

A New Interactive Hashing Theorem

Iftach Haitner*

Omer Reingold†

Abstract

Interactive hashing, introduced by Naor et al. [NOVY98], plays an important role in many cryptographic protocols. In particular, it is a major component in all known constructions of statistically-hiding commitment schemes and of zero-knowledge arguments based on general one-way permutations and on one-way functions. Interactive hashing with respect to a one-way permutations f is a two-party protocol that enables a sender that knows $y = f(x)$ to transfer a random hash $z = h(y)$ to a receiver. The receiver is guaranteed that the sender is committed to y (in the sense that it cannot come up with x and x' such that $f(x) \neq f(x')$ but $h(f(x)) = h(f(x')) = z$). The sender is guaranteed that the receiver does not learn any additional information on y . In particular, when h is a two-to-one hash function, the receiver does not learn which of the two preimages $\{y, y'\} = f^{-1}(z)$ is the one the sender can invert with respect to f .

This paper reexamines the notion of interactive hashing. We give an alternative proof for the NOVY-protocol [NOVY98], which seems to us significantly simpler and more intuitive than the original one. Moreover, the new proof achieves much better parameters (in terms of how security preserving the reduction is). Finally, our proof implies a more versatile interactive hashing theorem for a more general setting than that of [NOVY98]. One generalization relates to the selection of hash function h (allowing much more flexibility). More importantly, the theorem applies to the case where the underlying function f is hard-to-invert only on some given (possibly sparse) subset of the output strings. In other words, the theorem is tuned towards hashing of a value y that may be distributed over a sparse subset of the domain (rather than uniform on the entire domain as a random output of a one-way permutation is).

Our interest in interactive hashing is in part as a very appealing object (i.e., independent of any particular application). Furthermore, a major motivation for looking into interactive hashing is towards achieving a construction of statistical commitment schemes based on any one-way functions. Given the history of the problem, it seems that extending our understanding of interactive hashing is a natural approach towards relaxing the general assumptions needed for constructing statistical commitments. Indeed, our new theorem is already useful in easily and more directly implying constructions of statistical commitments based on the same assumptions as in [HHK⁺05]. In a subsequent work [HR06], we also show how to use the new theorem for constructions based on substantially weaker assumptions. That construction goes beyond the natural barrier of relying on “almost regular” one-way functions (and in particular imply a construction based on any exponentially-hard one-way function).

1 Introduction

Interactive hashing, introduced by Naor, Ostrovsky, Venkatesan and Yung [NOVY98], is a protocol that allows a sender \mathcal{S} to commit to a particular value while only revealing to a receiver \mathcal{R} some predefined information of this value. More specifically, \mathcal{S} commits to a value y while only revealing to \mathcal{R} the value $h, h(y)$, where h is some random hash function (we defer additional details on the choice of hash function). The two security properties of interactive hashing are *binding* (namely, \mathcal{S} is bounded by the protocol to at most one value of y) and *hiding* (namely, \mathcal{R} does not learn any impermissible information about y). As in [NOVY98], we will consider in this paper interactive hashing where the hiding property is statistical

*Dept. of Computer Science and Applied Math., Weizmann Institute of Science, Rehovot 76100, Israel. E-mail: iftach.haitner@weizmann.ac.il.

†Incumbent of the Walter and Elise Haas Career Development Chair, Department of Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel omer.reingold@weizmann.ac.il Research supported by grant no. 1300/05 from the Israel Science Foundation.

(i.e., the protocol preserves the secrecy of y even against an all-powerful \mathcal{R}), and the binding property is computational (i.e., it assumes that \mathcal{S} is computationally bounded).

Interactive hashing (in the flavor mentioned above) is closely related and to a large extent motivated by the fundamental notion of statistical commitment (i.e., statistically-hiding computationally-binding commitment schemes). In statistical commitments we again have a protocol between a sender \mathcal{S} and a receiver \mathcal{R} . However, here the sender \mathcal{S} commits to y without revealing *any* information about y . Statistical commitments can be used as a building block in constructions of statistical zero-knowledge arguments [BCC88, NOVY98] or certain coin-tossing protocols [Blu82, Lin03]. More generally, they have the following advantage over computationally hiding commitment schemes when used within protocols in which certain commitments are never revealed: in such a scenario, it need only be infeasible to violate the binding property *during the period of time the protocol is run*, whereas the committed values will remain hidden *forever* (i.e., regardless of how much time the receiver invests after completion of the protocol).

The relation between interactive hashing and statistical commitments goes beyond the similarity in definitions. On one hand, interactive hashing can easily be implemented using commitment schemes (simply commit to y using the commitment scheme and reveal whatever information needed on y in the clear). On the other hand, one of the main applications of interactive hashing protocols is for constructing statistical commitment schemes. Indeed, interactive hashing is a major component in all known constructions of statistical commitment that are based on the existence of one-way permutations or on (special forms of) one-way functions.

Naor et. al. [NOVY98] use their interactive-hashing protocol (from now on the NOVY protocol) in order to construct statistical commitments based on any one-way permutation. Haitner et. al. [HHK⁺05] make progress by using the NOVY protocol to construct statistical commitments based on regular one-way functions and also on the so called approximable-size one-way functions. Interactive hashing is also used by several other cryptographic protocols [OVY93, OY92]. In addition, interactive-hashing is used in “information theoretic setting” (i.e., no hardness assumptions are assumed) such as [CCM98, DHRS04, CS06, NV06]. Very recently, interactive hashing was used in the breakthrough result of Nguyen et al. [NOV06] to construct statistical zero-knowledge arguments for NP based on any one-way function. This result leaves open the question of constructing statistical commitments based on any one-way functions (a task that seems more feasible given the result of [NOV06]). Indeed, one of our major motivations for looking into interactive hashing is indeed the construction of statistical commitment schemes based on any one-way functions. However, before discussing our results and their applications let us look closer into interactive hashing.

1.1 Interactive hashing in the setting of one-way permutations

Consider the following two-party protocol between a sender \mathcal{S} and a receiver \mathcal{R} : The sender chooses a random element $x \in \{0,1\}^n$ and sets $y = f(x)$, where $f : \{0,1\}^n \rightarrow \{0,1\}^n$ is a one-way permutation. Next, the receiver selects a random pairwise independent two-to-one hash function $h : \{0,1\}^n \rightarrow \{0,1\}^{n-1}$ and sends its description to \mathcal{S} . Finally, \mathcal{S} sends $z = h(y)$ back to \mathcal{R} . Note that if both parties follow the protocol, then the following “binding” property is guaranteed: It is not feasible for \mathcal{S} to find a second element $x' \in \{0,1\}^n$ such that $f(x') \neq f(x)$ but $h(f(x')) = h(f(x)) = z$, although (exactly one) such element x' does exist. The reason is that the task of finding such x' can easily be shown to be equivalent in hardness to inverting f on a random output element (the latter task is assumed to be hard by the one-wayness of f). Furthermore, we are guaranteed to have the following “hiding” property: Let y_1 and y_2 be the two preimages of z w.r.t. h . Given \mathcal{R} ’s view of the communication (i.e., given the values h and z), it is indistinguishable whether the random element chosen by \mathcal{S} is $x_1 = f^{-1}(y_1)$ or $x_2 = f^{-1}(y_2)$. In this sense \mathcal{S} has committed to a bit (which indicates if it can produce the inverse of y_1 or that of y_2). This bit is statistically hidden from \mathcal{R} .

What happens, however, if \mathcal{S} selects x only *after* seeing h ? In such a case, it is quite plausible that \mathcal{S} would be able to “cheat” by producing $x, x' \in \{0,1\}^n$ such that $f(x) \neq f(x')$ but $h(f(x')) = h(f(x)) = z$.¹ The NOVY interactive hashing protocol prevents exactly such cheating. For that it employs a specific family

¹Assume for example that the one-way permutation equals the identity function on the set T of all strings that start with $n/4$ zeros (where n is the input length). Now given a hash function h all the cheating sender has to do is to find a collision $y_1 \neq y_2$, where $y_1, y_2 \in T$, such that $h(y_1) = h(y_2)$. Such a collision is likely to exist by the birthday paradox, and for many families of hash functions finding such a collision is very easy.

of hash functions such that each one of its functions h can be decomposed into $n - 1$ Boolean functions h_1, \dots, h_{n-1} (where $h(x) = h_1(x), \dots, h_{n-1}(x)$).² In the NOVY protocol, instead of sending h at once as described above, \mathcal{R} sends a single Boolean function h_i in each round. In return, the honest sender sends a bit $z_i = h_i(f(x))$. What about a cheating sender? Intuitively, a cheating sender has a significantly smaller leeway for cheating as it can no longer wait in selecting x till it receives the entire description of h . Still, it is highly non-trivial to argue (formally or even intuitively) that restricting the sender by adding interaction in this manner is sufficient in order to prevent the sender from cheating. Perhaps surprisingly, Naor et. al. [NOVY98] have shown that their protocol has the binding property even against a cheating verifier (namely, even a cheating verifier cannot produce $x, x' \in \{0, 1\}^n$ such that $f(x) \neq f(x')$ but $h(f(x')) = h(f(x)) = z$).

1.2 Interactive hashing in the sparse case

The NOVY interactive hashing protocol applies to one-way permutations and easily implies the existence of statistical commitments from any one-way permutation. How about constructing statistical commitments from, say, regular one-way functions (one-way functions where every output value has the same number of preimages)? In such a case we would like to interactively hash a value y (a random output of the one-way function) which is uniformly distributed in some subset L of $\{0, 1\}^n$ (rather than uniformly distributed in all of $\{0, 1\}^n$ as in the case of one-way permutations). What is the difficulty in directly hashing a value y that is taken from a set L that is sparse in $\{0, 1\}^n$? The NOVY-theorem guarantees that when hashing y with $h : \{0, 1\}^n \mapsto \{0, 1\}^{n-1}$ the sender is committed to a single value y (as shown in [NOV06] this holds even if the output of h is a bit shorter). However, when h outputs so many bits then most likely $h(y)$ completely determines y and statistical hiding is lost.

Facing the aforementioned difficulty, Haitner et. al. [HHK⁺05] first make the following observation: the NOVY protocol is still meaningful even when hashing a value y which is taken from a distribution that is “dense” in $\{0, 1\}^n$ (a bit more formally we would like the distribution to be sufficiently close to having min-entropy $n - O(\log n)$). In particular, if the one-way function is poly-to-one (i.e., each output has at most polynomial number of preimages in the image set of f), then the NOVY-protocol can be applied as is to give some weak form of statistical commitments that can later be amplified to full-fledge statistical commitments. To handle any regular one-way function, [HHK⁺05] applies additional layer of hashing to reduce to the dense case. This implies a construction of statistical commitments from any regular one-way function. [HHK⁺05] also gives a construction of statistical commitments from one-way functions where the size of the set of preimages of any output value y can be approximated (the so called approximable-size one-way functions). Intuitively this is obtained by a simple reduction of such functions to “almost regular” one-way functions. Interactive hashing in the sparse case arises in other works as well, most notably in the construction of statistical zero-knowledge arguments from any one-way function [NOV06].

1.3 Our Results and Applications

We introduce an alternative proof for the NOVY protocol, which relies in parts on the original proof due to [NOVY98] (the NOVY proof) but still seems to us significantly simpler. The proof follows a simple intuition that is sketched below in this section. Moreover, the parameters achieved by our proof are an improvement compared with the original ones. In our proof, given an algorithm A that breaks the binding property with probability ε_A we get an algorithm that inverts the one-way permutation in comparable time and with inverting probability $\varepsilon_A^2 \cdot \text{poly}(n)$. This is a substantial improvement and is much closer to natural limitations of the proof technique (see discussion in Section 5).³

In addition to being simpler and more security preserving, the new proof implies a more general interactive hashing theorem. The new theorem applies to every family of hash functions that is a product of Boolean families of pairwise independent hash functions (and not only to the special family of two-to-one hash functions used by [NOVY98]). More importantly, the new theorem directly applies to the “sparse case”: Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell(n)}$ be an efficiently computable function, let $L \subseteq \{0, 1\}^{\ell(n)}$. As mentioned above, when hashing a value $y \in L$, the NOVY proof only promises binding when using a hash function that outputs

²For more details on the definition of this family of hash functions see Section 4.

³We note that independently of our work, [NOV06] recently presented an $\varepsilon_A^3 \cdot \text{poly}(n)$ reduction. See discussion below for more detail on their work.

almost n bits. However, in such a case y is likely to be completely determined by $h(y)$ and statistical hiding cannot be guaranteed. Our theorem applies even when hashing to roughly $\lfloor \log(|L|) \rfloor$ bits. In particular, when h is taken from a family of hash functions $\overline{\mathcal{H}}: \{0, 1\}^{\ell(n)} \rightarrow \{0, 1\}^{\lfloor \log(|L|) \rfloor}$ that is a product of $\lfloor \log(|L|) \rfloor$ families of pairwise independent Boolean hash functions, we can show that a close variant of the NOVY protocol possesses the following binding property: If f is hard to invert on the uniform distribution over L , then any sender \mathcal{S}^* (even one which arbitrarily deviates from the protocol) cannot find two elements $x, x' \in f^{-1}(L)$ such that $f(x) \neq f(x')$ but $h(f(x)) = h(f(x')) = z$ (where z is the value determined by the protocol as $h(y)$).

