Let $(B_1, A_1, \dots, B_m, A_m)$ be a random variable denoting the transcript of the messages exchanged between A and B when both parties' coin tosses are chosen uniformly at random. Intuitively, the real entropy of A with respect to (A, B) is the entropy in A 's messages, where for each message A_i we take its entropy conditioned on the partial transcript (B_1, A_1, \dots, B_i) .

Consider now an adversary A^* which interacts with B . At each round, we ask what is the entropy of the next message of A^* conditioned not only on the partial transcript of previous messages but also on *the entire view of A^** (including previous coin flips). A^* is allowed to flip fresh random coins to generate its next message and this is indeed the source of entropy in the message (everything else in the view of A^* is fixed). We call this quantity the “accessible” entropy of A^* with respect to (A, B) . So that the definition is meaningful, we insist that the messages of A^* will be consistent with A and furthermore that A^* will be able to demonstrate this consistency. This is achieved by having A^* locally output (at each round) a string w such that when w is the random input of A the messages A would have sent are identical to those A^* did send so far.

It is interesting to note that if we put no computational restrictions on A^* then the entropy accessible to A^* can always be as high as the real entropy of (A, B) . Simply, at each round A^* can sample a new string w that is consistent with its messages so far and send a next message that is also consistent with w (i.e., send the string that A would have sent given the partial transcript if its random input was set to w). This strategy is not always possible for a computationally bounded A^* , and indeed the interesting protocols from the point of view of this work are protocols where a computationally bounded A^* can only access part of the real entropy (i.e., there is non-negligible inaccessible entropy).

Note that in the above informal definitions (which we formalize below), we only refer to an honest B . While we do so in this preliminary version for simplicity, natural analogues of these definitions for cheating B^* can be given as well.

3.1 Real Entropy

In this paper we will be interested in lower bounds on the real entropy. We will therefore define two variants — real Shannon entropy and real min-entropy (which is particularly suited for lower bounds on entropy). As we did in Section 2.2, we connect these two notions through the notion of real sample-entropy. In other words, for a fixed transcript we ask how surprising were the messages sent by A in this particular transcript. We then get real Shannon entropy by taking the expectation of this quantity over a random transcript and the min-entropy by taking the minimum (up to negligible statistical distance). An alternative approach would be to define the notions through sum of conditional entropies (as we do in the intuitive description in the introduction). This approach would yield closely related definitions, and in fact exactly the same definition in the case of Shannon entropy (see Lemma 3.3).

We say that a partial transcript $t = (b_1, a_1, \dots, b_j, a_j)$ and a sequence w of coin tosses is \mathbf{A} -consistent if \mathbf{A} would answer with a_1, \dots, a_j if its coins were w and it received messages b_1, \dots, b_j . We say that t is \mathbf{A} -consistent if there exists a w such that t and w are \mathbf{A} -consistent.

Definition 3.1 (real sample-entropy). *For an interactive algorithm \mathbf{A} and an \mathbf{A} -consistent partial transcript $t = (b_1, a_1, \dots, b_i)$, define random variables $W_i(t)$ and $A_i(t)$ as follows. Let $W_i(t)$ be selected uniformly at random from the set $\{w : t \text{ and } w \text{ are } \mathbf{A}\text{-consistent}\}$, and let $A_i(t) = \mathbf{A}(t; W_i(t))$. For a fixed message $a_i \in \text{Supp}(A_i)$ we define the real sample-entropy of a_i given t to be*

$$\text{RealH}_{\mathbf{A}}(a_i|t) := H_{A_i(t)}(a_i).$$

For a full transcript $t = (b_1, a_1, \dots, b_m, a_m)$ and a subset $I \subseteq [m]$ of rounds, we define the real sample-entropy of t in the rounds of I to be

$$\text{RealH}_{\mathbf{A}}^I(t) = \sum_{i \in I} \text{RealH}_{\mathbf{A}}(a_i|b_1, a_1, \dots, b_i).$$

Definition 3.2 (real entropy). *For an interactive protocol (\mathbf{A}, \mathbf{B}) as above and a subset $I \subseteq [m]$ of rounds, we say that \mathbf{A} has real Shannon entropy at least k in the rounds of I with respect to (\mathbf{A}, \mathbf{B}) , if*

$$\mathbb{E}_{t \leftarrow_{\mathbf{R}}(\mathbf{A}, \mathbf{B})} [\text{RealH}_{\mathbf{A}}^I(t)] \geq k.$$

In the case that $I = [m]$, we omit it from the above (and the following) notation.

We say that \mathbf{A} has real min-entropy at least k in the rounds of I with respect to (\mathbf{A}, \mathbf{B}) , if there is a negligible function $\varepsilon = n^{-\omega(1)}$ (where n is the security parameter) such that

$$\Pr_{t \leftarrow_{\mathbf{R}}(\mathbf{A}, \mathbf{B})} [\text{RealH}_{\mathbf{A}}^I(t) \geq k] \geq 1 - \varepsilon(n).$$

We observe that the real Shannon entropy simply amounts to measuring standard conditional Shannon entropy of \mathbf{A} 's messages when interacting with \mathbf{B} .

Lemma 3.3. *For an m -round interactive protocol (\mathbf{A}, \mathbf{B}) , let $(B_1, A_1, \dots, B_m, A_m)$ be a random variable denoting the transcript of the messages exchanged between \mathbf{A} and \mathbf{B} when both parties' coin tosses are chosen uniformly at random. Then*

$$\mathbb{E}_{t \leftarrow_{\mathbf{R}}(\mathbf{A}, \mathbf{B})} [\text{RealH}_{\mathbf{A}}^I(t)] = \sum_{i \in I} H(A_i|B_1, A_1, \dots, B_i).$$

Proof.

$$\begin{aligned} \mathbb{E}_{t \leftarrow_{\mathbf{R}}(\mathbf{A}, \mathbf{B})} [\text{RealH}_{\mathbf{A}}^I(t)] &= \sum_{i \in I} \mathbb{E}_{(b_1, a_1, \dots, b_m, a_m) \leftarrow_{\mathbf{R}}(\mathbf{A}, \mathbf{B})} [\text{RealH}_{\mathbf{A}}(a_i|b_1, a_1, \dots, b_i)] \\ &= \sum_{i \in I} \mathbb{E}_{(b_1, a_1, \dots, b_m, a_m) \leftarrow_{\mathbf{R}}(\mathbf{A}, \mathbf{B})} [H_{A_i(b_1, a_1, \dots, b_i)}(a_i)] \\ &= \sum_{i \in I} \mathbb{E}_{(b_1, a_1, \dots, b_m, a_m) \leftarrow_{\mathbf{R}}(\mathbf{A}, \mathbf{B})} [H_{A_i|B_1, A_1, \dots, B_i}(a_i|b_1, a_1, \dots, b_i)] \\ &= \sum_{i \in I} H(A_i|B_1, A_1, \dots, B_i), \end{aligned}$$

where in the third equality we use the fact that the conditional distribution of A_i given that $B_1 = b_1, A_1 = a_1, \dots, B_i = b_i$ is equal to the distribution $A_i(b_1, a_1, \dots, b_i)$ used in the definition of the real sample-entropy $\text{RealH}_A(a_i|b_1, a_1, \dots, b_i)$ (namely, sample coins w for A uniformly at random among those consistent with the history and output the corresponding message). \square

The next claim follows readily from [AH, PT, GV]:

Lemma 3.4. *Let A be an interactive algorithm that uses a random tape of length k , which it always sends as its last message. Then for every A -consistent transcript t , $\text{RealH}_A(t) = k$. In particular, for every interactive algorithm B the algorithm A achieves real entropy at least k with respect to (A, B) .*

3.2 Accessible Entropy

In this paper we will be interested in upper bounds on the accessible entropy. We will therefore define two variants - accessible Shannon entropy and accessible max-entropy (which is particularly suited for upper bounds on entropy). As in the case of real entropy, we connect these two notions through the notion of accessible sample-entropy. In other words, for a fixed view of the adversary A^* we ask how surprising were the messages sent by A^* . We then get accessible Shannon entropy by taking the expectation of this quantity over a random view and the max-entropy by taking the maximum (up to negligible statistical distance). Here too, the definitions obtained are closely related to the definitions one would obtain by considering a sum of conditional entropies (as we do in the intuitive description above). For the Shannon entropy, the definitions would in fact be identical (See Lemma 3.7).

Consider an adversarial strategy A^* that tosses its own fresh random coins s_i in each round before sending a_i , and then locally outputs a sequence w_i of coins for A as a “witness” to the fact that it is behaving consistently with A . So a partial view of A^* when interacting with B can be written in the form $v = (s_0, b_1, s_1, a_1, w_1, \dots, b_i, s_i, a_i, w_i)$. (Note that we also allow A^* some additional random coins s_0 at the start of the protocol.) For such a partial view v and a round $j \leq i$, define $\Gamma_j^A(v)$ to equal a_j if $(b_1, a_1, \dots, b_j, a_j)$ is A -consistent with w_j and to equal \perp otherwise. That is, we replace a message a_j sent by A^* with the failure symbol \perp if it is not accompanied with a consistent justification string w_j . Recall that in Section 2.3, we formalized notions that measure entropy (denoted H^*) in a way that discounts entropy that may come from failing.³

Definition 3.5 (accessible sample-entropy). *Let A^* be an interactive algorithm and let $v = (s_0, b_1, s_1, a_1, w_1, \dots, b_i, s_i, a_i, w_i)$ be an A^* -consistent partial view. Define random variables (S_i, A_i, W_i) by choosing S_i uniformly at random, and setting $(A_i, W_i) = A^*(s_0, b_1, s_1, a_1, w_1, \dots, b_i, S_i)$. For a fixed value $a_i \in \text{Supp}(A_i) \cup \{\perp\}$, we define the accessible sample-entropy of a_i given v as*

$$\text{AccH}_{A, A^*}(a_i|v) := H_{\Gamma_i^A(v, S_i, A_i, W_i)}^*(a_i).$$

For a view $v = (s_0, b_1, s_1, a_1, w_1, \dots, b_m, s_m, a_m, w_m)$ and a subset of rounds $I \subseteq [m]$, we define the accessible sample-entropy of v in the rounds of I to be

$$\text{AccH}_{A, A^*}^I(v) := \sum_{i \in I} \text{AccH}_{A, A^*}(\Gamma_i^A(v)|s_0, b_1, s_1, a_1, w_1, \dots, b_i).$$

³At a first reading of the following definition may be easiest to parse when considering A^* that never fails to supply a consistent witness. In such a case, $\text{AccH}_{A, A^*}(a_i|v) := H_{A_i}(a_i)$.

Definition 3.6 (accessible entropy). For an m -round interactive protocol (A, B) and $I \subseteq [m]$, we say that A has accessible entropy at most k in the rounds of I with respect to (A, B) , if for every PPT A^* ,

$$\mathbb{E}_{v \leftarrow \text{view}_{A^*}^R(A^*, B)} [\text{AccH}_{A, A^*}^I(v)] \leq k$$

We say that A has accessible max-entropy at most k in the rounds of I with respect to (A, B) , if for every PPT A^* , there is a negligible function $\varepsilon = \varepsilon(n)$ such that

$$\Pr_{v \leftarrow \text{view}_{A^*}^R(A^*, B)} [\text{AccH}_{A, A^*}^I(v) \leq k] \geq 1 - \varepsilon(n).$$

Accessible entropy can also be expressed in terms of standard conditional Shannon entropy.

Lemma 3.7. Let (A, B) be an m -round interactive protocol, and let A^* be an adversarial strategy as above. Define random variables $(S_0, B_1, S_1, A_1, W_1, \dots, B_m, S_m, A_m, W_m)$ denoting the view of A^* when interacting with B . Then

$$\mathbb{E}_{v \leftarrow \text{view}_{A^*}^R(A^*, B)} [\text{AccH}_{A, A^*}^I(v)] = \sum_{i \in I} H^*(\Gamma_i^A(V) | S_0, B_1, S_1, A_1, W_1, \dots, B_{i-1}, S_{i-1}, A_{i-1}, W_{i-1}, B_i).$$

Proof. Similar to the proof of Lemma 3.3. □

3.3 Manipulating Accessible and Real Entropy

In this section, we state two results on manipulating accessible and real entropy. The first tool, given by Proposition 3.8 below, deals with the affect of parallel repetition of a protocol on its real (Shannon) entropy and accessible (max) entropy. One effect of a t -fold parallel repetition (A^t, B^t) is that (for certain settings of parameters) the gap between real and accessible entropy can increase. The reason is that the real entropy is not much smaller than t times the real entropy of (A, B) and the accessible entropy is not much larger than t times the accessible entropy of (A, B) . Therefore, the difference between the quantities increases. A second useful effect of parallel repetition is in turning real Shannon entropy into real min-entropy. Note that the slight decrease in real entropy is due to this move from Shannon entropy to min-entropy (rather than from the parallel repetition itself).

Proposition 3.8 (gap amplification via parallel repetition). Let n be a security parameter and $\pi = (A, B)$ an m -round protocol. For $t \in \text{poly}(n) \cap \omega(\log^3 n)$, let $\pi^t = (A^t, B^t)$ be the t -fold parallel repetition of π . Then, π^t satisfies the following properties:

real entropy: For all $i \in [m]$, if the real Shannon entropy of A in round i with respect to π is at least k_{REAL} , then the real min-entropy of A^t in round i with respect to π^t is at least $t \cdot k_{\text{REAL}} - ut^{2/3}$, where u is an upper bound on the length of messages sent by A in π .

accessible max-entropy For any $I \subseteq [m]$ and any $s = \omega(\log n)$, if A has accessible max-entropy at most k_{ACC} in the rounds of I with respect to π , then A^t has accessible max-entropy at most $t \cdot k_{\text{ACC}} + s \cdot m$ in the rounds of I with respect to π^t .

