

ON THE COMPRESSIBILITY OF \mathcal{NP} INSTANCES AND CRYPTOGRAPHIC APPLICATIONS*

DANNY HARNIK[†] AND MONI NAOR[‡]

Abstract. We study compression that preserves the *solution* to an instance of a problem rather than preserving the instance itself. Our focus is on the compressibility of \mathcal{NP} decision problems. We consider \mathcal{NP} problems that have long instances but relatively short witnesses. The question is whether one can efficiently compress an instance and store a shorter representation that maintains the information of whether the original input is in the language or not. We want the length of the compressed instance to be polynomial in the length of the *witness* and polylog in the length of original input. Such compression enables succinctly storing instances until a future setting will allow solving them, either via a technological or algorithmic breakthrough or simply until enough time has elapsed. In this paper, we first develop the basic complexity theory of compression, including reducibility, completeness, and a stratification of \mathcal{NP} with respect to compression. We then show that compressibility (say, of SAT) would have vast implications for cryptography, including constructions of one-way functions and collision resistant hash functions from any hard-on-average problem in \mathcal{NP} and cryptanalysis of key agreement protocols in the “bounded storage model” when mixed with (time) complexity-based cryptography.

Key words. compression, NP problems, witness length, cryptography, collision resistant hash, bounded storage model

AMS subject classifications. 68Q15, 68Q17, 94A60

DOI. 10.1137/060668092

1. Introduction. In order to deal with difficult computational problems several well-established options were developed, including approximation algorithms, subexponential algorithms, parametric complexity, and average-case complexity. In this paper we explore our favorite approach for dealing with problems: *postpone* them (hopefully without cluttering our desk or disk). We initiate the study of the compressibility of \mathcal{NP} problems for their resolution in some future setting and in particular the cryptographic significance of such compression. Rather than solving a given instance, we ask whether a shorter instance with the same solution can be found efficiently. We emphasize that we are not interested in maintaining the information about the original instance (as is the case in typical notions of compression), but rather maintaining the solution only. The solution can possibly be much shorter than the input (as short as a yes/no answer); thus the potential of such a compression is high.

While the question of compressibility is interesting with respect to problems both inside and out of \mathcal{NP} , our focus is mostly on a special case, that of \mathcal{NP} problems that have long instances but relatively short witnesses. An \mathcal{NP} language L is defined by an efficiently computable relation R_L such that an input (or instance) x is in L if and only if there exists a witness w such that $R_L(x, w) = 1$. Throughout the paper, an \mathcal{NP} instance is characterized by two parameters m and n : the length of the instance

*Received by the editors August 23, 2006; accepted for publication (in revised form) August 19, 2009; published electronically January 15, 2010. This research was supported by a grant from the Israel Science Foundation. A short version of this paper appeared in *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, 2006.

<http://www.siam.org/journals/sicomp/39-5/66809.html>

[†]IBM Research Labs, Haifa 31905, Israel (danny.harnik@gmail.com). This research was conducted while at the Weizmann Institute of Science, Rehovot.

[‡]Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel (moni.naor@weizmann.ac.il).

x is denoted by m , and the length of the witness w is denoted by n . The problems of interest are those having relatively short witnesses, i.e., $n \ll m$, but not too short ($m \ll 2^n$). Traditionally, the study of \mathcal{NP} languages evolves around the ability or inability to efficiently decide if an instance is in the language or not, or to find a witness w for an instance $x \in L$ within polynomial time. We introduce the question of compressibility of such instances.

Example of compressing SAT instances. To illustrate the relevant setting, we use the well-known example of SAT. An instance Φ for SAT consists of a conjunctive normal form (CNF) formula over n variables, and we define that $\Phi \in \text{SAT}$ if there exists an assignment to the n variables that satisfies all the clauses of Φ . We begin with compressibility with respect to decision and discuss the search variant of compressibility later in the paper. In this example we consider the question of compressibility of SAT instances to shorter SAT instances.¹

Example 1.1 (compression of SAT instances). Does there exist an efficient algorithm and a polynomial $p(\cdot, \cdot)$ with the following input and output?

Input: A CNF formula Φ with m clauses over n variables.

Output: A formula Ψ of size $p(n, \log m)$ such that Ψ is satisfiable if and only if Φ is satisfiable.

The idea is that the length of Ψ should not be related to the original length m , but rather to the number of variables (or, in other words, to the size of the witness). Typically, we think of the parameters m and n as related by some function, and it is instructive (but not essential) to think of m as larger than any polynomial in n . So potentially, the length of Ψ can be significantly shorter than that of Φ .²

In general, one cannot expect to compress all the formulas, or else we would have an efficient algorithm for all \mathcal{NP} problems.³ However, once we restrict the attention to the case of a shorter witness, then compression becomes plausible. Note that if $\mathcal{P} = \mathcal{NP}$, then compression becomes trivial, simply by solving the satisfiability of Φ and outputting 1 if $\Phi \in \text{SAT}$ and 0 otherwise.

Motivation for compression. Compressing for the future is an appealing notion for various settings. There are numerous plausible scenarios that will give us more power to solve problems in the future. We could potentially find out that $\mathcal{P} = \mathcal{NP}$ and solve all our \mathcal{NP} problems then. We may have faster computers or better means of computing such as quantum computers or some other physical method for solving problems (see Aaronson [1] for a list of suggestions). Above all, the future entails lots and lots of time, a resource that is usually scarce in the present. Saving the problems of today as they are presented is wasteful, and compression of problems will allow us to store a far greater number of problems for better days.

Our interest in studying the issue of compression stems from the vast cryptographic implications of compressibility. We demonstrate three questions in cryptography that compression algorithms would resolve (see section 1.3). We are confident that the notion of compressibility will be found to have further applications both within and outside of cryptography. For example, in subsequent works, Dubrov and Ishai [26] show the relevance of the notion of compression to derandomization, and Dziembowski [28] shows that compression is related to the study of forward-secure storage

¹This example comes only as an illustration. We later consider the more general question of compression to instances that are not necessarily of the same language.

²Note that since our requirement for compression is only relevant for problems where $m \gg n$, an \mathcal{NP} -complete problem such as 3-SAT (where all clauses have exactly 3 literals) is irrelevant for compression, as m is already at most $O(n^3)$ in such formulas.

³Suppose that every formula can be compressed by a single bit; then sequentially reapplying compression to the input will result in a very short formula that may be solved by brute enumeration.

(see section 1.4 on related work). We note that a notion similar to compression has been useful (and well studied) in the context of parameterized complexity (see a comparison and discussion in section 1.4). The concept of compression of problems is also interesting beyond the confines of \mathcal{NP} problems and makes sense in any setting where the compression requires fewer resources than the actual solution of the problem.

1.1. Compression of \mathcal{NP} instances. We define the notion of compression with respect to an \mathcal{NP} language L . We associate with L a specific fixed \mathcal{NP} relation R_L that defines it (as mentioned above) as well as a function $n(x)$ that defines an upper bound on the length of a potential witness for an instance x .⁴ At times, for simplicity, we abuse notations and simply refer to the language L and omit the reference to the underlying relation R_L . In essence, a compression algorithm is a specialized Karp-reduction that also reduces the length of the instance.

DEFINITION 1.2 (compression algorithm for \mathcal{NP} instances). *Let $L = (R_L, n(\cdot))$ be an \mathcal{NP} language. Denote by m and n the instance length and the witness length, respectively. A compression algorithm for L is a polynomial-time machine Z along with a language L' and a polynomial $p(\cdot, \cdot)$ such that for all large enough m*

1. *for all $x \in \{0, 1\}^m$ with parameter n the length of $Z(x)$ is at most $p(n, \log m)$.*
2. *$Z(x) \in L'$ if and only if $x \in L$.*

The above definition is of an errorless compression. We also consider a probabilistic variant called ε -compression for some real function $\varepsilon : \mathbb{N} \rightarrow [0, 1]$. The probabilistic definition is identical to the errorless one except that Z is a probabilistic polynomial-time machine and the second property is augmented to

- 2'. *For large enough n , for all $x \in \{0, 1\}^m$ with parameter n it holds that*

$$\Pr[(Z(x) \in L') \Leftrightarrow (x \in L)] \geq 1 - \varepsilon(n),$$

where probability is over the internal randomness of Z . By default we require $\varepsilon(\cdot)$ to be negligible (i.e., $\varepsilon(n) = n^{-\omega(1)}$).⁵

The paper consists of two parts: Part I is a study of the concept of compression of \mathcal{NP} instances from a complexity point of view. Part II introduces the cryptographic applications of compression algorithms.

How much to compress. Definition 1.2 (of compression algorithms) requires a very strong compression, asking that the length of the compression be polynomial in n and $\log m$. For the purposes of Part I of the paper (the complexity study), it is essential that the length of the compression be at least subpolynomial in m in order to ensure that the reductions defined with respect to compressibility (see section 2.2) do compose. For clarity we choose a polynomial in $\log m$, although this may be replaced by any subpolynomial function $m'(\cdot)$ (i.e., $m' = m^{o(1)}$). We note that in the case of \mathcal{NP} problems one can assume that $n \geq \log m$, and then one can replace the polynomial $p(n, \log m)$ in Definition 1.2 by a polynomial in n alone (in sections 1.4 and 2.1 we compare this definition to the notion of *polynomial kernelization*). However, we choose not to restrict the scope of our discussion by making this assumption. Moreover, for Part II (the applications) Definition 1.2 may be significantly relaxed, where even a compression to length $m^{1-\varepsilon}$ for some constant ε suffices for some applications.

⁴Typically, the length n is part of the description of the problem (e.g., for Clique, SAT, Vertex cover, and others).

⁵Note that we can equivalently ask that the error be, say, $\varepsilon = \frac{1}{3}$. This is because the error can be reduced to negligible, albeit at the price of a worst compression rate (the polynomial $p(\cdot, \cdot)$ grows). See Claim 2.24.

The complexity of L' . In Definition 1.2 there is no restriction on the complexity of the language L' . All that is required is that there is enough information in $Z(x)$ to determine whether $x \in L$ or not. However, it is worth noting that if the compression is errorless, then the language L' must be in a class of nondeterministic-time $\text{poly}(m)$ that we denote $\mathcal{NP}(\text{poly}(m))$. That is, languages that are verifiable in time $\text{poly}(m)$ when given a nondeterministic hint (in order for $\text{poly}(m)$ to be well defined we assume that the parameter m is also encoded in the instance $Z(x)$). This fact follows simply from the definition of compression.⁶ In some cases it is essential to restrict L' to be in $\mathcal{NP}(\text{poly}(m))$, such as when defining the witness-retrievability property (Definition 6.1). Moreover, in some cases it is natural to further restrict L' to actually be in \mathcal{NP} (that is, in $\mathcal{NP}(\text{poly}(n, \log m))$). For instance, this is the case in the example for compression of SAT (Example 1.1).

Paper organization. In the rest of the introduction we survey the results of this paper, including Part I (the complexity study) and Part II (the cryptographic applications). In section 1.4 we discuss related and subsequent works. The main complexity study of the compressibility of \mathcal{NP} problems appears in section 2. The cryptographic applications are in sections 3.1, 5, and 6. In section 3.1 we describe the application of compression to constructing collision resistant hash functions (CRH) from any one-way function. Section 5 presents the implication to the hybrid bounded storage model (BSM), while section 6 discusses witness-retrievable compression and its application to the construction of oblivious transfer (OT) from any one-way function. We conclude with a discussion and some open problems (section 7).

1.2. Part I: Classifying \mathcal{NP} problems with respect to compression. In the first part of the paper, we develop the basic complexity theory of compression, including the following.

Examples of nontrivial compression. We demonstrate some nontrivial languages that do admit compression (section 2.1); however, these examples do not shed light on the general compression of other \mathcal{NP} problems. Moreover, it becomes clear that the traditional notions of reductions and completeness in \mathcal{NP} do not apply for the case of compression.

Reducibility and completeness. We define W-reductions, suitable for the study of compressibility, and define the corresponding notions of compression-complete and compression-hard languages for a class.

The \mathcal{VC} hierarchy. We introduce a classification of \mathcal{NP} problems with respect to compression. The classification presents a structured hierarchy of \mathcal{NP} problems that is quite different from the traditional view. We call this hierarchy \mathcal{VC} , short for “verification classes,” since the classification is closely related to the verification algorithm of \mathcal{NP} languages when allowed a preprocessing stage.

Compression of \mathcal{NP} search problems. For an \mathcal{NP} search problem, is there a compression algorithm that also maintains all the information about a witness for the problem? We show that the compression of a class of decision problems also implies compression for the corresponding class of search problems.

1.3. Part II: Implications to cryptography. In the second part of the paper, we show that compression has vast implications for cryptography. Specifically, various forms of compressibility for SAT (or even easier languages) would yield what follows.

⁶Suppose that there exists an errorless compression algorithm Z for L ; then define L' to be the language of all $Z(x)$ such that $x \in L$. Then, for every $y \in L'$ a verification algorithm takes as a nondeterministic witness a value x along with a witness to $x \in L$ and verifies that indeed $y = Z(x)$. Thus L' is in $\mathcal{NP}(\text{poly}(m))$.

One-way functions imply CRH. This construction requires compression that is *errorless*. This implication (as well as the next implication) is known to be impossible via “black-box” constructions, and thus our constructions necessarily make use of the code of the one-way function.

Hard-on-average problems in \mathcal{NP} imply one-way functions. More generally, given a language (not necessarily in \mathcal{NP}) that is hard on the average for nonuniform machines over a samplable distribution and a compression algorithm for a related language, one can construct a one-way function.

Impossibility of everlasting security via the “hybrid BSM.” This hybrid model mixes the standard “BSM” [62] with a (time) complexity-based assumption. This result shows that under compressibility of SAT, the hybrid model is no more powerful than the standard BSM.

OT from any one-way function. This requires a form of compression which preserves witnesses, which we show is unlikely to exist for general SAT instances.

1.4. Related work. The relationship between compression and complexity in general is a topic that has been investigated since the early days of complexity theory (i.e., Kolmogorov complexity [60]). However, the feature that we are studying in this work is compressibility with respect to the *solution* (witness) rather than the instance. This distinguishes our work from a line of seemingly related works about notions of compression (see [27, 78, 81] to name a few), all of which aim at eventually retrieving the input of the compression algorithm.

There are several examples of other relaxations of solving \mathcal{NP} problems in polynomial time. Each of these relaxations yields a corresponding classification of \mathcal{NP} . These classifications, like ours, are subtle and usually turn out to be different than the traditional \mathcal{NP} classification. For example, Papadimitriou and Yannakakis [75] introduce L-reductions and the classes MAX NP and MAX SNP, with respect to approximation algorithms. Impagliazzo, Paturi, and Zane [51] define reductions with respect to the solution in subexponential time.

The classification most related to ours is that of *parameterized complexity* (see the monographs on this subject by Downey and Fellows [24], Niedermeier [72], and Flum and Grohe [34]). Parameterized complexity studies the tractability of problems when one of the parameters is considered to be fixed or very small (this is called fixed parameter tractability). One of the basic techniques of acquiring efficient algorithms in this context is the method of “kernelization” that *may* yield natural compression algorithms (see examples in section 2.1). The kernelization method first shrinks the instance to a smaller instance whose size is only a function of the parameter and then solves it in brute force. However, in spite of the similarities between kernelization and compression, there are important differences. At a high level, kernelization is geared toward getting closer to a solution of the original instance. Our notion, on the other hand, requires compression per se, disregarding whether it is much harder to solve the compressed instance than the original one (in fact, in our main applications for constructing CRH and one-way functions in sections 3.1 and 4, the compressed instance never has to be solved). Indeed, we expect that new methods of compression that would resolve the problems we raise in this paper will utilize this property (that the compressed instance is harder to solve). That being said, a version of kernelization, namely, *polynomial kernelization*, is equivalent to deterministic compression to size $\text{poly}(n, \log m)$ if the witness size n is chosen to be the parameter. The question of polynomial kernelization has been raised independently from our work in the parameterized complexity community (e.g., see [34, Definition

9.1]). See a further discussion on kernelization in section 2.1. In addition, due to the above-mentioned similarities, the *Weft* hierarchy of parameterized complexity is reminiscent of the \mathcal{VC} -hierarchy: both being defined by reductions to circuits of bounded depth. However, as discussed above, our study of compression yields quite a different classification.

A related notion to parameterized complexity that is reminiscent of our work is *limited nondeterminism*, which started with the work of Kintala and Fischer [58]; see the survey by Goldsmith, Levy, and Mundheck [40]. This was further studied by Papadimitriou and Yannakakis [76], who defined a few syntactic classes within the class of polylog nondeterminism (*LOGNP* and *LOGSNP*). The interesting point is that several natural problems are complete for these classes. The notion of reduction used is the usual polynomial reduction, and hence the classifications arising from this study are very different from our \mathcal{VC} hierarchy. A related classification is the EW-hierarchy defined by Flum, Grohe, and Weyer [35]. This hierarchy is similar to the *Weft* classification of parameterized complexity but limits the running time to be only exponential in the witness length, thus being geared toward problems with polylogarithmic size parameters (as in *LOGNP*).

Subsequent works. Dubrov and Ishai [26] discussed the compression of problems and showed that a certain incompressibility assumption has an application to derandomization. Specifically they construct a pseudorandom generator that fools procedures that use more randomness than their output length. Their work was mostly conducted independently of ours, following conversations regarding an early phase of our work. In addition, inspired by our CRH construction, they prove that any one-way permutation can either be used for the above-mentioned derandomization or else can be used to construct a weak version of CRH.

Dziembowski [28] shows the relevance of our notion of witness-retrievable compression to achieving *forward-secure storage*. He shows a cryptanalytic result of such compression. Furthermore, following our approach for construction of OT from one-way functions, he shows that for every one-way function either a specific storage scheme is forward-secure, or there exists an (infinitely often) OT protocol based on this one-way function.

Recently some strong negative results about compression were shown. Fortnow and Santhanam [36] show that an errorless compression algorithm for SAT (or even for the class \mathcal{VC}_{OR}) entails the collapse of the polynomial hierarchy. Chen and Müller [14] notice that this generalizes to compression with a one-sided error. These results limit the application to constructing CRH functions (Theorem 3.1). The application may still be valid given a relaxed compression algorithm. For example, it suffices if the compression is successful only on instances that either have a witness to being satisfiable or have a witness to not being satisfiable. Note that the applications in sections 4 and 5 allow an error in the compression. Further results were found by Dell and van Melkebeek [19] who very recently showed hardness of compression of d -SAT beyond n^d and vertex cover beyond k^2 .

2. Part I: On the compression of \mathcal{NP} instances. Attempting to compress \mathcal{NP} instances requires a different approach than solving \mathcal{NP} problems. Intuitively, a solution for compression might arise while trying to solve the problem. While a full solution of an \mathcal{NP} problem may take exponential time, it is plausible that the first polynomial number of steps leaves us without an explicit solution but with a smaller instance. Indeed, some algorithms in the parameterized complexity world work like this (see some examples in the next section). On the other hand, we allow the possibil-

ity that the compressed version is actually harder to solve (computational timewise) than the original one (and may require a somewhat longer witness altogether).

2.1. Examples of compression algorithms for some hard problems. We start by showing several examples of compression algorithms for problems that are conjectured not to be in \mathcal{P} . Two of these example are \mathcal{NP} -complete problems, while the third is taken from cryptography.

Vertex cover. The well-studied \mathcal{NP} -complete problem of vertex cover receives as input a graph $G = (V, E)$ and asks whether there exists a subset of vertices $S \subseteq V$ of size at most k such that for every edge $(u, v) \in E$ either u or v are in S . The parameters are the instance length m , which is at most $O(|E| \log |V|)$, and the witness length $n = k \log |V|$.

CLAIM 2.1. *There exists a witness-retrievable compression algorithm for vertex cover.*

Proof. We are following the parameterized complexity algorithm for vertex cover (presented in [24] and attributed to Buss). If a vertex cover S of size k exists, then any vertex of degree greater than k must be inside the set S . The compression algorithm simply identifies all such vertices and lists them in the cover, while removing all their outgoing edges (that do not need to be covered by other vertices). The graph left after this process has maximal degree k , and furthermore all edges have at least one end in the cover. Thus, if the original graph has a k vertex cover, then the total number of edges left is at most k^2 (at most k vertices in the cover with at most k edges each). If there are more than k^2 edges, then the answer to the problem is NO; otherwise, such a graph can be represented by the list of all edges, which takes $k^2 \log k$ bits. The compression can be made witness-retrievable, since if we use the original labels of vertices to store the new graph, then the original cover is also a cover for the new compressed graph. \square

It is in fact possible to get the compressed instance to be a graph with $2k$ nodes, rather than k^2 nodes, as shown in [13] and [15] (see [72, Chapter 7]). It is interesting to note that we do not know of a compression algorithm for the Clique problem or the dominating set problem, which are strongly linked to the vertex cover problem in the traditional study of \mathcal{NP} , and, in fact, in Theorems 3.1, 5.2, and 6.2 we show strong implications of a compression algorithm for these languages.