Applications of The New Theorem The statistically-hiding commitment due to [HHK⁺05] can be described as follows: First, the output of the regular one-way function, f , is hashed (non-interactively) using an n -wise independent hash functions e into domain of size about $|Im(f)|$ to construct a new one-way function $g = e \circ f$. In the next step, the two parties invoke the NOVY interactive-hashing protocol w.r.t. g . Finally, \mathcal{S} sends $(b \oplus \langle g(x), r \rangle \bmod 2)$ to \mathcal{R} , where r is a random string chosen by \mathcal{R} . Using the properties of the NOVY protocol, [HHK⁺05] proved that this protocol is a “weak” statistical commitment scheme that can be amplified to full fledge statistical commitment. The new interactive hashing theorem can obtain a similar result more directly. Instead of applying the interactive hashing protocol to g one can now apply it directly to f and interactively hash $f(x)$ into a string of length $\lfloor \log(|Im(f)|) \rfloor - O(\log(n))$. Next, as before, \mathcal{S} sends $(b \oplus \langle f(x), r \rangle \bmod 2)$ to \mathcal{R} , where r is some random string chosen by \mathcal{R} . By the properties of the new interactive-hashing protocol w.r.t. $Im(f)$, it easily follows that \mathcal{S} is committed to b , whereas a semi-honest receiver does not learn “too much” about b (i.e., it has some noticeable uncertainty about the value of b). The weak hiding of this scheme can easily be amplified and the scheme can be protected against a malicious receiver (see [HKKM04, Theorem 4]). We therefore obtain a full-fledged statistical commitment scheme.

Haitner et. al. obtained statistical commitment based on regular one-way functions and on one-way functions where the size of the set of preimages of any output value y can be approximated (the so called approximable-size one-way functions). This is a natural milestone when extending results that were previously only known based on one-way permutations (see for example [GKL93, AGGM06]). In a subsequent work [HR06], we use the new interactive hashing theorem to go beyond this natural barrier and provide a construction of statistical commitments based one-way functions that seem very different than regular functions and do not seem to imply almost regular one-way functions by any simple transformation. Loosely, these are functions where the preimage size of an output value y can be non-trivially upper bounded. This means that (1) The bound is in fact a good approximation with some non-negligible probability (but can most of the times be a wild over-estimate), and (2) The bound may be false (i.e., an under estimate) only with some small (though still non-negligible) probability. This somewhat technical notion is interesting both because it substantially relaxes previous requirements, but in addition it implies some interesting corollaries. For example, based on this result we obtain a construction of statistical commitments from any *exponentially-hard* one-way function (i.e., a function that cannot be inverted in polynomial time with probability better than 2^{-Cn} for some constant $C > 0$). Note that this is a non-structural property (unlike being a permutation or being regular).

Related Work We note that independently of our work, Nguyen, Ong and Vadhan [NOV06] give a new proof for the NOVY protocol. Their proof follows the proof of [NOVY98] more closely than ours but still introduces various simplifications and parameter improvements. The main goal of the new proof is to generalize the protocol such that it allows hashing with a hash function that is poly-to-one rather than two-to-one as in [NOVY98]. In other words, they analyze the NOVY protocol with $n - \ell$ rounds where ℓ may be as large as $O(\log n)$ rather than $n - 1$ rounds in [NOVY98]. For a comparison between the parameters obtained by [NOVY98], [NOV06] and this paper, see Remark 3.10.

Relations Following [NOV06]) we state our protocol, and proof, in the more general setting of binary relations rather than functions. For instance, given a binary relation W that is hard to satisfy (i.e., given y it is hard to find x such that $(x, y) \in W$), we prove that following our interactive-hashing protocol, \mathcal{S} cannot find two pairs $(x_0, y_0), (x_1, y_1) \in W$ such that both y_0 and y_1 are consistent with the protocol, but

$y_0 \neq y_1$. Note that since every function, f , defines the natural binary relation $(x, f(x))$, any result w.r.t. binary relations implies an equivalent result w.r.t. functions.

Cheating Receiver In the new interactive hashing theorem, the hiding property of a cheating receiver is specified only with respect to a semi-honest receiver. A malicious receiver can learn $h(y)$ where h is not necessarily uniformly distributed. In fact, h can be determined adaptively based on partial knowledge of $h(y)$ (specifically, h_i is selected after learning the first $i - 1$ bits of $h(y)$). The NOVY protocol provides a cheating receiver exactly the same power in selecting h . Nevertheless, when y is selected uniformly in $\{0, 1\}^n$ and h is two-to-one, then regardless of how the receiver selects h there is one bit of knowledge on y that remains completely hidden from the receiver (i.e., given $z = h(y)$ there are two possibilities to the value of y , the hidden bit specifies which of these two values is the right one). In the general case, when y is distributed over a sparse subset, one should take more care in estimating the power of a cheating receiver. We note that in various settings, one can assume without loss of generality that receiver is semi-honest. In particular, this is the case for statistical commitments (see [HKKM04, Theorem 4]).

1.4 Proof Idea

We discuss our proof in the most basic setting where $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a one-way permutation and $L = \{0, 1\}^n$ (the general case is not much different). Our protocol consists of m rounds, where $n - \mathcal{O}(\log(n)) \leq m \leq n$. In each round, \mathcal{R} selects a random Boolean pairwise-independent hash function h_i and \mathcal{S} replies with $h_i(f(x))$. Let A be an algorithm that plays the sender's role in the protocol and at the end of the protocol outputs two elements $x_1, x_2 \in \{0, 1\}^n$. Assume that with some noticeable probability $\varepsilon > 0$, it holds that $f(x_1) \neq f(x_2)$ and both $f(x_1)$ and $f(x_2)$ are consistent with the protocol's transcript. It is easy to use A in order to construct an algorithm that inverts f with probability $\frac{\varepsilon}{2^n}$: Given input y , the algorithm chooses the hash functions at random and returns one of the two values A 's outputs.

Let's imagine that instead we are trying to invert f on the following distribution: The first $k \stackrel{\text{def}}{=} m - \log(\frac{1}{\varepsilon}) - C \log(n)$ (for some constant $C > 0$) Boolean hash functions, h_1, \dots, h_k , are chosen at random and only then a random element, y , is drawn from set of all the elements inside $\{0, 1\}^n$ that are consistent with A 's answers on h_1, \dots, h_k . We call the distribution induced on (y, h_1, \dots, h_k) by the above process D_{uni} . On the average, A has probability ε to cheat even when conditioned on h_1, \dots, h_k being selected. Therefore (since, with high enough probability, the number of y 's consistent with A 's answers on h_1, \dots, h_k is about $\frac{n^C}{\varepsilon}$), in this setting the naive algorithm, which selects the rest of the hash functions at random and returns one of A 's answers, inverts f with probability close to $\frac{\varepsilon^2}{n^C}$. In addition, a more careful analysis yields that the success probability of this inverting algorithm does not depend on inverting too few elements. More specifically, the subset of y 's that are consistent with A 's answers on h_1, \dots, h_k such that the naive algorithm inverts on them with "high enough" probability is of relative size $\sqrt{\varepsilon^2 n^C / \frac{n^C}{\varepsilon}} = \frac{\varepsilon^2}{n^{C/2}}$.⁴

Let's try to emulate the above setting on a random $y \in \{0, 1\}^n$. To do so, we choose the first k hash functions one by one, each time we keep sampling until we find a hash function that its value on y is consistent with A 's answer (if the answer is inconsistent, we "rewind" A to its state before it was asked the last "faulty" hash function). We call D_{src} the distribution the above process induces on (y, h_1, \dots, h_k) .

We would conclude the proof of the binding property if we could prove that the statistical difference between D_{uni} and D_{src} is smaller than $\frac{\varepsilon^2}{n^C}$ (recall that this is the inverting probability of the naive algorithm on D_{uni}). Unfortunately, we cannot prove such a strong bound. We manage to prove, however, that save but a small number of elements in the support of D_{uni} (which have total probability mass smaller than $\frac{\varepsilon^2}{n^{C/2}}$), the probability mass that each element has under D_{uni} is within a constant factor from its mass under D_{src} . It easily follows that we can invert y with noticeable probability over D_{src} , which directly implies that we can invert f (again, with noticeable probability) on the uniform distribution over $\{0, 1\}^n$.

⁴Loosely, let T be the set of y 's that A is likely to output their inverse (according to f). A random selection of h_1, \dots, h_k separates every two elements in T with probability $1 - 2^{-k}$. So unless the size of T is large enough (essentially $\sqrt{\varepsilon \cdot 2^k}$), one of the two values A output will be forced to be the inverse of an element outside of T . This will contradict the assumptions that values outside of T are only inverted with small probability.

1.5 Paper Organization

In Section 3, we generalize the definition of interactive-hashing, present our new construction and prove that it satisfies the new definition. In Section 4 we argue that the new proof can also be applied to the original NOVY protocol (that uses very specific hash functions). Discussion and further issues appear in Section 5.

2 Preliminaries

2.1 Notation

For $k \in \mathbb{N}$, we denote by $[k]$ the set $\{1, \dots, k\}$. We denote the concatenation of the strings x and y by $x \circ y$. Given a set L , we denote by $x \leftarrow L$ the experiment in which x is uniformly chosen from L . Let D be a distribution over the set L , the support of D is defined as: $\text{supp}(D) = \{x \in L : D(x) > 0\}$. We denote the probability of $L' \subseteq L$ w.r.t. D as $D(L') = \Pr_{x \leftarrow D}[x \in L']$. Given a function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ and a set $L \subseteq \{0, 1\}^*$, we denote the image of f on L as $f(L) = \{f(x) : x \in L\}$. We denote the running time of an algorithm A by T_A , where PPT stands for polynomial-time algorithm. Given two interactive Turing machines A and B , we denote the protocol they define by $\langle A, B \rangle$ and denote the following experiment by $(o_A \mid o_B) \leftarrow \langle A(i_A), B(i_B) \rangle$: The protocol $\langle A, B \rangle$ is invoked with inputs i_A and i_B and the outputs of the parties are assigned to o_A and o_B respectively.

2.2 Families of Pairwise-Independent Hash Functions

Definition 2.1 (*Efficient family of pairwise-independent hash functions*) Let \mathcal{H} be a family of functions mapping strings of length $\ell(n)$ to strings of length $m(n)$. We say that \mathcal{H} is an efficient family of pairwise independent hash functions (following [CW77]) if the following holds: ⁵

1. \mathcal{H} is polynomially samplable (in n).
2. There exists a polynomial-time algorithm that given $x \in \{0, 1\}^{\ell(n)}$ and a description of $h \in \mathcal{H}$ outputs $h(x)$.
3. For every distinct $x_1, x_2 \in \{0, 1\}^{\ell(n)}$ and every $y_1, y_2 \in \{0, 1\}^{m(n)}$, we have:

$$\Pr_h \leftarrow \mathcal{H}[h(x_1) = y_1 \wedge h(x_2) = y_2] = 2^{-2m(n)}.$$

It is well known ([CW77]) that there exists an efficient family of pairwise-independent hash functions for every choice of ℓ and m whose elements description size is $\mathcal{O}(\max\{\ell(n), m(n)\})$.

The following standard lemma (see for example, [Gol01, Lemma 4.3.1]) states that a random pairwise independent hash function partitions a given set into (almost) equal size subsets.

Theorem 2.2 Let \mathcal{H} be a family of pairwise independent hash functions mapping strings of length ℓ to strings of length m , let $L \subseteq \{0, 1\}^m$ and let $\ell = \frac{|L|}{2^m}$. Then for every $\alpha \in \{0, 1\}^m$ and $\delta > 0$

$$\Pr_h \leftarrow \mathcal{H}[|\{x \in L : h(x) = \alpha\}| - \ell| > \delta \ell] < \frac{1}{\delta^2 \ell}.$$

As in [NOVY98], our interactive-hashing protocol selects and evaluates an hash function incrementally. Therefore, the protocol is designed for hash functions that are product of hash functions defined next.

Definition 2.3 (*Product Hash Family*) Let \mathcal{H} be a family of functions mapping strings of length $\ell(n)$ to strings of length $m(n)$ and let $k(n) \in \mathbb{N}$. The k -product-family of \mathcal{H} , denoted $\mathcal{H}^{\times k(n)}$, is a family of functions mapping strings of length $\ell(n)$ to strings of length $k(n) \cdot m(n)$ which is defined as follows: The members of \mathcal{H} are all possible tuples \bar{h} of $k(n)$ functions from \mathcal{H} . For every such tuple $\bar{h} = (\bar{h}_1, \dots, \bar{h}_{k(n)})$ and every $x \in \{0, 1\}^{\ell(n)}$ we define $\bar{h}(x) = (\bar{h}_1(x), \dots, \bar{h}_{k(n)}(x))$.