Proof. The bound on real entropy follows readily from Lemma 2.3 and therefore we focus on establishing the bound on accessible max-entropy. Fix I and suppose on the contrary that there

exists an adversarial ppt \tilde{A}^* that upon interacting with B^t in π^t violates having accessible max-entropy at most $t \cdot k_{\text{ACC}} + s \cdot m$ in the rounds of I . That is, there exists a non-negligible function $\varepsilon = \varepsilon(n)$ such that with probability greater than ε over $\tilde{v} \stackrel{R}{\leftarrow} \text{view}_{\tilde{A}^*}(\tilde{A}^*, B^t)$, we have

$$\text{AccH}_{A^t, \tilde{A}^*}^I(\tilde{v}) > t \cdot k_{\text{ACC}} + s \cdot m.$$

By definition,

$$\begin{aligned} \text{AccH}_{A^t, \tilde{A}^*}^I(\tilde{v}) &= \sum_{i \in I} \text{AccH}_{A^t, \tilde{A}^*}(\Gamma_i^{A^t}(\tilde{v}) | \tilde{v}_{<i}) \\ &= \sum_{i \in I} H_{\Gamma_i^{A^t}(\tilde{v}_{<i}, S_i, \tilde{A}_i, \tilde{W}_i)}^*(\Gamma_i^{A^t}(\tilde{v})), \end{aligned}$$

where $\tilde{v}_{<i}$ denotes the portion of the view up to and including the i 'th message from B^t , S_i consists of the coins of \tilde{A} tossed just before sending its i 'th message, $\tilde{A}_i = (A_{i,1}, \dots, A_{i,t})$ is the message it sends, and $\tilde{W}_i = (W_{i,1}, \dots, W_{i,t})$ is its witness/justification string. Note that if we define random variables $\tilde{X}_i(\tilde{v}_{<i}) = \Gamma_i^{A^t}(\tilde{v}_{<i}, S_i, \tilde{A}_i, \tilde{W}_i)$, and for $j = 1, \dots, t$, define $X_{i,j}(\tilde{v}_{<i})$ to equal $A_{i,j}$ if the witness string $W_{i,j}$ is A -consistent for the j 'th execution and to equal \perp otherwise, then we are exactly in the situation of Lemma 2.6. (\tilde{X}_i equals \perp iff at least one $\tilde{X}_{i,j}$ is \perp and otherwise equals $(X_{i,1}, \dots, X_{i,t})$.) Thus, with probability at least $1 - 2^{-s} = 1 - \text{neg}(n)$ over \tilde{v} , we have

$$\sum_{j=1}^t H_{X_{i,j}(\tilde{v}_{<i})}^*(x_{i,j}) \geq H_{\tilde{X}_i(\tilde{v}_{<i})}^*(\tilde{x}_i) - s,$$

where the $x_{i,j}$'s and \tilde{x}_i are defined from the i 'th round of the view \tilde{v} analogously to how the $\tilde{X}_{i,j}$'s and \tilde{X}_i are defined from the random variables $(S_i, \tilde{A}_i, \tilde{W}_i)$. Summing over all $i \in I$, we have

$$\begin{aligned} \sum_{i \in I} \sum_{j=1}^t H_{X_{i,j}(\tilde{v}_{<i})}^*(x_{i,j}) &\geq \sum_i H_{\tilde{X}_i(\tilde{v}_{<i})}^*(\tilde{x}_i) - m \cdot s \\ &> t \cdot k_{\text{ACC}} \end{aligned}$$

with probability $\varepsilon(n) - \text{neg}(n) = 1/\text{poly}(n)$ over \tilde{v} . When this event occurs, there must exist a j such that $\sum_{i \in I} H_{X_{i,j}(\tilde{v}_{<i})}^*(x_{i,j}) > k_{\text{ACC}}$. Thus we can violate the accessible max-entropy of A with respect to π with the following ppt adversary A^* that interacts with B as follows:

1. Pick a random j in $[t]$ (using A^* 's initial coin tosses s_0) and simulate an execution of \tilde{A}^* with B^t .
2. We simulate B^t by using the external B for the j 'th execution of B and running $t - 1$ copies of B internally (the coin tosses for these copies of B are taken from A^* 's initial coin tosses s_0).
3. A^* will use the justification strings output by \tilde{A}^* for the j 'th execution of π .

Interactions of A^* with B correspond naturally to interactions of \tilde{A}^* and B^t . Moreover, it can be verified that for each fixed choice of j by A^* and every induced view \tilde{v} of \tilde{A}^* , the accessible sample-entropy achieved by A^* is precisely: $\sum_{i \in I} H_{X_{i,j}(\tilde{v}_{<i})}^*(x_{i,j})$. By the above, this is greater than k_{ACC} with non-negligible probability. \square

The second tool, given by Proposition 3.9 shows how to turn a protocol $\pi = (\mathbb{A}, \mathbb{B})$ for which a lower bound k_{REAL} on its real Shannon entropy is known to a different protocol (\mathbb{A}, \mathbb{B}) for which a lower bound k_{REAL}/m is known on the real Shannon entropy of (almost all of the) *individual messages*. The price of this transformation is in an increased round complexity (indeed the transformation essentially consists of sequential repetition of the original protocol). Since lower bounds for specific rounds are needed for our transformation of inaccessible entropy to statistically hiding commitments, Proposition 3.9 will indeed come useful. In cases where we will not want to pay the price of increased round complexity we will instead employ non-uniform advice (consisting of the individual bounds).

Proposition 3.9 (equalizing real entropy via sequential repetition). *Let n be a security parameter, $\pi = (\mathbb{A}, \mathbb{B})$ an m -round protocol. For every $t \in \text{poly}(n)$, there is a $(t + 1) \cdot m$ -round protocol $\pi' = (\mathbb{A}, \mathbb{B})$ satisfying the following properties:*

real entropy: *Let $I' = \{m + 1, \dots, tm\}$. Suppose the real Shannon entropy of \mathbb{A} with respect to π is at least k_{REAL} . Then, for all $i \in I'$, the real Shannon entropy of \mathbb{A} in round i with respect to π' is at least k_{REAL}/m .*

accessible max-entropy *If \mathbb{A} has accessible max-entropy at most k_{ACC} with respect to π , then \mathbb{A} has accessible max-entropy at most $t \cdot k_{\text{ACC}}$ with respect to π' .*

In particular, if we know that $k_{\text{ACC}} < (1 - 1/p) \cdot k_{\text{REAL}}$ for some polynomial p , then setting $t = 2p$, we obtain $\pi' = (\mathbb{A}', \mathbb{B}')$ where (1) for all $i \in I'$, the real Shannon entropy in round i w.r.t. π' is at least k_{REAL}/m , and (2) the accessible max-entropy of \mathbb{A}' w.r.t. π' is at most $t \cdot k_{\text{ACC}} < t \cdot k_{\text{REAL}} - 2k_{\text{REAL}} = k_{\text{REAL}}' - k_{\text{REAL}}$, where $k_{\text{REAL}}' = |I'| \cdot k_{\text{REAL}}/m$.

Proof. (of Prop 3.9) We begin by specifying the new protocol $\pi' = (\mathbb{A}, \mathbb{B})$:

1. \mathbb{B} chooses an offset $j \in [m]$ uniformly at random and sends j to \mathbb{A} .
2. For j rounds, the parties send to each other the special symbol \star .
3. The parties run $t - 1$ sequential repetitions of π , with \mathbb{A} and \mathbb{B} acting \mathbb{A} and \mathbb{B} respectively.
4. For $m - j$ rounds, the parties send to each other the special symbol \star .

To establish the statement on real Shannon entropy, observe that for every $i \in I'$, the i th message in a random execution in π' is identically distributed to a message in a random round of π . Hence,

$$\mathbb{E}_{t^{\text{R}}(\mathbb{A}, \mathbb{B})} \left[\text{RealH}_{\mathbb{A}}^{\{i\}}(t) \right] = \frac{1}{m} \sum_{j=1}^m \mathbb{E}_{t^{\text{R}}(\mathbb{A}, \mathbb{B})} \left[\text{RealH}_{\mathbb{A}}^{\{j\}}(t) \right] = \frac{1}{m} \mathbb{E}_{t^{\text{R}}(\mathbb{A}, \mathbb{B})} \left[\text{RealH}_{\mathbb{A}}^{[m]}(t) \right] = k_{\text{REAL}}/m$$

To obtain a bound the accessible max-entropy, consider any adversary \mathbb{A}^* in π' . Note that \mathbb{A}^* can only achieve zero accessible sample-entropy for the first j rounds and the last $m - j$ rounds. So if \mathbb{A}^* achieves accessible sample-entropy at least $t \cdot k_{\text{ACC}}$ in a particular view, it must have achieved accessible sample-entropy at least k_{ACC} in one of the blocks of m consecutive rounds corresponding to one of the interactions with \mathbb{B} . Based on this observation, we construct an adversary \mathbb{A}^* for π , which chooses an execution $i \in [t-1]$ and an offset $j \in [m]$ at random, internally simulates the first $i - 1$ interactions of \mathbb{A}^* with \mathbb{B} (using its initial coin tosses s_0), and then plays the i th interaction

of A^* against the external B . It can be verified that for every view in which A^* chooses execution i , the accessible sample-entropy achieved by A^* equals the accessible sample-entropy achieved by A^* in the rounds corresponding to the i 'th interaction with B . Thus, if A^* achieves accessible sample-entropy at least $t \cdot k_{\text{ACC}}$ with nonnegligible probability ε , then A^* achieves accessible sample-entropy at least k_{ACC} with probability at least ε/t , contradicting the hypothesis. \square

4 Entropy Gap to Commitment

In this section we present the main technical contribution of this paper, showing how any protocol with a noticeable gap between its real and accessible entropies can be converted into a statistically hiding and computationally binding commitment scheme. First we recall the definition of the latter:⁴

Definition 4.1. *A (bit) commitment scheme (S, R) is an efficient two-party protocol consisting of two stages. Throughout, both parties receive the security parameter 1^n as input.*

COMMIT. *The sender S has a private input $b \in \{0, 1\}$, which she wishes to commit to the receiver R , and a sequence of coin tosses σ . At the end of this stage, both parties receive as common output a commitment z .*

REVEAL. *Both parties receive as input a commitment z . S also receives the private input b and coin tosses σ used in the commit stage. This stage is non-interactive: S sends a single message to R , and R either outputs a bit (and accepts) or rejects.*

Definition 4.2. *A commitment scheme (S, R) is statistically hiding if*

COMPLETENESS. *If both parties are honest, then for any bit $b \in \{0, 1\}$ that S gets as private input, R accepts and outputs b at the end of the reveal stage.*

STATISTICAL HIDING. *For every unbounded strategy R^* , the distributions $\text{view}_{R^*}(S(0), R^*)$ and $\text{view}_{R^*}(S(1), R^*)$ are statistically indistinguishable.*

COMPUTATIONAL BINDING. *For every PPT S^* , S^* succeeds in the following game (breaks the commitment) with negligible probability in n :*

- S^* interacts with an honest R in the commit stage, which yields a commitment z .
- S^* outputs two messages τ_0, τ_1 such that for both $b = 0$ and $b = 1$, R on input (z, τ_b) accepts and outputs b .

The main theorem of this section is as follows:

Theorem 4.3 (restatement of Theorems 1.1, 1.2). *Assume that one-way functions exist. Then there exists an efficient transformation that takes as input a security parameter 1^n , an (efficient) m -round interactive protocol⁵ $\pi = (A, B)$, and unary parameters 1^m , 1^k and 1^p , and outputs a $O(mp)$ -round protocol Com with the following guarantee: if the real Shannon entropy of A with*

⁴We present the definition for bit commitment. To commit to multiple bits, we may simply run a bit commitment scheme in parallel.

⁵Given as a pair of circuits.

respect to π is at least k and the accessible max-entropy of \mathbf{A} with respect to π is at most $(1 - 1/p)k$, then \mathbf{Com} is a statistically hiding and computationally binding commitment scheme. Alternatively if $m = O(1)$, then \mathbf{Com} can also be made to have $O(1)$ rounds.⁶

The heart of the proof lies in following Lemma.

Lemma 4.4. *Assume that one-way functions exist. Then there exists an efficient transformation that takes as input a security parameter 1^n , an (efficient) m -round interactive protocol $\pi = (\mathbf{A}, \mathbf{B})$, a set of indices $I \subseteq [m]$ and a set of integers $\{\ell_i\}_{i \in I}$ where $\ell_i \geq 3n$ for all $i \in I$, and outputs an $O(m)$ -round commitment scheme $\mathbf{Com} = (\mathbf{S}, \mathbf{R})$ with the following properties:*

hiding: *If for all $i \in I$ the real min-entropy with respect to π in the i 'th round is at least ℓ_i , then \mathbf{Com} is statistically hiding.*

binding: *If the accessible max-entropy of \mathbf{A} in the rounds of I with respect to π is at most $\sum_{i \in I} \ell_i - 3n|I|$, then \mathbf{Com} is computationally binding.*

Informally (see Section 4.3 for the formal proof) Lemma 4.4 is used to prove Theorem 4.3 as follows: we start by applying the “equalizing real entropy transformation” (Proposition 3.9) on π to get an $O(pm)$ -round protocol π' for which the entropy gap still exists and (almost) all of π' 's rounds have the same *known* value of real entropy. Then we apply the gap amplification transformation (Proposition 3.8) on π' to get a protocol π'' where (almost) all of its rounds have large *known* value of *min-entropy*, and the accessible max-entropy of the protocol is much smaller than the sum of the rounds' min-entropy. Finally, we apply Lemma 4.4 on π'' to get an $O(mp)$ -round statistically hiding commitment.

In the case of a constant m , we skip the first “entropy equalizing” step and rather apply Proposition 3.8 directly on π , to get a protocol as π'' above, but for which we have no handle of the value of the (possibly different) min-entropies of each round. Since π and thus π'' is a constant round protocol, by applying Lemma 4.4 on π'' for polynomially many possible values for the min-entropies whose sum is “large enough” (this value is induced by the value of k), we get polynomially many commitments that are all binding and at least one of them is hiding. These commitment can be combined in a standard way to get a single scheme that is statistically hiding and computationally binding.

In order to prove Lemma 4.4 (again see Section 4.3 for the formal proof), we first show (in Section 4.1) show how to use a protocol with gap between its real min-entropy and its accessible max-entropy to get a secure “ m -phase commitment” (defined in Section 4.1). We then complete the proof Lemma 4.4, by showing (in Section 4.1) how to use such a secure m -phase commitments for constructing a full fledged statistically hiding commitment.

4.1 Entropy Gap to secure m -phase Commitment

We start by defining the notion of m -phase commitment, and then show how to construct such a commitment that is secure (in the sense given below) using a protocol with gap between its real min-entropy and its accessible max-entropy.

⁶By equipping the transformation with nonuniform advice, the number of rounds of \mathbf{Com} can be reduced to $O(m)$ also in the general case.

4.1.1 m -phase commitment

An m -phase commitment protocol consists of m sequentially executed commitment protocols such that each commit stage is statistically hiding and no polynomial-time strategy can break the binding in all m phases.

Definition 4.5. (*m -phase commitment scheme*) An m -phase commitment scheme, for $m \in \mathbb{N}$, is an m -phase efficient protocol between a sender S , getting as private input an m -bit string \bar{b} , and a receiver R , where both parties get a security parameter 1^n . Each phase consists of a possibly interactive commit-stage followed by a reveal-stage in which the sender sends a single message to the receiver. We put the following security requirements on the protocol:

hiding (against semi-honest receivers): After the i 'th commit stage (i.e., the commit stage of the i 'th phase), the value of $\bar{b}[i]$ is statistically hidden from R . Namely, for every $i \in [m]$ and every $\bar{b} \in \{0, 1\}^{i-1}$ and $\bar{b}' \in \{0, 1\}^{m-i}$ it holds that $\text{view}_R^i(S(\bar{b}, 0, \bar{b}'), R)(1^n)$ and $\text{view}_R^i(S(\bar{b}, 1, \bar{b}'), R)(1^n)$ are statistically indistinguishable, where view_R^i stands for R 's view after the i 'th commit stage.

