Inaccessible Entropy

Iftach Haitner
Microsoft Research
Cambridge, MA
iftach@microsoft.com

Omer Reingold*
Weizmann Institute of Science
Rehovot, Israel
omer.reingold@weizmann.ac.il

Salil Vadhan†
Harvard University
Cambridge, MA
salil@eecs.harvard.edu

Hoeteck Wee‡
Queens College, CUNY
Flushing, NY
hoeteck@cs.qc.cuny.edu

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).

Categories and Subject Descriptors:
F.0 [**Theory of Computation**]: General.

General Terms: Theory, Security.

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

*Supported by US-Israel BSF grant 2006060.

†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.

‡Part of this work was done while a post-doc at Columbia University, supported in part by NSF Grants CNS-0716245 and SBE-0245014.

STOC'09, May 31–June 2, 2009, Bethesda, Maryland, USA.

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* [11], which is the computational analogue of statistical distance, enabled bypassing Shannon's impossibility results on perfectly secure encryption [28], and provided the basis for the computational theory of pseudorandomness [3, 29]. A computational analogue of entropy, known as *pseudoentropy*, introduced by Håstad, Impagliazzo, Levin, and Luby [16], 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 $\mathrm{H}(X) := \mathrm{E}_{x \xleftarrow{R} X}[\log(1/\mathsf{Pr}[X = x])$, which measures the number of "bits of randomness" in X (on average). We will refer to $\mathrm{H}(X)$ as the *real entropy* of X to contrast with the computational analogues that we study. Håstad et al. [16] 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 $G : \{0,1\}^\ell \to \{0,1\}^n$ 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 $\mathrm{poly}(\log n)$. While several other computational analogues of entropy have been studied in the literature (cf., [2]), 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 $h : \{0,1\}^n \to \{0,1\}^m$ 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 $(B_1, A_1, \ldots, B_m, A_m)$ 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

$$\sum_i \mathrm{H}(A_i | B_1, A_1, \ldots, B_i),$$

where $\mathrm{H}(X|Y) = \mathrm{E}_{y \xleftarrow{R} Y}[\mathrm{H}(X|_{Y=y})]$ 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 $(B_1, S_1, A_1, W_1, \ldots, B_m, S_m, A_m, W_m)$ be random variables corresponding to the view of A^* when interacting with B. Then we define the *accessible entropy* achieved by A^* to be

$$\sum_i \mathrm{H}(A_i | B_1, S_1, A_1, W_1, \ldots, B_i).$$

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 $m + \mathrm{neg}(n)$, 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 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 B is replaced by a cheating strategy 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 [13].)

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* S and a *receiver* 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 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 [18] 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 [16], 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 $k - \mathrm{neg}(n)$. On the other hand, computational binding property says that the *accessible entropy* of m after the commit stage is at most $\mathrm{neg}(n)$.

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* (A,B) *in which the real entropy of* 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 [4] and universal one-way hash functions [20, 27]), 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 [21, 14].)

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 [14], 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 [14]. 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* $f : \{0,1\}^n \to \{0,1\}^n$*, 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 [19], which was designed to construct statistically hiding commitments from one-way *permutations*. A (different) variant of the [19] protocol was also used as the first step in the previous construction of statistically hiding commitments from one-way functions in [14]. 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" [12].

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 [16, 18]): 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 [9] 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 [16, 18]). 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 [5, 8]. However, we do know how to construct zero-knowledge proofs for NP that remain secure under parallel composition [7, 6] 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 [24], 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 [23], 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 [8], who considered 3-round protocols and (constant-round) public-coin protocols, settings in which perfect resampling is trivial.

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 [12]), 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 [12]).

We note that due to space limitations, many constructions and proofs are omitted in this extended abstract.

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, $\mathsf{Pr}[X = X]$ 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 $\mathrm{Supp}(X) := \{x : \mathsf{Pr}[X = x] > 0\}$. 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.