⁵The first two properties, regarding the efficiency of the family, implicitly assume an ensemble of families (one family for every value of n). For simplify of presentation, we only refer to a single family.

3 The New Interactive Hashing Theorem

In this section we present our extended definition for an interactive-hashing protocol and give a revised construction and new proof that match this definition.

3.1 Defining a New Notion of Interactive Hashing

We choose (following [NOV06]) to state our definitions in the setting of binary relations. This generalizes the original definition due to [NOVY98], which concentrates on the particular relations that are naturally defined by one-way permutations (see Corollary 3.14). In particular, the underlying relation is not necessarily efficiently computable or even not efficiently verifiable. Moreover, the relation is not necessarily defined over all strings of a given length, but might rather be defined over some small subset of the strings.

Notation: Let W be a binary relation over and let $y \in \{0, 1\}^*$, we denote the set $\{z \in \{0, 1\}^* : W(z, y) = 1\}$ by W_y .

Definition 3.1 (*Interactive-Hashing*) Let $\overline{\mathcal{H}}$ be a family of hash functions mapping strings of length $\ell(n)$ to strings of length $m(n)$. A polynomial-time protocol $\langle \mathcal{S}, \mathcal{R} \rangle$ is an $\overline{\mathcal{H}}$ -interactive-hashing protocol if the following hold:

- The inputs of \mathcal{S} are a string $y \in \{0, 1\}^{\ell(n)}$ and the security parameter 1^n . Following its interaction, \mathcal{S} outputs y .
- The input of \mathcal{R} is the security parameter 1^n . Following its interaction, \mathcal{R} outputs $(\overline{h}, \overline{z}) \in \overline{\mathcal{H}} \times \{0, 1\}^{m(n)}$.
- *Functionality* - For every $y \in \{0, 1\}^{\ell(n)}$, letting $(y, (\overline{h}, \overline{z})) = \langle \mathcal{S}(1^n, y), \mathcal{R}(1^n) \rangle$ it holds that $\overline{h}(y) = \overline{z}$.

The security of interactive-hashing protocol has two aspects. Binding the sender to y and concealing some information regarding y from \mathcal{R} . In this paper we focus on security w.r.t. polynomially-bounded sender and unbounded receiver, the complementary setting where the receiver is polynomially-bounded and the sender is unbounded, called Information Theoretic interactive-hashing, is not treated by this paper (for details on the information theoretic setting see for example [CCM98, DHRS04, CS06]).

We start by formalizing the binding property.

Definition 3.2 ($\text{BndBreak}^{L, W}$) Let $L \subseteq \{0, 1\}^{\ell(n)}$, let W be a binary relation and let $\overline{\mathcal{H}}$ be a family of hash functions mapping strings of length $\ell(n)$ to strings of length $m(n)$. The Boolean function $\text{BndBreak}^{L, W}$ is defined such that for any inputs $o_{\mathcal{S}} = ((x_0, y_0), (x_1, y_1))$ and $o_{\mathcal{R}} = (\overline{h}, \overline{z}) \in \overline{\mathcal{H}} \times \{0, 1\}^{m(n)}$, the value $\text{BndBreak}^{L, W}(o_{\mathcal{S}}, o_{\mathcal{R}})$ is one iff y_0 and y_1 are distinct elements inside L such that $\overline{h}(y_0) = \overline{h}(y_1) = \overline{z}$, $x_0 \in W_{y_0}$ and $x_1 \in W_{y_1}$.

Definition 3.3 (*Binding*) Let L , W and $\overline{\mathcal{H}}$ be as in Definition 3.2 and let $\langle \mathcal{S}, \mathcal{R} \rangle$ be an $\overline{\mathcal{H}}$ -interactive-hashing protocol. We say that $\langle \mathcal{S}, \mathcal{R} \rangle$ is binding w.r.t. L and W if for every PPT A that plays the role of \mathcal{S} in the protocol and outputs two pairs of elements, the following is negligible:

$$\Pr_{(o_A | o_{\mathcal{R}}) \leftarrow \langle A(1^n), \mathcal{R}(1^n) \rangle} [\text{BndBreak}^{L, W}(o_A, o_{\mathcal{R}})]$$

where the probability is taken over the random coins of A and \mathcal{R} .

The following definition states that the only information that an honest receiver acquires through the protocol about y is its hash value for a uniformly chosen hash function.

Definition 3.4 (*Secrecy preserving w.r.t. semi-honest receiver*) Let $\overline{\mathcal{H}}$ be a family of hash functions mapping strings of length $\ell(n)$ to strings of length $m(n)$ and let $\langle \mathcal{S}, \mathcal{R} \rangle$ be an $\overline{\mathcal{H}}$ -interactive-hashing protocol. For $y \in \{0, 1\}^{\ell(n)}$, we denote by $\text{view}_{\mathcal{R}}^{\mathcal{S}}(1^n, y)$ the distribution of \mathcal{R} 's view when interacting with $\mathcal{S}(y, 1^n)$ (this view simply consists of the sequence of messages \mathcal{R} receives from \mathcal{S} and its random coins), where this distribution is taken over the random coins of \mathcal{S} and \mathcal{R} . We say that $\langle \mathcal{S}, \mathcal{R} \rangle$ is secrecy-preserving (w.r.t. semi-honest receiver) if there exists a polynomial-time simulator Sim , such that for every $y \in \{0, 1\}^{\ell(n)}$ the distributions $\text{view}_{\mathcal{R}}^{\mathcal{S}}(1^n, y)$ and $(\text{Sim}(1^n, \overline{h}, \overline{h}(y)))_{\overline{h} \leftarrow \overline{\mathcal{H}}}$ are identical.

Remark 3.5 Some level of hiding can be guaranteed by our protocol even against malicious \mathcal{R} . Specifically, the protocol hides any information regarding the index of y among all the preimages of $\bar{z} = \bar{h}(y)$ w.r.t. \bar{h} . In the setting of [NOVY98] this information is quite meaningful and is also easy to construct. This is because y is chosen uniformly in $\{0,1\}^{\ell(n)}$ and regardless of the way the receiver selects \bar{h} , there are exactly two possible preimages of \bar{z} . The two preimages can be found easily and therefore the relative index of y is easy to construct. In the most general setting, however, we encounter two problems: Firstly, a malicious \mathcal{R} may be able to force the existence of only a single preimage of \bar{z} w.r.t. \bar{h} that lies in L . Secondly, it may be difficult to find the preimages of \bar{z} that lie in L .

We note that in several cryptographic applications of interactive hashing (e.g., statistically-hiding bit commitment, see [HKKM04, Theorem 4]), any protocol that is secure against an honest receiver can be compiled into a protocol that is secure against a malicious receiver.

3.2 The Construction

Construction 3.6 (Interactive Hashing) Let $m(n) \in \mathbb{N}$ and let \mathcal{H} be a family of efficiently computable Boolean functions defined over strings of length $\ell(n)$. The parties of the protocol are \mathcal{S} and \mathcal{R} , where the inputs of \mathcal{S} are 1^n and an $\ell(n)$ bits string y , and \mathcal{R} 's input is 1^n .

1. For $i = 1$ to $m(n)$:
 - (a) \mathcal{R} chooses uniformly at random $h_i \in \mathcal{H}$ and sends its description to \mathcal{S} .
 - (b) \mathcal{S} sends $z_i = h_i(y)$ back to \mathcal{R} .
2. \mathcal{S} locally outputs y .
3. \mathcal{R} outputs $(\bar{h}, \bar{z}) = (h_1, \dots, h_{m(n)}, z_1, \dots, z_{m(n)})$.

The following lemma is immediate from Definitions 3.1 and 3.4.

Lemma 3.7 Let $\bar{\mathcal{H}}$ be the $m(n)$ -product-family of \mathcal{H} , then $\langle \mathcal{S}, \mathcal{R} \rangle$ is a secrecy-preserving $\bar{\mathcal{H}}$ -interactive-hashing protocol.

3.3 The Main Theorem - Binding

Theorem 3.8 Let W be a binary relation and let $L \subseteq \{0,1\}^{\ell(n)}$. Let \mathcal{H} be an efficient family of pairwise independent Boolean hash functions defined over string of length $\ell(n)$ and let $\langle \mathcal{S}, \mathcal{R} \rangle$ be as in Construction 3.6. Then there exists an oracle algorithm $M^{(\cdot)}$ such that for any algorithm A , the running-time of M^A is $\mathcal{O}(\log(n)T_A(n) + m \log(n)T_{\mathcal{H}}(n))$ (where $T_{\mathcal{H}}(n)$ is an upper bound of the sampling and computing time of \mathcal{H}) and for large enough n ,

$$\Pr_y \leftarrow_L [M^A(y) \in W_y] \in \Omega\left(\frac{2^{m'(n)}}{|L|} \cdot \frac{\varepsilon_A(n)^2}{n^8}\right),$$

where $\varepsilon_A(n) \stackrel{\text{def}}{=} \Pr_{(o_A|o_{\mathcal{R}}) \leftarrow \langle A(1^n), \mathcal{R}(1^n) \rangle} [\text{BndBreak}^{L,W}(o_A, o_{\mathcal{R}})]$ and $m'(n) \stackrel{\text{def}}{=} \min\{m(n), \lfloor \log(|L|) \rfloor\}$.

Remark 3.9 We point out that $M^{(\cdot)}$ does not need to know W or ε_A .

Remark 3.10 (Comparing the parameters to [NOVY98] and [NOV06]) For $L = \{0,1\}^n$ and $m = n - 1$, the success probability of M^A is $\Omega(\frac{\varepsilon_A(n)^2}{n^8})$, where the running-time is still $\mathcal{O}(\log(n)T_A(n) + m \log(n)T_{\mathcal{H}}(n))$. We point that the same success probability and running-time apply also for the NOVY protocol (see Section 4 for details). This is an improvement in parameters compared with the analysis in [NOV06, Lemma B.2]. There the algorithm runs in time $\mathcal{O}(nT_A(n) + mnT_{\mathcal{H}}(n))$ and breaks f with probability $\Omega(\frac{\varepsilon_A(n)^3}{n^6})$. Finally, in the [NOVY98, [Lemma 2] analysis, the algorithm runs in time $\mathcal{O}(nT_A(n) + mnT_{\mathcal{H}}(n))$ (same as in [NOV06]) and only guarantees to break f with probability $\Omega(\frac{\varepsilon_A(n)^{10}}{n^8})$.

The following corollaries follow Lemma 3.7 and Theorem 3.8.

Definition 3.11 Let W be a relation and let $L \subseteq \{0, 1\}^{\ell(n)}$. We say that W is hard-to-satisfy on L if for any PPT A the probability $\Pr_{y \leftarrow L}[A(y) \in W_y]$ is negligible in n .

Corollary 3.12 Let L , m , $\overline{\mathcal{H}}$, W and $\langle \mathcal{S}, \mathcal{R} \rangle$ be as in Theorem 3.8. If W is hard-to-satisfy on L and $m > \log(|L|) - \mathcal{O}(\log(n))$, then the protocol $\langle \mathcal{S}, \mathcal{R} \rangle$ is a computationally-binding, secrecy-preserving (w.r.t. honest-receiver) $\overline{\mathcal{H}}$ -interactive hashing protocol w.r.t. L and W .

Definition 3.13 Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell(n)}$ be an efficiently computable function and let $L \subseteq \{0, 1\}^{\ell(n)}$. We say that f is hard to invert over L if for any PPT A the probability $\Pr_{y \leftarrow L}[A(y) \in f^{-1}(y)]$ is negligible in n .

Corollary 3.14 Let L , m , $\overline{\mathcal{H}}$ and $\langle \mathcal{S}, \mathcal{R} \rangle$ be as in Theorem 3.8. Let f be hard to invert over L , and let W be the binary relation defined by f (i.e., $x \in W_y$ iff $f(x) = y$). If $m > \log(|L|) - \mathcal{O}(\log(n))$, then the protocol $(\mathcal{S}, \mathcal{R})$ is a computationally-binding secrecy-preserving (w.r.t. honest-receiver) $\overline{\mathcal{H}}$ -interactive-hashing protocol w.r.t. L and W .