$\binom{m}{1}$ -binding: Let S^* be an algorithm interacting with R , and assume that after the i 'th reveal-stage, S^* locally outputs two strings $w_i^0 = (r^0, \bar{b}^0)$ and $w_i^1 = (r^1, \bar{b}^1)$. We say that S^* breaks the binding of the i 'th phase if $(w_i^0, (\text{trans}_i, a_i))$ and (w_i^1, trans_i) are S -consistent and $\bar{b}^0[i] \neq \bar{b}^1[i]$, where trans_i is the transcript of the protocol after i 'th commit-stage and a_i is the message of S^* in the i 'th reveal phase. We say that S^* breaks the $\binom{m}{1}$ -binding of Com in a given execution if it breaks the binding of all the phases simultaneously. Finally, Com is $\binom{m}{1}$ -binding, if no PPT breaks the $\binom{m}{1}$ -binding of Com with more than negligible probability.

An m -phase commitment scheme is **secure**, if it is both hiding (against semi-honest receivers) and $\binom{m}{1}$ binding.

Remark 4.6. We note that while the above primitive looks similar (and shares similar name) to the two-phase commitment schemes previously defined by Nguyen et al. [NV, HNO⁺], the two primitives enjoy very different security properties. The $\binom{2}{1}$ -binding property of the earlier notion only guarantees that the adversary cannot decommit the first-phase commitment to two values, and then break the binding also in the two second-phase commitments induced by both first-phase decommitments. This is in contrast to the above definition of $\binom{m}{1}$ binding, where it is infeasible for the adversary to be able break the subsequent phases induced by even one of the decommitments.

In Lemma 4.18, we take advantage of the latter stronger binding guarantee in order to convert any secure m -phase commitment into a full fledged statistically hiding commitment with the same round complexity. In contrast, the only transformation known from earlier form of two-phase commitment to statistically hiding commitment, due to [HNO⁺], incurs a multiplicative $\omega(\log n)$ factor to the round complexity.

4.1.2 Constructing m -phase commitment

Lemma 4.7. Assume that one-way functions exist. Then there exist an efficient transformation that takes as input parameter 1^n , 1^m and 1^t , an m -round interactive protocol $\pi = (A, B)$, a set of indices $I \subseteq [m]$ and a set of integers $\{\ell_i\}_{i \in I}$, where $\ell_i \geq 3n$ for all $i \in I$, and outputs an $O(m)$ -round $|I|$ -phase commitment scheme $\text{Com} = (S, R)$ with the following properties:

hiding: If for every $i \in I$ the real min-entropy of A in the i 'th round with respect to π is at least ℓ_i , then Com is hiding (against semi-honest receivers).

binding: If the accessible max-entropy of A in the rounds of I with respect to π is at most $\sum_{i \in I} \ell_i - 3n|I|$, then Com is $\binom{m}{1}$ binding.

public-coin: R is public-coin if B is.

Proof. (of Lemma 4.7) The high level idea is to use the rounds of (A, B) whose indices are inside I as the phases of (S, R) . More specifically, we use the $i \in I$ 'th message a_i of A (played by S) for masking the input bit $\bar{b}_{i'}$ of S , where i' is the index of i inside I . In order to do so, we add an additional step before S sends a_i , in which S sends the “hash value” of a_i such that the following hold:

- after sending the hash value, the real min-entropy a_i is still high (e.g., $\Omega(n)$), and
- if the accessible max-entropy of A in the i 'th round is lower than $\ell_i - 3n$ (i.e., given A 's view, the support size of a_i is smaller than $2^{\ell_i - 3n}$), then a_i is *determined* from the point of view of (even a cheating) S after sending the hash value.

We next extract (via inner product with a random string) a bit $\sigma_{i'}$ from a_i whose real min-entropy is close to one (i.e., its value is close to uniform from R 's point of view). Finally, we use $\sigma_{i'}$ for masking $\bar{b}_{i'}$. The hiding of the scheme follows from the guarantee about the min-entropy of A 's messages. The $\binom{m}{1}$ -binding of the scheme follows since the bound on the accessible max-entropy of A , yields that in (almost) every execution it holds that the accessible entropy of one of A 's messages is low.

As our basic building block we are using the following two-round information-theoretic interactive hashing protocol, taken from [DHRS].

Protocol 4.8. (*two-round interactive hashing protocol*) $(S_{\text{IH}}, R_{\text{IH}})$

Common input: security parameter 1^n , length parameter 1^{len} , an entropy threshold 1^ℓ , a family of len -wise independent hash functions $\mathcal{H}_{\text{len}} : \{0, 1\}^{\text{len}} \mapsto \{0, 1\}^\ell$ and a family of pairwise independent hash functions $\mathcal{H}_2 : \{0, 1\}^{\text{len}} \mapsto \{0, 1\}^n$.

S_{IH} 's private input: $x \in \{0, 1\}^{\text{len}}$.

1. R_{IH} selects uniformly at random $h_{\text{len}} \in \mathcal{H}_{\text{len}}$ and sends it to S_{IH} .
2. S_{IH} sends $y_1 = h_{\text{len}}(x)$ back to R_{IH} .
3. R_{IH} selects uniformly at random $h_2 \in \mathcal{H}_2$ and sends it to S_{IH} .
4. S_{IH} sends $y_2 = h_2(x)$ back to R_{IH} .

.....

We will use two properties of the above protocol. The first, which we will use for “hiding,” is that S_{IH} sends only $\ell + n$ bits to the R_{IH} . Thus, if S_{IH} 's input x comes from a distribution of min-entropy significantly larger than $\ell + n$, it will still have high min-entropy conditioned on R_{IH} 's view of the protocol (with high probability). On the other hand, the following “binding” property says that if x has max-entropy smaller than ℓ (i.e. is restricted to come from a set of size at most

2^ℓ), then it will have negligible entropy after the protocol (i.e. will be uniquely determined, except with negligible probability).

The following proposition readily follows from the proof of [DHRS, Theorem 5.6]

Proposition 4.9 (“statistical binding” property of $(S_{\text{IH}}, R_{\text{IH}})$). *Let $L \subseteq \{0, 1\}^{\text{len}}$ be a set of size at most 2^ℓ . Let S_{IH}^* be an (unbounded) adversary playing the role of S_{IH} in $(S_{\text{IH}}, R_{\text{IH}})$ and assume that following the protocol execution S_{IH}^* outputs two strings x_0 and x_1 . Then*

$$\Pr[x_0 \neq x_1 \in L \wedge \forall j \in \{0, 1\} \ h_{\text{len}}(x_j) = y_1 \wedge h_2(x_j) = y_2] < 2^{-\Omega(n)} .$$

The above proposition guarantees that after Protocol 4.8 ends, it is impossible (even for an unbounded party) to find two consistent inputs inside a “small set,” except with negligible probability. For our reduction to go through, we need the following stronger security guarantee: it should be infeasible (for polynomially bounded parties) to find two consistent inputs where even one of them lies inside a small set. Protocol 4.11 defined below uses universal one-way hash functions to achieve this strong requirement. We use the following definition of universal one-way hash functions.

Definition 4.10 (universal one-way hash functions [NY]). *Let $\mathcal{F} = \bigcup_n \mathcal{F}_n = \{f: \{0, 1\}^{\ell(n)} \mapsto \{0, 1\}^{m(n)}\}$ be an efficient family of functions. We say that \mathcal{F} is a family of universal one-way hash functions if the following conditions hold:*

Compression. $m(n) < \ell(n)$.

Target Collision Resistance. *The probability that A succeeds in the following game is negligible in n :*

1. $(x, \text{state}) \leftarrow A(1^n)$
2. $f \leftarrow \mathcal{F}_n$
3. $x' \leftarrow A(x, \text{state}, f)$ and A succeeds whenever $x' \neq x$ and $f(x') = f(x)$.

Adding these to the interactive hashing protocol, we obtain the following hashing protocol.

Protocol 4.11. (hashing protocol)⁷ $(S_{\text{H}}, R_{\text{H}})$.

Common input: security parameter 1^n , length parameter 1^{len} , entropy threshold 1^ℓ and a family of universal one-way hash function functions $\mathcal{F}_n = \{f: \{0, 1\}^{\text{len}} \rightarrow \{0, 1\}^n\}$.

S_{H} 's private input: $x \in \{0, 1\}^{\text{len}}$.

1. *The two parties run $(S_{\text{IH}}(x), R_{\text{IH}})(1^n, 1^{\text{len}}, 1^\ell)$, with S_{H} and R_{H} act S_{IH} and R_{IH} respectively.*
2. *R_{H} sends a random $f \in \mathcal{F}_n$ to S_{H} .*
3. *S_{H} sends $y = f(x)$ back to R_{H} .*

⁷Protocol 4.11 is of similar flavor (and indeed inspired by) to the protocol used by Haitner and Reingold [HR2] in their transformation of “two-phase” commitment to statistically hiding commitment. In fact, the protocol of [HR2] can be seen as a special case of Protocol 4.11, designed to work for singleton sets L_v (see Lemma 4.12).

Note that S_H now sends a total of at most $\ell + 2n$ bits to the receiver, and the following is the new binding property.

Lemma 4.12 (“computational binding” property of (S_H, R_H)). *Let $\{L_v \subseteq \{0, 1\}^{\text{len}}\}_{v \in \{0, 1\}^*}$ be a family of sets, each of size at most 2^ℓ . Let S_H^* be a PPT adversary playing the role of S_H in (S_H, R_H) , assume that before the protocol starts S_H^* outputs a string v and that, following the protocol execution, S_H^* outputs two strings x_0 and x_1 . Then*

$$\Pr[x_0 \neq x_1 \wedge (\exists j \in \{0, 1\} : x_j \in L_v) \wedge (\forall j \in \{0, 1\} h_{\text{len}}(x_j) = y_1 \wedge h_2(x_j) = y_2 \wedge f(x_j) = y)] = \text{neg}(n) ,$$

where y_1 and y_2 are the values of the corresponding messages in the embedded execution of (S_{IH}, R_{IH}) .

Proof. Assume towards a contradiction that there exists an PPT S_H^* that violates the “binding” of (S_H, R_H) with success probability at least $1/p$ for some $p \in \text{poly}$, and let the random variables V and Trans_{IH} , denote the value of the variables v and trans_{IH} induced by a random execution of (S_H^*, R_H) , where trans_{IH} denotes the transcript of the embedded execution of (S_{IH}, R_{IH}) . Proposition 4.9 yields that S_H^* that violates the binding of (S_H, R_H) with probability at least $1/2p$ also when we add the requirement that there exists a single element in L_V that is consistent with Trans_{IH} . We denote the value of this element by X . Let E be the event that after the interaction of (S_H^*, R_H) ends, there exists a single element in L_V that is consistent with Trans_{IH} the conditional probability that S_H^* cheats is at least $1/4p$. An averaging argument yields that $\Pr[E] \geq 1/2p$. Note that whenever E happens, X can be extracted efficiently with probability $1/4p$ - simulate a random continuation of (S_H^*, R_H) , and let x_0, x_1 be the output of S_H^* in the end of the execution. Output uniformly at random $x \in \{x_0, x_1\}$.

Consider the following efficient algorithm for violating the target collision resistance of \mathcal{F}_n .

Algorithm 4.13 (collision finder).

ColFinder

Input: security parameter 1^n .

Output x :

- Emulate a random execution of (S_H^*, R_H) till the execution of (S_{IH}, R_{IH}) ends, denote the state of the emulated protocol by **state**.
- Extract the value of X induced by the above emulation, where if the extraction fails abort.
- output (x, state) .

Finding Collision:

Input: $f \in \mathcal{F}_n$, $x \in \{0, 1\}^{\text{len}}$ and **state** $\in \{0, 1\}^*$.

Operation:

- Emulate a random execution of (S_H^*, R_H) conditioned on state and $F = f$, let x_0 and x_1 be the output of S_H^* in the end of the emulation.
- Output $\{x_0, x_1\} \setminus \{x\}$.

By the above observations, it follows that ColFinder violates target collision resistance \mathcal{F}_n with probability $\Omega(1/p^2)$. \square

Let (A, B) be an m -round protocol with “entropy gap”. We assume without loss of generality that all of A ’s messages are of the same length⁸ len . In the following we invoke Protocol 4.11 on each of the $|I|$ high min-entropy messages of A , to construct an $|I|$ -phase commitment scheme.

Protocol 4.14 ($|I|$ -phase commitment scheme).
Com = (S, R) .

Common input: security parameter 1^n , round complexity 1^m , message length 1^{len} , an m -round protocol (A, B) , a family of functions $\mathcal{F} = \bigcup_n \mathcal{F}_n = \{f: \{0, 1\}^{\text{len}} \rightarrow \{0, 1\}^n\}$, a set $I = \{i_1, \dots, i_{|I|}\}$ and a set of entropy thresholds $\{1^{\ell_i}\}_{i \in I}$.

S’s private input: $\bar{b} \in \{0, 1\}^{|I|}$.

The two parties initiate an execution of $(A(1^n), B(1^n))$, with S and R acting A and B respectively. The protocol continues through the following $|I|$ -phase protocol:

The j ’th commit stage:

1. The two parties continue the execution of $(A(1^n), B(1^n))$ until B sends its i_j ’th message to A .
2. The two parties run $(S_H(a_j), R_H)(1^n, 1^{\text{len}}, 1^{\ell_{i_j} - 3n})$ where S and R act as S_H and R_H respectively and a_j is the next message of A in the embedded execution of $(A(1^n), B(1^n))$.
3. S chooses a random $r \in \{0, 1\}^{\text{len}}$ and sends $(\langle r, a_j \rangle_2 \oplus \bar{b}[j], r)$ to R , where $\langle \cdot, \cdot \rangle_2$ denotes inner-product modulo 2.

The j ’th reveal stage:

S (acts as A) sends a_j to R (acts as B).

Claim 4.15 (hiding against semi-honest receivers). Assume that for all $i \in I$ the real min-entropy of A in the i ’th round is at least ℓ_i with respect to π , then **Com** is hiding (against honest receivers).

Proof. Let A_i denote the i ’th message of A in π and let V_i^B denote B ’s view before A_i is sent, where both random variables are with respect to a random execution of π . The assumption about ℓ_i yields that there exists a negligible function ε such that $\Pr_{t \leftarrow (A, B)}[\text{RealH}_A^i(t) \geq \ell_i] \geq 1 - \varepsilon(n)$. It follows that there exists a distribution (V_i^B, A_i') that is ε -close to (V_i^B, A_i) and for every value $v_i^B \in \text{Supp}(V_i^B)$ it holds that $A_i |_{V_i^B = v_i^B}$ has min-entropy at least ℓ_i .