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

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 $x \in \mathrm{Supp}(X)$, we define the sample-entropy of x with respect to X to be the quantity

$$\mathrm{H}_X(x) := \log(1/\mathsf{Pr}[X = x]).$$

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* $\mathrm{H}(X)$ and *min-entropy* $\mathrm{H}_\infty(X)$ as follows:

$$\begin{aligned} \mathrm{H}(X) &:= \mathop{\mathrm{E}}_{x \xleftarrow{\mathrm{R}} X} [\mathrm{H}_X(x)] \\ \mathrm{H}_\infty(X) &:= \min_{x \in \mathrm{Supp}(X)} \mathrm{H}_X(x) \end{aligned}$$

We will also discuss the *max-entropy* $\mathrm{H}_0(X) := \log(1/|\mathrm{Supp}(X)|)$. The term "max-entropy" and its relation to the sample-entropy will be made apparent below.

It can be shown that $\mathrm{H}_\infty(X) \leq \mathrm{H}(X) \leq \mathrm{H}_0(X)$ with equality if and only if X is flat. Thus, saying $\mathrm{H}_\infty(X) \geq k$ is a strong way

[1] The results of [23] 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).

of saying that X has "high entropy" and $\mathrm{H}_0(X) \leq k$ 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 $|\mathrm{H}(X) - \mathrm{H}(Y)| \leq \varepsilon \cdot \log |\mathcal{U}|$. 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 $\mathrm{H}_\infty(X) \geq k$ (resp., $\mathrm{H}_0(X) \leq k$) 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 $x \xleftarrow{R} X$, we have $\mathrm{H}_X(x) \geq k$. Then X is ε-close to a random variable X' such that $\mathrm{H}_\infty(X') \geq k$.*

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

Lemma 2.2. *1. Suppose that with probability at least $1-\varepsilon$ over $x \xleftarrow{R} X$, we have $\mathrm{H}_X(x) \leq k$. Then X is ε-close to a random variable X' such that $\mathrm{H}_0(X') \leq k$.*

2. Suppose that X is ε-close to a random variable X' such that $\mathrm{H}_0(X') \leq k$. Then with probability at least $1 - 2\varepsilon$ over $x \xleftarrow{R} X$, we have $\mathrm{H}_X(x) \leq k + \log(1/\varepsilon)$.

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.

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 $\mathrm{H}(X) = 1$ 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

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

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

$$\begin{aligned} \mathrm{H}^*(X) &= \mathop{\mathrm{E}}_{x \xleftarrow{R} X}[\mathrm{H}^*_X(x)] \\ &= \mathrm{H}(X | I(X = \perp)), \end{aligned}$$

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

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 $\mathrm{H}^*(X') \leq \mathrm{H}^*(X)$.

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 $x \xleftarrow{R} X$, we have $\mathrm{H}^*_X(x) \leq k$. For this notion, it can be shown that failing more often cannot increase entropy by much:

Lemma 2.4. *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 $x \xleftarrow{R} X'$, $\mathrm{H}^*_{X'}(x) \leq \mathrm{H}^*_X(x) + \log(1/\varepsilon)$.

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

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 $(\mathsf{A}, \mathsf{B})(1^n)$ 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 $(B_1, A_1, \ldots, B_m, A_m)$ 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 $(B_1, A_1, \ldots, B_i)$.

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).

[2]The term "smoothed entropy" was coined by Renner and Wolf [26] but the notion of smoothed min-entropy has commonly been used (without a name) in the literature on randomness extractors [22].

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

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

$$\mathrm{RealH}_{\mathsf{A}}(a_i|t) := \mathrm{H}_{A_i(t)}(a_i).$$

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

$$\mathrm{RealH}^I_{\mathsf{A}}(t) = \sum_{i \in I} \mathrm{RealH}_{\mathsf{A}}(a_i|b_1, a_1, \ldots, b_i).$$

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

$$\mathop{\mathrm{E}}_{t \xleftarrow{R} (\mathsf{A},\mathsf{B})} \left[\mathrm{RealH}^I_{\mathsf{A}}(t)\right] \geq k.$$

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

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

$$\mathsf{Pr}_{t \xleftarrow{R} (\mathsf{A},\mathsf{B})} \left[\mathrm{RealH}^I_{\mathsf{A}}(t) \geq k\right] \geq 1 - \varepsilon(n).$$

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

Lemma 3.3. *For an m-round interactive protocol (A, B), let $(B_1, A_1, \ldots, B_m, A_m)$ 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. Then*

$$\mathop{\mathrm{E}}_{t \xleftarrow{R} (\mathsf{A},\mathsf{B})} \left[\mathrm{RealH}^I_{\mathsf{A}}(t)\right] = \sum_{i \in I} \mathrm{H}(A_i|B_1, A_1, \ldots, B_i).$$