The Proof of Theorem 3.8. For simplicity we drop the dependency on n whenever it is clear from the context. We assume w.l.o.g. that $m' = m$, since any adversary that violates the binding of the m -round protocol, can violate with the same probability the binding of the protocol with $m' < m$ rounds. We denote by A^r the restriction of A to some fixed random coins $r \in \{0, 1\}^{T_A}$ and use throughout the proof the following random variables: For $k \in [m]$ and $(r, \overline{h^s}) \in \{0, 1\}^{T_A} \times \mathcal{H}^{\times k}$, let $A^{Com}(r, \overline{h^s}) \in \{0, 1\}^k$ be A^r 's answers when questioned by $\overline{h^s}$ and let $\text{Consist}(r, \overline{h^s}) = \{y \in L : \forall i \in [k] \ \overline{h^s}_i(y) = A^{Com}(r, \overline{h^s})_i\}$ (i.e., the set of y 's that are consistent with A 's answers). Finally, we assume w.l.o.g. that for any sequence of questions $\overline{h} \in \mathcal{H}^{\times m}$, A^r outputs two pairs of elements $((x_0, y_0), (x_1, y_1))$ and denote them by $A^{Dec}(r, \overline{h})$. For some value of $\text{ofs} \in \{0, \dots, m-1\}$ that will be specified below, we consider the following algorithm for satisfying W on L :

$M^A(y)$:

1. Choose uniformly at random $r \in \{0, 1\}^{T_A}$.
2. Let $\overline{h^s} \leftarrow \text{Searcher}(r, y)$.
3. Return $\text{Inverter}(r, \overline{h}, y_s)$.

where the algorithms *Searcher* and *Inverter* are defined as follows:

$\text{Searcher}(r, y)$:

1. Fix A 's random coins to r .
2. For $k = 1$ to $m - \text{ofs}$:
 - (a) Do the following $2 \log(n)$ times:
 - i. Set a value for h_k uniformly at random in \mathcal{H} .
 - ii. If $A^{Com}(r, (h_1, \dots, h_k))_k = h_k(y)$, break the inner loop.
3. Return $(h_1, \dots, h_{m-\text{ofs}})$.

$\text{Inverter}(r, \overline{h}, y_s)$:

1. Fix A 's random coins to r .
2. Choose uniformly at random $\overline{h^e} \in \mathcal{H}^{\times \text{ofs}}$.
3. Set $((x_0, y_0), (x_1, y_1)) \leftarrow A^{Dec}(r, (\overline{h^s}, \overline{h^e}))$.
4. Return x_0 with probability half and x_1 otherwise.

Remark 3.15 The value ofs will depend in our proof on ε_A . This seems to contradict Remark 3.9 that $M^{(\cdot)}$ does not need to know ε_A . Nevertheless, ofs can instead be selected at random with only a factor m

decrease in the success probability of M^A . More interestingly, setting $\mathbf{ofs} = 0$ will also guarantee M^A the success probability claimed in the theorem. The only affect of decreasing \mathbf{ofs} to zero is that \bar{h}^e will be selected by the rewinding method of *Searcher* rather than uniformly at random by *Inverter*. For every value \bar{h}^e that satisfies $y \in \text{Consist}(r, (\bar{h}^s, \bar{h}^e))$, we have that the probability of selecting it with the rewinding technique is only larger than the probability of uniformly selecting it. A value of \bar{h}^e such that $y \notin \text{Consist}(r, (\bar{h}^s, \bar{h}^e))$ will not contribute in our analysis to the success probability of M^A .

It follows that the distinction between *Searcher* and *Inverter* is not necessary for the proof. Still, following [NOVY98], we find this distinction very useful for pedagogical reasons.

Given that we use the proper data structure to support the rewinding action, it follows that the running time of M^A is $\mathcal{O}(\log(n)T_A(n) + m \log(n)T_{\mathcal{H}}(n))$, as stated in the theorem. As a first step in proving correctness, we show that $\Pr_{y \leftarrow L}[M^A(y) \in W_y] \in \Omega\left(\frac{2^{m(n)}}{|L|} \cdot \frac{\varepsilon_A(n)^3}{n^6}\right)$, since this proof has somewhat nicer abstraction than the one proving the stronger bound claimed in the theorem. In Section 3.7, we present the modifications needed for the stronger result.

We would like to set the value of \mathbf{ofs} to $\lceil 6 \log(n) + 2 \log(\frac{1}{\varepsilon_A}) \rceil + C$, where $C \in \mathbb{N}$ is some universal constant determined by the analysis. For that we need to assume that $m > \lceil 6 \log(n) + 2 \log(\frac{1}{\varepsilon_A}) \rceil + C$. If $m \leq 6 \log(n) + 2 \log(\frac{1}{\varepsilon_A}) + C$, then we can set $\mathbf{ofs} = m$ and conclude the proof of the theorem directly as follows:

$$\begin{aligned}
& \Pr_{y \leftarrow L}[M^A(y) \in W_y] \\
&= \frac{1}{|\{0,1\}^{T_A}| \cdot |L|} \sum_{y \in L, r \in \{0,1\}^{T_A}} \Pr_{\bar{h} \leftarrow \mathcal{H}^{\times m}}[\text{Inverter}(r, \bar{h}, y) \in W_y] \\
&\geq \frac{1}{|\{0,1\}^{T_A}| \cdot |L|} \cdot \frac{1}{2} \sum_{y \in L, r \in \{0,1\}^{T_A}} \Pr_{\bar{h} \leftarrow \mathcal{H}^{\times m}, ((x_0, y_0), (x_1, y_1)) \leftarrow A^{Dec}(r, \bar{h})} [x_0 \in W_y \vee x_1 \in W_y] \\
&\geq \frac{\varepsilon_A}{|L|} \\
&\in \Omega\left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^3}{n^6}\right).
\end{aligned} \tag{1}$$

We conclude that we can set $\mathbf{ofs} = \lceil 6 \log(n) + 2 \log(\frac{1}{\varepsilon_A}) \rceil + C$, and assume that $m > \mathbf{ofs}$.

We consider the following distributions:

Definition 3.16

- $D_{uni} \stackrel{\text{def}}{=} (r, \bar{h}, y)_{r \leftarrow \{0,1\}^{T_A}, \bar{h} \leftarrow \mathcal{H}^{\times(m-\mathbf{ofs})}, y \leftarrow \text{Consist}(r, \bar{h})}$,
- $D_{src} \stackrel{\text{def}}{=} (r, \bar{h}, y)_{r \leftarrow \{0,1\}^{T_A}, y \leftarrow L, \bar{h} \leftarrow \text{Searcher}(y, r)}$.

Given that y is uniformly chosen in L , then D_{src} is the distribution that *Inverter* is invoked upon through the execution of M^A . Thus, the probability that *Inverter* satisfies W over D_{src} equals to the success probability of M^A . On the other hand, it is rather easy to show that the probability that *Inverter* satisfies W over D_{uni} is noticeable (as a function of ε_A). Intuitively, this is because the distribution of \bar{h} in D_{uni} is uniform and this is also the distribution of \bar{h} that A encounters when interacting with \mathcal{R} . In fact, Lemma 3.19 states that the probability that *Inverter* satisfies W over D_{uni} is well spread: Even if we ignore the contribution to the success probability of some sufficiently small number of values in the support of D_{uni} , this success probability will remain noticeable. To sum up, we know that *Inverter* does well in satisfying W over D_{uni} and our goal is to show that it also does well over D_{src} . To this end, Lemma 3.26 will show that the distributions D_{uni} and D_{src} are “not too far” from each other (in a sense defined below). This will indeed allow us to complete our proof.

3.4 Most Pairs are Balanced

We start the formal proof by showing that in each step of the protocol the number of elements inside L that are consistent with the transcript so far is w.h.p. (regardless of A 's answers) not faraway from the expected value.

Definition 3.17 For $i \in \{0, \dots, m\}$, the pair $(r, \bar{h}) \in \{0, 1\}^{T_A} \times \mathcal{H}^{\times i}$ is balanced if

$$\frac{|L|}{3 \cdot 2^i} \leq |\text{Consist}(r, \bar{h})| \leq \frac{3 \cdot |L|}{2^i}.$$

Claim 3.18 For every $i \in \{0, \dots, m\}$,

$$\Pr_{(r, \bar{h}) \leftarrow \{0, 1\}^{T_A} \times \mathcal{H}^{\times i}}[(r, \bar{h}) \text{ is balanced}] \geq 1 - \frac{6n^2 2^i}{|L|}.$$

Proof: We say that $h \in \mathcal{H}$ is *good* w.r.t. the pair (r, \bar{h}) , if it partitions $\text{Consist}(r, \bar{h})$ into two (almost) equal size parts. That is, if $|\text{Consist}(r, (\bar{h}, h))| \in \left[\frac{|\text{Consist}(r, \bar{h})|}{2} (1 - \frac{1}{n}), \frac{|\text{Consist}(r, \bar{h})|}{2} (1 + \frac{1}{n}) \right]$. Note that if for a given pair $(r, \bar{h}) \in \{0, 1\}^{T_A} \times \mathcal{H}^{\times k}$ it holds that $\forall i \in [k-1]$ the function \bar{h}_i is good w.r.t. $(r, \bar{h}_{1, \dots, i-1})$, then (r, \bar{h}) is balanced. By Theorem 2.2 it follows that $\Pr_{h \in \mathcal{H}}[h \text{ is not good w.r.t. } (r, \bar{h})] < \frac{2|\text{Consist}(r, \bar{h})|}{n^2}$. Therefore, we can lower bound the probability that a pair is balanced as follows:

$$\begin{aligned} & \Pr_{(r, \bar{h}) \leftarrow \{0, 1\}^{T_A} \times \mathcal{H}^{\times i}}[(r, \bar{h}) \text{ is not balanced}] \\ & \leq \sum_{k=1}^i \Pr_{(r, \bar{h}) \leftarrow \{0, 1\}^{T_A} \times \mathcal{H}^{\times k}}[\bar{h}_k \text{ is not good w.r.t. } (r, \bar{h}_{1, \dots, k-1}) \mid (r, \bar{h}_{1, \dots, k-1}) \text{ is balanced}] \\ & \leq \sum_{k=0}^{i-1} \frac{3n^2 2^k}{|L|} \\ & \leq \frac{6n^2 2^i}{|L|}. \blacksquare \end{aligned}$$

3.5 Analyzing the Success Probability of *Inverter* on D_{uni} .

As mentioned above, it is rather easy to prove (much like the proof for the case that $m = \text{ofs}$) that the success probability of *Inverter* over D_{uni} is at least $\Omega(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^2}{n^6})$. Proving that, however, does not suffice to deduce that the success probability of *Inverter* over D_{src} is also high. The reason is that potentially the success probability of *Inverter* over D_{uni} could stem from a relatively few elements that have significantly smaller probability mass w.r.t. D_{src} than w.r.t. D_{uni} . To overcome this problem, we prove that the success of *Inverter* is sufficiently high even if we flatten this probability such that the contribution of any single element is small. Having that, we are guaranteed that the success probability of *Inverter* is high w.r.t. any distribution that assigns about the same mass to *most* elements in $\text{sup}(D_{uni})$ (and as we show later, D_{src} satisfies this property). More formally, for every $(r, \bar{h}^s) \in \{0, 1\}^{T_A} \times \mathcal{H}^{\times (m-\text{ofs})}$ we let

$$\varepsilon_{r, \bar{h}^s} \stackrel{\text{def}}{=} \Pr_{(o_A \mid o_{\mathcal{R}} = (\bar{h}, *)) \leftarrow \langle A^r(1^n), \mathcal{R}(1^n) \rangle} [\text{BndBreak}^{\text{L}, \text{W}}(o_A, o_{\mathcal{R}}) \mid \bar{h}_{1, \dots, m-\text{ofs}} = \bar{h}^s],$$

That is, $\varepsilon_{r, \bar{h}^s}$ is the cheating probability of A conditioned on (r, \bar{h}^s) . We define the *weight* of $y \in \text{Consist}(r, \bar{h}^s)$ by

$$\begin{aligned} w(r, \bar{h}^s, y) & \stackrel{\text{def}}{=} \\ & \frac{1}{2} \cdot \Pr_{(o_A = ((*, y_0), (*, y_1)) \mid o_{\mathcal{R}} = (\bar{h}, *)) \leftarrow \langle A^r(1^n), \mathcal{R}(1^n) \rangle} [\text{BndBreak}^{\text{L}, \text{W}}(o_A, o_{\mathcal{R}}) \bigwedge y \in \{y_0, y_1\} \mid \bar{h}_{1, \dots, m-\text{ofs}} = \bar{h}^s]. \end{aligned}$$

Note that $w(r, \overline{h^s}, y)$ is a lower bound on the probability that *Inverter* satisfies W on y conditioned on $(r, \overline{h^s})$. Also note that $\varepsilon_{r, \overline{h^s}} = \sum_{y \in \text{Consist}(r, \overline{h^s})} w(r, \overline{h^s}, y)$. Finally, we define the *decreased weight* of $(r, \overline{h^s})$ as

$$w_{dec}(r, \overline{h^s}, y) \stackrel{\text{def}}{=} \min \left\{ \frac{\varepsilon_A}{2^C/2n^3}, w(r, \overline{h^s}, y) \right\},$$

where $C > 0$ is the same universal constant that appears in the definition of **ofs**. The following lemma essentially implies that the success probability of *Inverter* over D_{uni} w.r.t. the decreased weight is noticeable.