⁸Using padding technique one can transform any protocol to one that all its messages are of the same length. It is easy to verify that such padding does not change the real/accessible entropy of the parties.

Let V_i^R be R's view in a random execution of Com just before the embedded A sent its i 'th message, and let V_i^B denote the part of V_i^R that is coming from the embedded execution of π and V_i^H the other part of v_i . Note that value of V_i^H is a probabilistic function of V_i^B , where V_i^B and A_i are distributed the same in $(S(U_m), R)$ and in (A, B) . Thus (V_i^R, A_i) is ε -close to a distribution (V_i^R, A'_i) such that for every value $v_i^R \in \text{Supp}(V_i^R)$ it holds that $A'_i |_{V_i^R=v_i^R}$ has min-entropy at least ℓ_i . Let H_i be the concatenation of the messages sent by S in the i 'th phase of $(S_H(U_m), R_H)$. Since $|H_i| = \ell_i - n$, it follows that (V_i^R, H_i, A_i) is $(\varepsilon + 2^{-n/2})$ -close to a distribution (V_i^R, H_i, A'_i) such that for every value $(v_i^R, h_i) \in \text{Supp}(V_i^R, H_i)$ it holds that $A'_i |_{(V_i^R, H_i)=(v_i^R, h_i)}$ has min-entropy at least $n/2$. Finally, by the Leftover Hash Lemma (cf., [ILL]), for every $\bar{b} \in \{0, 1\}^{i-1}$ and $\bar{b}' \in \{0, 1\}^{m-i}$ it holds that $\text{view}_i^{R'}(S(\bar{b}, 0, \bar{b}'), R)$ and $\text{view}_i^{R'}(S(\bar{b}, 1, \bar{b}'), R)$ are of statistical distance at most $\varepsilon + 2^{-\Omega(n)}$, where $V_i^{R'}$ stands for R's view after the i 'th commit stage. \square

Claim 4.16 ($\binom{m}{1}$ -binding). *If the accessible max-entropy of A with respect to I in π is at most $\sum_{i \in I} \ell_i - 3n |I|$, then Com is $\binom{m}{1}$ -binding.*

Proof. Let S^* be a PPT playing the role of S in Com and assume that it breaks the $\binom{m}{1}$ -binding of Com with non-negligible probability ε . For $i \in I$, let a_i denote the message that S^* sends in the i 'th reveal phase and let w_i^0 be the first justification string that S^* locally outputs after this stage. We define algorithm A^* for interacting with B as follows: A^* emulates a random execution of (S^*, R) while emulating R using the B it interacts with, and choosing the messages of R that are not of embedded B uniformly at random. Each time S^* send a message a_i to B (and in particular when $i \in I$), A^* locally outputs the string w_i^0 . Note that the view of S^* in a real execution of (S^*, R) and in the emulated execution induced by a random execution of (A^*, B) are the same.

Let V be distributed according to A^* 's in a random execution of (A^*, B) . Given $v \in \text{Supp}(V)$, let v_i be the prefix of v that reflects A^* 's view just after B sent its i 'th message and let $S_i(v) := \{a \in \{0, 1\}^*: \text{AccH}_{A, A^*}(a | v_i) \leq \ell_i - 3n\}$. Note that $|S_i(v)| \leq 2^{\ell_i - 3n}$. By Lemma 4.12, S^* cannot break (with nonnegligible probability) the binding of the i 'th phase of Com and in the same time have $a_i \in S_i(v)$. Thus, with nonnegligible probability for every $i \in I$ it holds that $a_i \notin S_i(v)$ and w_i^0 is a good justification string for a_i . Namely, A^* contradicts the bound on the accessible max-entropy of A in π . \square

By Rompel [Rom] (full proof in [KK]), it follows that assuming the existence of a one-way function, there exists a family of universal one-way hash functions from polynomial $\ell(n) \in \text{poly}(n)$ to $m(n) < \ell(n)$.⁹ By [NY, Lemma 2.1], this implies a construction for every $m(n)$ such that $\ell(n) = \text{poly}(m(n))$, which yields the following theorem.

Theorem 4.17 ([Rom, NY, KK]). *Assume that one way functions exist. Then for any positive polynomial $\ell(n) \geq n$, there exists a family of universal one-way hash functions mapping strings of length $\ell(n)$ to strings of length n .*

We conclude the proof of the lemma by applying the transformation guaranteed by the above theorem on the given one-way function to get a family of universal one-way hash functions $\mathcal{F} =$

⁹The Target Collision Resistance property of Definition 4.10 is somewhat stronger than the one given in [KK] (and somewhat weaker than the original definition in [NY]). The strengthening is in allowing A to transfer additional information, i.e., state, between the selection of x and finding the collision. We note that the proof in [KK] holds also w.r.t. to our stronger definition (and even w.r.t. the original definition of [NY]).

$\bigcup_n \mathcal{F}_n = \{f: \{0,1\}^{\text{len}} \mapsto \{0,1\}^n\}$, and output an instance of Protocol 4.14 with common inputs $(1^n, 1^m, \pi, \mathcal{F}, I$ and $\{\ell_i\}_{i \in I})$ as the m -phase commitment protocol. \square

4.2 Secure m -phase Commitment to Statistically Hiding Commitment

Lemma 4.18. *(converting m -phase commitments to standard commitments) There exist an efficient transformation that takes as input a security parameter 1^n , round parameter 1^m and a secure m -phase commitment scheme, and outputs an m -round commitment scheme, which is statistically hiding against honest receivers and computationally binding.*

Proof. We start by constructing a statistically hiding (against honest receivers) and weakly binding (i.e., the binding only holds with noticeable probability) standard commitment from an m -phase commitment (i.e., random value of \bar{b}), where the receiver chooses at random (and keeps private) an index $i \in [m]$ and a “mode bit” $M \in \{\text{“Com”}, \text{“Check”}\}$. In case of $M = \text{“Com”}$, following the i ’th commit stage the receiver asks the sender masks the secret bit with \bar{b}_i (i.e., the value of the i ’th commitment) and halts the execution. Otherwise ($M = \text{“Check”}$), following the i ’th reveal stage the receiver asks the sender to justify its reveal - to send random-coins and input that are consistent with the transcript (with respect to the honest sender). If the sender fails to justify, the receiver rejects, otherwise the execution halts (no additional value is sent in this case).

The hiding of the above commitment immediately follows from the hiding of the m -phase commitment, in the following we argue that it is also weakly binding. Given a sender that breaks the binding with too high probability (e.g., better than $1 - \frac{1}{4m}$), we use it for breaking the $\binom{m}{1}$ -binding of the underlying m -phase commitment as follows: following the i ’th commit stage, we send “Com” to the sender to get two openings of this commitment. We then “rewind” the last message, and continue the (real) execution of the commitment. Following the i ’th reveal stage, we send “Check” to the sender to get an opening of the i ’th commit that is consistent with the reveal stage. Then we again rewind the last message and continue the execution. It is easy to verify, that by doing the above we break the $\binom{m}{1}$ -binding with non-negligible probability. In order to get strongly binding commitment, we repeat the above weakly commitment in parallel. While in general the parallel repetition of a weakly computational binding commitment is not known to improve the binding, in the case of the above commitment the strongly binding of its parallel repetition follows via a rather straight forward argument (see detail below).

Our transformation sets $k = m \cdot \log^2 n$, and outputs the commitment $\text{Com}^{(k)} = (\mathcal{S}^{(k)}, \mathcal{R}^{(k)})$ defined below:

Protocol 4.19 (statistically hiding commitment scheme).

$\text{Com}^{(k)} = (\mathcal{S}^{(k)}, \mathcal{R}^{(k)})$.

Commit stage:

Common input: 1^n ,

$\mathcal{S}^{(k)}$ ’s **input:** $b \in \{0,1\}$.

$\mathcal{S}^{(k)}$ chooses uniformly at random $B^1, \dots, B^k \in \{0,1\}^m$, and $\mathcal{R}^{(k)}$ sets $J = [k]$. The two parties execute in parallel $(\mathcal{S}(B^1), \mathcal{R})(1^n), \dots, (\mathcal{S}(B^k), \mathcal{R})(1^n)$ with $\mathcal{S}^{(k)}$ and $\mathcal{R}^{(k)}$ act as \mathcal{S} and \mathcal{R} respectively in each execution, but with the following additional steps:

Real Commit: Following the i 'th commit stage, the parties engage in the following subprotocol:

1. $\mathsf{R}^{(k)}$ chooses a random subset $K_i \subseteq J$, where each $j \in J$ is independently chosen with probability $1/2m$, and sends its description to $\mathsf{S}^{(k)}$.
2. $\mathsf{S}^{(k)}$ sends $\sigma_i^j = b \oplus B_i^j$ for all $j \in K_i$ back to $\mathsf{R}^{(k)}$.
3. The parties halt the interactions of $(\mathsf{S}(B^j), \mathsf{R})(1^n)$ for all $j \in K_i$.

Consistency check: Following the i 'th reveal stage, the parties engage in the following subprotocol:

1. $\mathsf{R}^{(k)}$ chooses a random subset $L_i \subseteq (J \setminus K_i)$, where each $j \in (J \setminus K_i)$ is independently chosen with probability $1/2m$, and sends its description to $\mathsf{S}^{(k)}$.
2. $\mathsf{S}^{(k)}$ sends (r^j, B^j) for all $j \in L_i$ back to $\mathsf{R}^{(k)}$, where r^j are the random-coins used by S in the j 'th parallel execution.
3. $\mathsf{R}^{(k)}$ verifies for all $j \in L_i$ that $((r^j, B^j), \text{trans}_i^j)$ is S -consistent, where trans_i^j is the transcript of j 'th parallel execution. If one of the consistency checks fails, $\mathsf{R}^{(k)}$ aborts and rejects.
4. The parties halt the interactions of $(\mathsf{S}(B_j), \mathsf{R})(1^n)$ for all $j \in L_i$, and $\mathsf{R}^{(k)}$ sets $J = J \setminus (K_i \cup L_i)$.

Reveal stage: Generic. $\mathsf{S}^{(k)}$ sends its random-coins and the value of b to $\mathsf{R}^{(k)}$, and $\mathsf{R}^{(k)}$ verifies that they are $\mathsf{S}^{(k)}$ -consistent with the transcript.

It is easy to verify that $\text{Com}^{(k)}$ is correct. In the following claims, we prove that $\text{Com}^{(k)}$ is statistically hiding and computationally binding.

Claim 4.20. $\text{Com}^{(k)}$ is statistically hiding.

Proof. The hiding for case that $k = 1$ (i.e., $|J| = 1$), follows from the hiding of the m -phase commitment. Thus, the hiding for polynomial k (and in particular for $k = m \cdot \log^2 n$ as above) follows by a straight forward hybrid argument. □

Claim 4.21. $\text{Com}^{(k)}$ is computationally binding.

Proof. Let $\mathsf{S}^{(k)*}$ be a PPT that breaks the $\binom{m}{1}$ -binding of $\text{Com}^{(k)}$ with probability $\varepsilon > 1/\text{poly}$, and let V denote $\mathsf{S}^{(k)*}$'s view in a random execution of $(\mathsf{S}^*, \mathsf{R}^{(k)})$. Given $v \in \text{Supp}(V)$, let v_i be the view of S^* (induced by v) at the beginning of the i 'th Real commit stage, and let $J_i(v)$, $K_i(v)$ and $L_i(v)$ be the values of J and K_i and L_i in the i 'th Real Commit and Consistency check stages. We let $\gamma_{i,j}(v)$ [resp., $\delta_{i,j}(v)$] be the probability that S^* breaks the commitment in V conditioned on the event that V starts with v_i and $j \in K_i(V)$ [resp., $j \in L_i(V)$].

For $i \in [m]$ let $\text{Bad}_i^\gamma(v) = \{j \in J_i(v) : \gamma_{i,j}(v) \leq \varepsilon/2\}$ and let $\text{Bad}_i^\delta(v) = \{j \in J_i(v) : \delta_{i,j}(v) \leq \varepsilon/2\}$. We call v **bad** if there exists $i^*(v) \in [m]$ such that $\max \left\{ \left| \text{Bad}_{i^*(v)}^\gamma(v) \right|, \left| \text{Bad}_{i^*(v)}^\delta(v) \right| \right\} > k/2m$. For

$v, v' \in \text{Supp}(V)$, let $E_v(v') = 1$ iff $(K_{i^*}(v')(v') \cap \text{Bad}_{i^*}^\gamma(v) = \emptyset) \wedge (L_{i^*}(v') \cap \text{Bad}_{i^*}^\delta(v) = \emptyset)$. It follows that for every bad v it holds that

$$\Pr[E_v(V) = 1 \mid V \text{ starts with } v_{i^*}] \leq \left(1 - \frac{1}{2m}\right)^{k/2m} = \text{neg}.$$

Therefore, for every bad v and $j \in J_{i^*}(v)$ it holds that $\gamma'_{i,j}(v) \leq \gamma_{i,j}(v) + \text{neg}$ and $\delta'_{i,j}(v) \leq \delta_{i,j}(v) + \text{neg}$, where $\gamma'_{i,j}(v)$ [resp., $\delta'_{i,j}(v)$] is the probability that S^* breaks the commitment in V conditioned on the event that V starts with v_i , $j \in K_i(V)$ [resp., $j \in L_i(V)$] and $E_v(V) = 0$. It follows that

$$\begin{aligned} \Pr[S^* \text{ breaks the binding in } V \wedge V \text{ is bad}] & \quad (1) \\ & \leq \varepsilon/2 + \text{neg} < \frac{2\varepsilon}{3} \end{aligned}$$

Consequently, the event that S^* breaks the binding and V is not bad occurs with probability at least $\varepsilon/3$. In the following we use this for violating the $\binom{m}{1}$ -binding of Com . For $\ell \in [k]$, let S_ℓ^* be defined as follows:

Algorithm 4.22. S_ℓ^* .