On parameterized complexity and compression. The use of an algorithm from parameterized complexity for compression is not a coincidence, due to some similarities between the issues. In a parameterized problem, the instance is coupled with a parameter which is independent of the instance size (and typically considered to be much smaller). Recall that the objective is to solve such an instance in time that is an arbitrary function of the parameter, yet only polynomial in the instance length. The “problem kernel” method (see [24, Chapter 3] or [72, Chapter 7]) presents a two phase scheme for solving parameterized problems. Loosely speaking, the first phase, called “kernelization,” reduces the problem to a subinstance that is (i) small (a function of the parameter alone); and (ii) like compression, contains the answer to the original problem. Then the second phase solves the subinstance by brute force, running in time that is a function only of the subinstance length, e.g., exponential in this small instance size. Downey, Fellows, and Stege [25] (Lemma 4.7) show that kernelization precisely captures the problems that are fixed-parameter tractable. A *polynomial kernelization* (e.g., [34, Definition 9.1]) is a restriction of kernelization in which the output size is bounded by a polynomial in the parameter. In cases where the parameter is set to be the witness size, polynomial kernelization is equivalent to *deterministic*

compression to size $\text{poly}(n, \log m)$. (Recall that we may assume that $\log m \leq n$, since otherwise the instance can be solved by brute force in polynomial time.) Thus some kernelization algorithms yield compression algorithms and vice versa.

We note that taking the parameter to equal the witness size n is not very natural in the context of parameterized complexity, as the resulting parameterized problem is trivially fixed-parameter tractable by brute force search. Nevertheless, the existence of polynomial kernelizations for such problems has been considered in the parameterized complexity literature (independent of our work); see, e.g., [34].

Chapter 7 of the monograph of Niedermeier [72] contains several examples of polynomial size kernelizations (e.g., for the languages 3-hitting set and dominating set on planar graphs). These algorithms yield compression algorithms for the respective languages. We describe one additional example of a compression algorithm that is derived in this manner.

Minimum fill-in. The minimum fill-in problem is an \mathcal{NP} -hard problem that takes as input a graph G and a parameter k and asks whether there exist at most k edges that can be added to the graph that would turn it into a chordal graph, i.e., one with no induced cycles of length more than 3. This problem has applications in ordering a Gaussian elimination of a matrix.

CLAIM 2.2. *The minimum fill-in problem with parameter k has witness-retrievable compression.*

Proof. Kaplan, Shamir, and Tarjan [54] prove that this problem is fixed-parameter tractable. Their algorithm partitions the graph into two sets of nodes A and B , where A is of size k^3 and there are no chordless cycles (i.e., an induced cycle of length greater than three) in G that contain vertices in B . The complexity of this partition is $O(k^2|V||E|)$. They then prove that G has a k edge fill-in if and only if the graph induced by A has a k edge fill-in.

Thus the compression algorithm follows the same partitioning and stores only the graph induced by the small set A . The new graph has at most k^3 vertices and thus can be presented by only $\text{poly}(k) \log |k|$ bits. The fill-in for the new instance is exactly that of the original instance, and thus the compression can be witness-retrievable if the original labels of the vertices are used for the compressed graph as well. \square

2.1.1. Sparse languages. We call a language *sparse* if the language contains only of a small fraction of the words of any given length.

DEFINITION 2.3 (sparse language). *Let L be an \mathcal{NP} language with instance length m and parameter n , and denote $L_{m,n} = \{x \in \{0, 1\}^m \mid x \in L \text{ with witness of length } \leq n\}$; then L is sparse if there exists a polynomial $p(\cdot)$ such that for all sufficiently large m (with corresponding n) it holds that $|L_{m,n}| \leq 2^{p(n)}$.*

We show that all sparse languages can be compressed to a size that is dominated by the number of words that are actually in the language. This is shown by a generic compression algorithm for any sparse language. The idea is to apply a random (pairwise independent) hash function to the instance where the output of the hash is of length $2p(n)$ and thus substantially smaller than m . The new language contains all words that are hashed values of a word in the original language. We note that the compressed language L' lies in $\mathcal{NP}(\text{poly}(m))$ (recall that $\mathcal{NP}(\text{poly}(m))$ stands for nondeterministic-time $\text{poly}(m)$). In particular, it is not necessarily witness-retrievable.

Rather than formally presenting the method for a general sparse language, we describe the method via a sparse language that we call PRG-output. Note that for this language the method is witness-retrievable.

Example 2.4 (PRG-output). Let G be a pseudorandom generator (PRG) stretching an n bit seed to an m bit output (with m an arbitrary polynomial in n). Define

the language *PRG-output* over inputs $y \in \{0, 1\}^m$ as

$$L_G = \{y \mid \text{there exists an } x \text{ such that (s.t.) } G(x) = y\}.$$

As long as the underlying PRG is secure, then it is hard to decide whether an instance was taken randomly from $L(G)$ or from $\{0, 1\}^m$. Yet this language has a simple compression algorithm. Note that simply saving, say, the first $2n$ bits of the instance y is insufficient because if y differs only from $G(x)$ in one bit, then this bit may be anywhere in the m bits.

CLAIM 2.5. *There exists a witness-retrievable compression algorithm for PRG-output.*

Proof. Let \mathcal{H} be a family of almost pairwise independent hash functions from m bits to $2n$ bits. The compression algorithm simply chooses a random $h \in \mathcal{H}$ and outputs $(h(y), h)$. The new language is $L'_G = \{(z, h) \mid \text{there exists an } x \text{ s.t. } h(G(x)) = z\}$.

Naturally, if $y \in L_G$, then also $(h(y), h) \in L'_G$ with the same witness (and thus the witness-retrievability). On the other hand, if $y \notin L_G$, then by the properties of \mathcal{H} , for every seed x we have that $\Pr_h[h(G(x)) = h(y)] < O(2^{-2n})$, and by a union bound over all $x \in \{0, 1\}^n$, we get $\Pr_h[h(y) \in L'_G] < O(2^{-n})$. Finally, since there are almost pairwise independent hash functions whose description is of length $O(n) + \log m$ (for example, see [66]), then the algorithm is indeed compressing. Note that the compression algorithm described above is probabilistic and carries an error probability of $2^{-\Omega(n)}$ and also that the compressed language L' in this case is in $\mathcal{NP}(\text{poly}(m))$. \square

Sparse subset sum. We show another example of a compressible language called sparse subset sum that is sparse in a different sense than that of Definition 2.3. While the generic compression for sparse languages does not work for this language, it is compressible in its own right. Moreover, the compression algorithm for sparse subset sum is better than the generic algorithm in the sense that the compressed language in the specialized algorithm is in $\mathcal{NP}(\text{poly}(n, \log m))$ (or actually in \mathcal{NP}) rather than in $\mathcal{NP}(\text{poly}(m))$.

Example 2.6 (sparse subset sum). The language *sparse subset sum* takes as input n values x_1, \dots, x_n each in $\{0, 1\}^m$ (with $m \gg n$) and a target value $T \in \{0, 1\}^m$. An input is in the language if there is a subset $S \subseteq [n]$, where $\sum_{i \in S} x_i = T$ (the sum is taken modulo 2^m).

CLAIM 2.7. *There exists a witness-retrievable compression algorithm for sparse subset sum.*

Proof. To compress an instance of sparse subset sum, simply pick a large random prime $2^n < P < 2^{2n + \log m}$ and store the numbers $y_i = x_i \pmod P$ (for every $i \in [n]$), the target $T_P = T \pmod P$ and P (the idea of picking a prime P and working modulo P has been useful various applications, e.g., in the Karp–Rabin string matching algorithm [56]). The compressed instance is of length $O(n(n + \log m))$, and the compressed language is also subset sum (modulo P). If there exists a set S for which $\sum_{i \in S} x_i = T$, then also $\sum_{i \in S} y_i = T_P \pmod P$ (hence the witness-retrievability). On the other hand, we want that if the original instance was not in the language, then for any subset S it will hold that $\sum_{i \in S} y_i \neq T_P$. In order to get $\sum_{i \in S} y_i = T_P$, it is required that P is a divisor of $D = \sum_{i \in S} x_i - T$. However, D has at most m/n prime divisors that are greater than 2^n , while the prime P is taken from a range containing $O(2^{2n} m/n)$ primes (we assume $n \geq \log m$ in the calculations). Therefore, for every S it holds that $\Pr_P[\sum_{i \in S} y_i = T_P] \leq 2^{-2n}$, and by a union bound over all sets S , the probability of an error is bounded by 2^{-n} . \square

2.2. W-reductions and compression-completeness. The few examples of compression that we have showed clearly indicate that the study of \mathcal{NP} problems with respect to compression gives a distinct perspective, different from the traditional study of \mathcal{NP} . The reason is that the typical Karp-reduction between \mathcal{NP} problems does not distinguish between the length of the witness and the length of the instance. For example, when reducing SAT to the Clique problem, one builds a large graph from a CNF formula and asks whether or not it has a Clique of size k . However, in this new instance, the witness size⁷ is a polynomial in m (the length of the SAT formula) rather than n (the number of variables in the formula). Thus, it is not clear how to use a compression algorithm for Clique to get a compression algorithm for SAT.

W-reductions and compression-completeness. In order to show that a compression algorithm for L' implies a compression algorithm for L , a more restricted type of reduction is needed. We call this a *W-reduction*, and it is similar to a Karp-reduction but imposes an extra property on the length of the witness.

DEFINITION 2.8 (W-reduction). *For two \mathcal{NP} languages L and L' we say that L W-reduces to L' if there exist polynomials p_1 and p_2 and a polynomial-time computable function f that takes an instance x for L and outputs an instance $f(x)$ for L' such that*

1. $f(x) \in L'$ if and only if $x \in L$.
2. if x is of length m with witness length n , then $f(x)$ is of length at most $p_1(m)$ with witness length at most $p_2(n, \log m)$.

We first note that this reduction composes, namely what follows.

CLAIM 2.9. *If L W-reduces to L' and L' W-reduces to L'' , then L W-reduces to L'' .*

We next claim that W-reduction indeed fulfills its goal with respect to compression.

CLAIM 2.10. *Let L and L' be \mathcal{NP} languages such that L' W-reduces to L . Then given a compression algorithm for L , one can obtain a compression algorithm for L' .*

Proof. Suppose that x is an instance for language L' of length m with witness length n . The compression algorithm for L' runs as follows: First use the W-reduction to L and get an instance $f(x)$ for L , and then run the compression algorithm for L on $f(x)$. By the properties of the reduction, $f(x)$ is of length $m' \leq p_1(n, m)$ with witness length $n' \leq p_2(n, \log m)$. The outcome $Z(f(x))$ of the compression is therefore of length $\text{poly}(n', \log m') = \text{poly}(n, \log m)$. Furthermore, if L'' is the language that Z compresses to, then $Z(f(x)) \in L''$ if and only if $f(x) \in L$, which in turn happens if and only if $x \in L'$. Thus the combined process gives a compression algorithm for instances of L' . \square

We remark that in the complexity discussion of compression we choose to ignore the issue of witness-retrievability. Nevertheless, in order for the W-reduction to relay this property, the reduction itself must also have a witness-retrievability property. That is, given a witness w for $x \in L$, one can efficiently compute w' for $f(x) \in L'$ (without the knowledge of x). We define complete problems with respect to compression: these are defined similarly to the standard notion but with respect to W-reductions.

DEFINITION 2.11 (compression-complete). *A problem L is compression-complete for class \mathcal{C} if*

1. $L \in \mathcal{C}$;
2. for every $L' \in \mathcal{C}$ the language L' W-reduces to L .

A language is called compression-hard for class \mathcal{C} if requirement 2 holds (requirement 1 may or may not hold).

⁷The witness for Clique is a choice of k vertices from the graph.

The relevance of compression-complete problems is stated in the following simple claim.

CLAIM 2.12. *Let L be compression-complete for class \mathcal{C} ; then given a compression algorithm for L , one can obtain a compression algorithm for any $L' \in \mathcal{C}$.*

The proof follows directly from the definition of completeness and Claim 2.10.

2.3. The \mathcal{VC} classification. We now introduce the new classification arising from the study of compressibility of \mathcal{NP} problems. For this we define a series of \mathcal{NP} languages. Some notation: by a **circuit of depth k** , we mean a depth k alternating AND-OR circuit where the fan-in of the gates is bounded only by the size of the circuit and negations are only on the input variables (no NOT gates).

DEFINITION 2.13 (Depth $_k$ CircuitSAT). *For any $k \geq 2$ consider the \mathcal{NP} problem called Depth $_k$ CircuitSAT.*

Input: a circuit C of size m and depth at most k over n variables.

Membership: $C \in \text{Depth}_k\text{CircuitSAT}$ if there exists a satisfying assignment to C .

The next language, LocalCircuitSAT, is a less natural one. It is designed to capture computations that do not need to access the whole input, but can rather check only a sublinear fraction of the input (a good example is verifying that a set of vertices in a graph is indeed a Clique). Let x be a string of length m . If $I = (i_1, \dots, i_n)$ is a list of n locations in x , then we denote by $x(I)$ the values of x at these locations.

DEFINITION 2.14 (LocalCircuitSAT).

Input: A string x of length m and a circuit C over $n + n \cdot \log m$ variables and of size $(n + n \cdot \log m)$.⁸

Membership: $(x, C) \in \text{LocalCircuitSAT}$ if there exists a list I of n locations in x such that $C(x(I), I) = 1$.

We can now introduce our classification of \mathcal{NP} problems.

DEFINITION 2.15 (The \mathcal{VC} classification of \mathcal{NP} problems). *Consider \mathcal{NP} problems, where m denotes the instance size and n denotes the witness size. We define the class \mathcal{VC}_k for every $k \geq 0$. The definition is divided into three cases:*

- $k = 0$: *The class \mathcal{VC}_0 is the class of all languages that admit compression algorithms. There are two possible versions here, one considering errorless compression and the other allowing probabilistic compression with errors. We typically refer to the later, depending on the context.*
- $k = 1$: *The class \mathcal{VC}_1 is the class of all languages that W-reduce to LocalCircuitSAT.*
- $k \geq 2$: *The class \mathcal{VC}_k is the class of all languages that W-reduce to Depth $_k$ CircuitSAT.*

For any function $k(m, n)$ (where $k(m, n) \leq m$) also define $\mathcal{VC}_{k(m, n)}$ as the class of all languages that W-reduce to Depth $_{k(m, n)}$ CircuitSAT. Finally, define $\mathcal{VC} = \mathcal{VC}_m$ (the class for $k(m, n) = m$).

A first observation is that simply by definition, the languages LocalCircuitSAT and Depth $_k$ CircuitSAT are compression-complete for their respective classes. The most notable example of a complete language is for the class $\mathcal{VC} = \mathcal{NP}$, where the complete problem is CircuitSAT (satisfiability of a polynomial size circuit).

When discussing a W-reduction to a depth k circuit, we can actually assume without loss of generality that the top gate of this circuit is an AND gate (as we will show in the next claim). An immediate corollary is that SAT (that is, satisfiability of CNF formulas) is compression-complete for the class \mathcal{VC}_2 . Formally, let

⁸The choice of the circuit to be of size n' (over n' variables) is arbitrary, and other polynomial functions suffice as well. Furthermore, such a circuit of small size may be meaningful, since not all the variables have to be used and some might be just dummy variables.

$\text{Depth}_k\text{CircuitSAT}_{\text{AND}}$ denote the language $\text{Depth}_k\text{CircuitSAT}$ when restricted to circuits where the top gate is an AND gate.

CLAIM 2.16. *For any $k \geq 2$, we have that a language $L \in \mathcal{VC}_k$ if and only if L W-reduces to the language $\text{Depth}_k\text{CircuitSAT}_{\text{AND}}$.*

Proof. We show that any instance that contains a circuit where the top gate is an OR W-reduces to an instance with top gate AND. We prove this first for $k \geq 3$. Denote the input circuit by $C = \bigvee_j \bigwedge_t C_{j,t}$, where each $C_{j,t}$ is a top OR depth $(k - 2)$ circuit. If C is satisfiable, then $\bigwedge_t C_{j,t}$ is satisfiable for at least one choice of j . Add to the witness the index i of this satisfiable subcircuit (i is given by the boolean variables i_1, \dots, i_ℓ , where ℓ is logarithmic in $\text{poly}(m, n)$). For each j , denote $C'_{j,t} = C_{j,t} \vee i_1^{\bar{j}} \vee \dots \vee i_\ell^{\bar{j}}$, where $i^{\bar{j}}$ denotes $i \oplus j$. Notice that $C'_{j,t}$ is always satisfied for $j \neq i$, and for $j = i$ it is satisfied if and only if $C_{i,t}$ is satisfied. Thus the circuit can now be written as $C' = \bigwedge_{j,t} C'_{j,t}$ that is satisfiable if and only if the original circuit was. The top OR gate of C is therefore removed in the new instance C' while adding only a small number of variables; thus the input to the circuit witness remains of order $\text{poly}(n, \log m)$ as required.

In the case $k \geq 3$, the depth of the new instance becomes $k - 1$. In the case where $k = 2$, the bottom level that included only variables is transformed into an OR of variables; thus the new circuit is simply a CNF formula (and the depth remains $k = 2$). \square

The VC hierarchy. The \mathcal{VC} classification indeed defines a hierarchical structure. That is,

$$\mathcal{VC}_0 \subseteq \mathcal{VC}_1 \subseteq \mathcal{VC}_2 \subseteq \mathcal{VC}_3 \cdots \subseteq \mathcal{VC}.$$

And, in general, for every two polynomially bounded functions $k(n, m), \ell(n, m)$ such that for all n, m we have $k(n, m) \leq \ell(n, m)$, then $\mathcal{VC}_k(m, n) \subseteq \mathcal{VC}_\ell(m, n)$. Furthermore, $\mathcal{VC} = \mathcal{NP}$ by the definition of \mathcal{NP} . These observations follow trivially by the definitions, the only nontrivial part being the fact that $\mathcal{VC}_1 \subseteq \mathcal{VC}_2$, that is proved next.

LEMMA 2.17. $\mathcal{VC}_1 \subseteq \mathcal{VC}_2$.

Proof. We need to show a W-reduction from LocalCircuitSAT to SAT . The input is therefore a long string x and small circuit C on $n + n \log m$ variables. Let i_1, \dots, i_n denote the potential locations in the string that the circuit C receives as inputs. Also define the variables y_1, \dots, y_n to indicate the values of x in the corresponding locations (that is, $y_t = x_{i_t}$ for $t \in [n]$). Thus the circuit C runs on the variables y_1, \dots, y_n and the bits of i_1, \dots, i_n .

We first note that C is of size $p(n, \log m) = (n + n \log m)$ and may be reduced (via Cook's theorem [16]) to a CNF formula Φ_C over $O(p(n, \log m))$ variables and of size $O(p(n, \log m))$ that is satisfiable if and only if C is satisfiable.

Thus we have a CNF formula over the variables $y_1, \dots, y_n, i_1, \dots, i_n$, and some extra variables. This formula's satisfiability will be equivalent to the membership of the LocalCircuitSAT instance if we manage to force the variables of y to take the values $y_t = x_{i_t}$. This is done by adding additional clauses to the CNF in the following manner: For simplicity we describe this only for y_1 , where the same is repeated for every other y_t for $t \in [n]$. Define for each $j \in [m]$ a formula $\Phi_j = (y_1 = x_j) \vee (i_1 \neq j)$. Notice that $\Phi_{i_1} = 1$ if and only if $y_1 = x_{i_1}$. Denote the bits of i_1 by $i_{1,1}, \dots, i_{1,d}$, where $d = \lceil \log m \rceil$. An alternative way to write Φ_j is as the following CNF (recall that $i^{\bar{j}}$ denotes $i \oplus j$):

$$\Phi_j = \left(y_1 \vee \overline{x_j} \vee i_{1,1}^{\bar{j}} \vee \dots \vee i_{1,d}^{\bar{j}} \right) \wedge \left(\overline{y_1} \vee x_j \vee i_{1,1}^{\bar{j}} \vee \dots \vee i_{1,d}^{\bar{j}} \right).$$

Finally, to force $y_1 = x_{i_1}$ we simply take the new CNF to be $\Phi_C \wedge \bigwedge_{j \in [m]} \Phi_j$. The same is repeated to force $y_t = x_{i_t}$ for all $t \in [n]$. \square

2.4. The \mathcal{VC} classification and verification with preprocessing. We now discuss the \mathcal{VC} hierarchy from a different angle, that of the verification complexity of a language. This approach, though slightly more cumbersome than the definition via W-reductions, gives more intuition as to what it means to be in a class \mathcal{VC}_k . The new view defines the \mathcal{VC} hierarchy with respect to the verification algorithm for L , that is, the efficient procedure that takes a witness w for $x \in L$ and verifies that it is indeed a true witness. We point out that the nature of verification algorithms may vary when discussing different \mathcal{NP} problems. For example, in the k -Clique problem the verification algorithm needs to check only $O(k^2)$ edges in the graph and thus can read only a sublinear part of the instance. In SAT, on the other hand, all the clauses in the formula must be checked when verifying a witness.