Lemma 3.19

$$\mathbb{E}_{(r, \overline{h}, y_s) \leftarrow D_{uni}} [w_{dec}(r, \overline{h^s}, y)] \in \Omega \left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^3}{2^C n^6} \right).$$

Proof: Let $(r, \overline{h^s}) \in \{0, 1\}^{T_A} \times \mathcal{H}^{\times(m-\text{ofs})}$. We assume for simplicity a non-increasing order on the elements of $\text{Consist}(r, \overline{h^s})$ according to their wights and denote by $\text{Consist}(r, \overline{h^s})_i$ the i^{th} element of $\text{Consist}(r, \overline{h^s})$ by this order. The following claim states that the weight is not concentrated only on the first $\ell_{r, \overline{h^s}} \stackrel{\text{def}}{=} \lfloor \sqrt{2^{\text{ofs}-1} \varepsilon_{r, \overline{h^s}}} \rfloor$ heaviest elements of $\text{Consist}(r, \overline{h^s})$.

Claim 3.20 $\sum_{i=\ell_{r, \overline{h^s}}+1}^{|\text{Consist}(r, \overline{h^s})|} w(r, \overline{h^s}, \text{Consist}(r, \overline{h^s})_i) \geq \varepsilon_{r, \overline{h^s}}/4$.

Proof: Let $Z = \bigcup_{i=1}^{\ell_{r, \overline{h^s}}} \{\text{Consist}(r, \overline{h^s})_i\}$, by the pairwise independence of \mathcal{H} it follows that,

$$\begin{aligned} & \Pr_{(h_{m-\text{ofs}+1}, \dots, h_m) \leftarrow \mathcal{H}^{\times \text{ofs}}} [\exists y_0 \neq y_1 \in Z \text{ s.t. } \forall j \in \{m - \text{ofs} + 1, \dots, m\} \ h_j(y_0) = h_j(y_1)] \\ & \leq \frac{|Z|^2}{2^{\text{ofs}}} \leq \frac{2^{\text{ofs}} \varepsilon_{r, \overline{h^s}}}{2 \cdot 2^{\text{ofs}}} = \varepsilon_{r, \overline{h^s}}/2. \end{aligned}$$

Recall that A cheats with probability $\varepsilon_{r, \overline{h^s}}$. Since the probability that both y_0 and y_1 it returns are inside Z is at most $\varepsilon_{r, \overline{h^s}}/2$, it follows that the probability that A^r cheats successfully while at least one of y_0 and y_1 is outside Z is at least $\varepsilon_{r, \overline{h^s}}/2$. Note that each event where A^r cheats successfully and outputs an element $y_i = y$, contributes half its probability to the total weight of y . Thus, the sum of weights of the elements inside $\text{Consist}(r, \overline{h^s}) \setminus Z$ is at least $\varepsilon_{r, \overline{h^s}}/4$. ■

Since $\varepsilon_{r, \overline{h^s}} = \sum_{y \in \text{Consist}(r, \overline{h^s})} w(r, \overline{h^s}, y) \geq \sum_{i=1}^{\ell_{r, \overline{h^s}}} w(r, \overline{h^s}, \text{Consist}(r, \overline{h^s})_i)$, it follows that $w(r, \overline{h^s}, \text{Consist}(r, \overline{h^s})_{\ell_{r, \overline{h^s}}}) \leq \frac{\varepsilon_{r, \overline{h^s}}}{\ell_{r, \overline{h^s}}}$. We have set **ofs** to $\lceil 6 \log(n) + 2 \log(\frac{1}{\varepsilon_A}) \rceil + C$, therefore $\frac{2\varepsilon_{r, \overline{h^s}}}{\ell_{r, \overline{h^s}}} = \frac{2\varepsilon_{r, \overline{h^s}}}{\lfloor \sqrt{2^{\text{ofs}-1} \varepsilon_{r, \overline{h^s}}} \rfloor} \leq \frac{\varepsilon_A}{2^C/2n^3}$. Thus,

$$\sum_{y \in \text{Consist}(r, \overline{h^s})} w_{dec}(r, \overline{h^s}, y) \geq \sum_{i=\ell_{r, \overline{h^s}}+1}^{|\text{Consist}(r, \overline{h^s})|} w(r, \overline{h^s}, \text{Consist}(r, \overline{h^s})_i) \geq \varepsilon_{r, \overline{h^s}}/4. \quad (2)$$

Thus,

$$\begin{aligned} & \mathbb{E}_{(r, \overline{h}, y_s) \leftarrow D_{uni}} [w_{dec}(r, \overline{h^s}, y)] \\ & = \frac{1}{|\{0, 1\}^{T_A} \times \mathcal{H}^{\times \text{ofs}}|} \cdot \sum_{(r, \overline{h^s}, *) \in \text{sup}(D_{uni})} \frac{\sum_{y \in \text{Consist}(r, \overline{h^s})} w_{dec}(r, \overline{h^s}, y)}{|\text{Consist}(r, \overline{h^s})|} \\ & \geq \frac{1}{|\{0, 1\}^{T_A} \times \mathcal{H}^{\times \text{ofs}}|} \cdot \sum_{(r, \overline{h^s}, *) \in \text{sup}(D_{uni}) : |\text{Consist}(r, \overline{h^s})| \leq \frac{3|L|}{2^{m-\text{ofs}}}} \frac{\varepsilon_{r, \overline{h^s}}}{4} / \left(\frac{3|L|}{2^{m-\text{ofs}}} \right) \\ & = \frac{2^{m-\text{ofs}}}{12|L|} \cdot \frac{1}{|\{0, 1\}^{T_A} \times \mathcal{H}^{\times \text{ofs}}|} \cdot \sum_{(r, \overline{h^s}, *) \in \text{sup}(D_{uni}) : |\text{Consist}(r, \overline{h^s})| \leq \frac{3|L|}{2^{m-\text{ofs}}}} \varepsilon_{r, \overline{h^s}}. \end{aligned}$$

Claim 3.18 yields that $\Pr_{(r, \overline{h^s}, *) \leftarrow D_{uni}} \left[|\text{Consist}(r, \overline{h^s})| > \frac{3|L|}{2^{m-\text{ofs}}} \right] \in \mathcal{O} \left(\frac{n^2 2^{m-\text{ofs}}}{|L|} \right)$. Thus,

$$\mathbb{E}_{(r, \overline{h}, y_s) \leftarrow D_{uni}} [w_{dec}(r, \overline{h^s}, y)] \geq \frac{2^{m-\text{ofs}}}{12|L|} \cdot \left(1 - \mathcal{O} \left(\frac{n^2 2^{m-\text{ofs}}}{|L|} \right) \right) \cdot \varepsilon_A \in \Omega \left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^3}{2^C n^6} \right). \blacksquare$$

3.6 Analyzing the Success Probability of *Inverter* on D_{src} .

We would have liked to claim that D_{uni} is statistically close to D_{src} and thus, the proof would immediately follow Lemma 3.19. Unfortunately, we can only prove that the two distributions are at statistical distance that is much bigger than the success probability of *Inverter* on D_{uni} . Hence, we need to refine our approach by considering a different measure of distance. We call an element $y \in \text{sup}(D_{uni})$ “good”, if $\frac{D_{src}(y)}{D_{uni}(y)}$ is not too far from one. We prove that the total mass of the non-good elements in the support of D_{uni} is small. Thus, the probability that a good element is drawn from D_{uni} on which *Inverter* does well is noticeable. Having that, we deduce that *Inverter* also does well on D_{src} . Let us turn to a more formal discussion.

The Proximity Measure.

Definition 3.21 Let D_1 and D_2 be two distributions over a set Z , let $\varepsilon \in [0, 1]$ and let $a \geq 1$. We say that D_1 (ε, a) -approximates D_2 , if there exists a subset $Z' \subseteq Z$ such that the following holds:

1. $D_1(Z') \leq \varepsilon$.
2. For every $x \in \text{sup}(D_1) \setminus Z'$ it holds that $\frac{1}{a} \leq \frac{D_1(x)}{D_2(x)} \leq a$.⁶

The following propositions show that the proximity measure “behaves” similarly to the standard statistical distance measure. Since the proofs of following propositions are rather immediate, they are omitted from this version. The first proposition enables us to use hybrid arguments when proving the proximity between distributions.

Proposition 3.22 (Transitivity) Let D_1 , D_2 and D_3 be distributions over a set Z , let $\varepsilon_1, \varepsilon_2 \in [0, 1]$ and let $a_1, a_2 \geq 1$. Assuming that D_1 (ε_1, a_1) -approximates D_2 and that D_2 (ε_2, a_2) -approximates D_3 , then D_1 $(\varepsilon_1 + a_1 \varepsilon_2, a_1 a_2)$ -approximates D_3 .

Proposition 3.23 (Average) Let $\{D_1^i\}_{i=1}^m$ and $\{D_2^i\}_{i=1}^m$ be two distributions ensembles over some set Z such that for every $i \in [m]$ it holds that D_1^i (ε_i, a) -approximates D_2^i . Let P be some distribution over $[m]$ and for every $j \in \{0, 1\}$ let D_j be the distribution over Z defined as $D_j(x) = \sum_{i=1}^m P(i) D_j^i(x)$. Then D_1 (ε, a) -approximates D_2 , where $\varepsilon = \mathbb{E} x_i \leftarrow P[\varepsilon_i]$.

Proposition 3.24 (Extension) Let D_1 and D_2 be two distributions such that D_1 (ε, a) -approximates D_2 and let $P : (\text{sup}(D_1) \cup \text{sup}(D_2)) \rightarrow \{0, 1\}^*$ be a random process such that for every $x_0 \neq x_2 \in \text{sup}(D_1) \cup \text{sup}(D_2)$ it holds that $\Pr[P(x_0) = P(x_2)] = 0$, then $P(D_1)$ (ε, a) -approximates $P(D_2)$.

The following proposition is where the usefulness of the new proximity measure lies, since it implies that the expected value of any predicate over two distributions that are close to each other is similar.

Proposition 3.25 (Evaluation) Let D_1 and D_2 be two distributions such that D_1 (ε, a) -approximates D_2 . Let $\delta > 0$ and let $P : \text{sup}(D_1) \cup \text{sup}(D_2) \rightarrow [0, \delta]$, then $\mathbb{E} x_x \leftarrow D_2[P(x)] \geq \frac{1}{a} (\mathbb{E} x_x \leftarrow D_1[P(x)] - \varepsilon \delta)$.

Lemma 3.26 $D_{uni} \left(\mathcal{O} \left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^2}{2^C n^3} \right), 81 \right)$ -approximates D_{src} .

⁶Actually, for the purpose of this paper, it would suffice to require that $\frac{1}{a} \leq \frac{D_1(x)}{D_2(x)}$. We chose to use the symmetric definition since it seems more natural to us.

Proof: We “bridge” between D_{uni} and D_{src} using the following hybrid distributions. For every $\overline{h^s} \in \mathcal{H}^{\times k}$, we define the hybrid algorithm $Searcher^{\overline{h^s}}(r, y)$ that sets its first k hash functions to $\overline{h^s}$ and then continues as the original $Searcher$ algorithm does. For any $0 \leq k \leq m - \text{ofs}$ let

$$D^k \stackrel{\text{def}}{=} (r, (\overline{h^s}, \overline{h^e}), y)_{r \leftarrow \{0,1\}^{T_A}, \overline{h^s} \leftarrow \mathcal{H}^{\times k}, y \leftarrow \text{Consist}(r, \overline{h^s}), \overline{h^e} \leftarrow Searcher^{\overline{h^s}}(r, y)_{k+1, \dots, m-\text{ofs}}}.$$

Note that D^0 is equal to D_{src} and that $D^{m-\text{ofs}}$ is equal to D_{uni} . The following lemma states that every neighboring distributions are close to each other.

Lemma 3.27 *For every $0 \leq k < m - \text{ofs}$, let $\ell_k = \frac{|L|}{3 \cdot 2^k}$ and let $\delta_k = \Pr_{(r, \overline{h^s}) \leftarrow \{0,1\}^{T_A} \times \mathcal{H}^{\times k}}[(r, \overline{h^s}) \text{ is not balanced}]$. Then, $D^{k+1} \left(\delta_k + \frac{160 \cdot n^3}{\ell_k}, (1 + \frac{4}{n}) \right)$ -approximates D^k .*

Before proving Lemma 3.27, we use it to prove Lemma 3.26.

