

On Average-Case Hardness in TFNP from One-Way Functions*

Pavel Hubáček¹, Chethan Kamath², Karel Král¹, and Veronika Slívová¹

¹ Charles University

{hubacek,kralka,slivova}@iuuk.mff.cuni.cz

² Northeastern University

ckamath@protonmail.com

Abstract. The complexity class TFNP consists of all NP *search* problems that are *total* in the sense that a solution is guaranteed to exist for all instances. Over the years, this class has proved to illuminate surprising connections among several diverse subfields of mathematics like combinatorics, computational topology, and algorithmic game theory. More recently, we are starting to better understand its interplay with cryptography.

We know that certain cryptographic primitives (e.g. one-way permutations, collision-resistant hash functions, or indistinguishability obfuscation) imply average-case hardness in TFNP and its important subclasses. However, its relationship with the most basic cryptographic primitive – i.e., one-way functions (OWFs) – still remains unresolved. Under an additional complexity theoretic assumption, OWFs imply hardness in TFNP (Hubáček, Naor, and Yagev, ITCS 2017). It is also known that average-case hardness in most structured subclasses of TFNP does not imply any form of cryptographic hardness in a black-box way (Rosen, Segev, and Shahaf, TCC 2017) and, thus, one-way functions might be sufficient. Specifically, no negative result which would rule out basing average-case hardness in TFNP *solely* on OWFs is currently known. In this work, we further explore the interplay between TFNP and OWFs and give the first negative results.

As our main result, we show that there cannot exist constructions of average-case (and, in fact, even worst-case) hard TFNP problem from OWFs with a certain type of *simple* black-box security reductions. The class of reductions we rule out is, however, rich enough to capture many of the currently known cryptographic hardness results for TFNP. Our results are established using the framework of black-box separations (Impagliazzo and Rudich, STOC 1989) and involve a novel application of the reconstruction paradigm (Gennaro and Trevisan, FOCS 2000).

1 Introduction

The complexity class TFNP of *total* search problems [MP91], i.e., with syntactically guaranteed existence of a solution for all instances, holds a perplexing place in the hierarchy of computational complexity classes. The standard method for arguing computational hardness in TFNP is via clustering these problems into subclasses characterised by the existential argument guaranteeing their totality [Pap94]. This approach was particularly successful in illuminating the connections between search problems in seemingly distant domains such as combinatorics, computational topology, and algorithmic game theory (see, for example, [DGP09, KPR⁺13, FG18, FG19, GH19] and the references therein). However, all results of this type ultimately leave open the possibility of the existence of polynomial time algorithms for all of TFNP.

An orthogonal line of work, which can be traced to Papadimitriou [Pap94], shows the existence of hard problems in subclasses of TFNP under cryptographic assumptions. Such conditional lower bounds for structured subclasses of TFNP were recently given under increasingly more plausible cryptographic assumptions [Jeř16, BPR15, GPS16, HY17, KS17, CHK⁺19a, CHK⁺19b, EFKP20, BG20]. The end of the line

* This research was supported in part by the Grant Agency of the Czech Republic under the grant agreement no. 19-27871X, by the Charles University projects PRIMUS/17/SCI/9, UNCE/SCI/004, and GAUK 1568819, and by the Charles University grant SVV-2017-260452. The second author is supported by the IARPA grant IARPA/2019-19-020700009.

in this sequence of results would correspond to a “dream theorem” establishing average-case hardness in one of the lower classes in the TFNP hierarchy (e.g. CLS [DP11]) under some weak general cryptographic assumptions (e.g. the existence of one-way functions).

An informative parallel for the limits of this methodology can be drawn by considering average-case hardness of *decision* problems (i.e., languages) in $\text{NP} \cap \text{coNP}$. The existence of a hard-on-average decision problem in $\text{NP} \cap \text{coNP}$ follows from the existence of hard-core predicates for any one-way permutation [GL89]. However, the existence of injective one-way functions is insufficient for black-box constructions of hard-on-average distributions for languages in $\text{NP} \cap \text{coNP}$ even assuming indistinguishability obfuscation [BDV17] (in fact, [BDV17] ruled out even black-box constructions of worst-case hardness in $\text{NP} \cap \text{coNP}$ using these cryptographic primitives).

For total *search* problems, the existence of hard-on-average TFNP distributions is straightforward either from one-way permutations or collision-resistant hash functions. Moreover, there exist constructions of hard-on-average TFNP distributions either assuming indistinguishability obfuscation and one-way functions [BPR15, GPS16] or under derandomization-style assumptions and one-way functions [HNY17]. On the other hand, no analogue of the impossibility result for basing average-case hardness in $\text{NP} \cap \text{coNP}$ on (injective) one-way functions [BDV17] is currently known for TFNP. Rosen, Segev, and Shahaf [RSS17] showed that most of the known structured subclasses of TFNP do not imply (in a black-box way) any form of cryptographic hardness; thus, it is plausible that hard-on-average distributions in TFNP can be based *solely* on the existence of one-way functions.

Rosen et al. [RSS17] also provided some insight into the structure of hard-on-average distributions in TFNP. They showed that any hard-on-average distribution of a TFNP problem from any primitive which exists relative to a random injective trapdoor function oracle (e.g. one-way functions, injective trapdoor functions, or collision-resistant hash functions) must result in instances with a nearly exponential number of solutions. Even though the [RSS17] result restricts the structure of hard-on-average distributions in TFNP constructed from these cryptographic primitives, it certainly does not rule out their existence. Indeed, a collision-resistant hash function constitutes a hard-on-average TFNP distribution, albeit with an exponential number of solutions.

Motivated by the significant gap between negative and positive results, we revisit the problem of existence of average-case hardness in TFNP under weak general cryptographic assumptions and address the following question:

Are (injective) one-way functions sufficiently structured to imply hard-on-average total search problems?

Towards answering this question, we provide negative results and show that *simple* fully black-box constructions of hard-on-average TFNP distributions from injective one-way functions do not exist.

1.1 Our Results

We recall the details of the construction of a hard-on-average distribution in TFNP from one-way permutations to highlight the restrictions on the type of reductions considered in our results.

Consider the total search problem PIGEON, in which you are given a length-preserving n -bit function represented by a Boolean circuit C and are asked to find either a preimage of the all-zero string (i.e., $x \in \{0, 1\}^n : C(x) = 0^n$) or a non-trivial collision (i.e., $x \neq x' \in \{0, 1\}^n : C(x) = C(x')$). PIGEON is complete for a subclass of TFNP known as PPP, and Papadimitriou [Pap94] gave the following construction of a hard PIGEON problem from one-way permutations. Given a (one-way) permutation $f: \{0, 1\}^n \rightarrow \{0, 1\}^n$ and a challenge $y \in \{0, 1\}^n$ for inversion under f , the reduction algorithm defines an instance of PIGEON by the oracle-aided circuit C_y^f computing the function $C_y^f(x) = f(x) \oplus y$. It is not hard to see that the instance of PIGEON C_y^f has a unique solution corresponding to the preimage of y under f and, therefore, any algorithm solving it breaks the one-wayness of f .

Note that the above construction of a hard (on average) TFNP problem is extremely simple in various aspects:

- The construction is *fully black-box*, i.e., the PIGEON instance can be implemented via an oracle-aided circuit treating the one-way permutation as a black-box and the reduction inverts when given oracle access to an arbitrary solver for PIGEON.
- The reduction is *many-one*, i.e., a single call to a PIGEON-solving oracle suffices for finding the preimage of y .
- The reduction is *f -oblivious*, i.e., the oracle-aided circuit C_y^f defining the PIGEON instance depends only on the challenge y and does not depend on the one-way permutation f in the sense that C_y^f itself can be fully specified without querying f . In other words, given the challenge y , the instance C_y^f submitted to the PIGEON oracle by the reduction is, as an oracle-aided circuit, identical for all permutations f .
- The reduction is *deterministic*, i.e., it simply passes y to specify the PIGEON instance.

Such a fully black-box construction of PIGEON with a deterministic f -oblivious many-one reduction exists also assuming collision-resistant hash functions exist (folklore). Specifically, for any hash function $h: \{0, 1\}^n \rightarrow \{0, 1\}^{n-1}$ from the collision-resistant family, the PIGEON instance is defined as $C^h(x) = h(x) \| 1$, where $\|$ represents the operation of string concatenation. Since C^h concatenates the value $h(x)$ with 1 for any input x , it never maps to the all-zero string and, therefore, has the same collisions as h . Note that, unlike in the above construction from one-way permutations, the instances resulting from collision-resistant hash functions do not have a unique solution. In fact, there are always at least 2^{n-1} nontrivial collisions (even in two-to-one functions where each $y \in \{0, 1\}^{n-1}$ has exactly two preimages) and this structural property is inherent as shown by Rosen et al. [RSS17]. Importantly, the property of having nearly exponentially many solutions is not in contradiction with the resulting distribution being hard-on-average. Currently, there is no actual evidence suggesting that average-case hardness in TFNP cannot be based on the existence of injective one-way functions.

The above constructions motivate us to study whether there exist such “simple” constructions of an average-case hard TFNP problem under weaker cryptographic assumptions such as the existence of injective one-way functions, and we answer this question in negative (see Section 3.2 for the formal statement of Theorem 1).

Theorem 1 (Main Theorem - Informal). *There is no efficient fully black-box construction of a worst-case hard TFNP problem from injective one-way functions with a randomized f -oblivious non-adaptive reduction.*

Thus, we actually rule out a larger class of fully black-box constructions with reductions to injective one-way functions than the *deterministic f -oblivious many-one* reductions from the motivating examples of *average-case hardness* of PIGEON from one-way permutations, respectively collision-resistant hash functions. We rule out even constructions of *worst-case hard* TFNP problems using *randomized f -oblivious non-adaptive*³ reductions. The formal definitions of fully black-box constructions with f -oblivious non-adaptive reductions are given in Section 3.1 (see Definition 2 and Definition 3).