Simply looking at the verification algorithm of a language is not sufficient. For starters, classification according to verification does not distinguish between problems in \mathcal{P} that are trivially compressible and \mathcal{NP} -complete languages. Instead, we consider the notion of verification with preprocessing. This is the process for verifying that $x \in L$ when given a witness, that also allows a preprocessing stage to the instance.

DEFINITION 2.18 (verification with preprocessing). *Let L be an \mathcal{NP} language with instances of length m and witness length n . A pair of polynomial-time algorithms (P, V) are called a verification with preprocessing for L if the following two-step verification holds:*

1. Preprocessing: P gets an instance x and outputs a new instance $P(x)$.
2. Verification: There exists a polynomial $p(\cdot, \cdot)$ such that $x \in L$ if and only if there exists a witness w of length at most $p(n, \log m)$ such that $V(P(x), w) = 1$.

Notice that when allowing for preprocessing, all problems in \mathcal{P} have a pair (P, V) , where P solves the problem and stores the answer while V simply relays this answer. Thus when considering the complexity of V in this definition, easy problems indeed have very low complexity.

The \mathcal{VC} classification via verification with preprocessing. An alternative and equivalent way to view the classes in the \mathcal{VC} hierarchy is based on the verification algorithm V in a verification with preprocessing pair (P, V) . The problems are divided into two families:

- The class \mathcal{VC}_1 is the set of the languages that have very efficient verification (i.e., $poly(n, \log m)$ rather than $poly(n, m)$). We assume random access to the instance (suppose that the verification algorithm is a RAM); thus such a verification algorithm accesses only a sublinear fraction of the instance.
- The languages whose verification is not very efficient (run in time $poly(n, m)$). This family is further classified into subcategories. The class \mathcal{VC}_k is the class of languages where the verification algorithm V has a representation as a depth k polynomial size circuit (polynomial in n and m).

This definition is equivalent to the definition via W-reductions, since the W-reduction to the complete problem can simply be viewed as the preprocessing stage. In the other direction, every preprocessing stage is actually a W-reduction to the language defined by V .

It is interesting to note that Buss and Islam [8] give an alternative view with similar flavor to the *Weft* hierarchy of parameterized complexity. They call it “prepare,

guess and check” in which they essentially add a preprocessing phase to a previous approach of Cai and Chen [10].

2.5. Within \mathcal{VC}_1 —the class \mathcal{VC}_{OR} . Arguably, the most interesting class in the hierarchy is the bottom class \mathcal{VC}_1 . It contains many natural problems, such as Clique or small subset-sum,⁹ that test only local properties of the input. Furthermore, it is presumably the easiest class to find compression algorithms for. We further refine our hierarchy within the class \mathcal{VC}_1 by introducing another class, the class \mathcal{VC}_{OR} . Consider the language $OR(L)$ that takes a large OR of small instances of a language L .

DEFINITION 2.19 ($OR(L)$). *Let L be an \mathcal{NP} language. Define the language $OR(L)$ as follows:*

Input: m instances x_1, \dots, x_m to the language L , each of length n .

Membership: $(x_1, \dots, x_m) \in OR(L)$ if there exists $i \in [m]$ such that $x_i \in L$.

Specifically the language $OR(\text{CircuitSAT})$ is defined as follows:

Input: m circuits C_1, \dots, C_m where each circuit is of size n .

Membership: $(C_1, \dots, C_m) \in OR(\text{CircuitSAT})$ if at least one of the m circuits is satisfiable.

This language is used to define the following class.

DEFINITION 2.20. *The class \mathcal{VC}_{OR} is the class of all languages that W-reduce to $OR(\text{CircuitSAT})$.*

We first note that in each of the m small instances, the instance length and witness length are polynomially related. So unlike the general case where we focused only on short witness languages, when talking about $OR(L)$, any language $L \in \mathcal{NP} \setminus \mathcal{P}$ is interesting. For example, the language $OR(3 - \text{SAT})$ is not trivially compressible. Moreover, it is compression-complete for \mathcal{VC}_{OR} .

CLAIM 2.21. *Let L be any \mathcal{NP} -complete language; then $OR(L)$ is compression-complete for \mathcal{VC}_{OR} .*

Proof. The W-reduction from $OR(\text{CircuitSAT})$ to $OR(L)$ simply runs the standard Karp reduction from CircuitSAT to L for each of the m circuits independently. The witness for each circuit was of length at most n and is now of size $p(n)$ for some polynomial p . In addition the witness contains an index of the instance of L that is satisfied; thus the total witness length is $p(n) + \log m$. \square

For example, the problem $OR(\text{Clique})$ that gets m small graphs (over n vertices) and asks whether at least one of the graphs has k sized clique (where $k = O(n)$) is also compression-complete for \mathcal{VC}_{OR} .

CLAIM 2.22. $\mathcal{VC}_{OR} \subseteq \mathcal{VC}_1$.

Proof. This is best seen by W-reducing $OR(\text{Clique})$ to LocalCircuitSAT. Given graphs G_1, \dots, G_m for $OR(\text{Clique})$, generate the instance $x = G_1, \dots, G_m$ and a circuit C that receives the locations of a clique in one of the graphs and checks whether there are indeed edges in these locations and whether they form a clique (all belong to the same graph and are the edges induced by k vertices). The size of the circuit is $p(n, \log m)$ for some polynomial p , since it checks only locations in x that belong to one graph (of size n). Finally, add $p(n, \log m)$ dummy variables to the circuit so that the circuit C has size equal to the number of input variables (this is a technical requirement in the definition of LocalCircuitSAT). \square

Furthermore, $\mathcal{VC}_0 \subseteq \mathcal{VC}_{OR}$, since any compressible language can be W-reduced by the compression algorithm to a language with instance size $p(n, \log m)$, and this

⁹This problem takes m values and a target value and asks if there is a small (size n) subset of the values that adds up to the target.

instance can be reduced to CircuitSAT and viewed as an OR of a single small circuit and hence is in \mathcal{VC}_{OR} . Note that here too, one may need to add dummy variables to make the circuit of the same size as its input. Altogether we have that

$$\mathcal{VC}_0 \subseteq \mathcal{VC}_{OR} \subseteq \mathcal{VC}_1.$$

Finally, we show a language that is compression-hard for \mathcal{VC}_{OR} . This claim is also relevant to our cryptographic applications (in sections 3.1, 4, 5, and 6).

CLAIM 2.23. *Clique is compression-hard for \mathcal{VC}_{OR} .*

Proof. The language $OR(Clique)$ W-reduces to Clique simply by taking one graph that is the union of all the small graphs in the $OR(Clique)$ instance. Clearly there is a clique in the union if and only if there is a clique in at least one subgraph. \square

A similar claim is true for all problems involving searching for a connected subgraph of size n in a graph of size m as long as the problem of deciding whether a graph of size $p(n)$ contains such a subgraph is NP-hard for some polynomial $p(\cdot)$. This is true, for instance, for the problem of whether there is a path of length n .¹⁰

2.6. The \mathcal{VC} classification and some \mathcal{NP} problems. In general, most of the \mathcal{VC} classification focuses on W-reductions to depth k circuits. The reasoning for this is that there is a certain tradeoff between depth and the number of variables. More precisely, we can reduce the depth of a verification circuit, but only at the price of adding additional variables (this is done using methods from Cook’s theorem [16] and requires adding a variable for each gate in one intermediate level of the circuit). Since the number of variables is the focal point when discussing compression (as it coincides with the witness size), then depth turns out to be central in our classification.

Given our current state of knowledge, there are a few plausible views of the world. The two *end point* scenarios are (i) there is compression for every language in \mathcal{NP} (as would be implied by a compression algorithm for CircuitSAT), (ii) there is only compression for a few select problems, such as the examples in section 2.1. A third option is that there is a compression algorithm for some compression-complete problem in the hierarchy (say, for \mathcal{VC}_k), which would imply the collapse of all the classes below \mathcal{VC}_k to \mathcal{VC}_0 .

We will briefly go over a few key classes in the hierarchy and a few examples of natural \mathcal{NP} problems and their classification (as we know it) within the \mathcal{VC} hierarchy. We note that all the statements in this section apply also to compression with possible error (negligible in n).

The class \mathcal{VC}_0 . Currently we know that this class contains all the languages in \mathcal{P} , languages that are already compressed by definition (such as 3-SAT), and the languages that we showed compression algorithms to (vertex cover, PRG-output, and minimum-fill-in).

The class \mathcal{VC}_{OR} . This class contains all languages $OR(L)$ for an \mathcal{NP} language L . One natural example is the $OR(SAT)$ problem, which is actually a depth 3 circuit where the fan-in at the two bottom levels is bounded by n and only the top gate is allowed to be of greater fan-in. Some important languages in this class are those that need to be compressed in the cryptographic applications in sections 3.1, 5, and 6.

The class \mathcal{VC}_1 . Since we are only interested in problems where the witness size n is much smaller than the instance size m , then many natural problems with this

¹⁰It is interesting to note that whereas the problem of finding a path of length n is fixed parameter tractable [2], Feige and Kilian [32] give indications that the Clique problem is hard for small n (via subexponential simulations). This illustrates that such differences in parameterized complexity are not necessarily preserved in the classification of compression.

restriction are in \mathcal{VC}_1 . For example, graph problems that ask whether a small graph can be embedded in a large graph are all in \mathcal{VC}_1 . The Clique problem (with a clique of size n) or long-path (a path of length n that does not hit any vertex twice) are such small graph embedding problems. Small subset-sum is another natural language in \mathcal{VC}_1 . This language receives a set of m values and a target sum and asks whether there is a small (size n) subset for which the values add up exactly to the target sum (see also footnote in section 2.5).

Dominating Set. The problem asks, given a graph, whether there is a set of k vertices such that all the graph is in its neighbor set. A dominating set is in the class \mathcal{VC}_3 , as implied by the following verification: the witness is a set S , and the algorithm tests that \forall vertex $v \exists$ vertex $u \in S$ such that (u, v) is in the graph. The \forall translates to an AND gate and the \exists translates to an OR gate. Finally, testing that an edge is in the graph requires an AND over the $O(\log m)$ bits representing this edge. In total, this is a depth 3 circuit. Note that a straightforward verification of vertex cover will also yield a depth 3 circuit. However, while vertex cover is compressible and in \mathcal{VC}_0 , for a dominating set we are unaware of a better method. In addition, a dominating set is *compression-hard* for \mathcal{VC}_2 . This is seen by a standard reduction of SAT to a dominating set in which a SAT formula with n variables and m clauses is transformed into a graph with $m + 3n$ vertices with the property that the graph has a dominating set of size n if and only if the SAT formula is satisfiable.¹¹

Weighted-SAT. Given a CNF formula of length m , the problem asks if it has a satisfying assignment of weight at most k (at most k variables are assigned the value 1). Unlike our previous discussions of SAT, the number of variables here is only bounded by m , and the short witness simply consists of the list of all variables that receive the value 1 (that is, the witness is of length $n = k \log m$). This problem, with constant clause size, serves as the basic complete problem for the parameterized complexity class $W[2]$, which is at the bottom of the W-hierarchy (see [24]). However, with regard to compressibility, we only know how to place it in the class \mathcal{VC}_4 . This is shown by the following verification procedure (using the same logic as with dominating-set): For every witness (list) L , the algorithm tests that \forall clauses C either \exists a variable $x \in C$ such that $x \in L$ or \exists a negated variable $\bar{x} \in C$ such that $x \notin L$. The verification of $x \in L$ adds up to total depth 3 by testing that $\exists y \in L$ such that $x = y$ (where $x = y$ is tested by an AND over the bits of x and y). On the other hand, verifying that $x \notin L$ requires total depth 4 as it runs $\forall y \in L$, we have $x \neq y$. The overall depth is thus dominated by the negated variables and is thus 4.

OR of (large) instances. Consider the OR of CNF formulas over few variables (each CNF formula may be large, unlike in the language $OR(SAT)$, where the CNF formulas are considerably smaller than the fan-in of the OR gate). In other words, instances of this language are depth 3 circuits, where the top gate is an OR gate. Yet the language is actually in \mathcal{VC}_2 , as implied by Claim 2.16.

Integer programming (IP). An instance of IP consists of a list of m linear constraints on n integer variables with the goal of maximizing a linear target function over these n variables (under the list of constraints). Unlike its counterpart of linear programming, where the variables may take real values and is polynomial-time solvable, IP is \mathcal{NP} -hard even when the variables are restricted to taking only the values 0

¹¹In a nutshell, the reduction creates a triangle for each variable x_i of the formula. One of the nodes of the triangle is identified with the positive variable and another with its negation, while the third is connected only to the other two. In addition, a vertex is created for each clause in the formula. Now, each literal is connected with all of the clauses that it appears in. The generated graph has a dominating set of size n if and only if the formula is satisfiable.

and 1 (one of Karp’s original problems [55]). Thus, the decision variant of IP, where the number of constraints is much larger than the number of variables, is interesting with respect to compression. First, compressing it is at least as hard as compressing SAT: given an SAT instance with n variables and m constraints, it is simple to come up with a corresponding IP instance with $2n$ variables and m constraints, i.e., IP is \mathcal{VC}_2 -hard. On the other hand, a straightforward verification of a witness for this problem takes the proposed assignment for the n variables and checks if it satisfies each of the constraints. The verification of a linear constraint can be achieved in logarithmic depth (in n), placing IP in $\mathcal{VC}_k(n)$ for $k(n) = \Omega(\log n)$. We are unaware of a (significantly) better classification (of lower depth) for IP.

2.7. On reducing the error in compression algorithms. The error of a compression algorithm can be reduced substantially at the expense of a worse compression rate (the output length of the compression algorithm will be longer). The idea is simply to run and store the outcome of many executions of the compression, each time with a fresh and independent randomness. For example, by storing n independent executions and using a Chernoff bound we arrive at the following claim.

CLAIM 2.24. *Let Z be a compression algorithm for language L with outcome length $p(n, \log m)$ and $q, \delta > 0$ be such that (i) if $x \in L$, then $Z(x) \in L'$ with probability q , and (ii) if $x \notin L$, then $Z(x) \notin L'$ with probability $q + \delta$. Then there is a compression algorithm Z' with error $2^{-\Omega(\delta^2 n)}$ and outcome length $np(n, \log m)$.*

Note that this technique is limited by the growth of the output, and, in particular, one cannot use this method to achieve an error that is exponentially small in m (rather than n).

2.8. On compression of search problems. So far, the \mathcal{NP} problems that we discussed were all decision problems, that is, they ask if $x \in L$ and are answered by YES or NO. When considering a specific \mathcal{NP} relation R_L associated with L , the above decision problem has a natural search problem associated with it, which is to actually find a witness to $x \in L$ with respect to the relation R_L . A solution to such a problem is an n bit string rather than just a single bit.

Loosely speaking, a compression algorithm for the search instance should produce a shorter output that contains enough information about some witness for the original problem.

DEFINITION 2.25 (compression for search problem). *A compression algorithm for an \mathcal{NP} search problem L (with respect to R_L) is a pair of algorithms (Z, E) with a polynomial $p(\cdot, \cdot)$, where Z is a polynomial-time compression algorithm and E is an unbounded extraction algorithm. Given an instance x with witness parameter n , we should have that*

1. $Z(x)$ is of length at most $p(n, \log m)$;
2. if $x \in L$ and there is a witness of length n , then $E(Z(x)) = w$, where w is a witness to $x \in L$ with respect to R_L .

It is natural to consider the relationship between the difficulty of decision and search for a given problem, as was done in other settings, such as average-case complexity by Ben-David et al. [7]. We show that for any problem a compression for the decision variant also yields a compression for the search variant, *without an increase in the VC hierarchy.*

THEOREM 2.26. *For any $k \leq 1$, if the class \mathcal{VC}_k has a compression algorithm, then there is a compression algorithm for the search problem of a relation R_L of $L \in \mathcal{VC}_k$. This is true also for \mathcal{VC}_{OR} .*

Note that Theorem 2.26 holds also when a small error in the compression is allowed. The error in the resulting compression for search algorithm grows by a polynomial factor (by factor n^3) with respect to the error of the underlying compression for decision algorithm. This follows in a straightforward manner from the proof (by a union bound).

The technique of the proof below also comes in handy in proving Theorem 5.4, regarding the application of the ability to compress, say, SAT, to cryptanalysis in the hybrid BSM. In the following proof, a witness to $x \in L$ refers to a witness according to the specific relation R_L associated with L .

Proof. Given an instance x to a language L , for any $i \in [n]$, consider the \mathcal{NP} problem L_i that asks whether there exists an n bit witness w to $x \in L$ such that $w_i = 1$ (the i th bit of w is 1). The language L_i is also in \mathcal{VC}_k , since its verification circuit is the same as the one for L with an additional AND to the variable w_i (this AND gate is incorporated into the top level AND of the circuit; thus the depth remains k).

Our first attempt is to compress the instance x for every $i \in [n]$ with respect to the language L_i (denote such a compression by $Z_{L_i}(x)$). Thus we store $Z_{L_i}(x)$ for all $i \in [n]$, which amounts to $n \cdot p(n, \log m)$ bits for some polynomial $p(n, \log m)$ (this is also in $\text{poly}(n, \log m)$). Now suppose that there is only a *single* witness w to x ; then one can extract w bit by bit, by solving the compressed instance of each bit. However, this fails when w is not the only witness, and we might obtain inconsistent answers for the different bits.

The natural idea now is to use the reduction of Valiant and Vazirani [80] to a unique witness, as was done by Ben-David et al. [7] for showing that average NP being in BPP implies also a randomized search algorithm for average NP. The idea is to choose a pairwise-independent hash function h that is appropriately shrinking and add to the language the requirement that $h(w) = 0$. We use the following lemma.

LEMMA 2.27 (see [80]). *Let L be an \mathcal{NP} language, and for every x denote by W_x the set of all witnesses to $x \in L$. Let ℓ be such that $2^\ell \leq |W| \leq 2^{\ell+1}$. Let $\mathcal{H}_{\ell+2}$ be a family of pairwise independent hash functions with $h : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell+2}$ for all $h \in \mathcal{H}_{\ell+2}$. Then*

$$\Pr_{h \in \mathcal{H}_{\ell+2}}[|\{w : w \in W_x \text{ and } h(w) = 0\}| = 1] \geq \frac{1}{8}.$$

Let \mathcal{H} be a family of pairwise independent hash functions. Consider the \mathcal{NP} language $L^{\mathcal{H}}$ whose elements are of the form (x, h) , where $h \in \mathcal{H}$ maps strings of length n to some shorter length. We have that $(x, h) \in L^{\mathcal{H}}$ if there is a witness w for $x \in L$ and $h(w) = 0$. We note that this language is also in \mathcal{VC}_k , since the additional requirement that $h(w) = 0$ can be verified efficiently over n variables (the hash function h computation is efficient). By Cook's theorem this computation may be represented as a CNF formula ϕ_h over these variables plus only $\text{poly}(n)$ additional variables. Thus adding the requirement of the hash does not add to the depth of the verification circuit for L . This is easy too for \mathcal{VC}_k , and for \mathcal{VC}_{OR} note that we can add (conjunction) the CNF formula ϕ_h to each instance of CircuitSAT, while keeping the problem in \mathcal{VC}_{OR} .

Now, if we enumerate on all values of ℓ , then with probability at least $\frac{1}{8}$, for the correct ℓ we will get that $L^{\mathcal{H}}$ has a unique witness; storing $Z_{L_i^{\mathcal{H}}}(x, h)$ for all i suffices to maintain the information about this witness. This can be repeated sufficiently many times (say, $O(n)$ times) so that with overwhelming probability one of the attempts will indeed give a unique witness. However, this solution is also insufficient, since we have stored a list of $O(n^2)$ compressed values ($O(n)$ repetitions for each value of $\ell \in [n]$)

with the guarantee that with overwhelming probability one of them is a witness for x , but we do not know which one (recall that we cannot store the original instance and thus cannot verify that a witness is correct).

Our final attempt succeeds in reducing the list of potential witnesses into a unique and true witness. This compression is as follows: Denote by $L_{\bar{i}}$ the language that asks whether there exists an n bit witness w to $x \in L$ such that $w_i = 0$ (similar to L_i but with w_i negated). The compression of an instance x to the search problem L goes as follows.

For every $\ell \in [n]$, repeat the following n times:

- Choose $h \in_R \mathcal{H}_{\ell+2}$.
- For all $i \in [n]$ store $Z_{L_i^{\mathcal{H}}}(x, h)$ and $Z_{L_{\bar{i}}^{\mathcal{H}}}(x, h)$.