Operations: *Emulate a random execution of $(S^{(k)*}, R^{(k)})(1^n)$, while emulating $R^{(k)}$ using the real R in the ℓ 'th parallel executions and using $t - 1$ randomly emulated R 's for the other entries. In addition, do the following in the i 'th Real Commit and Consistency check phases (of $\text{Com}^{(k)}$):*

i 'th Real Commit phase: Do the following for n/ε times:

1. *Select (on behalf of $R^{(k)}$) a random subset $K_i \subseteq J$ conditioned on $\ell \in K_i$,¹⁰ and send K_i to $S^{(k)*}$.*
2. *Continue the emulation of $(S^{(k)*}, R^{(k)})$ till it ends.*
3. *If $S^{(k)*}$ broke the binding of $\text{Com}^{(k)}$:*
 - (a) *For $d \in \{0, 1\}$, set $w_d = (r^{d\ell}, \bar{b}^{d\ell})$, where $\{(r^{0j}, \bar{b}^{0j}), (r^{1j}, \bar{b}^{1j})\}_{j \in K_i}$ are the pairs of random-coins that $S^{(k)*}$ outputs when breaking the binding.*
 - (b) *Halt the loop.*
4. *Otherwise, rewind $(S^{(k)*}, R^{(k)})$ to the beginning of the Real Commit phase.*

If none of the above loops succeed (i.e., the emulated $S^{(k)}$ did not break the commitment), abort. Otherwise, rewind $(S^{(k)*}, R^{(k)})$ to its stage before the Real Commit phase. Select (on behalf of $R^{(k)}$) a random subset $K_i \subseteq J$ condition that $\ell \notin K_i$, and continue the emulation till the end of the Real Commit phase.*

¹⁰Where “random” stand for the same method done in a real execution of Com^k - each $j \in J$ is independently chosen with probability $1/2m$.

***i 'th Consistency check:** Do the following for n/ε times:*

1. *Select (on behalf of $\mathsf{R}^{(k)}$) a random subset $L_i \subseteq (J \setminus K_i)$ conditioned on $\ell \in L_i$, and send L_i to $\mathsf{S}^{(k)*}$.*
2. *If $\mathsf{S}^{(k)*}$ returns to $\mathsf{R}^{(k)}$ a valid justification string $\{\tilde{w}^j = (r^j, \bar{b}^j)\}_{j \in K_i}$:*
 - (a) *Locally output $(\tilde{w}^j = (r^j, \bar{b}^j), w_d = (r^{d^j}, \bar{b}^{d^j}))$, where $d \in \{0, 1\}$ is taken such that $\bar{b}^j[i] \neq \bar{b}^{d^j}[i]$.¹¹*
 - (b) *Halt the loop.*
3. *Rewind $(\mathsf{S}^{(k)*}, \mathsf{R}^{(k)})$ to the beginning of the Consistency check.*

If none of the above loops succeed (i.e., the emulated $\mathsf{S}^{(k)}$ failed to justify its output), abort. Otherwise, select (on behalf of $\mathsf{R}^{(k)}$) a random subset $L_i \subseteq (J \setminus K_i)$ condition that $\ell \notin L_i$, and send L_i to $\mathsf{S}^{(k)*}$, and continue the emulation till the end of the Consistency check.*

.....

We first note that $\mathsf{S}^{(k)*}$'s view in the beginning of the i 'th Real Commit phase, is distributed exactly as the view of $\mathsf{S}^{(k)*}$ in a random execution of $(\mathsf{S}^{(k)*}, \mathsf{R}^{(k)})$ conditioned that $j \in J_i$. Inequality (1) yields that with probability at least $\varepsilon/3m$ over a random choice of $j \in [k]$ and the random-coins of S^j and R , for all $i \in [m]$ it holds that $\delta_{i,j}(\tilde{V}) > \varepsilon/2$, where \tilde{V} denotes $\mathsf{S}^{(k)*}$'s view in the end of the emulation. Consider now the algorithm S^* that selects that selects $\ell \in [k]$ uniformly at random and acts in the interaction with R as S_ℓ^* . It follows that S^* breaks the $\binom{m}{1}$ -binding of Com with probability $(1 - O(2^{-n}))\varepsilon/3m$, and a contradiction is derived. \square

4.3 Putting it Together

Now start by proving Lemma 4.7 and then use Lemma 4.7 for proving Theorem 4.3.

Proof. (of Lemma 4.4) We start by applying Lemma 4.7 on $\pi = (\mathsf{A}, \mathsf{B})$, I and $\{\ell_i\}_{i \in I}$ to get an $O(m)$ -round m -phase commitment Com_m that is hiding and $\binom{m}{1}$ binding. Then we apply Lemma 4.18 on Com_m to get a standard an $O(m)$ -round commitment scheme $\text{Com} = (\mathsf{S}, \mathsf{R})$ with the following properties:

- If for all $i \in I$ it holds that A has real min-entropy at least ℓ_i in round $i \in I$ with respect to π , then Com is statistically hiding against honest receivers.
- If accessible max-entropy of A in the rounds of I with respect to π is at most $\sum_{i \in I} \ell_i - 3n |I|$, then Com is computationally binding.
- R is public-coin if B is.

¹¹Note that such value for d always exist - Since $w_0 = (r^{0^\ell}, \bar{b}^{0^\ell})$ and $w_1 = (r^{1^\ell}, \bar{b}^{1^\ell})$ (chosen in Line 3. of the i 'th "Real Commit phase" above) imply decommitment to different values of the i 'th commit stage (as otherwise $\mathsf{S}^{(k)*}$ has not broken the binding of Com^k), it follows that $\bar{b}^1[i] \neq \bar{b}^0[i]$. In particular either $\bar{b}^1[i]$ or $\bar{b}^0[i]$ is different than $\bar{b}^j[i]$.

So it is left to amplify the hiding to hold against arbitrary receivers, this is achieved by applying on Com the transformation guaranteed by the following fact.¹²

Fact 4.23. (*implicit in [HHK⁺]*) *Assume that one-way functions exist. Then there exists an efficient transformation that takes as input a security parameter 1^n , an m -round statistically hiding against honest receivers and computationally binding commitment scheme Com , and outputs an $O(m)$ -round statistically hiding and computationally binding commitment scheme Com' .*

Proof. Immediately follows from the proof of [HHK⁺, Theorem 6], by combining Construction 6.4 with Claims 6.5 and 6.6. □

Proof. (of Theorem 4.3) We first give a general reduction that works for nonconstant m , and then give a more efficient reduction for constant m 's. Our first step is to apply the equalizing real entropy transformation (Proposition 3.9) on π (with respect to $t = O(p)$) to get an $O(pm)$ -round protocol $\pi' = (A', B')$ such that the following hold:

- The real Shannon entropy of A' in round $i \in I$ with respect to π' is at least k/m , where $I := [m + 1, (m - 1) \cdot p]$.
- $k_{\text{ACC}\pi'}^I + 1 \leq k \cdot (p - 2)$, where $k_{\text{ACC}\pi'}^I$ is an upper bound on the accessible max-entropy of A in the rounds of I with respect to π' .

Our next step is to apply the gap amplification transformation (Proposition 3.8) on π' (with respect to $t \in O((npm \cdot \text{len})^3)$, where len is an upper bound on the length of A messages in π , to get a protocol $\pi'' = (A'', B'')$ with the following properties:

- The real *min-entropy* of A' round $i \in I$ with respect to π'' is at least $\ell = 3n + t \cdot (k(p - 2) - \frac{1}{2}) / |I|$
- $k_{\text{ACC}\pi''}^I < t(k \cdot (p - 2) - \frac{1}{2})$.

Finally, we apply Lemma 4.4 on π'' (with respect to I and $\{\ell_i = \ell\}_{i \in I}$), to get an $O(mp)$ -round statistically hiding and computationally binding commitment.

In the case of a constant m , we skip the first “entropy equalizing” step and rather apply Proposition 3.8 directly on π with $t = O((mnpk \cdot \text{len})^3)$ to get a protocol π' with the following properties:

- There exist $I \subseteq [m]$ and a set $\{\ell_i\}_{i \in I}$ (whose values might not be efficiently computable) such that the real *min-entropy* of A' in the round $i \in I$ with respect to π' is at least $\ell_i \in \{3n, \dots, t \cdot \text{len}\}$.
- $k_{\text{ACC}\pi'}^I + 3n |I| \leq t \cdot k \cdot (1 - 1/2p) \leq \sum_{i \in I} \ell_i$.

Consider all possible choices of I' and $\{\ell'_i\}_{i \in I'}$, where the values of the ℓ'_i 's are taken from the set $\{0\} \cup \{3n, \dots, t \cdot \text{len}\}$ such that $tk(1 - 1/2p) \leq \sum_{i \in I'} \ell'_i$ and $I' = \{i \in [m] : \ell'_i > 0\}$. By applying Lemma 4.4 to each of these *polynomially* many choices, we get a set of polynomially many commitments that are all binding and at least one of them is hiding. We achieve the statistically hiding and computationally binding commitment using any (round preserving) one-out-of-poly commitment combiner (cf., [HHK⁺, Lemma 5.3]). □

¹²Unfortunately, the following transformation does not preserve the public-coin property of the receiver.

5 Statistically Hiding Commitments from One-way Functions

Theorem 5.1 (restatement of Theorem 1.3). *Let $f: \{0, 1\}^n \mapsto \{0, 1\}^n$ be a one-way function, then there exists an efficient $O(n/\log n)$ -round protocol $\pi = (A, B)$ for which the following holds:*

1. (A, B) has real Shannon entropy n with respect to π .
2. A has accessible max-entropy at most $n - \omega(\log n)$ with respect to π .
3. B is public coin.

As an immediate corollary of Theorem 5.1 above and Theorem 4.3 we get an alternative and more round efficient construction to the one-way function based statistically hiding commitment of [HNO⁺].

Corollary 5.2. *Let $f: \{0, 1\}^n \mapsto \{0, 1\}^n$ be a one-way function, then there exists an $O(n^2/\log^2 n)$ -round statistically hiding and computationally binding commitment scheme.*

Remark 5.3. *We could also use Theorem 5.1 and Theorem 4.3 to get a $O(n/\log n)$ -round non-uniform (i.e., the parties get an additional non-uniform advice) statistically hiding and computationally binding commitment. Such a protocol matches the lower bound of Haitner et al. [HHRS] on the round complexity of fully-black-box reduction from statistically hiding commitment schemes to one-way functions.*

Proof. (of Theorem 5.1) We assume for simplicity that $n/\log n \in \mathbb{N}$, and let $(S_{\text{CIH}}, R_{\text{CIH}})$ be an instantiation of the *computational interactive hashing protocol* given by ([HR1, Protocol 3.6]) (building on [NOVY]) described next.

Protocol 5.4 (computational interactive hashing protocol).
 $(S_{\text{CIH}}, R_{\text{CIH}})$.

Common input: \mathcal{H} - a family pairwise independent hash functions from $\{0, 1\}^n$ to $\{0, 1\}^{\log n}$.

S_{CIH} 's input: $y \in \{0, 1\}^n$.

For $i = 1$ to $n/\log n$:

1. R_{CIH} selects uniformly at random $h_i \in \mathcal{H}$ and sends its description to S_{CIH} .
2. S_{CIH} sends $h_i(y)$ back to R_{CIH} .

.....

We use the following fact about Protocol $(S_{\text{CIH}}, R_{\text{CIH}})$:

Proposition 5.5. *Let $f: \{0, 1\}^n \mapsto \{0, 1\}^n$ be a one-way function, and for $j \in [n]$ let $T_j := \{y \in \{0, 1\}^n : 2^{j-1} \leq |f^{-1}(y)| < 2^j\}$. Then the following has negligible probability for every efficient S_{CIH}^* , $j \in [n]$ and a constant $c > 0$: after $\frac{n-j}{\log n} - c$ rounds, S_{CIH}^* outputs $x_0, x_1 \in \{0, 1\}^n$ such that $f(x_0) \neq f(x_1)$ and both $f(x_0)$ and $f(x_1)$ are in T_j and consistent with the protocol.¹³*

¹³We note that in order for Proposition 5.5 to hold, Protocol 5.4 could not be replaced with the (computational) interactive hashing of [NOVY]. The binding property guaranteed by the protocol of [NOVY] is only “meaningful” for $j > n - O(\log n)$, and thus cannot be used to imply the proposition.

Proof. Consider the function $f': \{0, 1\}^n \mapsto \{0, 1\}^n$, where $f'(x) = f(x)$ if $f(x) \in T_j$ and equals $y_0 \in (\{0, 1\}^n \setminus T_j)$ otherwise (where in case $T_j = \{0, 1\}^n$, we set $f' := f$). Note that one violates Proposition 5.5 with respect to f and j iff it violates Proposition 5.5 with respect to f' and j . Let $L \supseteq T_j$ be an arbitrary set of size $2^{n-(j-1)}$ inside $\{0, 1\}^n \setminus \{y_0\}$. Applying [HR1, Theorem 4.1] on f' and L , yields that if Proposition 5.5 does not hold with respect to f (and thus with respect to f'), then it is feasible to invert f' on the uniform distribution over L . In particular, there exists a PPT A such that $\Pr_{y \leftarrow L}[y \in T_j \wedge A(y) \in f^{-1}(y)] > \text{neg}$. Since for every $y' \in T_j$ it holds that $\Pr_{y \leftarrow f(U_n)}[y = y'] \leq \Pr_{y \leftarrow L}[y = y'] \leq 2 \cdot \Pr_{y \leftarrow f(U_n)}[y = y']$, it follows that $\Pr_{y \leftarrow f(U_n)}[A(y) \in f^{-1}(y)] \geq \Pr_{y \leftarrow f(U_n)}[y \in T_j \wedge A(y) \in f^{-1}(y)] > \text{neg}$, which contradicts the one-wayness of f (in the standard sense). \square

We let π be the following $m = (\frac{n}{\log n} + 2)$ -round protocol:

Protocol 5.6. (A, B).

Common input: 1^n .

1. A selects a random $x \in \{0, 1\}^n$ and set $y = f(x)$.
2. The two parties run $(S_{\text{CIH}}(y), R_{\text{CIH}})$, with A and B acting S_{CIH} and R_{CIH} respectively.
3. B sends a dummy message to A.¹⁴
4. A sends y to B.
5. B sends a dummy message to A.
6. A sends x to B.

Since the only random-coins of A are x , Lemma 3.4 yields that A has real Shannon entropy n with respect to π . In order to prove the Theorem, we need to show that the accessible max-entropy of A with respect to π is bounded by $n - \omega(\log n)$. Assume that a cheating A^* outputs $y \in T_j := \{y \in \{0, 1\}^n : 2^{j-1} \leq |f^{-1}(y)| < 2^j\}$. Proposition 5.5 yields that for any $c > 0$, the values of the last $(n-j)/\log n - c$ prior to last messages of A^* are determined given the first messages, and thus their accessible entropy is zero. We conclude the proof by showing that the other messages do not contribute too much accessible entropy to cover this loss.