Even though restricted, our results are the first step towards the full-fledged black-box separation of TFNP and (injective) one-way functions. We note that the full-fledged separation would necessarily subsume the known separation of collision-resistant hash functions and injective one-way functions [Sim98], for which, despite the recent progress, there are only non-trivial proofs [MM11, HRS15, BD19].

1.2 Our Techniques

Our results employ the framework of black-box separations [IR89, RTV04, Fis12]. The approach suggested in [IR89] for demonstrating that there is no fully black-box construction of a primitive P from another primitive Q is to come up with an oracle O relative to which Q exists, but every black-box implementation C^Q of P is broken. However, as explained in [RTV04, MM11], this approach actually rules out a larger class of constructions (so-called “relativized” constructions), and to rule out just fully black-box constructions it suffices to use the so-called two-oracle technique [HR04]. Here, the oracle O usually consists of two parts:

³ Reductions which can ask *multiple queries in parallel* to the TFNP oracle.

an idealised implementation of the primitive Q and a “breaker” oracle for primitive P . In our context, P corresponds to a TFNP problem and the oracle O comprises of a random injective function (which yields an injective one-way function) and a procedure SOLVE which provides a solution for any instance of a TFNP problem. To rule out the existence of fully black-box constructions of hard-on-average TFNP problems from injective one-way functions, one then has to argue that access to such a “breaker” oracle SOLVE for TFNP does not help any reduction R in inverting injective one-way functions. Designing such a SOLVE oracle and then arguing that it does not help inverting injective one-way function, which we carry out using the reconstruction paradigm of Gennaro and Trevisan [GT00], constitute the main technical challenges. Before giving an overview of these two steps, we explain the structural insight that is key to our separation, and guided us in the design of the two steps.

The existence of a “useless” solution. At the core of our negative result is a new structural insight about TFNP instances constructed from (injective) one-way functions. Observe that any one-way function gives rise to a search problem with a hard-on-average distribution which is total over its support (but all instances outside its support have no solution). Specifically, for any one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n+1}$, an instance is any $y \in \{0, 1\}^{n+1}$ and the solution for y is any $x \in \{0, 1\}^n$ such that $f(x) = y$. The hard-on-average distribution then corresponds to sampling x uniformly from $\{0, 1\}^n$ and outputting the instance $y = f(x)$ (as in the standard security experiment for one-way functions). When attempting to construct a hard search problem which is truly total and has a solution for all instances (not only for the support of the hard distribution), one has to face the frustrating obstacle in the form of “useless” solutions which do not help the reduction in inverting its challenge y . Note that, as the resulting TFNP problem must be total for all oracles f , there must exist a solution even for oracles with no preimage for the challenge y . By a simple probabilistic argument, it follows that for a random oracle f and a random challenge y , with overwhelming probability, there exists a solution to any TFNP instance which does not query a preimage of y , i.e., a “useless” solution from the perspective of the reduction.⁴

Thus, demonstrating a black-box separation would be straightforward if the TFNP solver knew which challenge y is the reduction attempting to invert. Our solver would simply output such a “useless” solution and we could argue via the reconstruction paradigm that no reduction can succeed in inverting y given access to our solver. In this work, we show that it is possible to construct a TFNP solver which returns such a “useless” solution with overwhelming probability even though the solver *does not know* the input challenge of the reduction.

Reduction-specific SOLVE. Note that a reduction in a fully black-box construction must succeed in breaking the primitive P when given access to *any* oracle SOLVE (see Definition 2). In other words, to rule out the existence of constructions with a fully black-box reduction, it is sufficient to show that for every reduction there exists a SOLVE which is not helpful in inverting; in particular, SOLVE may depend on the reduction. To enable SOLVE to answer the reduction’s query with a “useless” solution with overwhelming probability, we take exactly this approach and construct a reduction-specific SOLVE for any construction of a TFNP problem from injective one-way functions. We significantly differ in this respect from the previous works which relied on the reconstruction paradigm of Gennaro and Trevisan [GT00], e.g., the works which employed the collision-finding oracle of Simon [Sim98, HHR15, PV10, RS10, BKS11]. We note that the possibility of designing a breaker oracle which depends on the fully black-box construction was exploited already by Gertner, Malkin, and Reingold [GMR01], who considered SOLVE which depends on the implementation rather than the reduction algorithm (as in our case). That is, to rule out the construction of a primitive P from a primitive Q , they considered an oracle SOLVE that depends on the implementation C^Q of the primitive P , whereas in our case SOLVE depends on the reduction algorithm R that is supposed to break Q given access to

⁴ Note that the above argument fails in the case of one-way permutations, where the challenge $y \in \{0, 1\}^n$ is in the image for any permutation $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$. The construction of a TFNP problem then simply does not have to deal with the case when the challenge y is not in the image of f , and it can ensure that every solution is useful for inverting the challenge y . Indeed, the hard instances C_y^f of PIGEON from one-way permutations described in Section 1.1 have a unique solution on which C_y^f always queries a preimage of y under any f .

an algorithm that breaks C^Q . The possibility of proving black-box separations via reduction-specific oracles was also observed in the work of Hsiao and Reyzin [HR04] who, nevertheless, did not have to leverage this observation in their proofs.

On a high level, given that SOLVE can use the code of the reduction R , SOLVE can simulate R on all possible challenges y to identify the set of challenges on which R outputs the present instance that SOLVE needs to solve. As we show, the solution then can be chosen adversarially so that it avoids such solutions of interest to the reduction. To turn this intuition into a formal proof, one needs to show that our SOLVE indeed does not help in inverting injective one-way functions and we do so along the lines of the reconstruction paradigm of [GT00].

Applying the compression argument. Two important subtleties arise in the proof when we try to turn the reduction into a pair of compression and decompression algorithms, which we explain next. First, the reconstruction paradigm is conventionally applied to random permutations [GT00, HHS15], whereas the reduction R and the algorithm SOLVE are designed for random injective functions. The natural approach is to simply proceed with the same style of proof even in our setting. Specifically, one would presume that a similar incompressibility argument can be leveraged if we manage to somehow encode the image of the random injective function. While this intuition is correct in the sense that it allows *correct* compression and reconstruction, it turns out that the space required to encode the image is too prohibitive for reaching the desired contradiction with known information-theoretic lower bounds on the expected length of encoding for a random injective function. To resolve this issue, we construct compressor and decompressor algorithms for a random permutation, but we equip the algorithms with *shared randomness* in the form of a random injective function $h: \{0,1\}^n \rightarrow \{0,1\}^{n+1}$ independent of the random permutation $\pi: \{0,1\}^n \rightarrow \{0,1\}^n$ to be compressed. Whenever the compressor and decompressor need to provide the reduction or SOLVE with access to the injective function $f: \{0,1\}^n \rightarrow \{0,1\}^{n+1}$, they compose the permutation π with the shared injective function h and then pass off the composed injective function $f = h \circ \pi$ to the reduction. With this modification, we are able to show that any reduction which succeeds in inverting injective one-way functions given access to our SOLVE can be used to compress a random permutation on $\{0,1\}^n$ below a standard information-theoretic lower bound on the size of a prefix-free encoding of such random variable. We note that this is reminiscent of the approach used in [MM11] for separating *injective* one-way functions from one-way permutations.

Second, we cannot employ the actual oracle SOLVE in our compression and decompression algorithms: even though we can use the reduction when compressing and decompressing the random permutation, we must be able to *consistently* simulate SOLVE without accessing the whole permutation. In general, the choice of the “breaker” oracle that can be simulated efficiently without too many queries to the permutation (e.g., the collision finding oracle of Simon [Sim98, Bae14]) is a crucial part of the whole proof, and, a priori, it is unclear how to design a TFNP solver which also has such a property. Nevertheless, we show that there exists a SOLVE which can be efficiently simulated given only (sufficiently small) partial information about the permutation.

f-oblivious reductions. As our SOLVE simulates the reduction on possible challenges y , we need for technical reasons that the reduction is *f-oblivious* (namely, for correctness of our encoding and decoding algorithms). However, we believe that *f-obliviousness* is not overly restrictive as it is a natural property of security reductions. Besides the two fully black-box constructions of PIGEON with *f-oblivious* reductions described in section 1.1, *f-oblivious* security reductions appear also in the cryptographic literature – see for example the standard security reduction in the Goldreich-Levin theorem establishing the existence of hard-core predicate for any one-way function (note that this particular security reduction is also non-adaptive). An orthogonal notion of a π -oblivious construction appears in the work of Wee [Wee07]. However, it is the implementation of the constructed primitive which is “oblivious” to the one-way permutation π in his work.

1.3 Related Work

TFNP and its subclasses. The systematic study of total search problems was initiated by Megiddo and Papadimitriou [MP91] with the definition of complexity class TFNP. They observed that a “semantic” class such

as TFNP is unlikely to have complete problems, unless $\text{NP} = \text{coNP}$. As a resolution, Papadimitriou [Pap94] defined “syntactic” subclasses of TFNP with the goal of clustering search problems based on the various *existence theorems* used to argue their totality. Perhaps the best known such class is PPAD [Pap94], which captures the computational complexity of finding Nash equilibria in bimatrix games [DGP09, CDT09]. Other subclasses of TFNP include:

- PPA [Pap94], which captures computational problems related to the *parity argument* like Borsuk-Ulam theorem or fair division [FG19];
- PLS [JPY88], defined to capture the computational complexity of problems amenable to *local search* and its “continuous” counterpart $\text{CLS} \subseteq \text{PLS}$ [DP11] (in fact $\text{CLS} \subseteq \text{PLS} \cap \text{PPAD}$), which captures finding the computational complexity of finding (approximate) local optima of continuous functions and contains interesting problems from game theory (e.g., solving the simple stochastic games of Condon or Shapley); and
- PPP [Pap94] and PWPP \subseteq PPP [Jeř16], motivated by the *pigeonhole principle* and contain important problems related to finding collisions in functions.