The next claim follows readily from [1, 25, 10]:

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, $\mathrm{RealH}_{\mathsf{A}}(t) = k$. In particular, for every interactive algorithm B the algorithm A achives 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 $v = (s_0, b_1, s_1, a_1, w_1, \ldots, b_i, s_i, a_i, w_i)$. (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 $\Gamma^{\mathsf{A}}_j(v)$ to equal a_j if $(b_1, a_1, \ldots, b_j, a_j)$ 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]

[3] At 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, $\mathrm{AccH}_{\mathsf{A},\mathsf{A}^*}(a_i|v) := \mathrm{H}_{A_i}(a_i)$.

Definition 3.5 (accessible sample-entropy)**.** *Let* A^* *be an interactive algorithm and let* $v = (s_0, b_1, s_1, a_1, w_1, \ldots, b_i)$ *be an* A^**-consistent partial view. Define random variables* (S_i, A_i, W_i) *by choosing* S_i *uniformly at random, and setting*

$$(A_i, W_i) = \mathsf{A}^*(s_0, b_1, s_1, a_1, w_1, \ldots, b_i, S_i).$$

For a fixed value $a_i \in \mathrm{Supp}(A_i) \cup \{\perp\}$*, we define the* accessible sample-entropy of a_i given v *as*

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

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

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

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^**,*

$$\mathop{\mathrm{E}}_{v \xleftarrow{R} \mathrm{view}_{\mathsf{A}^*}(\mathsf{A}^*, \mathsf{B})} [\mathrm{AccH}^I_{\mathsf{A},\mathsf{A}^*}(v)] \leq k$$

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*

$$\mathsf{Pr}_{v \xleftarrow{R} \mathrm{view}_{\mathsf{A}^*}(\mathsf{A}^*, \mathsf{B})}[\mathrm{AccH}^I_{\mathsf{A},\mathsf{A}^*}(v) \leq k] \geq 1 - \varepsilon(n).$$

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* $(S_0, B_1, S_1, A_1, W_1, \ldots, B_m, S_m, A_m, W_m)$ *denoting the view of* A^* *when interacting with* B*. Then*

$$\mathop{\mathrm{E}}_{v \xleftarrow{R} \mathrm{view}_{\mathsf{A}^*}(\mathsf{A}^*, \mathsf{B})} [\mathrm{AccH}^I_{\mathsf{A},\mathsf{A}^*}(v)]$$
$$= \sum_{i \in I} \mathrm{H}^*(\Gamma_i^{\mathsf{A}}(V)|S_0, B_1, S_1, A_1, W_1, \ldots, W_{i-1}, B_i).$$

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 $(\mathsf{A}^t, \mathsf{B}^t)$ 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* $\pi = (\mathsf{A}, \mathsf{B})$ *an* m*-round protocol. For* $t \in \mathrm{poly}(n) \cap \omega(\log^3 n)$*, let* $\pi^t = (\mathsf{A}^t, \mathsf{B}^t)$ *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 *.*

The second tool, given by Proposition 3.9 shows how to turn a protocol $\pi = (\mathsf{A}, \mathsf{B})$ 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,* $\pi = (\mathsf{A}, \mathsf{B})$ *an* m*-round protocol. For every* $t \in \mathrm{poly}(n)$*, there is a* $(t+1) \cdot m$*-round protocol* $\pi' = (\mathbb{A}, \mathbb{B})$ *satisfying the following properties:*

real entropy: *Let* $I' = \{m+1, \ldots, tm\}$*. Suppose the real Shannon entropy of* 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* 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 the protocol $\pi' = (\mathbb{A}, \mathbb{B})$, $\mathbb{B}$ starts by picking a random offset and then runs $t-1$ *sequential* repetitions of π. Now, 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 $k_{\text{REAL}}{}' = |I'| \cdot k_{\text{REAL}}/m$.

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*.*

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

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 $\text{view}_{\mathsf{R}^*}(\mathsf{S}(0), \mathsf{R}^*)$ and $\text{view}_{\mathsf{R}^*}(\mathsf{S}(1), \mathsf{R}^*)$ 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 protocol[5] $\pi = (\mathsf{A}, \mathsf{B})$, 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 respect to π is at least k and the accessible max-entropy of A with respect to π is at most $(1-1/p)k$, then Com is a statistically hiding and computationally binding commitment scheme. Alternatively if $m = O(1)$, then 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 $\pi = (\mathsf{A}, \mathsf{B})$, 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 $\mathsf{Com} = (\mathsf{S}, \mathsf{R})$ 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 Com is statistically hiding.*

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 computationally binding.*

Informally, 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.