Proof: (of Lemma 3.26) By combining Lemma 3.27 and Proposition 3.22, we have that D_{uni} $\left(81 \left(\sum_{k=0}^{m-\text{ofs}-1} \left(\delta_k + \frac{160 \cdot n^3}{\ell_k} \right), 81 \right) \right)$ -approximates D_{src} . Claim 3.18 yields that $\delta_k < \frac{2n^2}{\ell_k}$ and therefore $\sum_{k=0}^{m-\text{ofs}-1} \left(\delta_k + \frac{160 \cdot n^3}{\ell_k} \right) \in \mathcal{O}(\frac{n^3}{\ell_{m-\text{ofs}}})$. Thus, $D_{uni} \left(\mathcal{O}(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^2}{2^{C n^3}}), 81 \right)$ -approximates D_{src} . ■

Proof: (of Lemma 3.27) Note that the only difference between D^k and D^{k+1} is their method of selecting y and $\overline{h_{k+1}}$. Therefore, in the proof we concentrate on the induced distributions on these values only. For any $(r, \overline{h^s}) \in \{0,1\}^{T_A} \times \mathcal{H}^{\times k}$, we define

- $D_{r, \overline{h^s}}^0 \stackrel{\text{def}}{=} (y, h)_{y \leftarrow \text{Consist}(r, \overline{h^s}), h \leftarrow Searcher^{\overline{h^s}}(r, y)_{k+1}}$
- $D_{r, \overline{h^s}}^1 \stackrel{\text{def}}{=} (y, h)_{h \leftarrow \mathcal{H}, y \leftarrow \text{Consist}(r, (\overline{h^s}, h))}$

The proof of Lemma 3.27 will follow from the next Lemma.

Lemma 3.28 *Let $(r, \overline{h^s}) \in \{0,1\}^{T_A} \times \mathcal{H}^{\times k}$ be balanced, then $D_{r, \overline{h^s}}^1 \left(\frac{160 \cdot n^3}{\ell_k}, (1 + \frac{4}{n}) \right)$ -approximates $D_{t, \overline{h^s}}^0$.*

By the extension and average properties of the proximity measure (Propositions 3.24 and 3.23), it follows from Lemma 3.28 that conditioned on any particular $(r, \overline{h^s})$ that is balanced, $D^{k+1} \left(\frac{160 \cdot n^3}{\ell_k}, (1 + \frac{4}{n}) \right)$ -approximates D^k . Since the probability that $(r, \overline{h^s})$ is not balanced is δ_k , the proof of Lemma 3.27 indeed follows. It therefore remains to prove Lemma 3.28.

Proof: (of Lemma 3.28) Consider the Boolean matrix $T^{|\text{Consist}(r, \overline{h^s})| \times |\mathcal{H}|}$, where $T(y, h) = 1$ iff $A^{Com}(r, (\overline{h^s}, h))_{k+1} = h(y)$ and zero otherwise. We identify the indices into T with the set $\text{Consist}(r, \overline{h^s}) \times \mathcal{H}$. The distribution $D_{r, \overline{h^s}}^1$ can be described in relation to T as follows: Choose a random column of T and draw the index of a random one entry from this column (where a “one entry” is simply an entry of the matrix that is assigned the value one). To argue about this process, let us compare the matrix T with the matrix $\hat{T}^{|\text{Consist}(r, \overline{h^s})| \times |\mathcal{H}|}$, where $\hat{T}(y, h) = h(y)$. Note that T can be derived from \hat{T} by flipping all values in some of its columns (where the column which corresponds to h is flipped whenever $A^{Com}(r, (\overline{h^s}, h))_{k+1} = 0$). By the pairwise independence of \mathcal{H} , it follows that most columns of \hat{T} are balanced (have about the same number of zeros and ones) and thus the same holds for T . Hence, the mass that $D_{r, \overline{h^s}}^1$ assigns to most one entries of T is close to $\frac{1}{|\mathcal{H}|} \cdot \frac{2}{|\text{Consist}(r, \overline{h^s})|}$.

The distribution $D_{t, \overline{h^s}}^0$ can also be described in relation to T as follows: Choose a random row of T and for $2 \log(2)$ times draw a random entry from this row. If a one entry is drawn, then choose its index and stop drawing, otherwise select the index of the last drawn entry. Using again the pairwise independent of \mathcal{H} , we can prove that most rows of T are balanced. It follows that w.h.p. a one entry is drawn from T . Hence, the mass that $D_{r, \overline{h^s}}^0$ assigns to most one entries in T is also close to $\frac{1}{|\mathcal{H}|} \cdot \frac{2}{|\text{Consist}(r, \overline{h^s})|}$. Since the support of $D_{r, \overline{h^s}}^1$ and the indices set of one entries in T are the same, the proof of the claim follows.

Let us turn to a more formal discussion. We define

$$\mathcal{H}^{Bad} \stackrel{\text{def}}{=} \left\{ h \in \mathcal{H} : \Pr_{y \leftarrow \text{Consist}(r, \overline{h^s})} [T(h, y) = 1] \notin \left[\frac{1}{2} \cdot \left(1 - \frac{1}{n} \right), \frac{1}{2} \cdot \left(1 + \frac{1}{n} \right) \right] \right\}.$$

The following claim, whose proof is immediate by the pairwise independence of \mathcal{H} (see Theorem 2.2), states that the relative size of \mathcal{H}^{Bad} is small.

Claim 3.29 $\Pr_{h \leftarrow \mathcal{H}}[h \in \mathcal{H}^{Bad}] \leq \frac{2n^2}{\ell_k}$.

Similarly, we define

$$Y^{Bad} \stackrel{\text{def}}{=} \left\{ y \in \text{Consist}(r, \overline{h^s}) : \Pr_{h \leftarrow \mathcal{H}}[T(h, y) = 1] \notin \left[\frac{1}{2} \cdot \left(1 - \frac{1}{n}\right), \frac{1}{2} \cdot \left(1 + \frac{1}{n}\right) \right] \right\}.$$

The following claim states the size of Y^{Bad} is small.

Claim 3.30 $|Y^{Bad}| < 54n^3$.

Proof: Let $Y_{Law}^{Bad} \stackrel{\text{def}}{=} \{y \in \text{Consist}(r, \overline{h^s}) : \Pr_{h \leftarrow \mathcal{H}}[T(h, y) = 1] < \frac{1}{2} \cdot (1 - \frac{1}{n})\}$. We assume that $|Y_{Law}^{Bad}| > 27n^3$ and derive a contradiction (the proof that $|Y_{Law}^{Bad}| < 27n^3$ is analogous). Consider the matrix $T|_{Y_{Law}^{Bad}}$, the restriction of T to the rows Y_{Law}^{Bad} . By definition, the rows of $T|_{Y_{Law}^{Bad}}$ have more zeros than ones. Hence, the matrix $T|_{Y_{Law}^{Bad}}$ itself has more zeros than ones. On the other hand, by the pairwise independence of \mathcal{H} it follows that most columns of $T|_{Y_{Law}^{Bad}}$ are balanced (have about the same number of zeros of ones). Therefore, $T|_{Y_{Law}^{Bad}}$ itself is balanced and a contradiction is derived. Formally, for $h \in \mathcal{H}$ let T_h be the number of ones in the h column, that is $T_h = \sum_{y \in Y_{Law}^{Bad}} T(y, h)$. We upper bound the expectation of T_h as follows,

$$\mathbb{E}_{h \leftarrow \mathcal{H}}[T_h] = \mathbb{E}_{h \leftarrow \mathcal{H}} \left[\sum_{y \in Y_{Law}^{Bad}} T(y, h) \right] = \sum_{y \in Y_{Law}^{Bad}} \mathbb{E}_{h \leftarrow \mathcal{H}}[T(y, h)] < |Y_{Law}^{Bad}| \left(\frac{1}{2} - \frac{1}{2n} \right). \quad (3)$$

Recall that $T(h, y) = 1$ if $A^{Com}(r, (\overline{h^s}, h))_{k+1} = h(y)$ and zero otherwise. Since the set Y^{Bad} is large enough, Theorem 2.2 yields that a random h splits w.h.p. the elements of Y^{Bad} into two almost equals size according to their consistency with A 's answer on h . That it, $\Pr_{h \leftarrow \mathcal{H}}[T_h < |Y_{Law}^{Bad}| \cdot (\frac{1}{2} - \frac{1}{3n})] < \frac{9n^2}{|Y_{Law}^{Bad}|} \leq \frac{1}{3n}$. Thus,

$$\begin{aligned} & \mathbb{E}_{h \leftarrow \mathcal{H}}[T_h] \\ & \geq \frac{1}{|H|} \cdot \sum_{h \in \mathcal{H}: T_h \geq |Y_{Law}^{Bad}| \cdot (\frac{1}{2} - \frac{1}{3n})} |Y_{Law}^{Bad}| \cdot \left(\frac{1}{2} - \frac{1}{3n} \right) \\ & > \left(1 - \frac{1}{3n} \right) \cdot |Y_{Law}^{Bad}| \left(\frac{1}{2} - \frac{1}{3n} \right) \\ & > |Y_{Law}^{Bad}| \left(\frac{1}{2} - \frac{1}{2n} \right), \end{aligned} \quad (4)$$

and a contradiction is derived. ■

The next claim concludes the proof of Lemma 3.28 by presenting a set that actualizes the stated proximity distance between $D_{r, \overline{h^s}}^1$ and $D_{r, \overline{h^s}}^0$.

Claim 3.31 Let $Z \stackrel{\text{def}}{=} \{(y, h) \in \text{const}(r, \overline{h^s}) \times \mathcal{H} : y \in Y^{Bad} \vee h \in \mathcal{H}^{Bad}\}$, then

1. $D_{r, \overline{h^s}}^1(Z) \leq \frac{160 \cdot n^3}{\ell_k}$.
2. For every $(y, h) \in \text{sup}(D_{r, \overline{h^s}}^1) \setminus Z$, it holds that $\frac{D_{r, \overline{h^s}}^1(y, h)}{D_{r, \overline{h^s}}^0(y, h)} \in [\frac{1}{1 + \frac{1}{n}}, 1 + \frac{1}{n}]$.

Proving Claim 3.31.1. Consider the partitioning of Z into $Z_1 \stackrel{\text{def}}{=} \text{const}(r, \overline{h^s}) \times \mathcal{H}^{Bad}$ and $Z_2 \stackrel{\text{def}}{=} Y^{Bad} \times (\mathcal{H} \setminus \mathcal{H}^{Bad})$. Since $\Pr_{D_{r, \overline{h^s}}^1}[Z_1] \leq \Pr_{h \leftarrow \mathcal{H}}[\mathcal{H}^{Bad}]$, Claim 3.29 yields that $\Pr_{D_{r, \overline{h^s}}^1}[Z_1] \leq \frac{2n^2}{\ell_k}$. Claim 3.30 yields that for any $h \in \mathcal{H}$, it holds that $\Pr_{(y, h') \leftarrow D_{r, \overline{h^s}}^1}[y \in Y^{Bad} \mid h' = h] \leq \frac{54n^3}{|\text{Consist}(r, (\overline{h^s}, h))|}$. Recall that for any $h \in \mathcal{H} \setminus \mathcal{H}^{Bad}$, it holds that $|\text{Consist}(r, (\overline{h^s}, h))| > (1 - \frac{1}{n}) \frac{|\text{Consist}(r, \overline{h^s})|}{2}$. Since $(r, \overline{h^s})$ is balanced, it

follows that $|\text{Consist}(r, (\overline{h^s}, h))| > (1 - \frac{1}{n}) \frac{\ell_k}{2}$. Thus, for any $h \in \mathcal{H} \setminus \mathcal{H}^{Bad}$, it holds that $\Pr_{(y, h') \leftarrow D_{r, \overline{h^s}}^1}[y \in Y^{Bad} \mid h' = h] \leq \frac{1}{(1 - \frac{1}{n})} \cdot \frac{2 \cdot 54n^3}{\ell_k} \leq \frac{150n^3}{\ell_k}$ and therefore, $\Pr_{D_{r, \overline{h^s}}^1}[Z_2] \leq \frac{150n^3}{\ell_k}$. We conclude that $\Pr_{D_{r, \overline{h^s}}^1}[Z] = \Pr_{D_{r, \overline{h^s}}^1}[Z_1] + \Pr_{D_{r, \overline{h^s}}^1}[Z_2] \leq \frac{160 \cdot n^3}{\ell_k}$. ■

Proving Claim 3.31.2. For $(y, h) \in \text{sup}(D_{r, \overline{h^s}}^1) \setminus Z$, it holds that $D_{r, \overline{h^s}}^1(y, h) = \Pr_{D_{r, \overline{h^s}}^1}[h] \cdot \Pr_{D_{r, \overline{h^s}}^1}[y \mid h] = \frac{1}{|\mathcal{H}|} \cdot \frac{1}{|\{y' \in \text{Consist}(r, \overline{h^s}) : T(y', h) = 1\}|}$. Since $h \notin \mathcal{H}^{Bad}$, it follows that $D_{r, \overline{h^s}}^1(y, h) \in [\frac{\gamma}{1 + \frac{1}{n}}, \frac{\gamma}{1 - \frac{1}{n}}]$, where $\gamma = \frac{1}{|\mathcal{H}|} \cdot \frac{2}{|\text{Consist}(r, \overline{h^s})|}$. Similarly, we have that $D_{r, \overline{h^s}}^0(y, h) = \frac{1}{|\text{Consist}(r, \overline{h^s})|} \cdot \Pr_{D_{r, \overline{h^s}}^0}[h \mid y]$. Calculating the value of $\Pr_{D_{r, \overline{h^s}}^0}[h \mid y]$, however, is a bit more subtle. Conditioned on y , it might be that the entry drawn from T is zero. Thus, the conditional probability of h is not the uniform one over $\{h' \in \mathcal{H} : T(y, h') = 1\}$. Still, since $y \notin Y^{Bad}$, we have that the conditional probability that $T(y, h') \neq 1$ is in $o(\frac{1}{n})$. (Recall the description $D_{r, \overline{h^s}}^0$: For $2 \log(n)$ rounds a random entry is selected from the y row of T , only if all the selected entries are zeros, then a zero entry is chosen). Therefore, the conditional probability of h is close to uniform over $\{h' \in \mathcal{H} : T(y, h') = 1\}$ and thus, $D_{r, \overline{h^s}}^0(y, h) = \frac{1}{|\text{Consist}(r, \overline{h^s})|} \cdot \frac{1}{|\{h' \in \mathcal{H} : T(y, h') = 1\}|} \cdot (1 \pm o(\frac{1}{n}))$. Using again the fact that $y \notin Y^{Bad}$, it follows that $D_{r, \overline{h^s}}^0(y, h) \in [\frac{\gamma - o(\frac{1}{n})}{1 + \frac{1}{n}}, \frac{\gamma + o(\frac{1}{n})}{1 - \frac{1}{n}}]$ and we conclude that $\frac{D_{r, \overline{h^s}}^1(y, h)}{D_{r, \overline{h^s}}^0(y, h)} \in [\frac{1}{1 + \frac{1}{n}}, 1 + \frac{1}{n}]$. ■