The extraction procedure is as follows: For all ℓ and $h \in \mathcal{H}_{\ell+2}$, solve all the compressed instance pairs. For every pair $Z_{L_i^{\mathcal{H}}}(x, h)$ and $Z_{L_{\bar{i}}^{\mathcal{H}}}(x, h)$, if both are negative or both are positive, then ignore all values that are compressed with this h . Only if for all i we have that exactly one of the instances being correct, then output the i th bit of w according to the result.

The above algorithm indeed compresses, since it adds only a factor of n^3 to the overall storage. With probability at least $1 - 2^{-\Omega(n)}$, at least one of the chosen h 's is successful in leaving exactly one witness to $x \in L_h$, and this witness will be extracted. Finally, if h did not leave exactly one witness, then this will be identified: If there are no witnesses, then $Z_{L_i^{\mathcal{H}}}(x, h)$ and $Z_{L_{\bar{i}}^{\mathcal{H}}}(x, h)$ will both be negative for all i . If there is more than one witness, then for at least one choice of i , both $Z_{L_i^{\mathcal{H}}}(x, h)$ and $Z_{L_{\bar{i}}^{\mathcal{H}}}(x, h)$ will be positive. \square

2.9. On maintaining other information. We have seen that compression may maintain much more than just a yes/no answer. A natural question to ask is, What other types of information may be maintained through compression algorithms? The following are some examples.

Number of witnesses. The compression described above actually maintains an approximation of the number of witnesses to $x \in L$ (with respect to R_L). Once the chosen k is too large, there will be a sharp drop in the probability of having a witness, and this can be observed when extracting the witnesses and indicate what is the right k .

An almost random witness. The compression above also outputs a witness that is almost uniformly distributed over W_x . Or, more accurately, the probability of getting each witness is bounded by a constant times $1/|W_x|$.

On maintaining all witnesses. As opposed to maintaining a single witness or the number of witnesses, a compressed instance cannot always maintain the information about *all* of the witnesses of an input instance. This is shown by the following simple information theoretic argument: encode an m bit string s with a disjunctive normal form (DNF) circuit C by constructing for each position $j \in [m]$ a formula C_j on $\log m$ variables. If $s[j] = 1$, then take C_j to be a circuit that is satisfied if and only if the variables encode the index j . If $s[j] = 0$, then C_j is the nonsatisfiable circuit $C_j = 0$. The circuit C is formed by taking an OR of all these circuits ($C = \bigvee_{j \in [m]} C_j$). The satisfying assignments of C correspond exactly to the 1's in s . Consider C as an input to the language as CircuitSAT.¹² Suppose that there exists a compression algorithm that maintains all of the witnesses of a circuit C . In particular, this means that the m bit string s may also be extracted from the compressed instance. But

¹²The circuit C is actually an instance for the language $OR(\text{CircuitSAT})$.

this is clearly impossible information theoretically, since m random bits may not be represented by $\text{poly}(n, \log m) < m$ bits. So we conclude that if our goal is to come up with a compression algorithm for SAT, then we must come up with a way of losing information about the witnesses.

In the examples of compression that we have seen in section 2.1, the compression algorithms for vertex cover, PRG-output, and minimum fill-in actually maintain all the witnesses. On the other hand, the compression for GapSAT (which we will see in section 2.10) does not necessarily maintain this information, as it is based on sampling.

2.10. Speculation on compression. We give two arguments that may be viewed as evidence to the existence and nonexistence of compression, respectively.

An optimistic view—compression of a promise problem and the PCP theorem. Consider the promise problem GapSAT that takes as input a CNF formula Φ of size m over n variables and the guarantee that either Φ is satisfiable or it is at most $(1 - \frac{1}{2n})$ -satisfiable (no assignment satisfies more than $(1 - \frac{1}{2n})$ of its clauses). The task is to decide if Φ is satisfiable or far from satisfiable.

Such a problem has a simple and witness-retrievable compression. The idea is to choose $O(n^2)$ random clauses from Φ and take the AND of these clauses to be the compressed formula Ψ . This compression works because if Φ is far from satisfiable, then for every assignment the formula Ψ is satisfied with probability at most 2^{-2n} (Ψ does not contain one of the $\frac{1}{2n}m$ unsatisfied clauses). Taking a union bound over all assignments, we get that with probability $(1 - 2^{-n})$ the formula Ψ has no satisfying assignment. On the other hand, if Φ is satisfiable, then the same assignment also satisfies Ψ (and hence the witness-retrievability). Note that our definition of GapSAT is robust in the sense that GapSAT is compressible whenever the gap is $(1 - \frac{1}{p(n)})$ for every choice of a polynomial $p(\cdot)$.

The above simple compression algorithm is especially interesting in light of the PCP theorem. One way to view the PCP theorem is as an efficient reduction from an instance of SAT to an instance of GapSAT. Thus one can hope to combine the PCP reduction with the above compression and get a compression for general SAT. However, reducing general SAT to GapSAT via the PCP is not a W-reduction, as the witness size grows to the order of the instance size. For starters, the PCP theorem is typically defined over 3-CNF formulas, and the reduction of a general size m CNF to a 3-CNF adds $O(m)$ variables. In order for this approach to achieve compression for SAT, we require a new PCP theorem that is actually a W-reduction.

GapSAT is just one example of a gap problem that admits compression. For instance, one can consider the promise problem GapClique, where a graph of size m either has a Clique of size m/n or contains no Clique of size n . As in the case of GapSAT, GapClique is compressible by sampling a subset of its vertices. Thus, coming up with a W-reduction from a general (n', m') -Clique problem (the graph of size m' either contains a clique of size n' or not) to (n, m) -GapClique would enable the compression of Clique. We view finding PCPs that are also W-reductions as a major research direction, especially in light of the recent new proof to the PCP theorem of Dinur [23].

This connection to succinct PCPs was subsequently studied by Fortnow and Santhanam [36]. They derive negative results on PCPs from the negative results on compression.

A pessimistic view—on oblivious compression. We have seen in section 2.9 that it is impossible to maintain all of the information in an instance when compressing it, and some information is necessarily lost (for example, the list of all witnesses cannot

be kept). On the other hand, we show that if compression exists, then it is not likely to lose too much information about the original instance. Such a result would entail the collapse of the polynomial hierarchy to its second level.

More formally, let Z be a compression algorithm for SAT. We consider it as a two-input algorithm taking a formula Φ and local randomness $r \in \{0, 1\}^\ell$. Denote by $Z(\Phi, U_\ell)$ the random variable taking the output of Z with fixed input Φ and random $r \in_R \{0, 1\}^\ell$. Let \mathbf{X} be a distribution over formulas. The random variable $Z(\mathbf{X}, U_\ell)$ denotes the output of Z under a choice of random r and a random Φ from the distribution \mathbf{X} .

The compression Z is said to be ε -oblivious if for every m, n , there exists a samplable distribution \mathbf{X} over satisfiable formulas of length m and with n variables such that for every satisfiable instance Φ (with parameters m and n) the distribution $Z(\Phi, U_\ell)$ and the distribution $Z(\mathbf{X}, U_\ell)$ are ε -statistically close.

CLAIM 2.28. *If there exists an ε -oblivious compression for SAT (with $\varepsilon \leq \frac{1}{3}$), then the polynomial hierarchy collapses to its second level.*

Proof. We show that if oblivious compression of SAT instances exists, then $\text{CoSAT} \in \mathcal{AM}$. Consider the following interactive proof that an instance $\Phi \notin \text{SAT}$. The verifier chooses a random satisfiable formula $\Psi \in \mathbf{X}$ randomness $r \in U_\ell$ and flips a random coin c . If $c = 0$, then the verifier sends $\xi = Z(\Phi, r)$ to the prover; if $c = 1$, he sends $\xi = Z(\Psi, r)$. The prover then answers 1 if the compressed instance is satisfiable and 0 otherwise. The verifier accepts if the prover's answer equals his bit c and rejects otherwise.

Completeness: If indeed $\Phi \notin \text{SAT}$, then the prover will be able to tell whether the verifier used a coin $c = 0$ or $c = 1$, simply by testing the satisfiability of ξ and replying correctly.

Soundness: Suppose that $\Phi \in \text{SAT}$; then by the obliviousness property of Z the message ξ is from nearly the same distribution whether $c = 0$ or $c = 1$, and the prover is bound to error with probability $\frac{1}{2} + \varepsilon$.

It should be noted also that the above impossibility result does not rely on the fact that the algorithm Z actually compresses but rather on the obliviousness property. \square

We note that the negative result of Fortnow and Santhanam [36] regarding deterministic compression of SAT can be viewed as a further development of these ideas.

3. Part II: Cryptographic applications.

3.1. Basing CRH functions on any one-way function. CRH functions are important cryptographic primitives with a wide range of applications, e.g., [70, 17, 57, 18, 65, 5]. Loosely speaking, a family of length-reducing functions \mathcal{H} is called CRH functions if no efficient algorithm can find collisions induced by a random member of the family. That is, no probabilistic polynomial-time Turing machine (PPTM) can find for a randomly chosen $h \in_R \mathcal{H}$ a pair of input strings x and x' such that $x \neq x'$ but $h(x) = h(x')$. In addition we want (i) an efficient algorithm for *sampling* from \mathcal{H} using (possibly secret) randomness (the secret coins approach is potentially more powerful than when only public coins are used [48]) and (ii) an efficient evaluation algorithm that given the description of $h \in \mathcal{H}$ and x produces $h(x)$. As mentioned in the introduction, CRHs have wide cryptographic applications; see discussion and formal definitions in, for example, [53]. We are interested in basing CRH on as general an assumption as possible. There is no known construction of CRH from general one-way functions or one-way permutations. Moreover, Simon [77] showed that basing

CRH on one-way permutations cannot be achieved using black-box reductions¹³. We show that compression can be used to bridge this gap.

THEOREM 3.1. *If there exists an errorless compression algorithm for SAT or for any problem that is compression-hard for \mathcal{VC}_{OR} , then there exists a construction of a family of CRH functions based on any one-way function.*

Proof. Let (COMMIT, VERIFY) be a statistically binding computationally hiding commitment scheme based on the one-way function f (see, for instance, [37] for formal definitions of commitments). Recall that the protocol COMMIT takes from the sender a string S and randomness r , and after an interaction the receiver gets a commitment σ . The polynomial-time algorithm VERIFY takes the commitment σ and a possible opening to value S' with randomness r' and verifies that S', r' are consistent with σ . One could take, for example, the commitment scheme of Naor [67] based on the one-way function f .¹⁴ In our setting we can work under the assumption that the sender (in the commitment) is honest, and in such a case, the commitment may be achieved without interaction at all.¹⁵

The CRH construction is inspired by the approach of Ishai, Kushilevitz, and Ostrovsky [53] for constructing CRH functions from private information retrieval (PIR). A high-level description is as follows: Choose a hash function from a naive hash family with no computational hardness guarantees; in the construction below we use the selection function, i.e., a random position i . The new hash function is defined by a computationally hiding commitment to the naive hash function, and the output of the new hash function is a compression maintaining the information of the committed naive hash function when applied to the input (i.e., compression of the formula that checks that the value is what it claimed to be). Intuitively, finding a collision would require guessing with nonnegligible advantage the naive hash function (the position i). The actual construction is given in Figure 1.

By the compressing properties of Z we get that h_{σ, r_Z} indeed shrinks its input (note that shrinkage by a single bit allows further shrinking by composition). We also have that sampling h_{σ, r_Z} from \mathcal{H} can be done efficiently (with secret coins).

As for collisions, let $x \neq x'$ be two strings in $\{0, 1\}^m$ that form a collision, i.e., $h_{\sigma, r_Z}(x) = h_{\sigma, r_Z}(x')$. This equality implies, by the property of the compression, that $\Phi_{\sigma, x}$ is satisfiable if and only if $\Phi_{\sigma, x'}$ is satisfiable (here we use the fact that the compression is errorless). Due to the binding property of the commitment, we have that any assignment satisfying Φ_{σ} must have $y = i$ (recall that i is the index that σ is a commitment to). Thus the first part of $\Phi_{\sigma, x}$ is only satisfied when $y = i$. But the second part is only satisfied if $x_y = 1$; thus $\Phi_{\sigma, x}$ is satisfied if and only if $x_i = 1$. We get that $\Phi_{\sigma, x}$ is satisfiable if and only if $x_i = 1$, and $\Phi_{\sigma, x'}$ is satisfiable if and only if $x'_i = 1$. Therefore it must be the case that $x_i = x'_i$, since otherwise one of them is 0 and the other one is 1, and the satisfiability of $\Phi_{\sigma, x}$ is different than that of $\Phi_{\sigma, x'}$. But for some j we have $x_j \neq x'_j$, and for that j we deduce that σ is not a commitment to j .

Suppose now that we have an efficient procedure that finds a collision x and x' for a given (σ, r_Z) with relatively high probability (an inverse polynomial in n). Whenever

¹³Simon's black-box impossibility result [77] is actually stated for the *public* coins version of CRH rather than the *secret* coins variant that we discuss. However, this separation also holds for the case of secret coins (as pointed out in [48]).

¹⁴To be more exact, the commitment of [67] can be based on the pseudorandom generator of Håstad et al. [47], which in turn can be based on the function f .

¹⁵In the scheme of Naor [67], the receiver is required to provide the sender with a (public) random string. Certainly, an honest sender can generate this string by himself without harming the properties of the commitment. Thus in such a setting, the sender can generate the commitment without interaction.

CRH family \mathcal{H}_f :

Description of the hash function: Let Z be a compression algorithm for SAT.

A function in the CRH collection is denoted h_{σ,r_Z} and defined by a commitment σ to a value $i \in [m]$ and randomness r_Z for Z . The commitment uses security parameter n .

Input to h_{σ,r_Z} : a string $x \in \{0,1\}^m$

The CNF formula $\Phi_{\sigma,x}$ is defined as follows:

- Denote by VERIFY_σ the algorithm VERIFY with the input σ fixed. That is, VERIFY_σ takes as inputs y and r and accepts if and only if they form a legal opening of the commitment σ (and in particular this means that $y = i$).
- Translate VERIFY_σ into a CNF formula Φ_σ (using Cook’s reduction) over the variables y_1, \dots, y_ℓ of y , the bits of r , and dummy variables added in the reduction.
- For every $j \in [m]$ define the clause $C_{j,x} = (y_1^{\bar{j}1} \vee y_2^{\bar{j}2} \vee \dots \vee y_\ell^{\bar{j}\ell})$ if $x_j = 0$ (where y^0 denotes \bar{y} and y^1 denotes y) and $C_{j,x} = 1$ if $x_j = 1$.
- Set

$$\Phi_{\sigma,x} = \Phi_\sigma \wedge \bigwedge_{j \in [m]} C_{j,x}.$$

The hash function:

$$h_{\sigma,r_Z}(x) = Z(\Phi_{\sigma,x}, r_Z).$$

FIG. 1. The construction of CRH from any one-way function.

the procedure indeed finds a collision, pick any j such that $x_j \neq x'_j$. For this j we can deduce that σ is *not* a commitment to j . This procedure can be used to break the hiding properties of the commitment scheme, since it yields an efficient method that distinguishes the commitment value from random with advantage $1/m$: Given (the real) i and a random one $i' \in [m]$ in a random order, run the above procedure to obtain j . If j equals one of the two values i or i' , then guess this one as the random one, and otherwise flip a coin. This contradicts our assumptions on building blocks (namely, the one-way function).

To prove the result when using compression for any language that is compression-hard for \mathcal{VC}_{OR} , a similar construction is defined based on the OR of small circuits rather than CNF formulas: For every $j \in [m]$ let $C_{\sigma,j}$ be the circuit that outputs one if and only if there exists randomness r such that σ is consistent with (j, r) (that is, σ is a possible commitment to the value j using randomness r). Let $C_{\sigma,x}$ be the circuit that takes the OR of all $C_{\sigma,j}$ such that $x_j = 1$, and let Z be a compression algorithm for the language $OR(\text{CircuitSAT})$. We define $h_{\sigma,r_Z}(x) = Z(C_{\sigma,x}, r_Z)$. The proof is identical to the case of SAT. \square

Note that instead of an errorless compression, we can do away with an error probability slightly smaller than 2^{-m} . That is, for all x we want the probability that $Z(\Phi_{\sigma,x}, r_Z)$ preserves the satisfiability of $\Phi_{\sigma,x}$ to be at least $1 - 2^{-m+u}$, where the probability is over σ and r_Z and $u \approx \log m$. In this case we can argue (using a union bound) that with probability at least $1 - 2^{-u}$, no x exists violating the preservation of satisfiability.

We also note that the construction is inherently nonblack box, as it uses the code of the one-way function (via the commitment) in the application of Cook's theorem. This is essential for the validity of the whole approach in light of the black-box impossibility of Simon [77]. Theorem 3.1 implies the following corollary.

COROLLARY 3.2. *If there exists an errorless compression algorithm for SAT or for any problem that is compression-hard for \mathcal{VC}_{OR} , then there exist statistically hiding, computationally binding commitment schemes based on any one-way function. The scheme requires two rounds of interaction.*

The corollary follows, since CRH imply statistically hiding bit commitment; see Naor and Yung [70] (and Damgård, Pederesen, and Pfitzmann [18] for commitment to many bits). Until recently, the known minimal assumptions for constructing statistically hiding bit commitments were the existence of one-way permutations [69] and the more general one-way functions with known preimage size [43]. Since the publication of the earlier version of this paper, statistically hiding bit commitments based on any one-way function were shown to exist [71, 44]. However, all of these protocols [69, 43, 44] require many rounds of interaction—at least linear in the security parameter (this was shown to be an inherent limitation of the technique [33, 42]). The commitments based on CRHs, on the other hand, are noninteractive, at least after the initial phase where the function $h \in \mathcal{H}$ is chosen. Such a noninteractive CRH also allows for commitment schemes with very low communication [57].

4. Basing one-way functions on hard instances. The next application shows that compression may be used in order to prove, in the terminology of [49], that **Pessiland** does not exist. Impagliazzo [49] summarizes five possibilities for how the world may look based on different computational assumptions. **Pessiland** is the option where it is easy to generate hard on the average instances, yet no one-way functions exist (or in other words one cannot efficiently generate *solved* hard instances). We show that compression may be used to overrule this possibility and place us in the setting of **Minicrypt** in which one-way functions do exist. More precisely, given a language (not necessarily in \mathcal{NP}) that is hard on the average for nonuniform machines over a samplable distribution and a compression algorithm for a related language, one can construct a one-way function. This result also employs nonblack-box techniques which are essential, as it was shown that there is no *black box* construction of a one-way function from any hard on the average language (over a samplable distribution). This was shown initially by Impagliazzo and Rudich (in unpublished work) and formally by Wee [83].

We start by defining the notion of hardness that we discuss. Denote by $(x \in L)$ the boolean value which corresponds to whether x is in L or not.

DEFINITION 4.1. *A language L is hard for polynomial size circuits over a distribution \mathcal{D} if for every family of polynomial size circuits $\{C_n\}$, for every polynomial $p(\cdot)$, and for all large enough n , it holds that*

$$\Pr_{x \leftarrow \mathcal{D}(1^n)}[C_n(x) = (x \in L)] \leq \frac{1}{2} + \frac{1}{p(n)}.$$

Let L be a language (not necessarily in \mathcal{NP}). Recall that the language $OR(L)$ with parameters m and n is defined as follows:

$$OR(L)_{m,n} = \{(x_1, \dots, x_m) \mid \forall i |x_i| \leq n \text{ and } \exists i \text{ such that } x_i \in L\}.$$

The following theorem demonstrates how compression of $OR(L)$ can be used to construct one-way functions.

THEOREM 4.2. *Given a language L that is hard for polynomial size circuits over a samplable distribution \mathcal{D} and a compression algorithm Z for $OR(L)$,*

1. *if Z is errorless, then there is a construction of CRH functions.*
2. *if Z allows a negligible error (negligible in n), there is a construction of a one-way function.*

Note that there is no restriction on the complexity of recognizing L , other than it being hard for circuits over a samplable distribution. In particular L need not be in \mathcal{NP} at all. If L does happen to be in \mathcal{NP} , then the above statement can use a general compression of a \mathcal{VC}_{OR} -complete language.

COROLLARY 4.3. *Let $L \in \mathcal{NP}$ be hard for polynomial size circuits over a samplable distribution \mathcal{D} (as in Definition 4.1). If there exists a compression algorithm for SAT or for any problem that is compression-hard for \mathcal{VC}_{OR} , then there is a construction of a one-way function. If the compression is errorless, then there is also a construction of CRH functions.*

Proof of Theorem 4.2. The proof follows by defining a family of hash functions h_S based on a compression algorithm. The claim is that, in the errorless case, h_S is a family of CRH functions (see section 3.1). If Z is error prone then we define a modified hash h'_S and prove that it is a family of *distributional* CRH functions. That is, it is hard to find a *random* collision for h'_S . This implies that h'_S naturally defines a distributional one-way function, which, in turn, implies the existence of one-way functions.