5. STATISTICALLY HIDING COMMITMENTS FROM ONE-WAY FUNCTIONS

Theorem 5.1 (restatement of Theorem 1.3)**.** *Let $f\colon \{0,1\}^n \mapsto \{0,1\}^n$ be a one-way function, then there exists an efficient $O(n/\log n)$-round protocol $\pi = (\mathsf{A}, \mathsf{B})$ 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 [14].

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

PROOF. (of Theorem 5.1) We assume for simplicity that $n/\log n \in \mathbb{N}$, and let $(\mathsf{S}_{\mathsf{CIH}}, \mathsf{R}_{\mathsf{CIH}})$ be an instantiation of the *computational interactive hashing protocol* given by [15, Protocol 3.6] (building on [19]) described next.

Protocol 5.3 (computational interactive hashing protocol)**.** $(\mathsf{S}_{\mathsf{CIH}}, \mathsf{R}_{\mathsf{CIH}})$**.**

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

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

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

1. *$\mathsf{R}_{\mathsf{CIH}}$ selects uniformly at random $h_i \in \mathcal{H}$ and sends its description to $\mathsf{S}_{\mathsf{CIH}}$.*
2. *$\mathsf{S}_{\mathsf{CIH}}$ sends $h_i(y)$ back to $\mathsf{R}_{\mathsf{CIH}}$.*

We use the following fact about Protocol $(\mathsf{S}_{\mathsf{CIH}}, \mathsf{R}_{\mathsf{CIH}})$, which can be deduced from [15, Theorem 4.1].

Proposition 5.4. *Let $f\colon \{0,1\}^n \mapsto \{0,1\}^n$ be a one-way function, and for $j \in [n]$ let*

$$T_j := \left\{ y \in \{0,1\}^n \colon 2^{j-1} \leq \left|f^{-1}(y)\right| < 2^j \right\}.$$

Then the following has negligible probability for every efficient $\mathsf{S}^*_{\mathsf{CIH}}$, $j \in [n]$ and a constant $c > 0$: after $\frac{n-j}{\log n} - c$ rounds, $\mathsf{S}^*_{\mathsf{CIH}}$ 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.

[5] Given as a pair of circuits.

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

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

Protocol 5.5. (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* $(\mathsf{S}_{\mathsf{CIH}}(y), \mathsf{R}_{\mathsf{CIH}})$*, with* A *and* B *acting* $\mathsf{S}_{\mathsf{CIH}}$ *and* $\mathsf{R}_{\mathsf{CIH}}$ *respectively.*
3. B *sends a dummy message to* A.
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 $y \in T_j := \left\{y \in \{0,1\}^n : 2^{j-1} \leq \left|f^{-1}(y)\right| < 2^j\right\}$. Proposition 5.4 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 negligble. We conclude the proof by showing that the other messages do not contribute too much accessible entropy to cover this loss.

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 [7, 6] 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 [8], 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 [8] (as well as more recent applications of the approach [17]), 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 [1, 25, 10].

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 [17, 8] 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 $\mathsf{success}$ 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 $\mathsf{success}$ 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 afore-mentioned 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 can start with a weaker security guarantee and end with the same conclusion.

Acknowledgements

We thank Oded Goldreich for helpful discussions.

7. REFERENCES

[1] AIELLO, W., AND HÅSTAD, J. Statistical zero-knowledge languages can be recognized in two rounds. *JCSS 42*, 3 (1991), 327–345.

[2] BARAK, B., SHALTIEL, R., AND WIGDERSON, A. Computational analogues of entropy. In *RANDOM-APPROX* (2003).

[3] BLUM, M., AND MICALI, S. How to generate cryptographically strong sequences of pseudo random bits. pp. 112–117.

[4] DING, Y. Z., HARNIK, D., ROSEN, A., AND SHALTIEL, R. Constant-round oblivious transfer in the bounded storage model. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004* (2004), pp. 446–472.

[5] FEIGE, U., AND SHAMIR, A. Witness indistinguishable and witness hiding protocols. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)* (1990), ACM Press, pp. 416–426.