Moving to the formal proof, we assume towards contradiction the existence of a PPT A^* , $p \in \text{poly}$ and $c > 0$ such that $\Pr[\sum_{i \in [m]} \text{AccH}_{A, A^*}^i(V) > n - c \log n] \geq 1/p(n)$, where V is A^* 's view induced by a random execution of (A^*, B) . Let V_i be A^* 's partial view after receiving the i 'th message of B, let A_i be A^* 's i 'th message and let X_i be A^* 's i 'th justification string, where all these random variables are taken with respect to a random execution of (A^*, B) . Let Γ_i take the value A_i if $(X_i, (V_i, A_i))$ is A-consistent (i.e., the transcript induced by V_i together with A_i are justified by X_i) and set it to \perp otherwise. In the following we assume with put lost of generality that there exists $i \in [m]$ for which $\Gamma_i \neq \perp$ (conditioned that no such i exists, the max-entropy of A^* 's messages is

¹⁴The current and the following two messages could be removed, as they are only added for simplifying the notations through the proof.

negligible), let $I^* = \max \{i \in [m]: \Gamma_i \neq \perp\}$ and let $Y = f(W_{I^*})$. Finally, for $j \in [n]$ let E_j be the event that $Y \in T_j$. It follows that there exists $j \in [n]$ such that

$$\Pr\left[\left(\sum_{i \in [m]} \text{AccH}_{\mathbf{A}, \mathbf{A}^*}^i(V)\right) > n - c \cdot \log n \wedge E_j\right] \geq 1/np \quad (2)$$

In the following we derive a contradiction by showing that the above event cannot happen with non-negligible probability. In order to do so we separate the sum in Inequality 2 into three different parts and bound each of them separately. We start by bounding the head part of the sum.

Claim 5.7. *For every $j \leq m - 2$ and $\varepsilon > 0$ it holds that*

$$\Pr\left[\left(\sum_{i=1}^j \text{AccH}_{\mathbf{A}, \mathbf{A}^*}^i(V)\right) > j \cdot \log n + \log(1/\varepsilon)\right] \leq \varepsilon .$$

Proof. Recall that $\text{AccH}_{\mathbf{A}, \mathbf{A}^*}^i(V) := \mathbf{H}_{(\Gamma_i | V_i)}^*(a_i(V))$, where $a_i(V)$ is \mathbf{A}^* 's i 'th message in V . Since for each $i \in [j]$ the variable Γ_i is taking values inside $\{0, 1\}^{\log n} \cup \{\perp\}$, Lemma 2.7 yields that

$$\begin{aligned} & \mathbb{E}\left[\prod_{i=1}^j 2^{\text{AccH}_{\mathbf{A}, \mathbf{A}^*}^i(V)}\right] \\ &= \mathbb{E}_{s_1 \leftarrow \text{AccH}_{\mathbf{A}, \mathbf{A}^*}^1(V)} \left[2^{s_1} \cdot \mathbb{E}_{s_2 \leftarrow \text{AccH}_{\mathbf{A}, \mathbf{A}^*}^2(V) | s_1} \left[2^{s_2} \cdots \mathbb{E}_{s_j \leftarrow \text{AccH}_{\mathbf{A}, \mathbf{A}^*}^j(V) | s_1, \dots, s_{j-1}} \left[2^{s_j} \right] \right] \right] \\ &\leq n^j , \end{aligned}$$

and the proof follows by a Markov bound. □

We next show that the accessible entropy of the “middle” part of the sum is very small.

Claim 5.8. *For every $j \in [n]$ and a constant $c > 0$ it holds that*

$$\Pr\left[\left(\sum_{i=\frac{n-j}{\log n}-c}^{m-1} \text{AccH}_{\mathbf{A}, \mathbf{A}^*}^i(V)\right) > 1 \wedge E_j\right] = \text{neg} .$$

Proof. Let $\ell = (n - j)/\log n - c$. Proposition 5.5 yields that after the first $\ell - 1$ rounds, there exists $y_0 \in T_j$ such that the the following event happens with only negligible probability: $E_j = 1$ and there exists $i \in \{\ell, \dots, m - 1\}$ such that Γ_i and is inconsistent with y_0 and not equal to \perp . It follows that

$$\begin{aligned} & \Pr\left[\left(\sum_{i=\ell}^{m-1} \text{AccH}_{\mathbf{A}, \mathbf{A}^*}^i(V)\right) > 1 \wedge E_j\right] \\ &\leq \Pr\left[\left(\exists i \in \{\ell, \dots, m - 1\} : \text{AccH}_{\mathbf{A}, \mathbf{A}^*}^i(V) > 1/m\right) \wedge E_j\right] \\ &= \text{neg} , \end{aligned}$$

where the equality follows by Lemma 2.8. □

Finally, we show that the entropy of the last message is not too large.

Claim 5.9. For every $j \in [n]$ and $\varepsilon > 0$ it holds that $\Pr[\text{AccH}_{\mathcal{A}, \mathcal{A}^*}^m(V) > j + \log(1/\varepsilon) \wedge E_j] \leq \varepsilon$.

Proof. Note that $\Gamma_m \in f^{-1}(A_{m-1}) \cup \{\perp\}$, and recall that $E_j = 1$ means that $A_{m-1} \in T_j$. Thus in case E_j holds, Lemma 2.7 yields that $\mathbb{E}[2^{\text{AccH}_{\mathcal{A}, \mathcal{A}^*}^m(V)}] \leq 2^j$ and the proof follows by a Markov bound. \square

Let $\varepsilon > 1/\text{poly}$ and $t = \frac{n-j}{\log n} - \frac{\log(1/\varepsilon)}{\log n}$. The above claims yield that

$$\begin{aligned} & \Pr\left[\left(\sum_{i \in m} \text{AccH}_{\mathcal{A}, \mathcal{A}^*}^i(V)\right) > n - 2 \cdot \log(1/\varepsilon) \wedge E_j\right] \\ & \leq \Pr\left[\left(\sum_{i=1}^{t-1} \text{AccH}_{\mathcal{A}, \mathcal{A}^*}^i(V)\right) > t \cdot \log n + \log(1/\varepsilon)\right] \\ & \quad + \Pr\left[\left(\sum_{i=t}^{m-1} \text{AccH}_{\mathcal{A}, \mathcal{A}^*}^i(V)\right) > 0 \wedge E_j\right] \\ & \quad + \Pr[\text{AccH}_{\mathcal{A}, \mathcal{A}^*}^m(V) > j + \log(1/\varepsilon) \wedge E_j] \\ & \leq 2\varepsilon + \text{neg} < 3\varepsilon . \end{aligned}$$

Taking $\varepsilon = \min\{1/3np, 1/n^{c/2}\}$, yields that $\Pr[(\sum_{i \in m} \text{AccH}_{\mathcal{A}, \mathcal{A}^*}^i(V)) > n - c \cdot \log n \wedge E_j] < 1/np$, which contradicts Equation 2. \square

6 Statistically Hiding Commitments from CZKP

In this section, we establish that constant-round statistically hiding commitments are necessary for constructing constant-round zero-knowledge proof systems for NP that remain secure under parallel composition (assuming the existence of one-way functions):

Theorem 6.1 (restatement of Theorem 1.4). *Suppose that nonuniformly secure one-way functions exist and that NP has constant-round (computational) zero-knowledge proofs that are black-box zero knowledge under parallel composition and that have an efficient prover. Then, there exist constant-round statistically hiding commitment schemes (with computational binding against nonuniform adversaries).*

We note that the converse is true, namely that constant-round statistically hiding commitment schemes imply constant-round black-box zero-knowledge proofs for NP that remain zero-knowledge under parallel composition [GK1, Gol] as well as the existence of one-way functions.

6.1 Proof overview

The proof of this theorem roughly proceeds by showing that the zero-knowledge protocol has gap between the real entropy of the verifier’s messages and the accessible entropy of the verifier’s messages, and then applying the construction of Theorem 4.3. The intuition for the accessible entropy of the verifier’s messages being small is that an adversary V^* achieving high accessible entropy should be hard to simulate. Indeed, the only advantage a black-box simulator has over a prover is its ability to “rewind” the verifier. But a verifier V^* achieving accessible high accessible entropy can “resample” new messages that are distributed similarly to the real verifier’s messages

every time it is rewound. Following Goldreich and Krawczyk [GK2], a simulator that successfully simulates accepting transcripts against such a “resampling” verifier can be turned into a prover strategy that convinces the real verifier to accept, which by soundness is possible only when $x \in L$. This enables us to distinguish YES and NO instances, contradicting the hardness of the language under consideration. We note that in [GK2] (as well as more recent applications of the approach [Kat]), V^* samples messages that are distributed identically to the real verifier’s messages. Here we argue instead need to argue that high accessible entropy implies that V^* ’s messages are distributed “similarly” to the real verifier’s messages; our analysis is inspired by [AH, PT, GV].

We now describe our approach in more detail:

Establishing an entropy gap. We want to make an argument of the following kind: if V^* achieves high accessible max-entropy while interacting with the honest prover, then it also achieves high accessible max-entropy while interacting with the black-box simulator. Once we prove such a statement, we may proceed as in [Kat, GK2] to construct a computationally unbounded “simulation-based cheating prover” to derive a contradiction to the soundness guarantee of the underlying proof system. However, formalizing such an argument presents two technical difficulties:

- First, “achieving high accessible max-entropy” is not an efficiently verifiable property, so it is not clear a-priori that the property is preserved under zero-knowledge simulation.
- Next, “achieving high accessible max-entropy” is an “online” property, whereas the black-box simulator does not interact with V^* in an online manner.

For these reasons, we will work with a weaker notion of accessible max-entropy, where we restrict attention to adversaries A^* that “know” when they have achieved high entropy as measured by some predicate `success` that is applied to its view, and for which the high entropy property holds in an arbitrary context (i.e. when interacting with an arbitrary strategy B^* , not just the honest B). We refer to this notion as “context-independent accessible max-entropy.” The predicate `success` will be the efficiently verifiable property used to address the first technical difficulty, and we will reason about whether V^* achieves high entropy while interacting with the “simulation-based cheating prover,” which will play the role of the aforementioned B^* . Unfortunately, we do not know how to achieve gap amplification (Proposition 3.8) for context-independent accessible max-entropy and as such, we are only able to construct commitment schemes starting from zero-knowledge proofs that remain secure under parallel composition.

From entropy gap to commitment scheme. Next, we show that an upper bound on context-independent accessible max-entropy is already sufficient to obtain a statistically hiding commitment via the transformation in Section 4; that is, we show that the transformation in Section 4 (specifically, Protocol 4.14) can start with a weaker security guarantee and end with the same conclusion.

6.2 Black-box zero knowledge

Definition 6.2 (zero knowledge). *An interactive proof system (P, V) for a language $L \in \text{NP}$ with relation R_L is zero knowledge if for every PPT V^* and polynomial p , there exists a PPT S such that for every $(x, \omega) \in R_L$ and every $z \in \{0, 1\}^{p(|x|)}$, the distributions $(P(\omega), V^*(z))(x)$ and $\text{Sim}(x, z)$ are computationally indistinguishable.*

Definition 6.3 (black-box zero knowledge). *An interactive proof system (P, V) for a language $L \in \text{NP}$ with relation R_L is black-box zero knowledge if there exists an oracle PPT Sim such that for every PPT V^* and polynomial p , and for every $(x, \omega) \in R_L$ and $z \in \{0, 1\}^{p(|x|)}$, the distributions $(P(\omega), V^*(z))(x)$ and $\text{Sim}^{V^*(x, z)}(x)$ are computationally indistinguishable.*

Definition 6.4 (black-box parallel zero knowledge). *An interactive proof system (P, V) for a language $L \in \text{NP}$ with relation R_L is black-box parallel zero knowledge if there exists an oracle PPT Sim such that for every PPT V^* and polynomials p, t , and for every n , $(x_1, \omega_1), \dots, (x_t, \omega_t) \in R_L$ with $t = t(n)$ and $|x_i| = n$ and $z \in \{0, 1\}^{p(n)}$, the distributions $(P^t(\omega_1, \dots, \omega_t), V^*(z))(x_1, \dots, x_t)$ and $\text{Sim}^{V^*(x_1, \dots, x_t, z)}(x_1, \dots, x_t)$, are computationally indistinguishable. Here, (P^t, V^t) denotes the t -fold parallel repetition of (P, V) .*

The notion of computational indistinguishability in the above definitions is the one from Section 2.4, which refers to *nonuniform* polynomial-time distinguishers (as is standard of treatments of zero knowledge). Moreover, following [GK2], we may make the following assumptions about Sim :

- It never asks the same query twice.
- Sim always queries V^* on a partial transcript (b_1, a_1, \dots, b_i) of the protocol. Moreover, whenever it makes such a query, it has previously queried V^* on all the proper prefixes, namely all sequences of the form (b_1, a_1, \dots, b_j) for $j < i$.
- If $(b_1, a_1, \dots, b_m, a_m)$ is the transcript that appears in the final output of Sim , then Sim has queried V^* on $(b_1, a_1, \dots, b_m, a_m)$.

6.3 Context-independent accessible entropy

We now define a variant of accessible max-entropy that only rules out adversaries who “know” when they have achieved high entropy (such that this holds even when they are interacting with an arbitrary strategy B^*). To capture the fact that adversaries A^* “know” when they have achieved high entropy, we will consider a “success predicate” success as applied to the view of A^* . An example of such a predicate in the context of a commitment scheme is to whether the commit phase transcript is accepting (i.e., the receiver does not abort), and whether the cheating sender locally outputs valid openings to two different values.

Now, consider the negation of the definition of accessible max-entropy (Definition 3.6). It says that there is a PPT adversary A^* who achieves a view v with accessible sample-entropy greater than k with probability at least $1/p(n)$ for some polynomial p . In the definition below, we require that this noticeable event can be recognized by some success predicate. That is, we require that A^* achieves high sample-entropy whenever $\text{success}(v)$ where v is A^* 's view (except for negligible probability), and this should hold even when A^* is interacting with an arbitrary strategy B^* .

Definition 6.5 (context-independent accessible entropy). *Let (A, B) be an m -round protocol, and $I \subseteq [m]$. We say that A has context-independent accessible max-entropy at most k in the rounds of I with respect to (A, B) , if there is no PPT A^* and an efficient computable predicate success satisfying the following conditions:*

- For any view v of A^* , $\text{success}(v)$ implies v is A -consistent.

- There is a polynomial p such that

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B)}[\text{success}(v)] \geq 1/p(n),$$

and

- There is a negligible function $\varepsilon(n)$ such that for every (possibly inefficient) strategy B^* ,

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B^*)}[\neg \text{success}(v) \text{ or } \text{AccH}_{A, A^*}^I(v) > k] \geq 1 - \varepsilon(n).$$

6.4 Main technical lemmas

Our first technical lemma states that we can construct a protocol with a large entropy gap for zero-knowledge proofs that remain secure under parallel composition.

Lemma 6.6 (CZKP to entropy gap). *Suppose that nonuniformly secure one-way functions exist and that every language in NP has a m -round (computational) zero-knowledge proof (P, V) that is black-box zero knowledge under parallel composition, where $m = \text{poly}(n)$. In addition, suppose (P, V) has an efficient prover. Then, there exists an m -round protocol $(A, B)(1^n)$ such that on security parameter n , there exist integers ℓ_1, \dots, ℓ_m for which:*

- For all $i \in [m]$, the real min-entropy of A in round i of (A, B) is at least ℓ_i .
- The context-independent accessible max-entropy of A in (A, B) is at most $\sum \ell_i - 3nm$.