The relative complexity of some of these classes was studied in [BCE⁺98] as it was shown using (worst-case) oracle separations that many of these classes are distinct.

Cryptographic hardness in TFNP. Hardness from standard cryptographic primitives was long known for the “higher” classes in TFNP like PPP and PPA. We have already mentioned that one-way permutations imply average-case hardness in PPP [Pap94] and existence of collision-resistant hashing (e.g. hardness of integer factoring or discrete-logarithm problem in prime-order groups) implies average-case hardness in PPP (as well as in PWPP). In addition, Jeřábek [Jeř16], building on the work of Buresh-Oppenheim [Bur06], showed that PPA is no easier than integer factoring.

However, it is only recently that we are better understanding the cryptographic hardness of the lower classes in TFNP. This was catalysed by the result of Bitansky et al. [BPR15] who reduced hardness in PPAD to indistinguishability obfuscation (and injective OWFs) expanding on Abbot, Kane, and Valiant [AKV04]; Hubáček and Yogev [HY17] extended this result to $\text{CLS} \subseteq \text{PLS} \cap \text{PPAD}$. The underlying assumption was relaxed further to cryptographic assumptions that are more plausible than indistinguishability obfuscation in [CHK⁺19a, CHK⁺19b, EFKP20]. Using similar ideas, Bitansky and Gerichter [BG20] presented a construction for hard-on-average distributions in the complexity class PLS in the random oracle model. Building on these results, a flurry of recent works [KPY20, KZ20, JK20, JKKZ20, LV20] further weakened the assumptions required for proving average-case hardness in CLS to sub-exponential hardness of learning with errors, bringing us closer to proving average-case hardness in CLS under a standard *concrete* cryptographic assumption.

One-way functions and TFNP. Hubáček et al. [HNY17] showed that average-case hardness in NP (which is implied by OWFs) implies average-case hardness in TFNP under complexity theoretical assumptions related to derandomization. Pass and Venkatasubramanian [PV19] recently complemented the [HNY17] result by showing that when OWFs *do not exist*, average-case hardness in NP implies average-case hardness in TFNP. However, a definitive relationship between OWFs and TFNP has remained elusive. This prompted Rosen et al. [RSS17] to explore impossibility of reducing TFNP hardness to OWFs. They gave a partial answer by showing that there do not exist hard-on-average distributions of TFNP instances over $\{0, 1\}^n$ with $2^{n^{o(1)}}$ solutions from any primitive which exists relative to a random injective trapdoor function oracle (e.g. one-way functions). Their main observation was that the argument in [Rud88], which separates one-way functions from one-way permutations, can be strengthened to separate other unstructured primitives from structured primitives (such as problems in TFNP). However, it seems that the [Rud88] argument has been exploited to its limits in [RSS17], and, therefore, it is not clear whether their approach can be extended to fully separate one-way functions and TFNP. Thus, the situation is contrasting to $\text{NP} \cap \text{coNP}$, the decision counterpart of TFNP, whose relationship with (injective) OWFs is much better studied. In particular, we know that hardness is implied by one way permutations but injective OWFs, even with indistinguishability obfuscation, (and, therefore, public-key encryption) cannot imply hardness in $\text{NP} \cap \text{coNP}$ in a black-box way [BDV17].

2 Preliminaries

Unless stated otherwise, all logarithms are base two. For $X \subseteq \{0, 1\}^*$, we use \overline{X} to denote the set $\{0, 1\}^* \setminus X$. For strings $x, y \in \{0, 1\}^*$, we use $x <_{\text{lex}} y$ or $y >_{\text{lex}} x$ to denote that x is lexicographically strictly smaller than y .

Notation 2 (Functions) Let $X, Y \subseteq \{0, 1\}^*$, $f: X \rightarrow Y$ be a function and $X' \subseteq X$ be a set.

1. $f \upharpoonright X'$ denotes the restriction of f on X' , i.e., the function $f': X' \rightarrow Y$ such that $\forall x \in X': f'(x) = f(x)$.
2. $\text{Dom}(f)$ denotes the domain of f , i.e., the set X .
3. $\text{Im}(f)$ denotes the image of f , i.e., the set $\{f(x) \mid x \in X\} \subseteq Y$.
4. $f[X']$ denotes the image of the restriction of f on X' , i.e., the set $\text{Im}(f \upharpoonright X')$.

Notation 3 (Injective functions) We denote by Inj_n^m the set of all injective functions from $\{0, 1\}^n$ to $\{0, 1\}^m$. For the special case when $n = m$ we get the set of all permutations on $\{0, 1\}^n$.

The set Inj is the set of all functions $f: \{0, 1\}^* \rightarrow \{0, 1\}^*$, such that f can be interpreted as a sequence $f = \left\{ f_n \mid f_n \in \text{Inj}_n^{m(n)} \right\}_{n \in \mathbb{N}}$ of injective functions, where $m: \mathbb{N} \rightarrow \mathbb{N}$ is an injective function such that for all $n \in \mathbb{N}$: $m(n) > n$ and $m(n) \leq 100n^{\log n}$.

We say that the function m is the type of f and we define the corresponding type operator $\tau: \text{Inj} \rightarrow (\mathbb{N} \rightarrow \mathbb{N})$ such that $\tau(f) = m$.

We denote the set of all possible types by \mathbb{T} , i.e.,

$$\mathbb{T} = \{\mu: \mathbb{N} \rightarrow \mathbb{N} \mid \exists f \in \text{Inj} \text{ such that } \tau(f) = \mu\}.$$

Through the paper f_n denotes the function $f \upharpoonright \{0, 1\}^n$ (i.e., restriction of f to the domain $\{0, 1\}^n$), where $f \in \text{Inj}$.

In our proofs, we often compose a function defined on all binary strings with a function defined only for binary strings of certain length; namely, we often want to compose a function from Inj with a permutation of n -bit strings. The desired resulting function should always be a function from all binary strings. For the ease of exposition, we extend the standard notation for function composition as follows.

Notation 4 (Function composition) Let X, Y, Z be any sets such that $X \subseteq Y$ and let $f: X \rightarrow Y$ and $g: Y \rightarrow Z$. We define the function $g \circ f: X \rightarrow Z$ as:

$$(g \circ f)(x) = \begin{cases} g(f(x)) & \text{if } x \in X, \\ g(x) & \text{if } x \in Y \setminus X. \end{cases}$$

Finally, we recall some basic information-theoretic results about prefix-free codes.

Definition 1 (Prefix-free code). A set of code-words $C \subseteq \{0, 1\}^*$ is a prefix-free code if there are no two distinct $c_1, c_2 \in C$ such that c_1 is a prefix (initial segment) of c_2 , i.e., for any two distinct $c_1, c_2 \in C$ there exists $0 \leq j < \min(|c_1|, |c_2|)$ such that $(c_1)_j \neq (c_2)_j$.

Proposition 1 (Theorem 5.3.1 in [CT12]). The expected length L of any prefix-free binary code for a random variable X is greater than or equal to the entropy $H(X)$.

Corollary 1. To encode a uniformly random permutation $\pi \in \text{Inj}_n^n$ using prefix-free encoding it takes at least $\log(2^n!)$ bits in expectation.

Proof. The entropy of a uniformly randomly chosen permutation from Inj_n^n is $\log(2^n!)$ as we choose uniformly at random from $2^n!$ distinct permutations. By Proposition 1, we get that the expected size of the encoding is at least $\log(2^n!)$. \square

3 Separating TFNP and Injective One-Way Functions

3.1 Fully Black-Box Construction of Hard TFNP Problem from iOWF

Below, we give a definition of fully black-box construction of a (worst-case) hard TFNP problem from an injective one-way function.

Definition 2 (Fully black-box construction of a worst-case hard TFNP problem from iOWF). A fully black-box construction of a worst-case hard TFNP problem from injective one-way function is a tuple (R, T_R, C, T_C, p) of oracle-aided algorithms R, C , functions T_R, T_C , and a polynomial p satisfying the following properties:

1. **R and C halt on all inputs:** For all $f \in \text{Inj}$, $n \in \mathbb{N}$, and $y, i, s \in \{0, 1\}^*$, the algorithm $R^f(1^n, y)$ halts in time $T_R(|y|)$, and the algorithm $C^f(i, s)$ halts in time $T_C(|i| + |s|)$.
2. **Correctness:** For all $f \in \text{Inj}$ and for all $i \in \{0, 1\}^*$, there exists $s \in \{0, 1\}^*$ such that $|s| \leq p(|i|)$ and $C^f(i, s) = 1$, i.e., for any instance of the TFNP problem there exists a solution of polynomial length.
3. **Fully black-box proof of security:** There exists a polynomial p' such that for all $f \in \text{Inj}$ and any oracle-aided algorithm SOLVE, if

$$\forall i \in \{0, 1\}^* : \text{SOLVE}^f(i) \text{ returns } s \text{ such that } C^f(i, s) = 1$$

then for infinitely many $n \in \mathbb{N}$,

$$\Pr_{x \leftarrow \{0, 1\}^n, R} [R^{f, \text{SOLVE}}(1^n, f(x)) = x] \geq \frac{1}{p'(n)} .$$

Definition 2 has the following semantics. The deterministic algorithm C specifies the TFNP problem and the algorithm R is the (security) reduction which, given access to any TFNP solver, breaks the security of any injective one-way function. For example in the case of the hard PIGEON problem from one-way permutations discussed in Section 1.1, C would be an algorithm which on input (C_y^f, x) , respectively (C_y^f, x, x') , outputs 1 if and only if $C_y^f(x) = 0^n$, respectively $C_y^f(x) = C_y^f(x')$. The reduction algorithm $R(1^n, y)$ simply queries the TFNP solver SOLVE with the instance $i = C_y^f$, i.e., a circuit computing the function $C_y^f(x) = f(x) \oplus y$, and outputs the solution s returned by SOLVE for which, by construction, $f(s) = y$.