[6] GOLDREICH, O. Concurrent zero-knowledge with timing, revisited. In *STOC* (2002), pp. 332–340.

[7] GOLDREICH, O., AND KAHAN, A. How to construct constant-round zero-knowledge proof systems for NP. *Journal of Cryptology 9*, 3 (1996), 167–190.

[8] GOLDREICH, O., AND KRAWCZYK, H. On the composition of zero-knowledge proof systems. *SIAM J. Comput. 25*, 1 (1996), 169–192. Preliminary version in *ICALP'90*.

[9] GOLDREICH, O., MICALI, S., AND WIGDERSON, A. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM 38*, 1 (1991), 691–729. Preliminary version in *FOCS'86*.

[10] GOLDREICH, O., AND VADHAN, S. P. Comparing entropies in statistical zero knowledge with applications to the structure of szk. In *IEEE Conference on Computational Complexity* (1999), pp. 54–.

[11] GOLDWASSER, S., AND MICALI, S. Probabilistic encryption. *Journal of Computer and System Sciences 28*, 2 (1984), 270–299.

[12] HAITNER, I., HOCH, J. J., REINGOLD, O., AND SEGEV, G. 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).

[13] HAITNER, I., HORVITZ, O., KATZ, J., KOO, C., MORSELLI, R., AND SHALTIEL, R. Reducing complexity assumptions for statistically-hiding commitment. In *Advances in Cryptology – EUROCRYPT 2005* (2005).

[14] HAITNER, I., NGUYEN, M., ONG, S. J., REINGOLD, O., AND VADHAN, S. 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*.

[15] HAITNER, I., AND REINGOLD, O. A new interactive hashing theorem. In *Proceedings of the 18th Annual IEEE Conference on Computational Complexity* (2007). Full version on authors' homepage.

[16] HÅSTAD, J., IMPAGLIAZZO, R., LEVIN, L. A., AND LUBY, M. A pseudorandom generator from any one-way function. *SIAM Journal on Computing 28*, 4 (1999), 1364–1396. Preliminary versions in *STOC'89* and *STOC'90*.

[17] KATZ, J. Which languages have 4-round zero-knowledge proofs? In *Theory of Cryptography, 5th Theory of Cryptography Conference, TCC 2008* (2008), pp. 73-88.

[18] NAOR, M. Bit commitment using pseudorandomness. *Journal of Cryptology 4*, 2 (1991), 151–158. Preliminary version in *CRYPTO'89*.

[19] NAOR, M., OSTROVSKY, R., VENKATESAN, R., AND YUNG, M. Perfect zero-knowledge arguments for NP using any one-way permutation. *Journal of Cryptology 11*, 2 (1998), 87–108. Preliminary version in *CRYPTO'92*.

[20] NAOR, M., AND YUNG, M. Universal one-way hash functions and their cryptographic applications. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)* (1989), ACM Press, pp. 33–43.

[21] NGUYEN, M., AND VADHAN, S. Zero knowledge with efficient provers. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)* (2006), ACM Press, pp. 287–295.

[22] NISAN, N., AND ZUCKERMAN, D. Randomness is linear in space. *Journal of Computer and System Sciences 52*, 1 (1996), 43–52.

[23] ONG, S. J., AND VADHAN, S. An equivalence between zero knowledge and commitments. In *Theory of Cryptography, 5th Theory of Cryptography Conference, TCC 2008* (2008), pp. 482-500.

[24] OSTROVSKY, R., AND WIGDERSON, A. One-way functions are essential for non-trivial zero-knowledge. In *Proceedings of the 2nd Israel Symposium on Theory of Computing Systems* (1993), IEEE Computer Society, pp. 3–17.

[25] PETRANK, E., AND TARDOS, G. On the knowledge complexity of np. In *FOCS* (1996), pp. 494–503.

[26] RENNER, R., AND WOLF, S. Smooth Renyi entropy and applications. In *IEEE International Symposium on Information Theory — ISIT 2004* (June 2004), IEEE, p. 233.

[27] ROMPEL, J. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)* (1990), pp. 387–394.

[28] SHANNON, C. Communication theory of secrecy systems. *Bell System Technical Journal 28*, 4 (1949), 656–715.

[29] YAO, A. C. Theory and applications of trapdoor functions. In *FOCS* (1982) pp. 80–91.