The next lemma (analogous to Lemma 4.7 in Section 4) states that we may exploit the entropy gap given by the previous lemma to obtain a $O(m)$ -phase commitment scheme. We note here that the transformation is exactly as before: using interactive hashing on the sender's messages, as specified in Protocol 4.14 (which in turn uses Protocol 4.11).

Lemma 6.7 (entropy gap to m -phase commitment scheme). *Suppose that nonuniformly secure one-way functions exist. Then, there exists an efficient transformation GapToComTransform that takes as input a security parameter 1^n and an efficient m -round interactive protocol $\pi = (A, B)$, a set of indices $I \subseteq [m]$ and a set of integers $\{\ell_i\}_{i \in I}$ where $\ell_i \geq 3n$ for all $i \in I$, and outputs an $O(m)$ -round $|I|$ -phase commitment scheme $\text{Com} = (S, R)$ with the following properties:*

hiding: *If for every $i \in I$, it holds that A has real min-entropy at least ℓ_i in round i with respect to π , then Com is statistically hiding.*

binding: *If the context-independent accessible max-entropy of A with respect to I in π is at most $\sum_{i \in I} \ell_i - 3n|I|$, then Com is $\binom{m}{1}$ binding.*

Proof of Theorem 6.1. Combining the two technical lemmas, we obtain a constant-round m -phase commitment scheme Com where m is also a constant. We may then proceed as in Sections 4.2 and 4.3 to transform Com into a constant-round standard commitment scheme. We stress here that since we are working with a constant m , we may try all possible settings of ℓ_1, \dots, ℓ_m as described at the end of Section 4.3.

6.5 Proof of Lemma 6.6

The protocol (A, B). Assuming the existence of one-way functions, there exists a language $L \in \text{NP}$ with corresponding relation R_L with distribution $(D_Y, W_Y)(1^n)$ on YES instance-witness pairs in R_L and a distribution D_N on NO instances such that

- $(D_Y, W_Y)(1^n)$ can be sampled in time $\text{poly}(n)$.
- D_Y and D_N are computationally indistinguishable.

Specifically, let $G : \{0, 1\}^{n/3} \rightarrow \{0, 1\}^n$ be a pseudorandom generator [HILL]. Then, $(D_Y, W_Y) = (G(U_{n/3}), U_{n/3})$ and D_N is the uniform distribution over $\{0, 1\}^n \setminus G(\{0, 1\}^{n/3})$. Let (P, V) be the zero-knowledge proof system for L with simulator Sim . We also assume that V sends its random tape as the last message. By parallel repetition, we may assume that (P, V) has completeness and soundness errors at most $1/100$, and let k_{REAL} denote the length of V 's random tape.

Now, consider the following protocol $(A, B)(1^n)$:

1. B samples t pairs $(x_1, \omega_1), \dots, (x_t, \omega_t)$ from (D_Y, W_Y) and sends (x_1, \dots, x_t) to A.
2. The two parties run $(P^t(\omega_1, \dots, \omega_t), V^t)(x_1, \dots, x_t)$, with B and A acting as P^t and V^t respectively.

We say that a transcript of the protocol (A, B) is *accepting* if A (acting as V) accepts for a majority of the instances x_1, \dots, x_t ; otherwise, we say that the transcript is *rejecting*.

Claim 6.8. *The protocol (A, B) satisfies the following properties:*

completeness: *A random transcript from the protocol (A, B) is rejecting with probability at most 2^{-t} .*

soundness: *For every (possibly inefficient) strategy B^* that sends $x_1, \dots, x_t \notin L$ in the first message, a random transcript from the protocol (A, B^*) is accepting with probability at most 2^{-t} .*

Proof. From the Chernoff bound, we know that a random transcript from the protocol (A, B) is rejecting with probability at most 2^{-t} (by the completeness of (P, V)). In addition, if $x_1, \dots, x_t \notin L$, then the probability that V accepts for a majority of the instances in a t -fold repetition of (P, V) on input (x_1, \dots, x_t) is at most 2^{-t} (by soundness to (P, V)). \square

Setting the parameters. Suppose the conditional Shannon entropy of the i 'th verifier message in (P, V) is k_i (for a random instance from D_Y); then, $k_1 + \dots + k_m = k_{\text{REAL}}$. Then,

1. By Proposition 2.3, the real min-entropy of A in round i is at least $\ell_i := tk_i - ut^{2/3}$ (where u is the length of the longest verifier message in (P, V)).
2. The real Shannon entropy of A in the rounds $[m]$ is tk_{REAL} .

We set $t = \text{poly}(n)$ so that $t > 4m \cdot (3n + ut^{2/3})$. As we shall show next, the context-independent accessible max-entropy of (A, B) is at most $tk_{\text{REAL}} - t/4 < tk_{\text{REAL}} - m(3n + ut^{2/3}) \leq \sum \ell_i - 3nm$.

Analysis for accessible max-entropy. Next, we shall establish the main technical claim of this section:

Claim 6.9. *The context-independent accessible max-entropy of A in (A, B) is at most $tk_{\text{REAL}} - t/4$.*

Suppose on the contrary that there exists a *ppt* A^* along with some predicate `success` that violates the guarantee on context-independent accessible max-entropy. This means in particular that there exists a non-negligible function $\varepsilon = \varepsilon(n)$ such that

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B)}[\text{success}(v)] > \varepsilon$$

Then, the high-level proof strategy, following the earlier outline (which is in turn based on [GK2]), is as follows:

Constructing a cheating verifier. Following [GK2], we will use A^* (with some modifications) as an adversarial verifier V^* wherein “rewinding doesn’t help” for black-box simulators. Then, by exploiting the small completeness error, the zero-knowledge property, and the fact that the distribution of V^* ’s messages is close to that of the honest verifier V , we may show that $\text{Sim}^{V^*}(D_Y^t)$ outputs accepting transcripts with high probability. Moreover, since D_Y and D_N are computationally indistinguishable, we also have that $\text{Sim}^{V^*}(D_N^t)$ outputs accepting transcripts with high probability.

Constructing a cheating prover. Again, following [GK2, Kat], we may construct from Sim^{V^*} a “simulation-based” stand-alone cheating prover B^* that sends a t -tuple of random `NO` instances drawn from D_N^t , interacts with A^* and convinces A^* to output accepting transcripts with noticeable probability. Now, by using the fact that A^* achieves high context-independent accessible max-entropy, we may argue that the distribution of A^* messages while interacting with B^* is close to that of the honest verifier V . This contradicts the soundness of (P, V) .

Step 1: constructing V^* . Recall that at each round i , A^* flips fresh random coins s_i to generate its next message a_i and a justification string w_i . Given $A^*(1^n)$, we construct a PPT cheating verifier V^* . On input $\bar{x} = (x_1, \dots, x_t)$ and auxiliary input h where h is a T -wise independent hash function where $T = \text{poly}(n)$ is an upper bound on the number of queries made by $\text{Sim}(\bar{x})$, $V_h^*(\bar{x}, h)$ does the following:

1. Pass the instances (x_1, \dots, x_t) to A^* as if coming from B .
2. Upon receiving a query $q_i = (b_1, a_1, \dots, b_i)$ (containing the first i prover messages), we compute $s_i = h(q_i)$ along with $(a_i, \omega_i) = A^*(q_i; s_0, s_1, \dots, s_i)$, where s_0, s_1, \dots, s_{i-1} are the randomness of A^* in the previous rounds, and respond with a_i .
3. Finally, V^* outputs the view of A^* .

Claim 6.10.

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B)}[\text{AccH}_{A, A^*}^{[m]}(v) > tk_{\text{REAL}} - t/4 \text{ and } v \text{ contains a rejecting } A\text{-consistent transcript}] \leq \text{neg}(n)$$

Proof. Let random variables $(S_0, B_1, S_1, A_1, W_1, \dots, B_m, S_m, A_m, W_m)$ denote the view of A^* when interacting with B . By Lemma 2.5, for each $i \in [m]$, and for all A^* -consistent views $v = (s_0, b_1, s_1, a_1, w_1, \dots)$, conditioned on $(S_0, B_1, A_1, W_1, \dots, B_i) = (s_0, b_1, a_w, w_1, \dots, b_i)$, with probability at least $1 - 2^{-t/8m}$ over $a_i \stackrel{R}{\leftarrow} A_i$, we have that

$$\text{AccH}_{A, A^*}(a_i | v) = H_{\Gamma_i^A(v, S_i, A_i, W_i)}^*(a_i) \leq H_{A_i}^*(a_i) + t/8m$$

This means that for all $i \in [m]$, with probability at least $1 - 2^{-t/8m} = 1 - \text{neg}(n)$ over $v = (s_0, b_1, s_1, a_1, w_1, \dots) \stackrel{R}{\leftarrow} \text{view}_{A^*}(A^*, B)$,

$$\begin{aligned} & H_{A_i | S_0, B_1, S_1, A_1, \dots, B_i}(a_i | s_0, b_1, s_1, a_1, \dots, b_i) \\ & \geq H_{A_i | S_0, B_1, S_1, A_1, \dots, B_i}^*(a_i | s_0, b_1, s_1, a_1, \dots, b_i) \\ & \geq \text{AccH}_{A, A^*}^i(a_i | v) - t/8m \end{aligned}$$

Taking a union bound and summing over $i \in [m]$, we have that with probability at least $1 - \text{neg}(n)$ over $v \stackrel{R}{\leftarrow} \text{view}_{A^*}(A^*, B)$,

$$\sum_{i=1}^m H_{A_i | S_0, B_1, S_1, A_1, \dots, B_i}(a_i | s_0, b_1, s_1, a_1, \dots, b_i) \geq \text{AccH}_{A, A^*}^{[m]}(v) - t/8$$

Now, let S denote the set of A^* -consistent views $(s_0, b_1, s_1, a_1, w_1, \dots)$ containing a rejecting A -consistent transcript and for which

$$\sum_{i=1}^m H_{A_i | S_0, B_1, S_1, A_1, \dots, B_i}(a_i | b_1, a_1, \dots, b_i) \geq tk_{\text{REAL}} - 3t/8$$

Then, it suffices to show that

$$\Pr[\text{view}_{A^*}(A^*, B) \in S] \leq \text{neg}(n)$$

Observe that for all $v = (s_0, b_1, s_1, a_1, w_1, \dots) \in S$:

$$\begin{aligned} & \Pr[\text{view}_{A^*}(A^*, B) = v] \\ & = \Pr[S_0 = s_0] \cdot \Pr[B_1 = b_1 | S_0 = s_0] \cdot \Pr[A_1 = a_1 | (S_0, B_1) = (s_0, b_1)] \cdot \\ & \quad \Pr[S_1 = s_1 | (S_0, B_1, A_1) = (s_0, b_1, a_1)] \cdots \\ & = \prod_{i=1}^m \Pr[B_i = b_i | (B_1, A_1, \dots, B_{i-1}, A_{i-1}) = (b_1, a_1, \dots, b_{i-1}, a_{i-1})] \cdot \prod_{i=1}^m 2^{-H_{A_i | S_0, B_1, S_1, A_1, \dots, B_i}(a_i | b_1, a_1, \dots, b_i)} . \\ & \quad \Pr[S_0 = s_0] \cdot \prod_{i=1}^m \Pr[S_i = s_i | (S_0, B_1, \dots, A_i) = (s_0, b_1, \dots, a_i)] \\ & \leq \prod_{i=1}^m \Pr[B_i = b_i | (B_1, A_1, \dots, B_{i-1}, A_{i-1}) = (b_1, a_1, \dots, b_{i-1}, a_{i-1})] \cdot 2^{-tk_{\text{REAL}} + 3t/8} . \\ & \quad \Pr[S_0 = s_0] \cdot \prod_{i=1}^m \Pr[S_i = s_i | (S_0, B_1, \dots, A_i) = (s_0, b_1, \dots, a_i)] \\ & = 2^{3t/8} \cdot \Pr[(A, B) = (b_1, a_1, \dots, b_m, a_m)] \cdot \Pr[S_0 = s_0] \cdot \prod_{i=1}^m \Pr[S_i = s_i | (S_0, B_1, \dots, A_i) = (s_0, b_1, \dots, a_i)], \end{aligned}$$

where the last equality holds because A sends its random coins as the last message. Summing over views in S , we have:

$$\begin{aligned}
& \Pr[\text{view}_{A^*}(A^*, B) \in S] \\
& \leq 2^{3t/8} \cdot \sum_{\substack{\text{rejecting transcripts} \\ (b_1, a_1, \dots, b_m, a_m)}} \left(\Pr[(A, B) = (b_1, a_1, \dots, b_m, a_m)] \cdot \right. \\
& \quad \left. \sum_{s_0, s_1, \dots, s_m} \Pr[S_0 = s_0] \cdot \prod_{i=1}^m \Pr[S_i = s_i \mid (S_0, B_1, \dots, A_i) = (s_0, b_1, \dots, a_i)] \right) \\
& \leq 2^{3t/8} \cdot \Pr[(A, B) \text{ outputs a rejecting transcript}] = \text{neg}(n)
\end{aligned}$$

Here, we use the fact (A, B) has completeness error 2^{-t} (Claim 6.8). \square

The following claim follows readily from the preceding claim and the fact that A^* violates the guarantee on context-independent accessible entropy (and $\text{success}(v)$ implies the transcript is A -consistent):

Claim 6.11.

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B)}[\text{success}(v) \text{ and } v \text{ is rejecting}] \leq \text{neg}(n)$$

Step 2: constructing B^* . Starting from A^* , we define a new (inefficient) strategy B^* which interacts with an external party A' (which could be A^* or A) as follows:

1. On input 1^n , sample and send to A' a t -tuple of instances $(x_1, \dots, x_t) \leftarrow D_N^t$ and pick a random subset $U = \{j_1, j_2, \dots, j_m\} \subset [T]$ of size m .
2. Internally simulate $\text{Sim}^{V^*}(x)$ step by step. We handle the j 'th query q_j , $j = 1, 2, \dots, T$, that Sim makes to V^* as follows. Suppose the query is of the form $q_j = (b_1, a_1, \dots, b_i)$.
 - if $j \notin U$: Pick and store a random string s_i associated with the query and look up the random strings s_1, \dots, s_{i-1} associated with the previous $i-1$ prefixes. Respond with $A^*(q_j; s_1, \dots, s_i)$.
 - if $j \in U$: If $i > 1$, output **fail** if we did not previously send b_1, \dots, b_{i-1} as the first $i-1$ messages to the external A' . Otherwise, forward b_i to the external A' as if the message comes from B^* and wait for the response a_i from A' . Next, look up the previous random strings s_1, \dots, s_{i-1} associated with the $i-1$ prefixes, and sample uniformly an s_i satisfying $A^*(q_j; s_1, \dots, s_i) = a_i$; output **fail** if no such s_i exists. Store s_i as the random string associated with q_j .
3. Output **fail** if Sim does not output the same transcript as occurred in the interaction with A' .