Below, we provide some additional remarks on important points in the above definition.

Reduction-specific SOLVE. Let us emphasize the order of quantifiers restricting the security reduction in Definition 2:

$$\exists(R, T_R, C, T_C, p) \forall f \forall \text{SOLVE} : \text{SOLVE}^f \text{ solves the TFNP problem } C \implies R^{f, \text{SOLVE}} \text{ inverts } f .$$

Importantly, the reduction must invert when given access to *any* oracle SOLVE. As a consequence, in order to establish a separation, the above statement is negated and it suffices to show that for every reduction there exists a solver (see proof of [HR04, Proposition 1] for more details). Thus, in the proof an oracle separation, the oracle SOLVE may even depend on the behaviour of the reduction R , and, in particular, SOLVE can simulate the security reduction R on an arbitrary input. We exploit these properties in establishing our results.

Constructions of worst-case vs. average-case hardness in TFNP. Our Definition 2 considers constructions of a *worst-case* hard TFNP problem – the reduction has access to SOLVE which is promised to return a solution to *any* instance of the TFNP problem. To capture constructions of *average-case* hardness in TFNP, we would need to extend the construction with an efficiently sampleable distribution D of instances of the TFNP problem and require the reduction to invert when given access to any SOLVE that returns solutions for instances coming from the specific distribution D (see Definition 5.1 in [RSS17]). However, given that we are proving a black-box separation, showing impossibility for worst-case hardness is a stronger result.

The quality of SOLVE. Note that we consider security reductions which invert given access to SOLVE which outputs a solution with probability 1, whereas some definitions in the literature allow the reduction to work only with some non-negligible probability. This also makes our negative result stronger – it is potentially easier to give a reduction when given access to SOLVE which is guaranteed to always return a solution.

Direct and indirect queries to f . The security reduction R can learn something about f in various ways. It may query f directly or the information might be deduced from the solution of some queried instance of the TFNP problem returned by SOLVE. We introduce the following notation in order to distinguish where queries originate, which allows us to reason about the view the security reduction has over the function f in our proof of Theorem 1.

Notation 5 (Query sets Q) *We distinguish the following sets of queries to oracles depending on where these queries originated and which oracle is queried.*

- Let $Q(C^f(i, s))$ denote the set of all preimages $x \in \{0, 1\}^*$ on which the oracle f has been queried by C running on an input (i, s) .
- Let $Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, y))$ denote the set of all instances $i \in \{0, 1\}^*$ on which the oracle SOLVE has been queried by R running on a security parameter n and challenge y .
- Let $Q_f^{\text{dir}}(R^{f, \text{SOLVE}}(1^n, y))$ denote the set of preimages $x \in \{0, 1\}^*$ on which the oracle f has been queried by R running on an input y and security parameter n .
- Let $Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y))$ denote the set of all preimages $x \in \{0, 1\}^*$ on which the oracle f has been queried indirectly, i.e., it has been queried by C running on an input (i, s) where $i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, y))$ and $s = \text{SOLVE}^f(i)$.
- Let $Q_f(R^{f, \text{SOLVE}}(1^n, y)) = Q_f^{\text{dir}}(R^{f, \text{SOLVE}}(1^n, y)) \cup Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y))$.

Note that these sets may not be disjoint. When f is a partial function (i.e., when f is not defined on all inputs) the query set contains all queries queried up to the point of the first undefined answer and the query with the undefined answer is included as well.

Restrictions on the power of the reduction. We consider various restricted classes of security reductions as defined below.

Definition 3 (Deterministic/randomized, many-one/non-adaptive, f -oblivious reductions). *Let (R, T_R, C, T_C, p) be a fully black-box construction of a hard TFNP problem from injective one-way functions.*

We distinguish deterministic / randomized reductions. For a randomized security reduction, we extend the input of R to a triple $(1^n, y, r)$, where the meaning of n , resp. y , remains unchanged (i.e., n is the security parameter, y is the challenge), and $r \in \{0, 1\}^$ is the randomness of the security reduction.*

The security reduction R is many-one if for all $f \in \text{Inj}$, for any oracle SOLVE and for all $y \in \{0, 1\}^$, $R^{f, \text{SOLVE}}(1^n, y)$ makes a single query to the oracle SOLVE.*

The security reduction R is non-adaptive if for all $f \in \text{Inj}$, for any oracle SOLVE and for all $y \in \{0, 1\}^$, all the queries of $R^{f, \text{SOLVE}}(1^n, y)$ to the oracle SOLVE are submitted in parallel (i.e., the queries to SOLVE do not depend on the answers received from SOLVE).*

The security reduction R is f -oblivious if for all $y \in \{0, 1\}^$, for any oracle SOLVE, and any pair of functions $f, f' \in \text{Inj}$, the distributions of queries $Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, y))$ and $Q_{\text{SOLVE}}(R^{f', \text{SOLVE}}(1^n, y))$ are identical (i.e., the queries to SOLVE depend only on the input y and are independent of the oracle f).*

3.2 Impossibility for a Deterministic f -Oblivious Many-One Reduction

In this section, we show that there is no fully black-box construction of a hard TFNP problem from injective one-way functions with a deterministic f -oblivious many-one reduction. The proof of this result is already non-trivial and highlights our main technical contributions. In Section 5, we explain how to extend this result to rule out fully black-box constructions even with a *randomized* f -oblivious *non-adaptive* reduction.

Algorithm 1: The oracle SOLVE.

<p>Hardwired : a fully black-box construction (R, T_R, C, T_C, p) of a hard TFNP problem from iOWF</p> <p>Oracle access: an injective function $f = \{f_n\}_{n \in \mathbb{N}} \in \text{Inj}$</p> <p>Input : an instance $i \in \{0, 1\}^*$</p> <p>Output : a solution $s \in \{0, 1\}^*$ such that $C^f(i, s) = 1$</p> <p>1 Compute $Y_i = \bigcup_{n=1}^{T_C(i +p(i))} \{y \in \text{Im}(f_n) \mid i \in Q_{\text{SOLVE}}(R^f, \text{SOLVE}(1^n, y))\}$</p> <p>2 Compute $N_i = \{n \in \mathbb{N} \mid Y_i \cap \text{Im}(f_n) \neq \emptyset\}$</p> <p>3 for $n \in N_i$ do</p> <p>4 Compute $Y_{i,n} = Y_i \cap \text{Im}(f_n)$</p> <p>5 end</p> <p>6 Compute $S_{i,f} = \{s \in \{0, 1\}^* \mid s \leq p(i) \text{ and } C^f(i, s) = 1\}$</p> <p>7 while <i>True</i> do</p> <p>8 $B_{i,f} = \{s \in S_{i,f} \mid f[Q(C^f(i, s))] \cap Y_i = \emptyset\}$</p> <p>9 if $B_{i,f} \neq \emptyset$ then</p> <p>10 return lexicographically smallest $s \in B_{i,f}$</p> <p>11 end</p> <p>12 Choose $n \in N_i$ such that $\frac{ Y_{i,n} }{2^n}$ is maximized.</p> <p>13 Set $N_i = N_i \setminus \{n\}$</p> <p>14 Set $Y_i = Y_i \setminus Y_{i,n}$</p> <p>15 end</p>
--

Theorem 1. *There is no fully black-box construction (R, T_R, C, T_C, p) of a worst-case hard TFNP problem from injective one-way functions with deterministic f -oblivious many-one reduction with success probability at least $2^{-0.1n}$ such that both running times $T_R, T_C \in O(n^{\text{polylog}(n)})$.*

In the above theorem, the running time of both R and C is restricted to quasi-polynomial. Note that the standard notion of cryptographic constructions requires R, C to run in polynomial time in order to be considered efficient. We are ruling out a broader class of potentially less efficient reductions.

The proof of Theorem 1 uses, on a high level, a similar template as other black-box separations in the literature. That is, we design an oracle O relative to which (injective) one-way functions exist but TFNP is broken (even in the worst case). We follow the two-oracle approach [HR04], and, therefore, our oracle $O = (f, \text{SOLVE})$ consist of:

1. $f \in \text{Inj}$: a sequence of random injective functions which implements injective one-way functions; and
2. SOLVE: a reduction-specific oracle that solves TFNP instances.

That a random injective function is one-way is a well-established result (see, e.g., Claim 5.3 in [RSS17] for the more general case of random functions). The bulk of technical work revolves around showing that f remains one-way even in the presence of SOLVE. For any fully black-box construction with a deterministic f -oblivious many-one reduction, we provide an oracle SOLVE which finds a solution for any TFNP instance (i.e., TFNP is easy in the presence of SOLVE) and argue that it does not help the reduction in inverting injective one-way functions. The description of our oracle SOLVE is given in Algorithm 1 and it is explained below.