We begin by proving the statement in the case of errorless compression. Define a family of hash functions h_S as follows. Each hash function is defined by $S = (\sigma_1^0, \sigma_1^1, \dots, \sigma_m^0, \sigma_m^1)$, a $2m$ -tuple of instances of length n from the domain of the distribution \mathcal{D} . Let Z be a compression algorithm for the language $OR(L)$. Define the hash function $h_S(x) = Z(\sigma_1^{x_1}, \dots, \sigma_n^{x_n})$. Suppose there exists an efficient procedure A that finds collisions for h_S over random $S \in \mathcal{D}^{2m}$. More precisely, there exists a polynomial $p(\cdot)$ such that for infinitely many n ,

$$\Pr_{S \in \mathcal{D}^{2m}} [A(S) = (x, x') \text{ such that } x \neq x' \text{ and } h_S(x) = h_S(x')] \geq \frac{1}{p(n)}.$$

Denote by \mathcal{D}_0 the restriction of the distribution \mathcal{D} to instances $\sigma \notin L$. Note that \mathcal{D}_0 is not necessarily samplable. We show that if there exists a procedure A that finds collisions over $S \in \mathcal{D}_0^{2m}$ (rather than \mathcal{D}^{2m}), then A can be used to break the hardness of the language L over \mathcal{D} . To complete the proof we then show that if A is successful over \mathcal{D}^{2m} , then it is also successful over \mathcal{D}_0^{2m} .

LEMMA 4.4. *Let A be an efficient algorithm and $p(\cdot)$ be a polynomial such that for infinitely many n ,*

$$\Pr_{S \in \mathcal{D}_0^{2m}} [A(S) = (x, x') \text{ such that } x \neq x' \text{ and } h_S(x) = h_S(x')] \geq \frac{1}{p(n)};$$

then there exists a family of polynomial size circuits C^A such that for infinitely many n ,

$$\Pr_{\sigma \in \mathcal{D}} [C^A(\sigma) = (\sigma \in L)] \geq \frac{1}{2} + \frac{1}{2np(n)}.$$

Proof of Lemma 4.4. By the assumption, the procedure A finds a collision with probability at least $\frac{1}{p(n)}$ (over \mathcal{D}_0^{2m}). Therefore, there exists an index $i \in [m]$ such that A finds a collision x, x' such that $x_i \neq x'_i$ (x and x' differed on the i th bit) with

probability at least $\frac{1}{np(n)}$ (since every collision must differ in at least one bit). This index i is used in the reduction described next.

The strategy of C^A for determining membership in L is as follows: Given an input σ drawn from the distribution \mathcal{D} , create a $2m$ -tuple S by putting σ in the i th pair in S (for example, define $\sigma_i^1 = \sigma$) and fill the other entries by random instances from the distribution \mathcal{D}_0 . The nonuniform hint is used to determine i and to supply the random samples from \mathcal{D}_0 . Now run the algorithm A on the tuple S and retrieve a collision x, x' (if A was successful). If $x_i \neq x'_i$, then answer $\sigma \notin L$. Otherwise, answer according to a random coin flip.

Under the restriction that $\sigma \notin L$, the tuple S is distributed precisely as the distribution \mathcal{D}_0^{2m} . Therefore, with probability at least $\frac{1}{np(n)}$ the algorithm A returns a collision with $x_i \neq x'_i$, and C^A answers correctly that $\sigma \notin L$.

On the other hand, under the restriction that $\sigma \in L$, the algorithm A cannot return a collision with $x_i \neq x'_i$. This is due to the fact that the outcome of $h_S(x)$ corresponds to whether $(\sigma_1^{x_1}, \dots, \sigma_n^{x_n})$ is in $OR(L)$ or not (by the correctness of the compression algorithm). But membership in $OR(L)$ is determined solely by the i th pair (all of the other pairs are not in L) and more precisely by the value of the bit x_i . Therefore, a collision can only occur if the i th bit is the same in x and x' . Thus, in this case the procedure C^A answers “not in L ” with probability exactly $\frac{1}{2}$.

Altogether, the procedure C^A answers correctly whenever $x_i \neq x'_i$ (happens with probability $\frac{1}{np(n)}$) and with probability $\frac{1}{2}$ otherwise. This amounts to a success probability of $\frac{1}{2} + \frac{1}{2np(n)}$. \square

It is left to show that A is as successful on \mathcal{D}^{2m} as it is on \mathcal{D}_0^{2m} . For this we define an event under which A is considered successful. In our case it is the cases where A running on S returns a collision under h_S (i.e., $A(S) = (x, x')$ such that $x \neq x'$ and $h_S(x) = h_S(x')$). We say that an algorithm’s success can be *efficiently verified* if there exists a polynomial-time computable relation R such that $R(A(S), S) = 1$ if and only if A was successful on S . This is clearly the case with collision finding, since one can verify efficiently whether the two outputs of A are distinct and collide under h_S . We conclude the first part of the theorem using the following claim.

CLAIM 4.5. *Let A be a polynomial time algorithm whose success can be verified efficiently, and let \mathcal{D} and \mathcal{D}_0 be defined as above. Then for every polynomial $p(\cdot)$ and all large enough n ,*

$$|\Pr_{S \leftarrow \mathcal{D}^{2m}}[A \text{ succeeds on } S] - \Pr_{S \leftarrow \mathcal{D}_0^{2m}}[A \text{ succeeds on } S]| < \frac{1}{p(n)}.$$

Proof sketch. Claim 4.5 is proved by a standard hybrid argument (see, e.g., [37, section 3.2.3]). Namely, one can use a distinguisher between \mathcal{D}^{2m} and \mathcal{D}_0^{2m} in order to distinguish between \mathcal{D} and \mathcal{D}_0 . This in turn is enough to break the hardness of L over \mathcal{D} . Note that nonuniformity is used in the reduction (for constructing hybrid distributions), and so this achieves only a contradiction if L is hard against *nonuniform* adversaries (circuits), even if the distinguisher between \mathcal{D}^{2m} and \mathcal{D}_0^{2m} is actually uniform. \square

This concludes the proof for the errorless case. We now turn to the case of error-prone compression. In this case we also incorporate the string r of random coins used by Z into the hash. Define

$$h'_S(x, r) = (Z_r(\sigma_1^{x_1}, \dots, \sigma_n^{x_n}), r).$$

Unlike the errorless case, we do not know that h'_S forms a CRH family (since the errors may form collisions that are easy to find). Rather, we first show that h'_S is a family of *distributional* CRH functions (DCRH) (a similar primitive was defined in [26]). Loosely speaking, this is a family such that for a randomly chosen hash in the family, no efficient algorithm can find a *random* collision of the hash. A DCRH is useful, since such a family translates to a *collection of distributional one-way functions* which in turn imply the existence of standard full-fledged one-way functions. A distributional one-way function is a function for which it is hard to find a *random* inverse of an output element (rather than just a single preimage as in standard one-way functions). This notion was defined by Impagliazzo and Luby [50], who showed that the existence of distributional one-way functions implies the existence of standard one-way functions. We use a straightforward generalization of distributional one-way functions to collections rather than a single function.

Note, however, that we show only that h'_S is a DCRH when S is sampled according to the distribution \mathcal{D}_0^{2m} . In particular, the key to the hash function cannot necessarily be sampled in an efficient manner. This eventually translates to a one-way function over a domain that might not be efficiently samplable. Unfortunately, one cannot apply Claim 4.5 to show that h'_S forms a DCRH also when S is taken from \mathcal{D}^{2m} , since the property of finding a *random* collision is not efficiently verifiable. Instead, we first construct a collection of one-way functions (via distributional one-way functions) in which the keys are chosen from \mathcal{D}_0^{2m} , and then apply Claim 4.5 to show that the one-wayness holds also for a collection chosen from \mathcal{D}^{2m} (using the fact that finding a single inverse is an efficiently verifiable property).

More formally, as in the case of CRH, a collection of functions consists of algorithms for sampling a key S and evaluating a hash function h'_S over the generated key (in our context we require only that the evaluation algorithm be efficient). For a fixed key S , suppose that h'_S takes inputs of length ℓ . For every such key S define the distribution \mathcal{C}_S over pairs (y, y') such that $y \in U_{\ell(n)}$ and y' is taken uniformly from the collection of the siblings of y (that is, from the set $\{y' \mid h'_S(y) = h'_S(y')\}$). A collection is said to be a *DCRH* if for every efficient algorithm A and every negligible function $\varepsilon(\cdot)$, the probability over the keys $\Pr_S[A(S)$ is $\varepsilon(n)$ -close to \mathcal{C}_S] is negligibly small (i.e., $n^{-o(1)}$). We will first show that h'_S as defined above is a DCRH when S is sampled from \mathcal{D}_0^{2m} . This is implied directly from the following lemma (proof appears after the proof of Theorem 4.2).

LEMMA 4.6. *Let A be an efficient algorithm, $\varepsilon_A(\cdot)$ be a negligible function, and $p(\cdot)$ be a polynomial such that for infinitely many n*

$$\Pr_{S \in \mathcal{D}_0^{2m}}[A(S) \text{ is } \varepsilon_A(n)\text{-close to } \mathcal{C}_S] \geq \frac{1}{p(n)};$$

then there exists an efficient circuit C^A such that for infinitely many n ,

$$\Pr_{\sigma \in \mathcal{D}}[C^A(\sigma) = (\sigma \in L)] \geq \frac{1}{2} + \frac{1}{3np(n)}.$$

We now show that a DCRH implies a collection of distributional one-way functions and start by defining this notion. As before, a collection of functions consists of algorithms for sampling a key S (given security parameter) and evaluating a keyed function f_S over the generated key (where the sampling algorithm is not necessarily efficient in our case). A collection is said to be *distributional one-way* if the probability $\Pr_S[(A(f_S(U_\ell), S), f_S(U_\ell)) \text{ is } \varepsilon(n)\text{-close to } (U_\ell, f_S(U_\ell))]$ is negligibly small (i.e.,

$n^{-o(1)}$). The distributions are taken over the choice of the input in U_ℓ and the random coins of A .

CLAIM 4.7. *Any DCRH also forms a collection of distributional one-way functions.*

Proof sketch. This is shown by demonstrating that a procedure A for breaking the distributional one-wayness of f_S can be used to break the distributional collision resistance of this function. Define the procedure B^A as follows: (i) choose a random $x \in U_\ell$, (ii) compute $x' = A(f_S(x), S)$, and (iii) if $x \neq x'$, then output (x, x') ; otherwise repeat from (i). If, for a given S , the procedure A is such that $(A(f_S(U_\ell), S), f_S(U_\ell))$ is ε -close to $(U_\ell, f_S(U_\ell))$, then the output of B^A is ε -close to C_S . \square

We now use the result of [50] that constructs a standard one-way function from a distributional one-way function. The same transformation holds also for collections of functions (the notion that we use), since the proof holds separately for each function in the family. Thus we derive standard collections of one-way functions (for definition, see, e.g., [37]).

LEMMA 4.8 (from [50, Lemma 1]). *If there is a collection of distributional one-way functions, then there is a collection of one-way functions.*

At this point we have a collection of one-way functions f_S in which the key S is sampled from the distribution \mathcal{D}_0^{2m} (which is not necessary efficiently samplable). We can now apply Claim 4.5 to show that this holds also when S is sampled from the distribution \mathcal{D}^{2m} (which is efficiently samplable). We use the fact that the success of an adversary in finding an inverse of $f_S(x)$ is efficiently verifiable (unlike the success in finding a *random* inverse). The final step is a standard transformation from a collection of a one-way function to a single one-way function (e.g., see [37, section 2.7.4, Exercise 18]). This concludes the proof of Theorem 4.2. \square

Proof of Lemma 4.6. The proof resembles that of the errorless case (Lemma 4.4), and in fact the circuit C^A is essentially the same circuit (barring the minor technicality of ignoring the r part of the inputs).

Recall that the construction in Lemma 4.4 identifies an index i for which a collision with $x_i \neq x'_i$ is found with probability at least $\frac{1}{np(n)}$. Given an instance $\sigma \in \mathcal{D}$, it generates a $2m$ -tuple S with σ in the i th pair and the rest filled with random instances from \mathcal{D}_0 . In Lemma 4.4 when one was given a collision with $x_i \neq x'_i$, we could immediately deduce that $\sigma \notin L$. This is not the case when an error is allowed, since for all we know, the algorithm A might always return an x, x', r such that Z with randomness r errs on either x or x' . What we show is that if A returns a collision according to the required distribution C_S , then with all but negligible probability this collision is a “good” collision (good in the sense that Z_r errs on neither), in which case we can safely deduce that if $x_i \neq x'_i$, then $\sigma \notin L$.

CLAIM 4.9. *Let Z be a compression algorithm for $OR(L)$ with error probability ε_Z ; then for any $S \in \mathcal{D}^{2m}$, $\Pr_{(x, x', r) \leftarrow C_S}[Z_r \text{ errs on either } x \text{ or } x'] < 2\varepsilon_Z$.*

By the assumption on A we get that with probability at least $\frac{1}{np(n)}$, the algorithm A returns a collision with $x_i \neq x'_i$, and by Claim 4.9 we have that with all but probability $2\varepsilon_Z + \varepsilon_A$ (a negligible probability), this collision implies that $\sigma \notin L$ (recall ε_A is the statistical distance of the output of a successful A from C_S). Thus the circuit C^A distinguishes between $\sigma \in L$ and $\sigma \notin L$ with advantage at least $\frac{1}{2np(n)} - \varepsilon_Z - \frac{\varepsilon_A}{2}$ (and in particular with advantage $\frac{1}{3np(n)}$). This concludes the proof of Lemma 4.6. \square

Proof of Claim 4.9. When sampling from \mathcal{C}_S , the first value (x, r) in the collision is simply taken according to the uniform distribution. In particular r is sampled independently of x , and by the definition of compression, for every x , at most an ε_Z fraction of the r 's yield an error. Moreover, when ignoring the first pair, the second value (x', r) is also uniformly distributed. This is because the probability of getting a value (x', r) as the second element in a collision is the probability of hitting a sibling of (x', r) (according to h'_S) as the first element and then the probability of choosing it out of all siblings. Denote the sibling set of (x', r) by $Sib_{(x', r)}$ and the combined length $|x'| + |r|$ by ℓ . Then the probability of getting (x', r) is $\frac{|Sib_{(x', r)}|}{2^\ell}$ for hitting $Sib_{(x', r)}$ times $\frac{1}{|Sib_{(x', r)}|}$ for hitting (x', r) within the set. Thus each element appears as the second element with probability $\frac{1}{2^\ell}$. Therefore, the probability of Z_r having an error on at least one of the values in the collision is at most $2\varepsilon_Z$ (by a union bound). \square

5. On everlasting security and the hybrid BSM. The *BSM*, introduced by Maurer [62], bounds the *space* (memory size) of dishonest players rather than their running time. The model is based on a long random string \mathcal{R} of length m that is publicly transmitted and accessible to all parties. Security relies on the assumption that an adversary cannot possibly store all of the string \mathcal{R} in his memory. The requirement is that the honest parties Alice and Bob can interact using a small local storage of size n (where n is significantly smaller than m), while security is guaranteed against an eavesdropper Eve with a much larger, yet bounded storage space.

This model has enjoyed much success for the task of private key encryption. It has been shown that Alice and Bob who share a short private key can exchange messages secretly using only a very small amount of storage,¹⁶ while an eavesdropper who can store up to a constant fraction of \mathcal{R} (e.g., $\frac{1}{2}m$ bits) learns essentially nothing about the messages (this was shown initially by Aumann and Rabin [4] and improved in [3, 22, 30, 61] and ultimately in Vadhan [79]). These encryption schemes have the important property called *everlasting security* (put forward in [3, 22]). Once the broadcast is over and \mathcal{R} is no longer accessible, then the message remains secure even if the private key is exposed and Eve gains larger storage capacity.

In contrast, the situation is less desirable when Alice and Bob do not share any secret information in advance. The solution of Cachin and Maurer [9] for this task requires Alice and Bob to use storage of size at least $n = \Omega(\sqrt{m})$, which is not so appealing in this setting. Dziembowski and Maurer [29] proved that this is also the best one can do.

The hybrid BSM. The inability to achieve secure encryption in the BSM with memory requirements smaller than $n = \sqrt{m}$ has led to the following suggestion that we call the *hybrid BSM*: Let Alice and Bob agree on their secret key using a computationally secure key agreement protocol (e.g., the Diffie–Hellman protocol [21]). The rationale being that while an unbounded eavesdropper will eventually break the key, if this happens after the broadcast had already occurred, then the knowledge of the shared key would be useless by then (this should be expected from the everlasting security property, where getting the shared key after the broadcast has ended is useless). This hybrid model is very appealing, as it attempts to achieve everlasting security by adding assumptions on the ability of an adversary that has a *strict time*

¹⁶Alice and Bob require only $n = O(\ell + \log m + \log \frac{1}{\varepsilon})$ bits of memory to exchange an ℓ bit message with error ε .

limit. Assumptions of this sort are generally very reasonable, since all that we require is that the computational protocol is not broken in the short time period between its execution and the transmission of \mathcal{R} . For instance, an assumption such as the Diffie–Hellman key agreement [21] cannot be broken within half an hour and can be made with far greater degree of trust than actually assuming the long-term security of this protocol.

Somewhat surprisingly, Dziembowski and Maurer [29] showed that this rationale may fail. They introduce a specific computationally secure key agreement protocol (containing a nonnatural modification based on PIR protocols). If this key agreement protocol is used in the hybrid BSM setting with a specific private key scheme, then the eavesdropper can completely decrypt the encrypted message. However, their result does not rule out the possibility that the hybrid idea will work with some other key agreement protocol. For instance, using the plain Diffie–Hellman key agreement may still work.

In this work we show that if compression of SAT exists, then there exists an attack on the everlasting security of *any* hybrid BSM scheme.

5.1. Two possible models. The notation we use for the storage bounds of the honest parties is n_A and n_B (respectively) and for Eve’s bound is m_E . For simplicity we take $n_A = n_B = n$ and use an abuse of notation by setting $m_E = m$ (where actually it should be that $m_E = \frac{1}{2}m$).

We define the hybrid BSM¹⁷ as a setting, where the running time of the eavesdropper Eve is polynomially bounded up until and during the broadcast of \mathcal{R} and unbounded after that. We discuss two variants of a BSM scheme. We first discuss these in the standard BSM where the eavesdropper is unbounded over time, and then we compare them to the hybrid setting where computational restrictions are imposed.

- *The basic BSM scheme.* The basic scheme allows interaction only up to the start of the broadcast of \mathcal{R} (after that only the encrypted message is sent). Thus the key is fully determined by the time the broadcast has ended. Such a scheme is fully breakable in the standard (nonhybrid) BSM (without an initial secret key), since the unbounded adversary can find some randomness consistent with Alice’s view and simulates Alice’s actions and thus recover the encryption key.¹⁸ Basic schemes in the hybrid BSM are interesting, as they include any combination of a key agreement protocol with a private key scheme (such as the one described by [29] and [45]). We show that if sufficiently strong compression exists, then there exist attacks on any such scheme.
- *The general BSM scheme.* Alice and Bob interact both before *and* after the broadcast of \mathcal{R} . Dziembowski and Maurer [29] show that such a scheme is breakable unless $n > \Omega(\sqrt{m})$ (without initial secret keys). For the hybrid BSM, we show that if compression exists, then there exists an attack on any such scheme as long as $n > \Omega(\sqrt{m/p(n, \log m)})$ for some polynomial p (related to the polynomial of the compression algorithm and to the running time of the protocol that Alice and Bob use).

Thus we prove that if compression of SAT (or of any \mathcal{VC}_{OR} -hard language) is feasible, then the hybrid BSM is essentially no more powerful than the standard BSM.

¹⁷The hybrid BSM model and notions of everlasting security in this model are formally defined in [45].

¹⁸Since Alice must be able to decrypt the message, simulating Alice with any randomness that is consistent with the transcript must output the same key.

On random oracles. The authors of this paper show in [45] that if all parties are given access to a random oracle, then there actually exists everlasting security in the hybrid BSM without an initial key and with low storage requirements from Alice and Bob.¹⁹ Therefore, finding a compression algorithm for SAT would show an example of a task that is achievable with random oracles but altogether impossible without them.²⁰ This would constitute an argument against relying (blindly) on random oracles to determine whether a task is feasible at all. This is different than previous results such as [5, 11, 41, 64, 6], which show specific protocols that becomes insecure if the random oracle is replaced by a function with a small representation. Model separation results were discussed by Nielsen [73, 74] (for noninteractive noncommitting encryption) and Wee [82] (for obfuscating point functions), but the separations there are between the programmable and nonprogrammable random oracle models. In contrast, the hybrid BSM result in [45] holds also if the oracle is nonprogrammable.