Putting It All Together

Recall that $\Pr[\text{Inverter}(r, \overline{h}, y) \in W_y] \geq w(r, \overline{h}, y)$. We therefore have that $\Pr_{y \leftarrow L}[M^A(y) \in W_y] = \Pr_{(r, \overline{h}, y) \leftarrow D_{src}}[\text{Inverter}(r, \overline{h}, y) \in W_y] \geq \mathbb{E}_{(r, \overline{h}, y) \leftarrow D_{src}}[w_{dec}(r, \overline{h}, y)]$. We can now relate this expectation to an expectation over the distribution D_{uni} (on which we have a better handle). For that we use Lemma 3.26 regarding the proximity of the distributions D_{src} and D_{uni} and the evaluation property of the proximity measure (Proposition 3.25). Recalling that by its definition $w_{dec}(r, \overline{h}, y) \leq \frac{\varepsilon_A}{2^{C/2}n^3}$, we can deduce that

$$\Pr_{y \leftarrow L}[M^A(y) \in W_y] \geq \frac{1}{81} \left(\mathbb{E}_{(r, \overline{h}, y) \leftarrow D_{uni}}[w_{dec}(r, \overline{h}, y)] - \mathcal{O}\left(\left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^2}{2^C n^3}\right) \cdot \frac{\varepsilon_A}{2^{C/2}n^3}\right) \right).$$

Finally, Lemma 3.19 yields that,

$$\Pr_{y \leftarrow L}[M^A(y) \in W_y] \geq \frac{1}{81} \left(\Omega\left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^3}{2^C n^6}\right) - \mathcal{O}\left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^3}{2^{\frac{3}{2}C} n^6}\right) \right) \in \Omega\left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^3}{n^6}\right),$$

for large enough C . ■

3.7 Achieving $\Omega(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^2}{n^8})$.

Since the success probability of M^A is at most the success probability of *Inverter* over D_{uni} (at least by our proof's method) and since the latter is at most $\frac{\varepsilon_A}{2^{\text{ofs}}}$, in order to get a better bound on M^A 's success probability, we have to consider smaller values for **ofs**. Following the same lines as the proof of Lemma 3.26, we could prove that for any value of **ofs**, it holds that D_{uni} ($\mathcal{O}(\frac{2^m}{|L|} \cdot n^3 2^{-\text{ofs}}, 81)$)-approximates D_{src} (where D_{uni} and D_{src} are redefined w.r.t. the new value of **ofs**). The problem is, however, that the value we have previously chosen for **ofs** is the smallest value for which the “additive part” (i.e., the $\mathcal{O}(\frac{2^m}{|L|} \cdot n^3 2^{-\text{ofs}})$ part) in the proximity gap between D_{uni} and D_{src} does not overwhelm the success probability of *Inverter* on D_{uni} . Fortunately, by reexamining the proof of Lemma 3.26 w.r.t. **ofs** = $\lceil 8 \log(n) + \log(\frac{1}{\varepsilon_A}) \rceil + 13$, we can prove that the set that actualizes the proximity between D_{uni} and D_{src} is rather evenly spread along the possible pairs of $(r, \overline{h^s})$. That is, we have the following lemma.

Lemma 3.32 (Stronger version of Lemma 3.26) *There exist a set $T \subseteq \text{sup}(D_{uni})$ such that the following holds:*

1. For any $(r, \overline{h}, y) \in \text{sup}(D_{uni}) \setminus T$ it holds that $\frac{1}{81} \leq \frac{D_{src}(r, \overline{h}, y)}{D_{uni}(r, \overline{h}, y)} \leq 81$.

$$2. \Pr_{(r, \bar{h}, *) \leftarrow D_{uni}} [|\text{Consist}(r, \bar{h}^s) \cap T| > 54n^4] \in \mathcal{O}(\frac{\varepsilon_A}{n^5}).$$

The advantage of the above equation over Lemma 3.26, is that it guarantees that conditioned on almost every choice of (r, \bar{h}^s) , the distribution D_{uni} approximates D_{src} well. In contrast to Lemma 3.26 that only guarantees that D_{uni} approximates D_{src} well on the average over the choice of (r, \bar{h}^s) . In particular, the above lemma allows us to relate the success probability of M^A over D_{src} to the success probability of M^A over D_{uni} conditioned on, almost, all values of (r, \bar{h}^s) .

Before using Lemma 3.26 to derive Theorem 3.8 (rather than the weaker version of the previous section), we first need to prove a different version of Lemma 3.19. Recall that Lemma 3.19 states that the success probability of *Inverter* over D_{uni} is noticeable, even if we ignore elements on which *Inverter*'s success probability is higher than some threshold. The next lemma states that the success probability of *Inverter* over D_{uni} is noticeable, even if we ignore elements of some given set, T , whose relative part in all but $\frac{\varepsilon_A}{2}$ of the pairs (r, \bar{h}^s) is not too big (keep in mind that the set of Lemma 3.26 is such a set).

Lemma 3.33 (*New version of Lemma 3.19*) *Let $T \subseteq \text{sup}(D_{uni})$ and let $w_T(x) = w(x)$ if $x \notin T$ and zero otherwise. If $\Pr_{(r, \bar{h}, *) \leftarrow D_{uni}} [|\text{Consist}(r, \bar{h}^s) \cap T| > 54n^4] < \varepsilon_A/2$, then*

$$\mathbb{E}_{(r, \bar{h}, y) \leftarrow D_{uni}} [w_T(r, \bar{h}^s, y)] \in \Omega(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^2}{n^8}).$$

Proof: By a Markov argument it follows that the success probability of A is at least $\frac{\varepsilon_A}{2}$, even if it is “forced” to fail on every $(r, \bar{h}^s) \in \{0, 1\}^{T_A} \times \mathcal{H}^{\times(m-\text{ofs})}$ such that $\varepsilon_{r, \bar{h}^s} < \frac{\varepsilon_A}{2}$. Therefore, we assume w.l.o.g. that either $\varepsilon_{r, \bar{h}^s} = 0$ or $\varepsilon_{r, \bar{h}^s} \geq \frac{\varepsilon_A}{2}$. The proof of the lemma follows by the next claim.

Claim 3.34 *Let $(r, \bar{h}^s) \in \{0, 1\}^{T_A} \times \mathcal{H}^{\times(m-\text{ofs})}$ and let $V \subseteq \text{Consist}(r, \bar{h}^s)$ be a subset of size at most $54n^4$. Then,*

$$\sum_{y \in \text{Consist}(r, \bar{h}^s) \setminus V} [w(r, \bar{h}^s, y)] \geq \frac{\varepsilon_{r, \bar{h}^s}}{4}.$$

The above Claim, whose proof we defer for now, yields that for any pair (r, \bar{h}) such that $|\text{Consist}(r, \bar{h}^s) \cap T| \leq 54n^4$, it holds that

$$\sum_{y \in \text{Consist}(r, \bar{h})} [w_T(r, \bar{h}, y)] \geq \frac{\varepsilon_{r, \bar{h}}}{4}. \quad (5)$$

Hence,

$$\begin{aligned} & \mathbb{E}_{(r, \bar{h}, y) \leftarrow D_{uni}} [w_T(r, \bar{h}^s, y)] \\ &= \frac{1}{|\{0, 1\}^{T_A} \times \mathcal{H}^{\times \text{ofs}}|} \cdot \sum_{(r, \bar{h}^s, *) \in \text{sup}(D_{uni})} \frac{\sum_{y \in \text{Consist}(r, \bar{h}^s)} w_T(r, \bar{h}^s, y)}{|\text{Consist}(r, \bar{h}^s)|} \\ &\geq \frac{1}{|\{0, 1\}^{T_A} \times \mathcal{H}^{\times \text{ofs}}|} \cdot \frac{2^{m-\text{ofs}}}{3|L|} \sum_{(r, \bar{h}^s, *) \in \text{sup}(D_{uni}): |\text{Consist}(r, \bar{h}^s) \cap T| \leq 54n^4, |\text{Consist}(r, \bar{h}^s)| \leq \frac{3|L|}{2^{m-\text{ofs}}}} \left(\sum_{y \in \text{Consist}(r, \bar{h}^s)} w_T(r, \bar{h}^s, y) \right) \\ &\geq \frac{2^{m-\text{ofs}}}{12|L|} \cdot \frac{1}{|\{0, 1\}^{T_A} \times \mathcal{H}^{\times \text{ofs}}|} \cdot \sum_{(r, \bar{h}^s, *) \in \text{sup}(D_{uni}): |\text{Consist}(r, \bar{h}^s) \cap T| \leq 54n^4, |\text{Consist}(r, \bar{h}^s)| \leq \frac{3|L|}{2^{m-\text{ofs}}}} \varepsilon_{r, \bar{h}^s}. \end{aligned}$$

Recall that Lemma 3.28 states that $\Pr_{(r, \bar{h}, *) \leftarrow D_{uni}} [|\text{Consist}(r, \bar{h}^s) \cap T| > 54n^4] \in \mathcal{O}(\frac{\varepsilon_A}{n^5})$ and that Claim 3.18 yields that $\Pr_{(r, \bar{h}^s, *) \leftarrow D_{uni}} [|\text{Consist}(r, \bar{h}^s)| > \frac{3|L|}{2^{m-\text{ofs}}}] \in \mathcal{O}(\frac{n^2 2^{m-\text{ofs}}}{|L|})$. Thus,

$$\mathbb{E}_{(r, \bar{h}, y) \leftarrow D_{uni}} [w_T(r, \bar{h}^s, y)] \geq \frac{2^{m-\text{ofs}}}{12|L|} \cdot \left(1 - \mathcal{O}(\frac{\varepsilon_A}{n^5}) - \mathcal{O}(\frac{n^2 2^{m-\text{ofs}}}{|L|}) \right) \cdot \varepsilon_A \in \Omega(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^2}{n^8}). \blacksquare$$

Proof: (of Claim 3.34) Recall that $w(r, \bar{h}, y_s) = \frac{1}{2} \cdot \Pr_{(o_A=((*, y_0), (*, y_1)) | o_{\mathcal{R}}=(\bar{h}, *)) \leftarrow \langle A^r(1^n), \mathcal{R}(1^n) \rangle} [\text{BndBreak}^{\text{L}, \text{W}}(o_A, o_{\mathcal{R}}) \wedge y \in \{y_0, y_1\} \mid \bar{h}_1, \dots, m-\text{ofs} = \bar{h}^s]$ and assume for simplicity some non-increasing order on the elements of $\text{Consist}(r, \bar{h}^s)$. Claim 3.20 states that $\sum_{i=\lfloor \sqrt{2^{\text{ofs}-1} \varepsilon_{r, \bar{h}^s}} \rfloor + 1}^{|\text{Consist}(r, \bar{h}^s)|} w(r, \bar{h}^s, \text{Consist}(r, \bar{h}^s)_i) \geq \varepsilon_{r, \bar{h}^s}/2$. We assume w.l.o.g. that $\varepsilon_{r, \bar{h}^s} \geq \frac{\varepsilon_A}{2}$. Therefore, $\sqrt{2^{\text{ofs}-1} \varepsilon_{r, \bar{h}^s}} > \sqrt{2^{\text{ofs}-2} \varepsilon_A} > 54n^4$ and we conclude that $\sum_{y \in \text{Consist}(r, \bar{h}^s) \setminus V} w(r, \bar{h}^s, \text{Consist}(r, \bar{h}^s)_i) \geq \frac{\varepsilon_{r, \bar{h}^s}}{4}$. ■

Putting It All Together

Let T be the set whose existence is guaranteed by Lemma 3.32. Recall that $\Pr[\text{Inverter}(r, \bar{h}, y) \in W_y] \geq w(r, \bar{h}, y)$. We therefore have that $\Pr_{y \leftarrow L}[M^A(y) \in W_y] = \Pr_{X(r, \bar{h}, y) \leftarrow D_{\text{src}}}[\text{Inverter}(r, \bar{h}, y) \in W_y] \geq \mathbb{E}_{X(r, \bar{h}, y) \leftarrow D_{\text{src}}}[w_T(r, \bar{h}, y)]$. Since D_{uni} approximates D_{src} well on any element in T , it follows that

$$\Pr_{y \leftarrow L}[M^A(y) \in W_y] \geq \frac{1}{81} \mathbb{E}_{(r, \bar{h}, y) \leftarrow D_{\text{uni}}} [w_T(r, \bar{h}, y)].$$

Finally, Lemma 3.33 yields that,

$$\Pr_{y \leftarrow L}[M^A(y) \in W_y] \in \Omega\left(\frac{2^m}{|L|} \cdot \frac{\varepsilon_A^2}{n^8}\right). \blacksquare$$

4 Applying Our New Proof to NOVY

The NOVY protocol is basically an instance of the protocol given in Construction 3.6, where the number of rounds is set to $n-1$ and W is naturally defined by a one-way permutation (i.e., given a one-way permutation f over strings of length n , then W consists on all the pairs $(x, f(x))$ where $x \in \{0, 1\}^n$). The main difference is that rather than using the same family of Boolean pairwise independent hash functions in each round, the NOVY protocol uses a different family for each round. Specifically, the protocol's i^{th} family \mathcal{H}^i is defined by the set of all strings of the form $0^{i-1}1\{0, 1\}^{n-i}$, where for $h \in \mathcal{H}^i$ and $x \in \{0, 1\}^n$ the hash value $h(x)$ is defined as $\langle h, x \rangle \bmod 2$.