Oracle SOLVE: Let (R, T_R, C, T_C, p) be the construction of a hard TFNP problem from injective one-way function with a deterministic f -oblivious many-one security reduction which is hardwired in the oracle SOLVE. Ideally, SOLVE should output a solution i which gives the reduction R no information about the inversion of its challenge y . Unfortunately, SOLVE is unaware of the particular challenge y on which $R^f(1^n, y)$ queried SOLVE with the instance i . Nevertheless, SOLVE can compute the set Y_i of all challenges y on which the reduction would query the instance i .⁵ The challenges in Y_i become “protected” and SOLVE will attempt to

⁵ Here we crucially rely on f -obliviousness of the reduction algorithm R which ensures that Y_i depends only on the image of f , which allows SOLVESIM to simulate SOLVE without querying f on too many inputs.

provide a solution which does not reveal a preimage of any $y \in Y_i$, i.e., s such that $C^f(i, s)$ does not make an f -query on a preimage of any $y \in Y_i$.

Note that we could run into a potential technical issue when defining Y_i , as the set of *all* challenges y on which R queries the instance i might be infinite. However when the instance i is queried by the security reduction R on some very long challenge y then C contributes no indirect query to $f^{-1}(y)$ as the running time of C depends only on the length of the instance i . More formally: the running time of C is bounded by $T_C(|i| + p(|i|))$ thus C cannot query f on longer inputs. Therefore, we can consider only possible challenges y from $\text{Im}(f_n)$ for $n \leq T_C(|i| + p(|i|))$.

On lines 2–6, SOLVE computes the following auxiliary sets N_i , $Y_{i,n}$, and $S_{i,f}$. The set N_i contains all the input lengths for the preimages x such that the reduction $R^{f, \text{SOLVE}}(1^n, f(x))$ queries the instance i . SOLVE then splits Y_i into subsets $Y_{i,n}$ using the input lengths of interest in N_i . Finally, SOLVE computes the set $S_{i,f}$ which is the set of *all* possible solutions for the instance i .

The strategy of SOLVE is to return a solution from the set of “benign” solutions $B_{i,f}$, which do not induce any query to preimages of the protected challenges in Y_i . If there is any such benign solution then SOLVE simply halts and returns the lexicographically smallest one. Unfortunately, it might be the case that every solution queries a preimage of some $y \in Y_i$, e.g., if the instance i is queried for all challenges y of a given preimage length n and on each solution s at least one x of length n is queried (i.e., $B_{i,f} = \emptyset$ unless we remove $Y_{i,n}$ from Y_i). Since SOLVE in general cannot output a solution while protecting the whole set Y_i , it will proceed to gradually relax the condition on the set of protected challenges.

Note that we might allow SOLVE to return a solution even though it induces queries to preimages of protected challenges, as long as the reduction queries the instance i on the corresponding image length often enough, as any fixed solution induces only a bounded number of queries to f (bounded by T_C). Therefore, if the set of challenges on which R queries i is dense enough w.r.t. some image length then, with overwhelming probability, an arbitrary solution will be benign for the random challenge y given to the reduction. Thus, we allow SOLVE to return a solution revealing preimages of challenges from the auxiliary set $Y_{i,n}$ maximizing $\frac{|Y_{i,n}|}{2^n}$. If the fraction $\frac{|Y_{i,n}|}{2^n}$ is small then SOLVE is able to find a benign solution which protects the preimages of length n (see Claim 13). Whereas, if the fraction $\frac{|Y_{i,n}|}{2^n}$ is large enough then any fixed solution will be benign w.r.t. the actual challenge of R with overwhelming probability as each solution can induce queries to only a small number of preimages of challenges from the set $Y_{i,n}$ (see Claim 12).

In order to show formally that SOLVE does not help in inverting the injective one-way function, we employ an incompressibility argument similar to [GT00]. Specifically, we present algorithms ENCODE _{n} (given in Algorithm 2) and DECODE _{n} (given in Algorithm 3) which utilize the reduction R to allow compression of a random permutation more succinctly than what is information-theoretically possible. When compressing the random permutation by ENCODE _{n} , we have access to the whole permutation and we can effectively provide the reduction with access to SOLVE. However, to be able to use the reduction also in the DECODE _{n} , we have to be able to simulate access to our SOLVE oracle given access only to a partially defined oracle f (as we are reconstructing f). For the description of the algorithm SOLVESIM, which simulates the SOLVE oracle for the purpose of decoding in DECODE _{n} , see Algorithm 4.

ENCODE _{n} algorithm: The algorithm ENCODE _{n} (Algorithm 2) uses the reduction R to compress a random permutation π on bit strings of length n . Note that even though R succeeds in inverting an injective function, for technical reasons, we leverage its power in order to compress a permutation. One particular issue we would run into when trying to compress an injective function f which is not surjective is that the encoding would have to comprise also of the encoding of the image of f which might render the encoding inefficient.

Nevertheless, in order to use the reduction for compressing, we must provide it with oracle access to an injective function which is *not a bijection*. Thus, we equip ENCODE _{n} (as well as DECODE _{n}) with an injective function h . ENCODE _{n} then computes the function f as a composition of the functions $h \circ \pi$ and uses the reduction with respect to the composed oracle f . We emphasize that h is independent of π , therefore it cannot be used in order to compress π on its own.

First, ENCODE _{n} computes the set INV_f which is the set of all challenges y on which the reduction successfully inverts (i.e., the reduction returns $f^{-1}(y)$). Then ENCODE _{n} computes the set G_f , which is the

Algorithm 2: The algorithm ENCODE_n .	
Hardwired	: a fully black-box construction (R, T_R, C, T_C, p) of a hard TFNP problem from iOWF
Common Input:	an injective function $h \in \text{Inj}$ shared with DECODE_n
Input	: a permutation $\pi \in \text{Inj}_n^n$ on $\{0, 1\}^n$
Output	: an encoding \mathcal{M} of π
1	$f = h \circ \pi$, i.e., $f(x) = \begin{cases} h(\pi(x)) & \text{for all } x \text{ of length } n \\ h(x) & \text{otherwise} \end{cases}$
2	$\text{INV}_f = \{y \in \text{Im}(h_n) \mid R^{\text{SOLVE}, f}(1^n, y) = f^{-1}(y)\}$
3	$G_f = \{y \in \text{INV}_f \mid f^{-1}(y) \notin Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y))\}$
4	$Y_f = \emptyset$ and $X_f = \emptyset$
5	while $G_f \neq \emptyset$ do
6	Pick lexicographically smallest $y \in G_f$
7	$G_f = G_f \setminus (f[Q_f(R^{f, \text{SOLVE}}(1^n, y))] \cup \{y\})$
8	$Y_f = Y_f \cup \{y\}$ and $X_f = X_f \cup \{f^{-1}(y)\}$
9	end
10	if $ X_f < 2^{0.6n}$ then
11	return $\mathcal{M} = (0, \pi)$
12	end
13	else
14	return $\mathcal{M} = (1, X_f , Y_f, X_f, \sigma = f \upharpoonright (\{0, 1\}^n \setminus X_f)) \in \{0, 1\}^{1+n + \lceil \log \binom{2^n}{ Y_f } \rceil + \lceil \log \binom{2^n}{ X_f } \rceil + \lceil \log(\{0, 1\}^n \setminus X_f !) \rceil}$
15	end

set of “good” challenges y , on which the reduction successfully inverts even though SOLVE returns a solution which does not induce a query to any preimage of y . This set is used to reduce the size of the trivial encoding of f – the part of f corresponding to the challenges in G_f will be algorithmically reconstructed by DECODE_n using the security reduction R .

Specifically, ENCODE_n computes Y_f , the subset of G_f for which the preimages will be algorithmically reconstructed, as follows: ENCODE_n processes the challenges y in G_f one by one in lexicographically increasing order and stores all f -queries needed for reconstruction by R (i.e, for any x such that there was an f -query x , the element $f(x)$ is removed from the “good” set G_f as we cannot reconstruct the preimage of y using R without knowing the image of x under f).

ENCODE_n outputs an encoding \mathcal{M} which describes the size of X_f , the sets Y_f and X_f (where X_f is the set of preimages corresponding to Y_f), and the partial function representing the function f on inputs of length n outside of X_f . Thus, the encoding saves bits by not revealing the bijection between X_f and Y_f which is algorithmically reconstructed instead (Lemma 4). Specifically, the size of X_f (equal to the size of Y_f) can be encoded using $\log 2^n = n$ bits. Y_f is a subset of $\text{Im}(f_n) = \text{Im}(h_n)$ and it is encoded using $\lceil \log \binom{2^n}{|Y_f|} \rceil$ bits as the index of the corresponding subset of size $|Y_f|$ (the set X_f is encoded in a similar manner). Finally, the bijection between $\{0, 1\}^n \setminus X_f$ and $\text{Im}(f) \setminus Y_f$ is encoded as the index of the corresponding permutation on a set of size $|\{0, 1\}^n \setminus X_f|$ using $\lceil \log(|\{0, 1\}^n \setminus X_f|!) \rceil$ bits.

A small technicality arises when the set X_f , respectively the set Y_f , is not large enough, the above mentioned encoding would be inefficient as the trivial encoding outputting the whole description of the permutation π would use fewer bits. Thus, ENCODE_n simply outputs the trivial encoding when X_f is too small. The first bit of the encoding distinguishes between the two cases.

DECODE_n algorithm: The encoding returned by ENCODE_n is uniquely decodable by DECODE_n given in Algorithm 3 (see Section 4.2). When the output of ENCODE_n starts with 0, the rest of the encoding is an explicit encoding of π and we are immediately done with its reconstruction. If the output starts with bit 1, the following n bits represent $|X_f| = |Y_f|$. DECODE_n then reads the following $\lceil \log \binom{2^n}{|X_f|} \rceil$ bits of the encoding to reconstruct the set Y_f (as the j -th subset of 2^n of size $|X_f|$). Similarly, DECODE_n reconstructs the set X_f using

Algorithm 3: The algorithm DECODE_n .