5.2. The basic hybrid BSM.

DEFINITION 5.1 (basic hybrid BSM scheme). *A basic hybrid BSM scheme consists of the following: Alice and Bob with storage bound n run a protocol Π that is polynomial in n (this could be a key agreement scheme with security parameter n). Denote by T the transcript of this protocol. Alice and Bob use their respective views of the protocol Π (i.e., the transcript T and their local randomness) to agree on at most n locations of bits from the broadcast string \mathcal{R} that they should store. They store these bits and then use the stored bits to generate an encryption key K (the scheme requires that they agree on the same key).²¹*

We show that sufficiently strong compression of SAT can be used to break any hybrid BSM scheme. For such a scheme to be secure it is required that the key K remains secret in the presence of an eavesdropper that runs in polynomial time up until and during the broadcast, but is unbounded after it. We refer the reader to [46] for rigorous definitions of security (the attack presented below is not sensitive to the actual definition).

For the discussion here, take K to be a one bit key. The general idea is that while the eavesdropper may not figure out in time what locations to store, he can use this transcript to save a relatively short (compressed) CNF formula whose satisfiability coincides with the value of the key K . Later, when he is given unbounded computational power, he will be able to extract this bit from the compressed formula.

THEOREM 5.2. *If there exists a compression algorithm for SAT or for any compression-hard language for \mathcal{VC}_{OR} , with polynomial p_1 , then any basic hybrid BSM scheme can be broken using memory $p_2(n, \log m)$ (where p_2 is a polynomial related to p_1 and the running time of the protocol Π).*

Proof. Denote the locations of the bits that Alice and Bob store by i_1, \dots, i_n . Consider the algorithm V that takes the transcript T_Π and the broadcast string \mathcal{R} as inputs and Alice's local randomness, and locations i_1, \dots, i_n as a witness. The algorithm should check if the witness and inputs are indeed consistent with one another (for example, V should verify that a key agreement with the randomness of Alice, the

¹⁹This does not contradict the compressibility of SAT, since the cryptanalytic result in the hybrid BSM model is not black-box and thus is not preserved in the presence of a random oracle.

²⁰Note that finding an algorithm that actually solves SAT would render more natural tasks (e.g., symmetric encryption) possible in the random oracle model and impossible without it. Of course, finding a compression algorithm seems more likely and does not rule out (most of) cryptography.

²¹The discussion is also valid if the parties are required to reach an agreement with all but negligible probability.

transcript T indeed chooses the indices i_1, \dots, i_n to store) and output 1 if and only if they are consistent and generate an encryption key $K = 1$. The main observation is that the \mathcal{NP} language defined by this relation V is in \mathcal{VC}_1 . Thus, if SAT has a compression algorithm, then there is also a compression algorithm for all of \mathcal{VC}_1 (from Lemma 2.17), including the language defined by V .

The attack of the eavesdropper Eve is as follows: Eve generates the verification program V and feeds the instance (T, \mathcal{R}) to the compression algorithm for the language V . By the properties of the compression, the output is a CNF formula that is satisfiable if and only if $K = 1$. The length of the output is of some polynomial length $p_2(n, \log m)$. If the polynomial p_2 is sufficiently small, then the compressed instance is shorter than Eve's space bound $\frac{1}{2}m$, and she stores this output. Finally, at a later stage, Eve can use her unbounded powers to solve the compressed problem and retrieve the bit K .

We note that a slightly more involved argument works also with compression for \mathcal{VC}_{OR} . The idea is to use independent compression for the bit $\mathcal{R}(i_j)$ for every $j \in [n]$. Every such $\mathcal{R}(i_j)$ may be presented as the OR of m circuits of size $p(n)$ each for some polynomial p . \square

5.3. The general hybrid BSM. The general scheme is like the basic one, but the encryption key K is not necessarily fully defined by the end of the broadcast. In addition, the parties are allowed to interact after the broadcast is over. We note that the bounded storage key exchange scheme of Cachin and Maurer [9] requires such late interaction.

DEFINITION 5.3 (general hybrid BSM scheme). *The general hybrid BSM scheme consists of the following: Alice and Bob with storage bound n engage in a protocol Π_1 that runs in time polynomial in n . Denote by T_1 the transcript of this protocol. Each of the two parties, Alice and Bob, uses its respective view of the protocol Π_1 to determine at most n locations in the broadcast string \mathcal{R} and stores the bits in these locations. After the broadcast they interact in a second protocol Π_2 (with transcript T_2), at the end of which they both agree on encryption key K (with all but negligible error probability).*

THEOREM 5.4. *If there exists a compression algorithm for SAT or for any compression-hard language for \mathcal{VC}_{OR} with compression $p_1(n, \log m)$, then there exists an attack on any general hybrid BSM scheme where $n^2 > m/p_2(n, \log m)$ (where p_2 is a polynomial related to p_1 and the running time of the protocol Π_1).*

Proof. Denote by A_{T_1} the set of all possible random strings r_A of Alice that are consistent with the transcript T_1 (recall that T_1 is executed in full before the string \mathcal{R} is broadcast, and therefore A_{T_1} is fully determined by T_1). Let $s_A = S_A(T_1, \mathcal{R}, r_A)$ denote the bits that Alice stores at the end of the broadcast when running with randomness r_A , transcript T_1 , and broadcast string \mathcal{R} . Finally, denote by $\mathbf{S}_A(T_1, \mathcal{R})$ the random variable that is $S_A(T_1, \mathcal{R}, r_A)$ for a uniform choice of $r_A \in A_{T_1}$. That is, $\mathbf{S}_A(T_1, \mathcal{R})$ is distributed over all possible s_A 's that Alice may store when running with transcript T_1 and broadcast string \mathcal{R} . Similarly we denote by $\mathbf{S}_B(T_1, \mathcal{R})$ the corresponding possible view of Bob.

The proposed strategy for Eve is to store n independent samples from the random variable $\mathbf{S}_A(T_1, \mathcal{R})$. For this purpose we denote by $\mathbf{S}_E(T_1, \mathcal{R})$ (for any \mathcal{R} and T_1) the random variable that consists of n independent samples of $\mathbf{S}_A(T_1, \mathcal{R})$. An important observation due to Maurer [63] is that the uncertainty of Eve regarding the underlying key is upper bounded by the mutual information between the views of Alice and Bob given Eve's view. Formally, the relevant quantity is $I(\mathbf{S}_A(T_1, \mathcal{R}); \mathbf{S}_B(T_1, \mathcal{R}) |$

$\mathbf{S}_E(T_1, \mathcal{R})$). The success of Eve’s strategy follows from the two lemmata below, the first due to Dziembowski and Maurer [29] and the second due to Maurer [63].

LEMMA 5.5 (see [29]). *Let $\mathbf{S}_A(T_1, \mathcal{R})$, $\mathbf{S}_B(T_1, \mathcal{R})$, and $\mathbf{S}_E(T_1, \mathcal{R})$ be defined as above. Then*

$$I(\mathbf{S}_A(T_1, \mathcal{R}); \mathbf{S}_B(T_1, \mathcal{R}) \mid \mathbf{S}_E(T_1, \mathcal{R})) \leq n^2/m.$$

LEMMA 5.6 (see [63, Theorem 3]). *Let \mathbf{V}_A , \mathbf{V}_B , and \mathbf{V}_E be random variables denoting the respective views of Alice, Bob, and Eve. Let $\mathbf{K}_A = K_A(\mathbf{V}_A)$ and $\mathbf{K}_B = K_B(\mathbf{V}_B)$ be procedures of Alice and Bob to extract a mutual secret key from their respective views such that $\mathbf{K} = \mathbf{K}_A = \mathbf{K}_B$ with all but negligible probability. Then $H(\mathbf{K}) \leq I(\mathbf{V}_A; \mathbf{V}_B \mid \mathbf{V}_E)$.*

A strategy for an eavesdropper is therefore to store n independent samples of the random variable $\mathbf{S}_A(T_1, \mathcal{R})$. Lemmata 5.5 and 5.6 assert that Eve’s entropy of the encryption key K is at most n^2/m in such a case. A crucial point is that an encryption key that has entropy significantly lower than 1 (from Eve’s point of view) can be predicted with high probability by an unbounded Eve, rendering the scheme insecure. Thus if an eavesdropper has $O(m)$ storage capacity, then the scheme is insecure as long as $n^2 = O(m)$.²²

Lemma 5.5 was used in [29] in a setting where the eavesdropper is unbounded and can hence sample the random variable $\mathbf{S}_A(T_1, \mathcal{R})$. However, in our setting the eavesdropper is computationally bounded and does not have the power to generate this distribution. Instead, we use compression to store information about samples of $\mathbf{S}_A(T_1, \mathcal{R})$ to be extracted *after* the broadcast is over (when the eavesdropper is computationally unbounded).

The main idea is to use compression for search problems, as was discussed in section 2.8. Define the \mathcal{NP} language L_A as follows:

$$L_A = \{(T_1, \mathcal{R}) \mid \exists \text{ witness } w = (r_A, s_A) \text{ such that } r_A \in A_{T_1} \text{ and } s_A = S_A(T_1, \mathcal{R}, r_A)\}.$$

The first thing to note is that L_A is in \mathcal{VC}_{OR} . This is shown once more by the same argument as in Theorems 5.2 or 3.1 and based on the fact that the protocol Π_1 is polynomial-time in n . Once this is established, then given a compression algorithm for \mathcal{VC}_{OR} we invoke Theorem 2.26 to get a compression algorithm to the search problem associated with L_A . Running this compression once allows us to extract a witness to L_A and in particular to get one sample s_A of a consistent view of Alice. Running this n times supposedly gives n samples of such a view, which suffices to break the scheme by Lemma 5.5.

However, in order to invoke Lemma 5.5, we need the samples to be taken according to the distribution $\mathbf{S}_A(T_1, \mathcal{R})$, which is defined by a uniform distribution over $r_A \in A_{T_1}$. We will show that while sampling via the compression of search problems does not give the desired distribution exactly, it is still sufficiently close to be useful.

A closer inspection of our compression for search technique from section 2.8 shows that we do not necessarily sample uniformly from A_{T_1} . However, we do sample close to uniformly, in the sense that no element in A_{T_1} gets more than double the probability of another element in A_{T_1} . We then show that taking a constant times many samples as was originally needed guarantees that amongst the stored bits we have n random

²²When considering n_A and n_B that are not necessarily identical, the actual requirement is for Eve to store n_B samples of $\mathbf{S}_A(T_1, \mathcal{R})$ (each sample is of length n_A). Subsequently the scheme is insecure as long as $n_A \cdot n_B < O(m_E)$.

samples of the random variable $\mathbf{S}_A(T_1, \mathcal{R})$, and thus we have stored enough bits from \mathcal{R} to break the scheme.

Recall from section 2.8 that the compression algorithm for search problems chooses a random pairwise-independent hash function h and saves only a witness (r_A, s_A) that is *uniquely* hashed to the value 0 by h . Since r_A fully determines s_A (when given T_1 and \mathcal{R}), then without loss of generality we view the witness simply as r_A . Furthermore, assume without loss of generality that r_A is of length n . Suppose that $\ell \in [n]$ is such that $2^\ell < |A_{T_1}| \leq 2^{\ell+1}$. Let $\mathcal{H}_{\ell+2}$ be a family of pairwise independent hash functions with $h : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell+2}$ for all $h \in \mathcal{H}_{\ell+2}$. Then for every $r_A \in A_{T_1}$ the probability that a random $h \in \mathcal{H}_{\ell+2}$ uniquely maps r_A to zero is at most $2^{-(\ell+2)}$ (since $\Pr_{h \in \mathcal{H}_{\ell+2}}[h(r_A) = 0] = 2^{-(\ell+2)}$). By the pairwise independence of \mathcal{H} , it holds that for all other $r'_A \in A_{T_1}$ with $r'_A \neq r_A$ we have that $\Pr_{h \in \mathcal{H}_{\ell+2}}[h(r'_A) \neq 0 | h(r_A) = 0] = 1 - 2^{-(\ell+2)}$. By a union bound over all $r'_A \in A_{T_1}$ with $r'_A \neq r_A$, combined with the probability that $h(r_A) = 0$, we get

$$\Pr_{h \in \mathcal{H}_{\ell+2}}[h \text{ uniquely maps } r_A \text{ to } 0] \geq 2^{-(\ell+2)} \cdot \frac{1}{2} = 2^{-(\ell+3)}.$$

Altogether, for all $r_A \in A_{T_1}$ it holds that

$$2^{-(\ell+2)} \geq \Pr_{h \in \mathcal{H}_{\ell+2}}[h \text{ uniquely maps } r_A \text{ to } 0] \geq 2^{-(\ell+3)}.$$

Thus whenever the output of h is indeed of length $\ell + 2$, the probability of sampling $r_A \in A_{T_1}$ is almost uniform (up to a factor of 2 for each element).²³ Since we repeat the compression for every choice of $\ell \in [n]$, then in particular samples are stored for the correct ℓ .

By Lemma 2.27 we know that at least $\frac{1}{8}$ of the repeated compressions indeed store information about a valid witness (a sample of $r_A \in A_{T_1}$). Thus, choosing, say, $9n$ independent $h \in \mathcal{H}_{\ell+2}$ guarantees at least n samples (by a Chernoff bound, as the choices are independent). But as mentioned above, these samples are just close to uniform over A_{T_1} rather than truly uniform. The solution is to simply run more instances of this process, say, for $25n$ independent choices of $h \in \mathcal{H}_{\ell+2}$. This would guarantee that with overwhelming probability, at least $3n$ of these choices have a valid witness. We show that from these slightly biased samples we can extract n truly uniform samples of witnesses.

This last argument follows by a method for generating uniformly distributed samples from A_{T_1} . At a first stage, $3n$ samples are taken using the unique hashing method. Now a diluting second stage is run in order to extract the actual samples: Suppose that the least likely element to be sampled gets probability p_{min} . For any element r_A that is sampled with probability p_{r_A} , keep the sample with probability $\frac{p_{min}}{p_{r_A}}$ and delete it otherwise. Thus every element is eventually chosen with the same probability p_{min} , and since $\frac{p_{min}}{p_{r_A}} \geq \frac{1}{2}$, then at least n samples are eventually chosen (with overwhelming probability). Note that the diluting stage is not necessarily efficiently computable. However, the probability p_{r_A} can be computed using the adversaries unbounded running time, since these probabilities are fully defined by the transcript T_1 , which can be stored in its entirety (as it is of length polynomial in n). Therefore an unbounded eavesdropper may indeed extract n uniform samples from her view. \square

²³Note that the almost uniformity of the samples actually holds for every choice of the parameter ℓ . Therefore, this property can be relied on even if the correct choice of ℓ is unknown.

Note: In the two models that we consider we limit the honest parties to access and store at most n actual bits from the broadcast string \mathcal{R} . This is in contrast to storing some function of \mathcal{R} with a bound on the function's output length (an ability that the adversary is entitled to). This is a legitimate requirement, as the honest parties should run algorithms that are considerably more efficient than the adversary's. It should be noted, however, that our Theorems (5.2 and 5.4) hold also if the honest players can store functions, albeit they then call for a compression algorithm for all of \mathcal{NP} (rather than just for the lowest class \mathcal{VCR}).

6. On witness retrievable compression and public key cryptography based on any one-way function.

6.1. On OT from any one-way function. The two top worlds that Impagliazzo considers in his survey [49] are *Minicrypt*, where one-way functions exist but OT protocols do not exist (in this world some interesting cryptographic applications are possible, and in particular *shared* key cryptography exists) and *Cryptomania*, where OT protocols do exist (and hence also a wide range of cryptographic applications like secure computation and *public* key cryptography). The last application we discuss is an attempt to use compression in order to prove that *Minicrypt*=*Cryptomania*. Whether OT can be constructed from any one-way function is a major open problem in cryptography. Impagliazzo and Rudich [52] addressed this problem and proved that key agreement protocols (and hence also OT) cannot be constructed from any one-way function using “black-box” reductions. Many researchers view this impossibility result as an indication that general one-way functions are insufficient for public key cryptography.

We explore an approach of using compression in order to bridge the gap between the two worlds. In order to do so we introduce an additional requirement of a compression algorithm.

DEFINITION 6.1 (witness-retrievable compression). *Let Z, L , and L' define a compression algorithm as in Definition 1.2, and let R_L and $R_{L'}$ be \mathcal{NP} relations for L and L' , respectively. The compression is said to be witness-retrievable with respect to R_L and $R_{L'}$ if there exists a probabilistic polynomial-time machine W such that for every input x , if $x \in L$, then for every witness w_x for x with respect to R_L it holds that $w_y = W(w_x, Z(x))$ is a witness for $Z(x) \in L'$ with respect to $R_{L'}$. We allow a negligible error in the success of W (where probability is over the internal randomness of Z and W).*

We describe an approach to bridging this gap using witness-retrievable compression of a specific language. More precisely, we demonstrate a construction of an OT protocol (see definition in, for instance, [38]) from any one-way function using such a compression algorithm.

THEOREM 6.2. *If SAT has a witness-retrievable compression algorithm (even for the specific type of CNF formulas in the proof below), then one can construct an OT from any one-way function.*

As in the CRH construction (Theorem 3.1), the conditional construction of OT in Theorem 6.2 is inherently nonblack-box. Unfortunately we show that this approach cannot work with a compression algorithm for the *general* SAT problem, due to the following theorem.²⁴

²⁴The first version of this paper [46] (dated Feb 17, 2006) did not contain this theorem and was hence more optimistic on the possibility of finding a witness-preserving compression algorithm for SAT.

THEOREM 6.3. *If one-way functions exist, then there is no witness-retrievable compression of SAT.*

Furthermore, we also rule out the possibility of other types of witness-retrievable compression that may be sufficient for Theorem 6.2. More precisely, the impossibility of witness-retrievable compression does not change when allowing an error in the retrieval or when dealing with a case where there is a unique witness (see Corollary 6.9). These developments rule out basing the approach of Theorem 6.2 on the compression of a general and standard language. The approach may still work out with a witness-retrievable compression algorithm for the specific CNF formulas as constructed in the proof of Theorem 6.2.

Finally, we point out that almost all of the examples of compression algorithms in this paper (in sections 2.1 and 2.10) are in fact witness-retrievable. This demonstrates that these examples fall short of the general compression that we are seeking. In fact a major obstacle in achieving compression for problems such as SAT seems to be that most natural approaches would be witness-retrievable.

Proof of Theorem 6.2. The construction actually builds a PIR protocol and then uses the construction of Di Crescenzo, Malkin, and Ostrovsky [20] to build an OT protocol from the PIR protocol. Recall that a PIR protocol has a sender with a database of size m and a receiver that chooses to learn one entry from the database (see precise definition in, e.g., [20]). It is required that the receiver learns the bit of his choice, but a computationally bounded sender learns essentially nothing about this choice. In addition, the total communication should be strictly smaller than m .

Let f be a one-way function and take (COMMIT, VERIFY) to be a commitment based on the one-way function f (as in section 3.1). In this proof we work under the assumption that the parties are semihonest (that is, the parties follow the protocol as prescribed and are only allowed to try and infer extra information from the transcript of the protocol). The semihonest assumption is justified by the compiler of Goldreich, Micali, and Wigderson [39] that showed how to transform a semihonest protocol into one against malicious parties (again, the only needed cryptographic assumption is the existence of a one-way function). Consider the protocol in Figure 2.

It remains to show that the protocol PIR_f is indeed a PIR protocol. Due to the fact that the commitment is binding (up to a negligible error), an assignment satisfying Φ_σ must have $x = i$ (recall that i is the index that Bob committed to). Thus the first part of Φ is only satisfied when $x = i$. But the second part is only satisfied if $D[x] = 1$; thus Φ is satisfied if and only if $D[i] = 1$. By the property of the compression algorithm, also Ψ is satisfiable if and only if $D[i] = 1$. Hence, using the witness-retrievable properties of the compression, Bob figures out whether or not Ψ is satisfiable and learns the bit $D[i]$ (up to a negligible error).

The second property is that the sender Alice learns no computational information about Bob's choice. This follows directly from the guarantees of the commitment scheme (note that Bob does not send any information outside of the commitment). The third and final requirement regards the length of the communication. But the length of the communication is a fixed polynomial in $p(n)$ (depending on the commitment protocol and the parameter of the compression algorithm). So choosing a large enough databases with $m > p(n)$ guarantees a nontrivial PIR protocol and hence an OT protocol. \square

Note that the OT protocol derived in Theorem 6.2 is a one-round protocol (that is, one message sent from the receiver followed by one message from the sender). This follows from the construction of the PIR protocol and the construction of [20] that preserves the number of rounds. One implication of this fact is that such an OT

Protocol PIR_f :

Alice's input: database D of m bits. Let $D[i]$ denote the i th bit in D .

Bob's input: index $i \in [m]$. Denote the bits of i by i_1, \dots, i_ℓ .