The following claim will be useful later:

Claim 6.12. For all views v of A^* : $\Pr[\text{view}_{A^*}(A^*, B^*) = v] \geq \frac{1}{T^m} \Pr[\text{Sim}^{V^*}(D_N^t) = v]$.

Proof. The claim follows readily from the following observations (c.f. [GK2]):

- $\{\text{Sim}^{V_H}(D_N^t)\}$ (where H is a randomly chosen t -wise independent hash function) and $\{(A^*, B^*)\}_{B^* \neq \text{fail}}$ are identically distributed.
- $\Pr[B^* \neq \text{fail in } (A^*, B^*)] \geq 1/T^m$.

□

Claim 6.13.

$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B^*)}[\text{AccH}_{A, A^*}^{[m]}(v) > tk_{\text{REAL}} - t/4 \text{ and } v \text{ contains an accepting } A\text{-consistent transcript}] \leq \text{neg}(n)$

Proof. Let random variables $(S_0, B_1, S_1, A_1, W_1, \dots, B_m, S_m, A_m, W_m)$ denote the view of A^* when interacting with B^* . Following the proof of Claim 6.10, let S denote the set of A^* -consistent views $(s_0, b_1, s_1, a_1, w_1, \dots)$ containing an accepting A -consistent transcript and for which

$$\sum_{i=1}^m H_{A_i | S_0, B_1, S_1, A_1, \dots, B_i}(a_i | b_1, a_1, \dots, b_i) \geq tk_{\text{REAL}} - 3t/8$$

Then, it suffices (as before) to show that

$$\Pr[\text{view}_{A^*}(A^*, B^*) \in S] \leq \text{neg}(n)$$

As before, for all $v = (s_0, b_1, s_1, a_1, w_1, \dots) \in S$:

$$\begin{aligned} & \Pr[\text{view}_{A^*}(A^*, B^*) = v] \\ &= 2^{3t/8} \cdot \Pr[(A, B) = (b_1, a_1, \dots, b_m, a_m)] \cdot \Pr[S_0 = s_0] \cdot \prod_{i=1}^m \Pr[S_i = s_i | (S_0, B_1, \dots, A_i) = (s_0, b_1, \dots, a_i)], \end{aligned}$$

Summing over views in S , we have (as before):

$$\Pr[\text{view}_{A^*}(A^*, B^*) \in S] \leq 2^{3t/8} \cdot \Pr[(A, B^*) \text{ outputs an accepting transcript}] \leq 2^{-5t/8} = \text{neg}(n)$$

Here, we use the fact A has soundness error 2^{-t} against strategies B^* that send $x_1, \dots, x_t \notin L$ in the first message (Claim 6.8). □

The following claim follows readily from the preceding claim and the fact that A^* violates the guarantee on context-independent accessible entropy:

Claim 6.14.

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B^*)}[\text{success}(v) \text{ and } v \text{ is accepting}] \leq \text{neg}(n)$$

Step 3: deriving a contradiction. It follows from Claim 6.11 that

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B)}[\text{success}(v) \text{ and } v \text{ is accepting}] \geq \varepsilon - \text{neg}(n)$$

By the zero-knowledge property and the indistinguishability of D_Y, D_N , we have:

$$\Pr_{v \leftarrow \text{Sim}^{V^*}(D_N^t)}[\text{success}(v) \text{ and } v \text{ is accepting}] \geq \varepsilon - \text{neg}(n)$$

By Claim 6.12, this implies:

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B^*)}[\text{success}(v) \text{ and } v \text{ is accepting}] \geq \varepsilon/2T^m$$

which contradicts Claim 6.14.

6.6 Proof of Lemma 6.7

As noted earlier, we use Protocol 4.14 (using hashing to select the sender's messages) to transform (A, B) into a m -phase commitment scheme; the only difference being that we instantiate the protocol with a family \mathcal{F} of universal one-way hash functions with *non-uniform* security. The proof of Lemma 6.7 is analogous to that of Lemma 4.7. Hiding is exactly as before, and follows immediately from Claim 4.15. As such, it suffices to establish the following strengthening of Claim 4.16 used to establish $\binom{m}{1}$ -binding.

Claim 6.15 ($\binom{m}{1}$ -binding). *Assume that \mathcal{F} is a family of universal one-way hash functions with non-uniform security and that the context-independent accessible max-entropy of A with respect to I in (A, B) is at most $\sum_{i \in I} \ell_i - 3n |I|$. Then, Com is $\binom{m}{1}$ -binding.*

Proof. Suppose there exists a PPT S^* playing the role of S in Com that breaks the $\binom{m}{1}$ -binding of Com with non-negligible probability ε . We need to show that there exists a A^* along with a predicate success that breaks the context-independent accessible max-entropy of A in (A, B) . We proceed in several steps:

Step 1: constructing A^* . Consider the same algorithm A^* for interacting with any B^* (where B^* may be B) as in the proof of Claim 4.16, that is: A^* emulates a random execution of (S^*, R) while emulating R using the B^* it interacts with, and choosing the messages of R that are not of embedded B^* (specifically, those of R_H), uniformly at random. In more detail, A^* proceeds as follows:

1. A^* uses s_0 to pick a random tape for S^* .
2. For $i = 1, 2, \dots, m$, upon receiving b_i from B^* , A^* proceeds as follows:

If $i \in I$, A^* uses s_i to pick (h_{len}, h_2, f) corresponding to the messages of R_H in the i 'th commit phase of Com and sends a_i , where a_i is the message S^* sends in the i 'th reveal phase of Com . In addition, when S^* locally outputs (w_i^0, w_i^1) , A^* locally outputs w_i^0 .

If $i \notin I$, A^* forwards whatever message a_i sent by S^* upon receiving b_i .

Step 2: the predicate success . Let success denote the predicate on a view v from $\text{Supp}(\text{view}_{A^*}(A^*, B^*))$ which evaluates to true iff v is A -consistent and S^* breaks the $\binom{m}{1}$ -binding of Com for all the phases in I . We claim that A^* and success violates the guarantee on context-independent accessible max-entropy. By definition, success satisfies the first condition. Next, observe that the view of S^* in a real execution of (S^*, R) and in the emulated execution (A^*, B) are identically distributed. It follows from the definition of success and the assumption about S^* that

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B)}[\text{success}(v)] \geq \varepsilon,$$

Step 3: quantifying over B^* . It remains to show that for every (possibly inefficient) strategy B^* ,

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B^*)} \left[\neg \text{success}(v) \text{ or } \text{AccH}_{A, A^*}^I(v) > \sum_{i \in I} \ell_i - 3n |I| \right] \geq 1 - \text{neg}(n).$$

Suppose otherwise; that is, there exists a strategy B^* and a non-negligible function ε_0 such that

$$\Pr_{v \leftarrow \text{view}_{A^*}(A^*, B^*)} \left[\text{success}(v) \text{ and } \text{AccH}_{A, A^*}^I(v) \leq \sum_{i \in I} \ell_i - 3n |I| \right] \geq \varepsilon_0.$$

Then, by an averaging argument, there exists an $i^* \in I$ and a partial view $v_{i^*}^* = (s_0, b_1, a_1, w_1, \dots, b_{i^*})$ such that

$$\Pr_{v \leftarrow \text{view}_{A^*}((A^*, B^*)|v_{i^*}^*)} \left[\text{success}(v) \text{ and } \text{AccH}_{A, A^*}(a_{i^*} \mid v_{i^*}^*) \leq \ell_{i^*} - 3n \right] \geq \varepsilon_0/m.$$

We define $L_{v_{i^*}^*}$ to be the set $\{a_{i^*} \neq \perp \mid \text{AccH}_{A, A^*}(a_{i^*} \mid v_{i^*}^*) \leq \ell_{i^*} - 3n\}$. It is clear that $|L_{v_{i^*}^*}| \leq 2^{\ell_{i^*} - 3n}$. Then,

$$\Pr_{v \leftarrow \text{view}_{A^*}((A^*, B^*)|v_{i^*}^*)} \left[\text{success}(v) \text{ and } a_{i^*} \in L_n \right] \geq \varepsilon_0/m.$$

Now, we can construct a nonuniform S_H^* that contradicts the binding property of (S_H, R_H) as formalized in Lemma 4.12. (Here, we use the fact that \mathcal{F} has non-uniform security, and therefore the binding property of (S_H, R_H) holds even against non-uniform adversaries.) Specifically, since S_H^* is non-uniform, it can have $v_{i^*}^*$ hardwired in and thus need not generate it efficiently. \square

Acknowledgements

We thank Oded Goldreich for helpful discussions.

References

- [AH] W. Aiello and J. Håstad. Statistical Zero-Knowledge Languages can be Recognized in Two Rounds. *JCSS*, 42(3):327–345, 1991.
- [BSW] B. Barak, R. Shaltiel, and A. Wigderson. Computational Analogues of Entropy. In *RANDOM-APPROX*, 2003.
- [BM] M. Blum and S. Micali. How to Generate Cryptographically Strong Sequences of Pseudo Random Bits. pages 112–117, 1982.
- [DHRS] Y. Z. Ding, D. Harnik, A. Rosen, and R. Shaltiel. Constant-Round Oblivious Transfer in the Bounded Storage Model. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004*, pages 446–472, 2004.
- [FS] U. Feige and A. Shamir. Witness Indistinguishable and Witness Hiding Protocols. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 416–426. ACM Press, 1990.
- [Gol] O. Goldreich. Concurrent zero-knowledge with timing, revisited. In *STOC*, pages 332–340, 2002.
- [GK1] O. Goldreich and A. Kahan. How to Construct Constant-Round Zero-Knowledge Proof Systems for NP. *Journal of Cryptology*, 9(3):167–190, 1996.

- [GK2] O. Goldreich and H. Krawczyk. On the Composition of Zero-Knowledge Proof Systems. *SIAM J. Comput.*, 25(1):169–192, 1996. Preliminary version in *ICALP’90*.
- [GMW] O. Goldreich, S. Micali, and A. Wigderson. Proofs that Yield Nothing But Their Validity or All Languages in NP Have Zero-Knowledge Proof Systems. *Journal of the ACM*, 38(1):691–729, 1991. Preliminary version in *FOCS’86*.
- [GV] O. Goldreich and S. P. Vadhan. Comparing Entropies in Statistical Zero Knowledge with Applications to the Structure of SZK. In *IEEE Conference on Computational Complexity*, pages 54–, 1999.
- [GM] S. Goldwasser and S. Micali. Probabilistic Encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
- [HHRS] I. Haitner, J. J. Hoch, O. Reingold, and G. Segev. Finding Collisions in Interactive Protocols – A Tight Lower Bound on the Round Complexity of Statistically-Hiding Commitments. In *Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS)*, 2007.
- [HHK⁺] I. Haitner, O. Horvitz, J. Katz, C. Koo, R. Morselli, and R. Shaltiel. Reducing Complexity Assumptions for Statistically-Hiding Commitment. In *Advances in Cryptology – EUROCRYPT 2005*, 2005.
- [HNO⁺] I. Haitner, M. Nguyen, S. J. Ong, O. Reingold, and S. Vadhan. Statistically Hiding Commitments and Statistical Zero-Knowledge Arguments from any One-Way Function. *SIAM Journal on Computing*, 2009. To appear. Preliminary versions in *FOCS ’06* and *STOC ’07*.
- [HR1] I. Haitner and O. Reingold. A New Interactive Hashing Theorem. In *Proceedings of the 18th Annual IEEE Conference on Computational Complexity*, 2007. Full version on authors’ homepage.
- [HR2] I. Haitner and O. Reingold. Statistically-Hiding Commitment from Any One-Way Function. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*. ACM Press, 2007.
- [HRVW] I. Haitner, O. Reingold, S. Vadhan, and H. Wee. Inaccessible Entropy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC ’09)*, 31 May–2 June 2009. To appear.
- [HILL] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. Preliminary versions in *STOC’89* and *STOC’90*.
- [ILL] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random Generation from one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 12–24. ACM Press, 1989.
- [Kat] J. Katz. Which Languages Have 4-Round Zero-Knowledge Proofs? In *TCC*, pages 73–88, 2008.

- [KK] J. Katz and C. Koo. On Constructing Universal One-Way Hash Functions from Arbitrary One-Way Functions. Technical Report 2005/328, Cryptology ePrint Archive, 2005.
- [Nao] M. Naor. Bit Commitment Using Pseudorandomness. *Journal of Cryptology*, 4(2):151–158, 1991. Preliminary version in *CRYPTO'89*.
- [NOVY] M. Naor, R. Ostrovsky, R. Venkatesan, and M. Yung. Perfect Zero-Knowledge Arguments for NP Using Any One-Way Permutation. *Journal of Cryptology*, 11(2):87–108, 1998. Preliminary version in *CRYPTO'92*.
- [NY] M. Naor and M. Yung. Universal One-Way Hash Functions and their Cryptographic Applications. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 33–43. ACM Press, 1989.
- [NV] M. Nguyen and S. Vadhan. Zero Knowledge with Efficient Provers. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 287–295. ACM Press, 2006.
- [NZ] N. Nisan and D. Zuckerman. Randomness is Linear in Space. *Journal of Computer and System Sciences*, 52(1):43–52, 1996.
- [OV] S. J. Ong and S. P. Vadhan. An Equivalence Between Zero Knowledge and Commitments. In *TCC*, pages 482–500, 2008.
- [OW] R. Ostrovsky and A. Wigderson. One-Way Functions are Essential for Non-Trivial Zero-Knowledge. In *Proceedings of the 2nd Israel Symposium on Theory of Computing Systems*, pages 3–17. IEEE Computer Society, 1993.
- [Pas] R. Pass. Parallel Repetition of Zero-Knowledge Proofs and the Possibility of Basing Cryptography on NP-Hardness. In *IEEE Conference on Computational Complexity*, pages 96–110, 2006.
- [PT] E. Petrank and G. Tardos. On the Knowledge Complexity of NP. In *FOCS*, pages 494–503, 1996.
- [RW] R. Renner and S. Wolf. Smooth Renyi Entropy and Applications. In *IEEE International Symposium on Information Theory — ISIT 2004*, page 233. IEEE, June 2004.
- [Rom] J. Rompel. One-Way Functions are Necessary and Sufficient for Secure Signatures. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 387–394, 1990.
- [Sha] C. Shannon. Communication Theory of Secrecy Systems. *Bell System Technical Journal*, 28(4):656–715, 1949.
- [Yao] A. C. Yao. Theory and Applications of Trapdoor Functions. In *FOCS*, pages 80–91, 1982.