<p>Hardwired : a fully black-box construction (R, T_R, C, T_C, p) of a hard TFNP problem from iOWF</p> <p>Common Input: an injective function $h \in \text{Inj}$ shared with ENCODE_n</p> <p>Input : an encoding \mathcal{M}</p> <p>Output : a permutation $\pi \in \text{Inj}_n^n$</p> <p>1 Parse $\mathcal{M} = (b, \mathcal{M}')$, where $b \in \{0, 1\}$</p> <p>2 if $b = 0$ then</p> <p>3 Decode π from \mathcal{M}'</p> <p>4 return π</p> <p>5 end</p> <p>6 Parse $\mathcal{M}' = (X_f , Y_f, X_f, \sigma)$</p> <p>7 Set partial function $f' = \begin{cases} \sigma & \text{for inputs of length } n \\ h & \text{otherwise} \end{cases}$ // f' is defined only outside X_f</p> <p>8 while $Y_f \neq \emptyset$ do</p> <p>9 Pick lexicographically smallest $y \in Y_f$</p> <p>10 Let $f''(x) = \begin{cases} y & \text{for all } x \in \text{Dom}(h) \setminus \text{Dom}(f') \\ f'(x) & \text{otherwise} \end{cases}$</p> <p>11 $x = R^{f'', \text{SOLVE}_{\text{SIM}}(h, f', \cdot)}(1^n, y)$</p> <p>12 Let $f'(x) = y$</p> <p>13 Set $Y_f = Y_f \setminus \{y\}$</p> <p>14 end</p> <p>15 return $\pi = (h^{-1} \circ f') \upharpoonright \{0, 1\}^n$</p>

the following $\lceil \log \binom{2^n}{|X_f|} \rceil$ bits. The remaining bits represent σ , a restriction of f on all the n -bit inputs outside of X_f , given by the index of the corresponding bijection between $\{0, 1\}^n \setminus X_f$ and $\text{Im}(f) \setminus Y_f$. Note that such encoding of σ does preserve the structure of the restriction but it loses the information about the domain and image of σ . However, both are easy to reconstruct. The domain is simply $\{0, 1\}^n \setminus X_f$ and the image of σ can be computed from Y_f and the common input h as $\text{Im}(\sigma) = \text{Im}(f) \setminus Y_f = \text{Im}(h \circ \pi) \setminus Y_f = \text{Im}(h) \setminus Y_f$.

DECODE_n then computes the remaining preimages one by one in lexicographic order using the security reduction R , adding the reconstructed mapping into a partial function f' . Note that during the computation of the preimage of y , the reduction might make an f -query on x which has no defined output. But as DECODE_n takes $y \in Y_f$ in the same order as ENCODE_n added them to the set Y_f , this happens if and only if the preimage of y is being queried. Thus, we answer any such query by y (it is crucial that this happens only for f -queries made directly by R) which is captured in the definition of the auxiliary function f'' defined by DECODE_n and used as the oracle for the security reduction R .

Once ENCODE_n finds the preimages of all challenges y from Y_f , the function f' is defined everywhere. To reconstruct the permutation π on $\{0, 1\}^n$, DECODE_n can simply compose the inverse of h with the reconstructed function f' .

SOLVE_{SIM} algorithm: For the ease of presentation we usually do not explicitly mention the oracle h as it is given by context (we run DECODE_n and $\text{SOLVE}_{\text{SIM}}$) with respect to only one h at a time.

The computation of the algorithm $\text{SOLVE}_{\text{SIM}}$ (Algorithm 4) is similar to the computation of SOLVE (Algorithm 1). First, $\text{SOLVE}_{\text{SIM}}$ computes the sets Y_i . There is one big difference between SOLVE and $\text{SOLVE}_{\text{SIM}}$. As $\text{SOLVE}_{\text{SIM}}$ does not have access to the whole f it uses h or the partial knowledge of f , namely the partial function f' everywhere f is used in the SOLVE algorithm.

- We use h whenever we need to determine the image of f_n for some n . As $\forall n \in \mathbb{N}: \text{Im}(h_n) = \text{Im}(f_n)$ using $\text{Im}(h_n)$ instead of $\text{Im}(f_n)$ makes no difference to the computation.

Algorithm 4: The algorithm SOLVESIM.

<p>Hardwired : a fully black-box construction (R, T_R, C, T_C, p) of a hard TFNP problem from iOWF</p> <p>Input : A function $h \in \text{Inj}$, partial injective function $f' \in \text{Inj}$, and an instance $i \in \{0, 1\}^*$</p> <p>Output : Solution, i.e., $s \in \{0, 1\}^*$ such that $C^{f'}(i, s) = 1$</p> <p>1 Compute $Y_i = \bigcup_{n=1}^{T_C(i +p(i))} \{y \in \text{Im}(h_n) \mid i \in Q_{\text{SOLVE}}(R^{h, \text{SOLVE}}(1^n, y))\}$</p> <p>2 Compute $N_i = \{n \in \mathbb{N} \mid Y_i \cap \text{Im}(h_n) \neq \emptyset\}$</p> <p>3 for $n \in N_i$ do</p> <p>4 Compute $Y_{i,n} = Y_i \cap \text{Im}(h_n)$</p> <p>5 end</p> <p>6 Compute $S_{i,f'} = \{s \in \{0, 1\}^* \mid s \leq p(i) \text{ and } Q(C^{f'}(i, s)) \subseteq \text{Dom}(f') \text{ and } C^{f'}(i, s) = 1\}$</p> <p>7 while <i>True</i> do</p> <p>8 $B_{i,f'} = \{s \in S_{i,f'} \mid f[Q(C^{f'}(i, s))] \cap Y_i = \emptyset\}$ // "benign" solutions</p> <p>9 if $B_{i,f'} \neq \emptyset$ then</p> <p>10 return lexicographically smallest $s \in B_{i,f'}$</p> <p>11 end</p> <p>12 Choose $n \in N_i$ such that $\frac{ Y_{i,n} }{2^n}$ is maximized.</p> <p>13 Set $N_i = N_i \setminus \{n\}$</p> <p>14 Set $Y_i = Y_i \setminus Y_{i,n}$</p> <p>15 end</p>
--

- The second place where h is used instead of f is when SOLVESIM computes the set Y_i . Specifically, when determining images y for which the security reduction R queries the given instance i , the algorithm SOLVESIM computes the same Y_i as if it used f by the f -obliviousness of the security reduction.
- In all other places, SOLVESIM uses the partial knowledge of f (i.e., the partial function f'). This causes a real difference in the computation. In particular, the set $S_{i,f'}$ (as computed by SOLVESIM) may differ a lot from $S_{i,f}$ (as computed by SOLVE) as some solutions from $S_{i,f}$ potentially query some unknown parts of f . Thus, the set $S_{i,f'}$ computed by SOLVESIM is just a subset of the whole $S_{i,f}$. The set $S_{i,f'}$ contains only the solutions SOLVESIM is “aware of” (f' is defined for all queries and thus SOLVESIM may verify the solution). The rest of the computation is practically the same, except that SOLVESIM is restricted just to the set of solutions $S_{i,f'}$. The main trick is that we make sure that SOLVESIM is aware of the solution which should be returned and it does not matter that it ignores other solutions of the instance.

Structure of the proof of Theorem 1. For ease of presentation and understanding, we divide the proof into four lemmata, Lemma 1 through 4. Lemma 1 shows that given an instance i of the TFNP problem represented by the algorithm C^f , our SOLVE always returns a solution, i.e., an s such that $C^f(i, s) = 1$ (formal proof is given in Section 4.1). Thus, any distribution of instances of the TFNP problem is easy in the presence of SOLVE.

Lemma 1. *For any instance $i \in \{0, 1\}^*$ and any $f \in \text{Inj}$, the algorithm $\text{SOLVE}^f(i)$ halts and returns a solution, i.e., it returns a string $s \in \{0, 1\}^*$ such that $|s| \leq p(|i|)$ and $C^f(i, s) = 1$.*

To argue that SOLVE does not help in inverting injective functions, we analyze the joint properties of the algorithms ENCODE_n and DECODE_n . First, we show that DECODE_n always returns the correct permutation encoded by ENCODE_n (see Section 4.2 for the formal proof).

Lemma 2. *For all $n \in \mathbb{N}$, $\pi \in \text{Inj}_n^n$, and $h \in \text{Inj}$,*

$$\text{DECODE}_n^h(\text{ENCODE}_n^h(\pi)) = \pi,$$

where ENCODE_n , respectively DECODE_n , is given in Algorithm 2, respectively Algorithm 3.

We crucially rely on f -obliviousness of the reduction for the proof of Lemma 2. It is the property which allows us to simulate the algorithm SOLVE during the decoding phase, as SOLVESIM needs to be able to compute the same set Y_i as SOLVE does. Moreover, SOLVESIM cannot query f on all preimages as SOLVE does when computing Y_i . Due to f -obliviousness of the reduction, we may substitute f by h in the computation of Y_i in SOLVESIM as the resulting set depends only on the image of the function given to R as an oracle (and $\text{Im}(f) = \text{Im}(h)$).

Second, we show that the encoding output by ENCODE_n is prefix-free (see Section 4.3 for the formal proof).

Lemma 3. *Let $h \in \text{Inj}$ be any injective function and $n \in \mathbb{N}$, then the encoding given by the algorithm ENCODE_n (Algorithm 2) is prefix-free, i.e.,*

$$\forall \pi, \pi' \in \text{Inj}_n^n \text{ such that } \pi \neq \pi': \text{ENCODE}_n^h(\pi) \text{ is not a prefix of } \text{ENCODE}_n^h(\pi').$$

Finally, we bound the expected size of the encoding given by ENCODE_n (see Section 4.4) which contradicts the information-theoretic bound implied by Corollary 1.