1. **Bob commits to i :** Bob commits to i with randomness r_B , Alice receives $\sigma = \text{COMMIT}(i, r_B)$.
2. **Alice computes Φ :** The CNF formula Φ is defined as follows:
 - Denote by VERIFY_σ the algorithm VERIFY with the input σ fixed. That is, VERIFY_σ takes as inputs x and r and accepts if and only if they form a legal opening of the commitment σ (and in particular this means that $x = i$).
 - Translate VERIFY_σ into a CNF formula Φ_σ (using Cook's reduction) over the variables x_1, \dots, x_ℓ of x , the bits of r , and dummy variables added in the reduction.
 - For every $j \in [m]$ define the clause $C_j = (x_1^{j_1} \vee x_2^{j_2} \vee \dots \vee x_\ell^{j_\ell})$ if $D[j] = 0$ (where x^0 denotes \bar{x} and x^1 denotes x) and $C_j = 1$ if $D[j] = 1$.
 - Set

$$\Phi = \Phi_\sigma \wedge \bigwedge_{j \in [m]} C_j.$$

3. **Alice Compresses Φ :** Let (Z, W) be a witness-retrievable compression algorithm for CNF formulas of the form of Φ . Alice runs $\Psi = Z(\Phi)$ and sends Ψ to Bob.
4. **Bob checks witness:** Note that Bob knows the witness to VERIFY_σ and can compute a witness w for Φ_σ . Bob checks if $W(w, \Psi)$ is a satisfying assignment for Ψ . If it is, Bob outputs 1; otherwise he outputs 0.

FIG. 2. The construction of a PIR protocol from any one-way function.

protocol may be used to construct a two-round key agreement scheme, that in turn maybe used to construct a public key encryption. In general, this is achieved by fixing the first message of the protocol to be as the public key.

COROLLARY 6.4. *If there exists a witness-retrievable compression algorithm for the type of SAT instances that occur in in the proof of Theorem 6.2, then based on any one-way function one can construct a public key encryption scheme (PKE) that is semantically secure against chosen plain text attacks.*

6.2. On the limitation of the witness retrievability property. Witness-retrievable compression is defined (Definition 6.1) as a compression with an additional probabilistic polynomial time algorithm W such that for every witness w_x for R_L it holds that $w_y = W(w_x, Z(x))$ is a witness for $Z(x) \in L'$. Recall that nearly all of the examples of compression algorithms (in sections 2.1 and 2.10) are in fact witness-retrievable (the exception being compression of general sparse languages, Definition 2.3). This property is essential to the success of the construction of the OT protocol in Theorem 6.2 (without it the receiver would have to run in time that is super-polynomial). In this section we show that if one-way functions exist, then a compression algorithm for SAT cannot be witness-retrievable (this regards the general

language SAT rather than a specific distribution of instances as generated in Theorem 6.2). Moreover, this statement also holds for other general languages mentioned in Theorem 6.2 (that are potentially easier to compress than SAT). In particular, there is no witness-retrievable compression for the Clique language or for the language $OR(SAT)$ (that is complete for \mathcal{VC}_{OR}). We give the formal statements below with respect to the language $OR(SAT)$ and deduce the statements for SAT and Clique as corollaries.

We also rule out other natural definitions of witness-retrievability that would have been sufficient for the proof of Theorem 6.2 to go through. Suppose we relax the witness-retrievability requirement to hold only with some probability ε ; then we show that if one-way functions exist, then this probability ε has to be very low, at most an inverse polynomial in m . Such a low probability of success is *not* sufficient for the OT construction in Theorem 6.2 to follow (we note though, that witness-retrievability with this low success probability is still sufficient for the cryptanalytic result in [28]). We then show that the same situation also holds for languages that are guaranteed to have *unique witnesses* (i.e., unique-SAT and unique- $OR(SAT)$). This is of relevance, since the instances being compressed in the proof of Theorem 6.2 all have at most a single witness.²⁵

We emphasize again that the OT construction may still be successful under the compression of formulas of the specific type that are generated in the proof. However, we cannot generalize this method to work with compression of a more standard language.

On the impossibility of perfect witness retrieval. Recall that the language $OR(SAT)$ takes as an input a list of m CNF formulas (each of length n) and accepts if at least one of the formulas is satisfiable. Consider the following way of generating an instance of $OR(SAT)$. Take m bit commitments $\sigma_1, \dots, \sigma_m$, each with security parameter n (see proof of Theorem 3.1 for a definition and discussion of commitments in our context). For each commitment σ_i , generate using Cook's theorem a CNF formula ϕ_{σ_i} that is satisfiable if and only if σ_i is a commitment to 1. As an instance of $OR(SAT)$, we take the OR of the m CNF formulas $\phi_{\sigma_1}, \dots, \phi_{\sigma_m}$. We denote this instance by $\phi(\sigma_1, \dots, \sigma_m)$. Denote by w_{σ_i} a satisfying assignment for ϕ_{σ_i} (such an assignment can be generated by an opening of σ_i to the value 1). The assignment w_{σ_i} also serves as a witness for $\phi(\sigma_1, \dots, \sigma_m) \in OR(SAT)$. Our first impossibility result is for compression of $OR(SAT)$ with errorless witness-retrievability.

LEMMA 6.5. *If one-way functions exist, then there is no witness-retrievable compression for $OR(SAT)$ with perfect witness-retrieval.*

Proof. The proof follows by showing that a witness-retrievable compression Z for $OR(SAT)$ can be used to transmit an m bit string between two parties with sublinear communication. As a setup stage, the receiver generates m random commitments to 1 and m random commitments to 0 and sends them to the sender. Denoted these by $(\sigma_1^1, \dots, \sigma_m^1)$ and $(\sigma_1^0, \dots, \sigma_m^0)$, respectively.

For every string $x \in \{0, 1\}^m$ denote $\phi_x = \phi(\sigma_1^{x_1}, \dots, \sigma_m^{x_m})$ (where x_i denotes the i^{th} bit of x). In order to send string $x \in \{0, 1\}^m$ the sender sends $Z(\phi_x)$ to the receiver. We claim that the receiver can, with overwhelming probability, learn the string x , thus contradicting the fact that the message sent is significantly shorter

²⁵The relevant instances in Theorem 6.2 actually have a unique witness only if there exists a commitment scheme that has only a *unique opening*. As this is not necessarily the case when given any one-way function, we consider for simplicity the case of one-way permutations (that guarantee a unique opening commitment scheme).

than m . Note that the receiver knows witnesses $w_{\sigma_i^1}$ for all i and that a witness for $\phi_x \in OR(SAT)$ consists of a witness $w_{\sigma_i^1}$ of a $\phi_{\sigma_i^1}$ that is included in ϕ_x . The receiver extracts x as follows:

Procedure Rec on input $Z(\phi_x)$.

- For every $i \in [m]$
 1. run $w = W(Z(\phi_x), w_{\sigma_i^1})$;
 2. if w is a witness for $Z(\phi_x)$, then set $y_i = 1$; otherwise, set $y_i = 0$.
- Output $y = y_1, \dots, y_m$.

Denote by X_i the random variable of the i th bit of x and by Y_i the random variable of the corresponding output of *Rec*. We view the process as a channel between a sender who holds the random variables $X = X_1, \dots, X_m$ to a receiver who gets the random variables $Y = Y_1, \dots, Y_m$ and claim that with overwhelming probability $Y = X$.

If $X_i = 1$, then the opening of σ_i^1 should yield a witness for $Z(\phi_x)$, from the perfect witness-retrievability, and thus $Y_i = 1$. We should show that if $X_i = 0$, then indeed $Y_i = 0$ (up to a negligible error). Note that X is uniformly distributed over $\{0, 1\}^m$, whereas Y is determined by the random choice of commitments $(\sigma_1^1, \dots, \sigma_m^1)$ and $(\sigma_1^0, \dots, \sigma_m^0)$, the random coins of Z and W , and the random variable X .

CLAIM 6.6. *Let X and Y be the random variables described above. Then for every $i \in [m]$ (possibly related to m, n) and every polynomial $q(\cdot)$ and all sufficiently large n ,*

$$\Pr[Y_i = 1 | X_i = 0] < \frac{1}{q(n)}.$$

Note that the Claim 6.6 holds also if the underlying witness-retrieval algorithm is nonperfect. This will be used in the proof of Lemma 6.7.

Proof. Suppose that the claim is false, that is, for some $q(\cdot)$, for infinitely many n , and some i (possibly related to n), $\Pr[Y_i = 1 | X_i = 0] \geq 1/q(n)$. For simplicity we first deal with the case where $\Pr[Y_i = 1 | X_i = 0] = 1$. In other words, $W(Z(\phi_x), w_{\sigma_i^1})$ always outputs a witness for $Z(\phi_x)$. Consider the two distributions \mathcal{L}_0 and \mathcal{L}_1 on lists of $m - 1$ commitments:

- Distribution \mathcal{L}_0 is defined by a random and independent choice of $m - 1$ commitments to 0.
- Distribution \mathcal{L}_1 is defined by first choosing at random a string $V_1, V_2, \dots, V_{m-1} \in \{0, 1\}^{m-1}$ and then generating $m - 1$ independent commitments to V_1, V_2, \dots, V_{m-1} .

From the hiding property of commitment schemes, it holds that these two distributions are indistinguishable, i.e., given a list L of $m - 1$ commitments, no computationally bounded distinguisher can tell with nonnegligible bias whether L was generated by \mathcal{L}_0 or \mathcal{L}_1 . We will show that if the premise of the claim is false, it is possible to distinguish the two distributions (without knowledge of the openings to any of the commitments in the list).

Given a list L of $m - 1$ commitments, the distinguisher generates σ_i^0 and σ_i^1 and the corresponding witnesses. He then generates a formula ϕ by adding σ_i^0 to the i th position in the list L and runs the compression on ϕ . The distinguisher then runs $w = W(Z(\phi), w_{\sigma_i^1})$ and checks whether w is a witness to $Z(\phi)$. By the assumption, w will indeed be a witness every time that ϕ is satisfiable. On the other hand, w cannot be a witness if ϕ is *not* satisfiable, simply by the properties of the compression. Thus if w is indeed a witness for $Z(\phi)$, then it must be that $\phi \in OR(SAT)$, and there is some commitment to 1 in the list, and thus L was generated from \mathcal{L}_1 . Otherwise, it

means that $\phi \notin OR(SAT)$, and the original list was from \mathcal{L}_0 (ignoring the negligible probability that \mathcal{L}_1 generates a list containing only commitments to 0).

Now if $\Pr[Y_i = 1 | X_i = 0] \geq \frac{1}{q(n)}$ for some polynomial $q(\cdot)$, then the distinguisher follows the same procedure with the difference that

- if $w = W(Z(\phi), w_{\sigma_i^1})$ is a witness for $Z(\phi)$, then output \mathcal{L}_1 .
- if w is not a witness, flip a coin and output either \mathcal{L}_0 or \mathcal{L}_1 accordingly.

In case w was indeed a witness, the distinguisher is guaranteed to be correct. Therefore, the above procedure gives an advantage $\frac{1}{2q(n)}$ in distinguishing between \mathcal{L}_0 and \mathcal{L}_1 , contradicting the hiding properties of the commitment scheme. \square

Note that the distributions \mathcal{L}_0 and \mathcal{L}_1 will be useful also in the discussion of the unique witnesses case (Lemma 6.8). \square

On nonperfect witness retrievability. We now show that the witness-retrieval procedure is possible only if its success probability is sufficiently low (we denote the success probability by $\frac{1}{q(n,m)}$). We upper bound the success probability by a function of the rate of compression that the algorithm Z achieves (we denote by $p(n,m)$ the polynomial that bounds the length of the output of Z , i.e., the compressed instance).

LEMMA 6.7. *Suppose one-way functions exist, and suppose that (Z,W) is a witness-retrievable compression for $OR(SAT)$ such that for every ϕ with parameters m,n the following holds:*

1. *The compression parameter $|Z(\phi)| \leq p(n,m)$.*
2. *The success probability of W is at least $\frac{1}{q(n,m)}$, where probability is over the random coins of Z and W as well as the choice of the witness.*

Then $q(n,m) \geq \Omega(\frac{m}{p(n,m)})$.

Proof. The proof uses the same setting as in the proof of Lemma 6.5. Once more, the sender sends a compressed value $Z(\phi_x)$ to the receiver that runs the procedure Rec , and we view this process as a channel between a sender who holds the random variables $X = X_1, \dots, X_m$ to a receiver who gets the random variables $Y = Y_1, \dots, Y_m$. Only this time if $X_i = 1$, it is not guaranteed that also $Y_i = 1$ (since the witness-retrievability is no longer perfect). Instead, our assumption on the success probability of W translates to $\Pr[Y_i = 1 | X_i = 1] \geq \frac{1}{q(n,m)}$ for a random i . Since X_i is a uniformly distributed bit, then $\Pr[Y_i = 1] \geq \frac{1}{2q(n,m)}$ for a random i .

In addition, Claim 6.6 states that for every i it holds that $\Pr[Y_i = 1 | X_i = 0] \in neg(n)$. Thus, if $Y_i = 1$, then $X_i = 1$ with overwhelming probability, and therefore $H(X_i | Y_i = 1) \in neg(n)$ for every i (where H denotes the Shannon entropy). We use the above-mentioned facts to provide an upper bound on the average entropy of X_i (average over i) when given Y :

$$\begin{aligned} \mathbb{E}_i[H(X_i | Y)] &= \mathbb{E}_i[Pr(Y_i = 1)H(X_i | Y_i = 1) + Pr(Y_i = 0)H(X_i | Y_i = 0)] \\ &\leq \frac{1}{2q(n,m)} \cdot neg(n) + \left(1 - \frac{1}{2q(n,m)}\right) \cdot 1 \\ &\leq 1 - \frac{1}{2q(n,m)} + neg(n). \end{aligned}$$

The first inequality is true, since $H(X_i | Y_i = 0) \leq 1$ for every i . We deduce an upper bound on the entropy of X when given Y :

$$H(X|Y) \leq \sum_i H(X_i | Y) = m\mathbb{E}_i[H(X_i | Y)] \leq m \left(1 - \frac{1}{2q(n,m)} + neg(n)\right).$$

Hence, when the receiver gets $Z(\phi_x)$ (and can generate Y), the receiver's entropy of

X deteriorates by

$$H(X) - H(X | Y) \geq \Omega\left(\frac{m}{q(n, m)}\right).$$

This can only happen if the sender sent at least $\Omega(\frac{m}{q(n, m)})$ bits to the receiver, and thus $p(n, m) \geq \Omega(\frac{m}{q(n, m)})$ as required. \square

Note that the construction of OT protocols from one-way functions in Theorem 6.2 requires that the compression rate $p(n, m) \leq O(m^{1-\epsilon})$ for some constant $\epsilon > 0$. Thus, when put in the context of constructing OT protocols, the above lemma states that a useful compression algorithm for $OR(SAT)$ cannot have witness-retrievability with probability that is better than $O(\frac{1}{m^\epsilon})$. In order to achieve nontrivial PIR protocols (via Theorem 6.2), one would require witness-retrievability with a better success probability.

On witness retrieval with a unique witness. The limitations on witness-retrievability hold also when there is only a single witness, which is the case in our cryptographic applications. For this we consider the *promise problem* $OR(SAT)^U$ that is $OR(SAT)$ with a guarantee that every instance has at most one satisfying assignment. We generate the interesting instances of $OR(SAT)^U$ as above, from sets of commitments. In this case the set of commitments should be such that at most one of the commitments is to the value 1. For simplicity we also assume that each commitment has a unique opening (this may be achieved using one-way permutation), so overall such instances have the unique witness property.

LEMMA 6.8. *Suppose one-way permutations exist, and suppose that (Z, W) is a witness-retrievable compression for $OR(SAT)^U$ such that for every input ϕ with parameters m, n the following holds:*

1. *The compression parameter is $|Z(\phi)| \leq p(n, m)$.*
2. *The success probability of W is at least $\frac{1}{q(n, m)}$ for a polynomial $q(\cdot, \cdot)$, where probability is over the random coins of Z and W .*

Then $\frac{1}{q(n, m)} - \frac{p(n, m)}{m} \in \text{neg}(n)$.

Proof. Suppose that there is a witness-retrievable compression (Z, W) for $OR(SAT)^U$ that succeeds with probability $\frac{1}{q(n, m)}$. In similar fashion to the proof of Claim 6.6 we will show that in such a case one can efficiently distinguish if a list of $m - 1$ commitments was generated by the distribution \mathcal{L}_0 or by the distribution \mathcal{L}_1 . Recall that the distribution \mathcal{L}_0 is a random choice of $m - 1$ commitments to 0, while the distribution \mathcal{L}_1 is a choice of $m - 1$ random commitments (commitments to either 0 or 1). The distinguisher works without knowledge of the openings to any of the commitments, thus contradicting the hiding properties of the commitment scheme.

The distinguisher generates a random commitment σ^1 to 1, along with its witness w_{σ^1} . Now, given a list L of $m - 1$ commitments, the distinguisher creates an instance ϕ by adding σ^1 in a random position in the list L and runs the compression on ϕ . The distinguisher then tries to retrieve a witness to $Z(\phi)$ using the opening w_{σ^1} . In the case where $L \in \mathcal{L}_0$, ϕ is an instance of $OR(SAT)^U$, and thus, by the assumption, the distinguisher will retrieve a witness with probability at least $\frac{1}{q(n, m)}$. On the other hand, if $L \in \mathcal{L}_1$, then the instance ϕ is a general instance of $OR(SAT)$ (without the promise of the unique witness). Lemma 6.7 states that there exists a ϕ for which the witness-retrieval succeeds with probability at most $\frac{p(n, m)}{m}$. A more careful inspection of the proof of Lemma 6.7 shows that this statement also holds for a randomly chosen ϕ (generated by choosing m random commitments, not all of which are to 0). Thus, if

$L \in \mathcal{L}_1$, then the witness-retrieval succeeds on ϕ with probability at most $\frac{p(n,m)}{m}$ (with probability taken over the choice of $L \in \mathcal{L}_1$ and the randomness of the distinguisher). Overall, the distinguisher accepts with probability at least $\frac{1}{q(n,m)}$ when L is from \mathcal{L}_0 and at most $\frac{p(n,m)}{m}$ when L is from \mathcal{L}_1 . So if $\frac{1}{q(n,m)} - \frac{p(n,m)}{m}$ is larger than a polynomial fraction in n , then this procedure has a distinguishing advantage between \mathcal{L}_0 and \mathcal{L}_1 , contradicting the security of the commitment scheme. \square

All our results have been stated for the language $OR(SAT)$. However, they may be applied for other languages such as SAT and Clique. In particular, we get the statement with respect to SAT as a corollary (since a compression for SAT can be used as a compression for $OR(SAT)$ via the same reduction as in Lemma 2.17).

COROLLARY 6.9. *Suppose one-way functions exist, and let (Z, W) be a witness-retrievable compression for SAT (or for Unique-SAT) such that for every input ϕ with parameters m, n the following holds:*

1. *The compression parameter $|Z(\phi)| \leq p(n, m)$.*
2. *The success probability of W is at least $\frac{1}{q(n,m)}$, where probability is over the random coins of Z and W as well as the choice of the witness.*

Then $q(n, m) \geq \Omega(\frac{m}{p(n,m)})$.

7. Discussion and open problems.

7.1. Discussion—A unified perspective of the applications. In sections 3.1, 4, and 6 we presented three separate applications of compression that have a similar flavor: A CRH from one-way functions using perfect compression (section 3.1), a CRH/one-way function from hard on average language using perfect/imperfect compression (section 4), and PIR/OT from a one-way function using witness-retrievable compression (section 6). These constructions have a common underlying principle and can be viewed as variants on this main theme. The basic observation is that compression of $OR(L)$, where L is a “hard on average” language, can be used to construct PIR protocols in which the receiver is unbounded. This construction follows by generalizing a standard approach in the design of PIR protocols (e.g., [59]). In this method the receiver generates a sequence of n commitments hiding the characteristic vector of its selection, and the server computes an encoding of the XOR (alternatively, OR) of all of the committed values which correspond to the 1-entries of the database. When decoded, this value amounts to the bit that the receiver was seeking. The nontriviality in the PIR protocol stems from the fact that the length of the latter encoding can be made shorter than the length of the database. Typically this is achieved by using homomorphic properties of specific commitment schemes. In our case, this is achieved via the compression of $OR(L)$ (where L is the language defined by the commitment scheme). Thus the use of compression here can be viewed as a relaxation of the traditional use of homomorphic commitments.

The result of section 3.1 follows from this general scheme combined with the observation that PIR with an unbounded receiver implies CRH (via the reduction of [53]). Section 6 observes that the receiver in the PIR protocol can be made efficient if the underlying compression is witness-retrievable. The results of section 4 follow by further observing that the CRH construction doesn’t require the committed vector to be known to anyone, and moreover this construction remains collision resistant even if the committed vector is uniformly random (otherwise one could break the semantic security of the commitment). Thus the commitments can be replaced by random instances of a hard on average language. When compression is imperfect, the CRH is relaxed to a distributional variant which still implies a one-way function.