We would like to apply Theorem 3.8 also to the NOVY protocol. We could do so directly if the families $\{\mathcal{H}^i\}_{i=1}^{n-1}$ would be guaranteed to be pairwise independent w.r.t. $\{0, 1\}^n$.⁷ The latter, however, does not hold and therefore we have to refine our approach. Fortunately, the proof of the theorem does not require that the families of Boolean hash function to be pairwise independent w.r.t. the initial set of inputs L (in the NOVY case w.r.t. $\{0, 1\}^n$), but rather to be pairwise independent w.r.t. the elements of the initial set that are consistent with the protocol so far. It turns out that given that the initial set is $\{0, 1\}^n$, the families of Boolean hash functions used by NOVY are “enough” pairwise independent on the relevant set. Thus, the proof of Theorem 3.8 can be also applied to the NOVY setting.

Let's us turn to a more formal discussion. For every $0 \leq k \leq n-2$, $\bar{h} \in (\mathcal{H}^1 \times \dots \times \mathcal{H}^k)$ and $\bar{z} \in \{0, 1\}^k$, let $\text{Consist}(\bar{h}, \bar{z})$ be the set of elements inside $\{0, 1\}^n$ that are consistent with \bar{h} and \bar{z} (i.e., $\{y \in \{0, 1\}^n : \forall i \in [k] \bar{h}_i(y) = \bar{z}_i\}$). By induction, it follows that for any possible pair (\bar{h}, \bar{z}) and element $y_2 \in \{0, 1\}^{n-k}$, there exists a *single* element $y_1 \in \{0, 1\}^k$ (which depends on y_2 and on (\bar{h}, \bar{z})) such that $y_1 y_2 \in \text{Consist}(\bar{h}, \bar{z})$. Hence, for any $y \in \text{Consist}(\bar{h}, \bar{z})$ there exists exactly one other element $y' \in \text{Consist}(\bar{h}, \bar{z})$ for which $y_{k+2\dots n} = y'_{k+2\dots n}$. Thus, for any other element $y'' \in \text{Consist}(\bar{h}, \bar{z})$, which is different than y and y' , it holds that the random variables $h(y)$ and $h(y'')$ (and also $h(y')$ and $h(y'')$), where h is a random function from \mathcal{H}^{k+1} , are independent. Therefore, every subset $Z \subseteq \text{Consist}(\bar{h}, \bar{z})$ can be partitioned into two almost equal size subsets (i.e., of difference in size at most one) such that \mathcal{H}^{k+1} is pairwise independent w.r.t. both subsets. Through the proof of Theorem 3.8, we use the pairwise independence property of the hash functions to prove that the following holds: Every fixed large enough subset $Z \subseteq \text{Consist}(\bar{h}, \bar{z})$ is partitioned w.h.p. by a random Boolean hash function into two parts of almost the same size.⁸ By the latter observation it holds

⁷Note that the proof of Theorem 3.8 does not require that the same family is used in each round.

⁸Actually, save but the proof of Claim 3.30, we only need this property w.r.t. $Z = \text{Consist}(\bar{h}, \bar{z})$. Note that by the above observation about the structure of $\text{Consist}(\bar{h}, \bar{z})$, every $h \in \mathcal{H}^{k+1}$ *always* partitions $\text{Consist}(\bar{h}, \bar{z})$ into two equal parts.

that with high enough probability such a partition also happens w.r.t. the family \mathcal{H}^{k+1} . Thus, the proof of Theorem 3.8 holds for the NOVY protocol as well.

5 Discussion and Further Research

One interesting question is to come with a reduction from interactive hashing to one-way permutation that is even more security preserving. Particularly, is there such a reduction that is linearly-preserving [HL92] (i.e., where the time-success ratio of an adversary inverting the one-way permutation is only larger by a multiplicative polynomial factor than the time-success ratio of an adversary breaking the interactive-hashing protocol). There are three possible directions for an improvement: (1) Presenting a more secure protocol than the NOVY protocol (or our variant), (2) Giving a better reduction from an adversary that breaks the interactive hashing to one that breaks the one-way permutations, or (3) Improving the analysis of the reduction mentioned in 2. Note that our improvement in parameters over the NOVY proof is mainly in the third item (i.e., the analysis of the reduction). As we now argue, it seems that improving the analysis in itself cannot imply a linear-preserving reduction.

Consider an algorithm M for inverting a one-way permutation that uses an adversary A of the NOVY protocol in the following black-box manner: On $y \in \{0, 1\}^n$, it keeps sampling random hash functions and rewinding A , until it finds a series of $n - 1$ hash functions on which A 's answers is consistent with y . Then, it returns one of A 's outputs as the candidate preimage of y . Assume that A operates as follows: For $\varepsilon > 0$, it replies with random answers on the first $n - \log(\frac{1}{\varepsilon})$ questions (hash functions) and then randomly selects two distinct elements, $y_1, y_2 \in \{0, 1\}^n$, that are consistent with the protocol so far. For the remaining hash functions A does the following: if both y_1 and y_2 yield the same answer then it answers with this value, otherwise, it selects randomly one of the elements and from now on answers according to this element. At the end of the protocol A checks whether both y_1 and y_2 are consistent with the protocol. If the answer is positive, it inverts f on both y_1 and y_2 and outputs the result (recall that the reduction does not assume that A is efficient and therefore it is allowed for example to invert f using exhaustive search), otherwise it outputs \perp . Since \mathcal{H} is a family of pairwise independent hash function, the random variables $h(y_1)$ and $h(y_2)$, for a randomly chosen hash function h , are independent.⁹ Thus, the probability that A breaks the NOVY protocol is exactly ε . On the other hand, in order for M to succeed, y has to be selected by A as one of the elements in $\{y_1, y_2\}$. Since the number of elements that are consistent with the protocol after $n - \log(\frac{1}{\varepsilon})$ steps is $1/\varepsilon$, it follows that this happens with probability 2ε . Given that $y \in \{y_1, y_2\}$, say that $y = y_1$, M has to choose in each step an hash function h for which $A(h) = h(y) = h(y_2)$. By the independence of $h(y)$ and $h(y_2)$, it follows that the probability that $A(h) = h(y) \neq h(y_2)$ is exactly $\frac{1}{4}$. Therefore, the probability that in all the last $\log(\frac{1}{\varepsilon})$ steps it holds that $A(h) = h(y) = h(y_2)$, is at most $(\frac{3}{4})^{\log(\frac{1}{\varepsilon})} < \varepsilon^{0.4}$. We conclude that the overall success probability of M^A is at most $2 \cdot \varepsilon^{1.4}$.

We note that for the above case, it is easy to present an algorithm that inverts f , using black-box access to A , with probability that is very close to ε . Nevertheless, it is possible that one can generalize and strengthen the above argument to preclude any linearly-preserving black-box reduction from interactive hashing to one-way permutation. Such a separation would be quite informative (an easier task would be to rule out any black-box proof that the NOVY protocol is linearly preserving).

Presenting an interactive hashing protocol with fewer rounds is another challenging task. The number of rounds in the NOVY and our protocols is quite high (i.e., $\theta(n)$).¹⁰ Any substantial improvement over this number would be very interesting. The complementary task would be to present a lower-bound on the number rounds for any black-box reduction.

⁹As mentioned in Section 4, the hash functions used by the NOVY protocol are not exactly pairwise independent. However, almost the same argument holds for the NOVY hash functions.

¹⁰It is easy to extend our analysis to a variant of our protocol that uses $\theta(\log(n))$ bits output hash functions rather than Boolean hash functions. Therefore, the number of rounds in this modified protocol would be $\theta(n/\log(n))$. This was also recently shown for a modification of the NOVY analysis [KS06]

References

- [AGGM06] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on NP-hardness. In ACM, editor, *38th IEEE FOCS*, 2006.
- [BCC88] G. Brassard, D. Chaum, and C. Crépeau. Minimum disclosure proofs of knowledge. *J. Computer and System Sciences*, 37(2):156–189, 1988.
- [Blu82] M. Blum. Coin flipping by phone. In *IEEE COMPCOM*, 1982.
- [CCM98] Cachin, Crepeau, and Marcil. Oblivious transfer with a memory-bounded receiver. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 1998.
- [CS06] C. Crpeau and G. Savvides. Optimal reductions between oblivious transfers using interactive hashing. In *Advances in Cryptology - Eurocrypt '06, Lecture Notes in Computer Science*, 2006.
- [CW77] I. Carter and M. Wegman. Universal classes of hash functions. In *9th ACM Symposium on Theory of Computing*, pages 106–112, 1977.
- [DHRS04] Ding, Harnik, Rosen, and Shaltiel. Constant-round oblivious transfer in the bounded storage model. In *Theory of Cryptography Conference (TCC), LNCS*, volume 1, 2004.
- [GKL93] O. Goldreich, H. Krawczyk, and M. Luby. On the existence of pseudorandom generators. *SIAM Journal of Computing*, 22(6):1163–1175, 1993.
- [Gol01] O. Goldreich. Randomized methods in computation - lecture notes. 2001.
- [HHK⁺05] Haitner, Horvitz, Katz, Koo, Morselli, and Shaltiel. Reducing complexity assumptions for statistically-hiding commitment. In *EUROCRYPT: Advances in Cryptology: Proceedings of EUROCRYPT*, 2005.
- [HKKM04] Omer Horvitz, Jonathan Katz, Chiu-Yuen Koo, and Ruggero Morselli. Reducing complexity assumptions for statistically-hiding commitment. Cryptology ePrint Archive, Report 2004/341, 2004. <http://eprint.iacr.org>.
- [HL92] A. Herzberg and M. Luby. Pubic randomness in cryptography. In *Advances in Cryptology - CRYPTO '92, Lecture Notes in Computer Science*, volume 740, pages 421–432. Springer, 1992.
- [HR06] I. Haitner and O. Reingold. Statistically-hiding bit-commitment from exponentially hard one-way functions. Manuscript in preparation, 2006.
- [KS06] T. Koshiba and Y. Seri. Round-efficient one-way permutation based perfectly concealing bit commitment scheme. ECCS, TR06-093, 2006.
- [Lin03] Y. Lindell. Parallel coin-tossing and constant-round secure two-party computation. *Journal of Cryptology*, 16(3):143–184, 2003.
- [NOV06] M. Nguyen, S. Ong, and S. Vadhan. Statistical zero-knowledge arguments for NP from any one-way function. Electronic Colloquium on Computational Complexity (ECCC), TR06-075, 2006.
- [NOVY98] 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.
- [NV06] M. Nguyen and S. Vadhan. Zero knowledge with efficient provers. In *Proceedings of the 38th ACM Symposium on Theory of Computing*, 2006.
- [OVY92] Rafail Ostrovsky, Ramarathnam Venkatesan, and Moti Yung. Secure commitment against A powerful adversary. In *9th Annual Symposium on Theoretical Aspects of Computer Science*, volume 577 of *lnes*, pages 439–448, Cachan, France, 13–15 February 1992. Springer.

- [OY93] Ostrovsky, Venkatesan, and Yung. Interactive hashing simplifies zero-knowledge protocol design. In *EUROCRYPT: Advances in Cryptology: Proceedings of EUROCRYPT*, 1993.