Lemma 4. *Let (R, T_R, C, T_C, p) be a fully black-box construction of a hard TFNP problem from an injective one-way function. Assume $n \in \mathbb{N}$ is large enough so that $n \geq 50$ and $2q(n) + 1 \leq 2^{0.2n}$, where $q(n)$ is the maximal number of f -queries made by C on the queried instance (see Definition 4). Let the success probability of R be $\beta \geq 2^{-0.1n}$, i.e., for any f we have*

$$\Pr_{x \leftarrow \{0,1\}^n} [R^{f, \text{SOLVE}}(1^n, f(x)) = x] = \beta \geq 2^{-0.1n}.$$

Then

$$\exists h \in \text{Inj}: \mathbb{E}_{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj}} [|\text{ENCODE}_n^h(\pi)|] \leq \log 2^{n!} - \frac{8}{10} n 2^{0.1n}.$$

We claim (see Claim 10) that the upper bound $2q(n) + 1 \leq 2^{0.2n}$ used in the statement of the lemma is without loss of generality for large enough n and for all quasi-polynomial (and, hence, also for efficient) algorithms R, C . We use this fact again in the proof of the main theorem (Theorem 1), and refer the readers to Section 4.4 for the precise statement and its proof.

Equipped with the above lemmata, we next prove Theorem 1.

Proof (of Theorem 1). Suppose to the contrary that there is such a reduction (R, T_R, C, T_C, p) . By Lemma 1, the algorithm SOLVE (Algorithm 1) returns a valid solution with probability one. Thus, the reduction R must invert f with high probability when given access to any oracle $f \in \text{Inj}$ and our oracle SOLVE, i.e.,

$$\Pr_{x \leftarrow \{0,1\}^n} [R^{f, \text{SOLVE}}(1^n, f(x)) = x] \geq \frac{1}{p'(n)}$$

for some polynomial p' and infinitely many $n \in \mathbb{N}$.

Let $n \in \mathbb{N}$ be large enough such that

1. $2q(n) + 1 \leq 2^{0.2n}$,
2. $\Pr_{x \leftarrow \{0,1\}^n} [R^{f, \text{SOLVE}}(1^n, f(x)) \in f^{-1}(f(x))] \geq \frac{1}{p'(n)}$,

where $q(n)$ is the maximal number of f -queries made by C on the queried instance (see Definition 4). As already pointed out, the quasi-polynomial bounds on running times $T_C, T_R \in O(n^{\text{polylog}(n)})$ imply that $q(n) \in o(2^{0.2n})$ (see Claim 10). Thus, for large enough n , the upper bound $2q(n) + 1 \leq 2^{0.2n}$ holds without loss of generality.

For any $h \in \text{Inj}$, we can use the algorithm ENCODE_n^h (Algorithm 2) to encode a given permutation $\pi \in \text{Inj}_n^n$. Decodability of the encoding follows from Lemma 2. Moreover, by Lemma 3, the encoding is a prefix-free code. By Lemma 4, there is a function $h \in \text{Inj}$ such that the pair of algorithms ENCODE_n^h and DECODE_n^h defines an encoding of $\pi \leftarrow \text{Inj}_n^n$ with expected length at most $\log(2^{n!}) - \frac{8}{10} n 2^{0.1n}$. This contradicts the information-theoretic bound on the expected length of any prefix-free encoding of a random permutation on $\{0, 1\}^n$ given by Corollary 1. \square

4 Proofs of the Supporting Lemmata

The first lemma we need to prove is that the algorithm SOLVE halts and returns a valid solution on any input TFNP instance i . Section 4.2 is devoted to proving the correctness of our encoding, i.e., the fact that the algorithm DECODE_n uniquely and correctly decodes ENCODE_n 's input permutation. We show that the encoding itself is prefix-free in Section 4.3.

Section 4.4 is the core of the incompressibility argument. We show that the expected size of our encoding is smaller than the information-theoretic bound.

4.1 SOLVE Always Returns a Solution

In this section we show that algorithm SOLVE halts and returns a proper solution. When the algorithm SOLVE halts it is easy to show that the returned solution is a valid solution of the given TFNP problem and that it is of the correct length (such a solution has to exist by correctness of the reduction). Thus the proof essentially reduces to proving that SOLVE eventually halts, at latest when the set Y_i becomes empty.

Lemma 1. *For any instance $i \in \{0, 1\}^*$ and any $f \in \text{Inj}$ the algorithm $\text{SOLVE}^f(i)$ halts and returns a solution, i.e., it returns a string $s \in \{0, 1\}^*$ such that $|s| \leq p(|i|)$ and $C^f(i, s) = 1$.*

Proof. First observe that all sets $Y_i, N_i, Y_{i,n}$ and $S_{i,f}$ are well defined and can be computed in finite time. Notice that in each iteration of the while loop we remove one n from the set N_i . Thus after $|N_i|$ iterations the set N_i is empty. During the whole algorithm we keep the invariant that

$$Y_i \subseteq \bigcup_{n \in N_i} \{0, 1\}^n$$

(see lines 1, 2 and 14, 13). After $|N_i|$ iterations of the while loop the set Y_i is empty as well. Once $Y_i = \emptyset$, the set $B_i = S_{i,f}$. As the set $S_{i,f}$ is nonempty by correctness of the TFNP problem (see Definition 2 - Correctness) the algorithm SOLVE always halts.

Finally when the algorithm returns (only line 10 of Algorithm 1), it returns a string s from the set $B_i \subseteq S_{i,f}$. Thus $|s| \leq p(|i|)$ and $C^f(i, s) = 1$ by the definition of the set $S_{i,f}$ (see line 6 of Algorithm 1). \square

4.2 Encoding Is Uniquely Decodable

In this section, we show that the encoding given by ENCODE_n (Algorithm 2) is uniquely decodable by DECODE_n (Algorithm 3), i.e., we prove the following lemma.

Lemma 2. *For any $n \in \mathbb{N}$, any $\pi \in \text{Inj}_n^n$ and for any $h \in \text{Inj}$*

$$\text{DECODE}_n^h(\text{ENCODE}_n^h(\pi)) = \pi,$$

where $\text{ENCODE}_n, \text{DECODE}_n$ are as described in Algorithms 2 and 3.

To prove this lemma, we work with a partial function f' computed during the decoding. We show by induction on the number of steps of DECODE_n that during the whole computation f' agrees with the function f . During decoding we fill in the undefined values of f' . To find the missing values we find preimages of challenges y from image of f one by one using the reduction algorithm R .

We need to show that the missing mapping is computed correctly. To this end, we show that the reduction as run during the decoding phase gives the same output as it gave when run during the encoding phase. To prove this we show that all f -queries (both direct and indirect) made by the reduction (when run in the decoding phase) are answered exactly as if f would be queried.

The first step is to show that all direct f -queries are answered correctly (Claim 6).

Claim 6. Let $n \in \mathbb{N}$ be any number, $h \in \text{Inj}$ be any function, $\pi \in \text{Inj}_n^n$ be any permutation and let $f = h \circ \pi$. Let \mathcal{M} be the output of $\text{ENCODE}_n^h(\pi)$ (see Algorithm 2) and assume that $\mathcal{M} = (1, |X_f|, Y_f, X_f, \sigma)$. Let $j \in \mathbb{N}$ be such that $1 \leq j \leq |X_f|$.

Let y be the j -th lexicographically smallest string from Y_f (i.e., y is picked in the j -th iteration of the while loop, see line 8 of Algorithm 3). Let

$$f' = f \upharpoonright (\{x \in X_f \mid f(x) \text{ is the } k\text{-th lexicographically smallest in } Y_f \text{ for } k < j\} \cup \overline{X_f})$$

and

$$f'' : \{0, 1\}^* \rightarrow \{0, 1\}^* \text{ be such that } f''(x) = \begin{cases} f'(x) & \text{for } x \in \text{Dom}(f') \\ y & \text{otherwise} \end{cases}$$

Then

$$\forall x' \in Q_f^{\text{dir}}(R^{f, \text{SOLVE}}(1^n, y)) \quad f''(x') = f(x').$$

Note that if DECODE_n computed the first $j - 1$ preimages then f', f'' from Algorithm 3 correspond to their definition given in Claim 6. This is formalized in the proof of Claim 8.

Proof (of Claim 6). Let x' be any query from $Q_f^{\text{dir}}(R^{f, \text{SOLVE}}(1^n, y))$. We distinguish the following four cases:

$x' \notin X_f$: In this case $f'(x')$ is defined (it is set before the first iteration of the while loop). Thus $f''(x') = f'(x') = f(x')$, where the last equality follows from the assumptions.

$x' \in X_f \wedge f(x') <_{\text{lex}} y$: Then $f'(x')$ is defined in the j -th iteration, because DECODE_n (Algorithm 3) computes preimage of $f(x')$ before y . Thus $f''(x') = f'(x') = f(x')$, where the last inequality follows from the assumptions.

$x' \in X_f \wedge f(x') = y$: In this case $f'(x')$ is undefined (we are computing the preimage of y in j -th iteration of while loop), thus by definition of f'' , $f''(x') = y = f(x')$.