7.2. Future directions and open problems. The issue of compressibility and the corresponding classification introduced in this work raise many open problems and directions. The obvious one is to come up with a compression algorithm for a problem like SAT or Clique (or some \mathcal{VC}_{OR} -complete or hard problem). Note that the new impossibility results of Fortnow and Santhanam [36] do not rule out the possibility of error-prone compression for these languages. We have seen compressibility of some interesting \mathcal{NP} languages, and hence the question is, Where exactly is the boundary between compressibility and incompressibility? We tend to conjecture that it is in the low levels of the \mathcal{VC} hierarchy. We view PCP amplification methods such as the recent result of Dinur [23] as potential leads toward achieving compression. This is because these results show a natural amplification of properties on a graph and could potentially be combined with a simple compression of promise problems (such as the example for GapSAT in section 2.10). The main issue is doing the PCP amplification without introducing many new variables. Due to the recent results of [36] and [14], the underlying PCP in such an approach must also introduce some level of errors.

In particular, the following task would suffice for achieving nontrivial compression: given CNF formulae ϕ_1 and ϕ_2 (not necessarily with short witnesses), come up with a CNF formula ϕ that (1) satisfiability of the new formula coincides with very high probability with the satisfiability of $\phi_1 \vee \phi_2$, and (2) shorter than the combined lengths of ϕ_1 and ϕ_2 . By shorter, we mean of length $(1 - \epsilon)(|\phi_1| + |\phi_2|)$. The reason this is sufficient is that we can apply it recursively and obtain nontrivial compression for $OR(SAT)$, which implies the cryptographic applications.

Short of showing a compression for general complexity classes, it would be interesting to come up with further interesting compression algorithms as well as to obtain more hardness results. For instance, is Clique or any other embedding problem complete for \mathcal{VC}_1 ? Is there a natural and simple complete problem for \mathcal{VC}_1 ? Also, the \mathcal{VC} hierarchy is by no means the ultimate classification with respect to compressibility. One can hope to further refine this classification, especially within the confines of \mathcal{VC}_1 . Moreover, it would be interesting to find connections of the \mathcal{VC} hierarchy to other classifications (e.g., in the style of Feige [31] for average case complexity and approximation algorithms and Chen et al. [12] for parameterized complexity and subexponential algorithms).

Regarding the cryptographic application of getting a CRH from one-way functions (Theorem 3.1), one issue, Is how essential is the requirement that the compression will be errorless (this question is even more interesting due to the new impossibility results of [36])? We know that this requirement can be relaxed to hold with an error that is exponentially small in m . However, it is unknown whether a CRH can be constructed from any one-way function using a compression algorithm that errs with probability that is, say, exponentially small in n and $\log m$. Note that using typical amplification techniques for CRH are unsuccessful. For example, taking a concatenation of several independently chosen hash functions on the same input fails, since reducing the adversary's success probability to allow using the a union bound requires using too many ($\Omega(m)$) independent functions for the overall hash to still be shrinking. Another question in this regard is whether compression of languages outside of \mathcal{NP} is possible. For example, applications such as the construction of a CRH (in sections 3.1 and 4) can work also with compression of the language $AND(L)$ (which may not have a short witness) or $XOR(L)$ (not in \mathcal{NP}) rather than $OR(L)$.

Especially in light of the apparent hardness of compression, it is valuable to understand what are the implications of *incompressibility*. We have seen the necessity of incompressibility for the security of schemes in the hybrid BSM. Other examples

are the previously mentioned works of Dubrov and Ishai [26] regarding derandomization and Dziembowski [28] with respect to forward-secure storage. In order to gain confidence in an incompressibility assumption when used in a cryptographic setting, it is important to come up with an *efficiently falsifiable* assumption²⁶ of this nature (see [68]).

Finally we feel that we have just scratched the surface of an important topic, and in the future there will be other implications of compressibility or the impossibility of compression, whether in cryptography or in other areas.

Acknowledgments. We thank Yuval Ishai for many helpful comments and specifically for pointing out that the CRH construction does not require witness-retrievability. We are also grateful to Alon Rosen, Ronen Shaltiel, and Gillat Kol for their comments on the presentation and Hovav Shacham for conversations regarding witness-retrievable compression. Finally we thank the anonymous referees for FOCS and SICOMP, Salil Vadhan and Mike Langston for their helpful comments and suggestions, and Mike Fellows for pointing out some references.

REFERENCES

- [1] S. AARONSON, *NP-complete problems and physical reality*, SIGACT News, 36 (2005), pp. 30–52.
- [2] N. ALON, R. YUSTER, AND U. ZWICK, *Color-coding*, J. ACM, 42 (1995), pp. 844–856.
- [3] Y. AUMANN, Y.Z. DING, AND M.O. RABIN, *Everlasting security in the bounded storage model*, IEEE Trans. Inform. Theory, 48 (2002), pp. 1668–1680.
- [4] Y. AUMANN AND M.O. RABIN, *Information theoretically secure communication in the limited storage space model*, in Advances in Cryptology–CRYPTO ’99, Lecture Notes in Comput. Sci. 1666, Springer, New York, 1999, pp. 65–79.
- [5] B. BARAK, *How to go beyond the black-box simulation barrier*, in Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, 2001, pp. 106–115.
- [6] M. BELLARE, A. BOLDYREVA, AND A. PALACIO, *An uninstantiable random-oracle-model scheme for a hybrid-encryption problem*, in Advances in Cryptology – EUROCRYPT ’2004, Lecture Notes in Comput. Sci. 3027, 2004, Springer, New York, pp. 171–188.
- [7] S. BEN-DAVID, B. CHOR, O. GOLDBREICH, AND M. LUBY, *On the theory of average case complexity*, J. Comput. System Sci., 44 (1992), pp. 193–219.
- [8] J. BUSS AND T. ISLAM, *Simplifying the weft hierarchy*, Theoret. Comput. Sci., 351 (2006), pp. 303–313.
- [9] C. CACHIN AND U. MAURER, *Unconditional security against memory-bound adversaries*, in Advances in Cryptology – CRYPTO ’97, Lecture Notes in Comput. Sci. 1294, Springer, New York, 1997, pp. 292–306.
- [10] L. CAI AND J. CHEN, *On the amount of nondeterminism and the power of verifying*, SIAM J. Comput., 26 (1997), pp. 733–750.
- [11] R. CANETTI, O. GOLDBREICH, AND S. HALEVI, *The random oracle methodology, revisited*, J. ACM, 51 (2004), pp. 557–594.
- [12] J. CHEN, B. CHOR, M. FELLOWS, X. HUANG, D. JUEDES, I. KANJ, AND G. XIA, *Tight lower bounds for certain parameterized NP-hard problems*, Inform. and Comput., 201 (2005), pp. 216–231.
- [13] J. CHEN, I. KANJ, AND W. JIA, *Vertex cover: Further observations and further improvements*, J. Algorithms, 41 (2001), pp. 280–301.
- [14] Y. CHEN AND M. MÜLLER, *SAT is unlikely to be compressible*, manuscript, 2007.
- [15] B. CHOR, M. FELLOWS, AND D. JUEDES, *Linear kernels in linear time, or how to save k colors in $O(n^2)$ steps*, in WG 04, Lecture Notes in Comput. Sci. 3353, Springer-Verlag, New York, 2004, pp. 257–269.
- [16] S.A. COOK, *The complexity of theorem-proving procedures*, in Proceedings of the 3rd ACM Symposium on the Theory of Computing, 1971, pp. 151–158.
- [17] I. DAMGÅRD, *A design principle for hash functions*, in Advances in Cryptology - CRYPTO ’89, Lecture Notes in Comput. Sci. 435, Springer, New York, 1989, pp. 416–427.

²⁶An efficiently falsifiable assumption is one for which it is possible to create verifiable challenges so that if the assumption is false, then the challenges can be solved.

- [18] I. DAMGÅRD, T. PEDERSEN, AND B. PFITZMANN, *On the existence of statistically hiding bit commitment schemes and fail-stop signatures*, in Advances in Cryptology - CRYPTO '93, Lecture Notes in Comput. Sci. 773, Springer, New York, 1993, pp. 250–265.
- [19] H. DELL AND D. VAN MELKEBEEK, *Satisfiability Allows No Nontrivial Sparsification Unless the Polynomial-time Hierarchy Collapses*, manuscript.
- [20] G. DI CRESCENZO, T. MALKIN, AND R. OSTROVSKY, *Single database private information retrieval implies oblivious transfer*, in Advances in Cryptology – EUROCRYPT '2000, Lecture Notes in Comput. Sci. 1807, Springer, New York, 2000, pp. 122–138.
- [21] W. DIFFIE AND M.E. HELLMAN, *New directions in cryptography*, IEEE Trans. Inform. Theory, 22 (1976), pp. 644–654.
- [22] Y.Z. DING AND M.O. RABIN, *Hyper-encryption and everlasting security*, in Annual Symposium on Theoretical Aspects of Computer Science (STACS), Lecture Notes in Comput. Sci. 2285, Springer, New York, 2002, pp. 1–26.
- [23] I. DINUR, *The PCP theorem by gap amplification*, J. ACM, 54 (2007), article 12.
- [24] R. DOWNEY AND M. FELLOWS, *Parameterized Complexity*, Springer, New York, 1999.
- [25] R. DOWNEY, M. FELLOWS, AND U. STEGE, *Parameterized complexity: A systematic method for confronting computational intractability*, in Contemporary Trends in Discrete Mathematics, AMS DIMACS Proc. Ser. 49, AMS, Providence, RI, 1999, pp. 49–100.
- [26] B. DUBROV AND Y. ISHAI, *On the randomness complexity of efficient sampling*, in Proceedings of the 38th ACM Symposium on the Theory of Computing, 2006, pp. 711–720.
- [27] C. DWORK, J. LOTSPIECH, AND M. NAOR, *Digital signets: Self-enforcing protection of digital information*, in Proceedings of the 28th ACM Symposium on the Theory of Computing, 1996, pp. 489–498.
- [28] S. DZIEMBOWSKI, *On forward-secure storage*, in Advances in Cryptology – CRYPTO '06, Lecture Notes in Comput. Sci. 4117, Springer, New York, 2006, pp. 251–270.
- [29] S. DZIEMBOWSKI AND U. MAURER, *On generating the initial key in the bounded-storage model*, in Advances in Cryptology – EUROCRYPT 2004, Lecture Notes in Comput. Sci. 3027, Springer, New York, 2004, pp. 126–137.
- [30] S. DZIEMBOWSKI AND U. MAURER, *Optimal randomizer efficiency in the bounded-storage model*, J. Cryptology, 17 (2004), pp. 5–26.
- [31] U. FEIGE, *Relations between average case complexity and approximation complexity*, in Proceedings of the 34th ACM Symposium on the Theory of Computing, 2002, pp. 534–543.
- [32] U. FEIGE AND J. KILIAN, *On limited versus polynomial nondeterminism*, Chicago J. Theoret. Comput. Sci., 1997 (1997), pp. 1–20.
- [33] M. FISCHLIN, *On the impossibility of constructing non-interactive statistically-secret protocols from any trapdoor one-way function*, in Topics in Cryptology–CT-RSA 2002, The Cryptographer's Track at the RSA Conference, 2002, pp. 79–95.
- [34] J. FLUM AND M. GROHE, *Parameterized Complexity Theory*, Springer, New York, 2006.
- [35] J. FLUM, M. GROHE, AND M. WEYER, *Bounded fixed-parameter tractability and $\log^2 n$ nondeterministic bits*, in Proceedings of the 31st International Colloquium on Automata, Languages and Programming (ICALP) 2004, Lecture Notes in Comput. Sci. 3142, Springer, New York, 2004, pp. 555–567.
- [36] L. FORTNOW AND R. SANTHANAM, *Infeasibility of instance compression and succinct PCPs for NP*, in Proceedings of the 40th ACM Symposium on the Theory of Computing, ACM, New York, 2008, pp. 133–142.
- [37] O. GOLDREICH, *Foundations of Cryptography*, Cambridge University Press, London, 2001.
- [38] O. GOLDREICH, *Foundations of Cryptography - Volume 2*, Cambridge University Press, London, 2004.
- [39] O. GOLDREICH, S. MICALI, AND A. WIGDERSON, *Proofs that yield nothing but their validity, or all languages in NP have zero-knowledge proof systems*, J. ACM, 38 (1991), pp. 691–729.
- [40] J. GOLDSMITH, M. LEVY, AND M. MUNDHENK, *Limited nondeterminism*, SIGACT News, 27 (1996), pp. 20–29.
- [41] S. GOLDWASSER AND Y. TAUMAN KALAI, *On the (in)security of the Fiat-Shamir paradigm*, in Proceedings of the 44th IEEE Symposium on Foundations of Computer Science, 2003, pp. 102–111.
- [42] I. HAITNER, J. HOCH, O. REINGOLD, AND G. SEGEV, *Finding collisions in interactive protocols – a tight lower bound on the round complexity of statistically-hiding commitments*, in Proceedings of the 48th IEEE Symposium on Foundations of Computer Science, 2007, pp. 669–679.
- [43] I. HAITNER, O. HORVITZ, J. KATZ, C. KOO, R. MORSELLI, AND R. SHALTIEL, *Reducing complexity assumptions for statistically-hiding commitment*, in Advances in Cryptology – EUROCRYPT '2005, Lecture Notes in Comput. Sci. 3494, Springer, New York, 2005, pp. 58–77.

- [44] I. HAITNER AND O. REINGOLD, *Statistically-hiding commitment from any one-way function*, in Proceedings of the 39th ACM Symposium on the Theory of Computing, 2007, pp. 1–10.
- [45] D. HARNIK AND M. NAOR, *On everlasting security in the hybrid bounded storage model*, in Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP) 2006, Part II, Lecture Notes in Comput. Sci. 4052, Springer, New York, 2006, pp. 192–203.
- [46] D. HARNIK AND M. NAOR, *On the compressibility of NP instances and cryptographic applications*, in Electron. Colloq. Comput. Complex., TR06-022, 2006.
- [47] J. HÅSTAD, R. IMPAGLIAZZO, L.A. LEVIN, AND M. LUBY, *A pseudorandom generator from any one-way function*, SIAM J. Comput., 29 (1999), pp. 1364–1396.
- [48] C. HSIAO AND L. REYZIN, *Finding collisions on a public road, or do secure hash functions need secret coins?*, in Advances in Cryptology – CRYPTO '04, Lecture Notes Comput. Sci. 3152, Springer, New York, 2004, pp. 92–105.
- [49] R. IMPAGLIAZZO, *A personal view of average-case complexity*, in Proceedings of the 10th Annual Structure in Complexity Theory Conference, 1995, IEEE Computer Society Press, pp. 134–147.
- [50] R. IMPAGLIAZZO AND M. LUBY, *One-way functions are essential for complexity based cryptography*, in Proceedings of the 30th IEEE Symposium on Foundations of Computer Science, 1989, pp. 230–235.
- [51] R. IMPAGLIAZZO, R. PATURI, AND F. ZANE, *Which problems have strongly exponential complexity?*, in Proceedings of the 39th IEEE Symposium on Foundations of Computer Science, 1998, pp. 653–663.
- [52] R. IMPAGLIAZZO AND S. RUDICH, *Limits on the provable consequences of one-way permutations*, in Proceedings of the 21st ACM Symposium on the Theory of Computing, 1989, pp. 44–61.
- [53] Y. ISHAI, E. KUSHILEVITZ, AND R. OSTROVSKY, *Sufficient conditions for collision-resistant hashing*, in 2nd Theory of Cryptography Conference (TCC '05), Lecture Notes in Comput. Sci. 3378, Springer, New York, 2005, pp. 445–456.
- [54] H. KAPLAN, R. SHAMIR, AND R. TARJAN, *Tractability of parameterized completion problems on chordal, strongly chordal, and proper interval graphs*, SIAM J. Comput., 28 (1999), pp. 1906–1922.
- [55] R. KARP, *Reducibility among combinatorial problems*, in Complexity of Computer Computations, R. Miller and J. Thatcher, eds., Plenum Press, New York, 1972, pp. 85–103.
- [56] R. KARP AND M. RABIN, *Efficient randomized pattern-matching algorithms*, IBM J. Res. Dev., 31 (1987), pp. 249–260.
- [57] J. KILIAN, *A note on efficient zero-knowledge proofs and arguments*, in Proceedings of the 24th ACM Symposium on the Theory of Computing, 1992, pp. 723–732.
- [58] C. KINTALA AND P. FISCHER, *Refining nondeterminism in relativized polynomial-time bounded computations*, SIAM J. Comput., 9 (1980), pp. 46–53.
- [59] E. KUSHILEVITZ AND R. OSTROVSKY, *Replication is not needed: Single database, computationally-private information retrieval*, in Proceedings of the 38th IEEE Symposium on Foundations of Computer Science, 1997, pp. 364–373.
- [60] M. LI AND P. VITÁNYI, *An Introduction to Kolmogorov Complexity and Its Applications*, 2nd ed., Springer, New York, 1997.
- [61] C. LU, *Encryption against space-bounded adversaries from on-line strong extractors*, J. Cryptology, 17 (2004), pp. 27–42.
- [62] U. MAURER, *Conditionally-perfect secrecy and a provably-secure randomized cipher*, J. Cryptology, 5 (1992), pp. 53–66.
- [63] U. MAURER, *Secret key agreement by public discussion*, IEEE Trans. Inform. Theory, 39 (1993), pp. 733–742.
- [64] U. MAURER, R. RENNER, AND C. HOLENSTEIN, *Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology*, in Proceedings of the 1st Theory of Cryptography Conference (TCC '04), Lecture Notes in Comput. Sci. 2951, Springer, New York, 2004, pp. 21–39.
- [65] S. MICALI, *CS proofs*, in Proceedings of the 35th IEEE Symposium on Foundations of Computer Science, 1994, pp. 436–453.
- [66] J. NAOR AND M. NAOR, *Small-bias probability spaces: Efficient constructions and applications*, SIAM J. Comput., 22 (1993), pp. 838–856.
- [67] M. NAOR, *Bit commitment using pseudorandomness*, J. Cryptology, 4 (1991), pp. 151–158.
- [68] M. NAOR, *On cryptographic assumptions and challenges*, in Advances in Cryptology – CRYPTO '03, Lecture Notes in Comput. Sci. 2729, Springer, New York, 2003, pp. 96–109.

- [69] M. NAOR, R. OSTROVSKY, R. VENKATESAN, AND M. YUNG, *Perfect zero-knowledge arguments for NP using any one-way permutation*, J. Cryptology, 11 (1998), pp. 87–108.
- [70] M. NAOR AND M. YUNG, *Universal one-way hash functions and their cryptographic applications*, in Proceedings of the 21st ACM Symposium on the Theory of Computing, 1989, pp. 33–43.
- [71] M. NGUYEN, S. ONG, AND S. VADHAN, *Statistical zero-knowledge arguments for NP from any one-way function*, in Proceedings of the 47th IEEE Symposium on Foundations of Computer Science, 2006, pp. 3–14.
- [72] R. NIEDERMEIER, *Invitation to Fixed Parameter Algorithms*, Oxford University Press, London, 2006.
- [73] J.B. NIELSEN, *Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case*, in Advances in Cryptology – CRYPTO '02, Lecture Notes in Comput. Sci. 2442, Springer, New York, 2002, pp. 111–126.
- [74] J.B. NIELSEN, *On protocol security in the cryptographic model*, BRICS Diss. Ser. DS-03-8, University of Aarhus, Denmark, 2003.
- [75] C. PAPADIMITRIOU AND M. YANNAKAKIS, *Optimization, approximation, and complexity classes*, in Proceedings of the 20th ACM Symposium on the Theory of Computing, 1988, pp. 229–234.
- [76] C. PAPADIMITRIOU AND M. YANNAKAKIS, *On limited nondeterminism and the complexity of the V-C dimension*, J. Comput. System Sci., 53 (1996), pp. 161–170.
- [77] D. SIMON, *Finding collisions on a one-way street: Can secure hash functions be based on general assumptions?*, in Advances in Cryptology – EUROCRYPT 1998, Lecture Notes in Comput. Sci. 1403, Springer, New York, 1998, pp. 334–345.
- [78] L. TREVISAN, S. VADHAN, AND D. ZUCKERMAN, *Compression of samplable sources*, Comput. Complexity, 14 (2005), pp. 186–227.
- [79] S. VADHAN, *Constructing locally computable extractors and cryptosystems in the bounded storage model*, J. Cryptology, 17 (2004), pp. 43–77.
- [80] L. VALIANT AND V. VAZIRANI, *NP is as easy as detecting unique solutions*, Theoret. Comput. Sci., 47 (1986), pp. 85–93.
- [81] H. WEE, *On pseudoentropy versus compressibility*, in Proceedings of the IEEE Conference on Computational Complexity, 2004, pp. 29–41.
- [82] H. WEE, *On obfuscating point functions*, in Proceedings of the 37th ACM Symposium on Theory of Computing, 2005, pp. 523–532.
- [83] H. WEE, *Finding pessiland*, in Theory of Cryptography Conference (TCC '06), Lecture Notes in Comput. Sci. 3876, Springer, New York, 2006, pp. 429–442.