$x' \in X_f \wedge f(x') >_{\text{lex}} y$: As $f(x') >_{\text{lex}} y$, the ENCODE_n (Algorithm 2) processes y before $f(x')$. Then if x' is in $Q_f^{\text{dir}}(R^{f, \text{SOLVE}}(1^n, y))$, ENCODE_n removes $f(x')$ from G_f (line 7, Algorithm 2). Thus $f(x')$ is never added into Y_f which is a contradiction. □

Claim 6 shows that each direct oracle query to f (that is a query done by the reduction itself) is answered correctly by f'' . Next we show that both algorithms SOLVE (Algorithm 1) and SOLVESIM (Algorithm 4) answer the TFNP query with the same solution whenever the algorithm DECODE_n queries SOLVESIM during the simulation of the security reduction R .

The main observation is that the algorithm SOLVESIM is aware of the solution returned by SOLVE (i.e., f' is defined for all f -queries made during the computation). Then we show that the procedure which picks the solution out of all solutions in algorithm SOLVE gives the same result in algorithm SOLVESIM , which is possibly aware only about few solutions of the instance. Here we utilize the fact that the algorithm SOLVE always returns lexicographically smallest solution.

Claim 7. Let $n \in \mathbb{N}$ be any number, $h \in \text{Inj}$ be any function, $\pi \in \text{Inj}_n^n$ be any permutation and let $f = h \circ \pi$. Let \mathcal{M} be the output of $\text{ENCODE}_n^h(\pi)$ (see Algorithm 2) and assume that $\mathcal{M} = (1, |X_f|, Y_f, X_f, \sigma)$. Let $j \in \mathbb{N}$ be such that $1 \leq j \leq |X_f|$. Let y be the j -th lexicographically smallest string from Y_f (i.e., y is picked in the j -th iteration of the while loop). Let

$$f' = f \upharpoonright (\{x \in X_f \mid f(x) \text{ is the } k\text{-th lexicographically smallest in } Y_f \text{ for } k < j\} \cup \overline{X_f}).$$

Then the query to SOLVE is answered correctly, i.e.,

$$\forall i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, y)) \quad \text{SOLVE}^f(i) = \text{SOLVESIM}(h, f', i).$$

Proof. We fix $i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, y))$. Let s_{SOLVE} be the solution returned by $\text{SOLVE}^f(i)$ and s_{SOLVESIM} be the solution returned by $\text{SOLVESIM}(h, f, i)$. To prove $s_{\text{SOLVE}} = s_{\text{SOLVESIM}}$, we examine the sets $S_{f,i}, B_{f,i}, S_{f',i}$ and $B_{f',i}$ used in the algorithms SOLVE and SOLVESIM . For the ease of presentation we denote the sets as follows:

- $S_{\text{SOLVE},f}$ Let $S_{\text{SOLVE},f}$ denote the set $S_{f,i}$ as computed by algorithm SOLVE (Line 1, line 6).
- $S_{\text{SOLVESIM},f'}$ Let $S_{\text{SOLVESIM},f'}$ denote the set $S_{f',i}$ as computed by algorithm SOLVESIM (Algorithm 4, line 6).
- $B_{\text{SOLVE},f,k}$ Let $B_{\text{SOLVE},f,k}$ denote the set $B_{f,i}$ as computed in the k -th iteration of the while loop by the algorithm SOLVE (Algorithm 1, line 8).
- $B_{\text{SOLVESIM},f',k}$ Let $B_{\text{SOLVESIM},f',k}$ denote the set $B_{f',i}$ as computed in the k -th iteration of the while loop of the algorithm SOLVESIM (Algorithm 4, line 8).

Moreover by f -obliviousness the sets $N_i, Y_i, Y_{i,n}$ depend only on the image of f which is identical to the image of h and h is known to SOLVESIM . Thus all sets used in algorithms SOLVE and SOLVESIM except for $S_{\text{SOLVE},f}, S_{\text{SOLVESIM},f'}, B_{\text{SOLVE},f,k}, B_{\text{SOLVESIM},f',k}$ are independent of f (just dependent on its image) and thus can be computed by both SOLVE and SOLVESIM .

First recall that $y \in Y_f$ thus by the definition of Y_f , respectively the definition of $G_f \supseteq Y_f$ (Algorithm 2, lines 3, 4, 6, and 8) we get that $f^{-1}(y) \notin Q(C^f(i, s))$ as images of all $x \in Q_f(R^{f, \text{SOLVE}}(1^n, y)) \supseteq Q(C^f(i, s))$ are removed from G_f (Algorithm 2, line 7). This means that for any $x \in Q(C^f(i, s))$ either $x \notin X_f$ or $f(x) <_{\text{lex}} y$, in both cases $f'(x)$ is defined and $f'(x) = f(x)$ by the assumption of the claim. Thus

$$s_{\text{SOLVE}} \in S_{\text{SOLVESIM},f'}. \quad (1)$$

On the other hand by the assumptions of the claim $f' \subseteq f$, thus every $s \in S_{\text{SOLVESIM},f'}$ is a solution also with respect to f . In other words $s_{\text{SOLVESIM}} \in S_{\text{SOLVE},f}$ and more generally the following equation holds:

$$S_{\text{SOLVESIM},f'} \subseteq S_{\text{SOLVE},f}. \quad (2)$$

During the entire computation both SOLVE and SOLVESIM hold the exact same set Y_i , as that is computed and updated independently of f , respectively f' . Thus by the definition of $B_{\text{SOLVESIM},f',k}$ and $B_{\text{SOLVE},f,k}$ and the equation 2 we have that

$$\forall k \in \mathbb{N} \quad B_{\text{SOLVESIM},f',k} = B_{\text{SOLVE},f,k} \cap S_{\text{SOLVESIM},f'}. \quad (3)$$

Fix $k \in \mathbb{N}$ such that $B_{\text{SOLVE},f,k}$ or $B_{\text{SOLVESIM},f',k}$ is nonempty. We show that both sets $B_{\text{SOLVE},f,k}$ and $B_{\text{SOLVESIM},f',k}$ are non-empty. Moreover we show the following:

$$\{s_{\text{SOLVE}}, s_{\text{SOLVESIM}}\} \subseteq B_{\text{SOLVE},f,k} \cap B_{\text{SOLVESIM},f',k} \quad (4)$$

We distinguish the following cases:

- $B_{\text{SOLVE},f,k} = \emptyset$ By the assumption $B_{\text{SOLVESIM},f',k} \neq \emptyset$. Thus SOLVESIM returns in the k -th iteration of the while loop and the returned solution s_{SOLVESIM} has to be contained in $B_{\text{SOLVESIM},f',k}$. By the equation 3 we get that $s_{\text{SOLVESIM}} \in B_{\text{SOLVE},f,k}$. Which is a contradiction as we assumed that $B_{\text{SOLVE},f,k}$ is empty.
- $B_{\text{SOLVE},f,k} \neq \emptyset$ As $B_{\text{SOLVE},f,k} \neq \emptyset$, the SOLVE algorithm returns in the k -th iteration of the while loop and thus $s_{\text{SOLVE}} \in B_{\text{SOLVE},f,k}$. By the equation 1 s_{SOLVE} is contained in $S_{\text{SOLVESIM},f'}$ too. Thus by Equation 3 $s_{\text{SOLVE}} \in B_{\text{SOLVESIM},f',k}$. Now using the fact that $B_{\text{SOLVESIM},f',k} \neq \emptyset$ and the same argumentation as in the case above we get that $s_{\text{SOLVESIM}} \in B_{\text{SOLVESIM},f',k}$ and $s_{\text{SOLVESIM}} \in B_{\text{SOLVE},f,k}$. Which proves that both $s_{\text{SOLVE}}, s_{\text{SOLVESIM}} \in B_{\text{SOLVE},f,k} \cap B_{\text{SOLVESIM},f',k}$.

Thus we proved that both algorithms SOLVE and SOLVESIM return in the k -th iteration of their while loop. Recall that the lexicographically smallest solution from $B_{\text{SOLVE},f,k}$, respectively $B_{\text{SOLVESIM},f',k}$ is returned. Thus by Equation 4 both algorithms SOLVE and SOLVESIM return the same string $s_{\text{SOLVE}} = s_{\text{SOLVESIM}}$. \square

We combine the Claims 6 and 7 to show that the function f' is defined consistently with f during the whole computation.

Claim 8. Let $n \in \mathbb{N}$ be any number, $h \in \text{Inj}$ be any function, $\pi \in \text{Inj}_n^n$ be any permutation and let $f = h \circ \pi$. Let \mathcal{M} be the output of $\text{ENCODE}_n^h(\pi)$ (see Algorithm 2) and assume that $\mathcal{M} = (1, |X_f|, Y_f, X_f, \sigma)$. Let $j \in \mathbb{N}$ be such that $0 \leq j \leq |X_f|$. Let f' be the partial function computed by $\text{DECODE}_n^h(\mathcal{M})$ after j iterations of the while loop (line 8, Algorithm 3). Then

$$f' \subseteq f$$

i.e., for any $x \in \text{Dom}(f')$ we have $f'(x) = f(x)$.

Proof. We proceed by induction on the number of iterations of the while loop of the algorithm DECODE_n (Algorithm 3). Before the first iteration, we set $f' = \sigma$ for inputs of length n (see line 11 of Algorithm 2). Recall that $\sigma = f \upharpoonright (\{0, 1\}^n \setminus X_f)$ on inputs of length n . For inputs of length different from n we set f' according to h , which is the same as f as $f(x) = h(x)$ for all inputs x of length different from n . Thus, we immediately get $f' \subseteq f$.

In each iteration of the while loop, the algorithm DECODE_n (Algorithm 3) sets the preimage of exactly one $y \in \text{Im}(f)$. We show that the preimage is computed correctly. Suppose that we are in the j -th iteration of the while loop (line 8, Algorithm 3) and y is chosen in the j -th iteration as the lexicographically smallest string in Y_f (i.e., for all $y' \in \text{Im}(f)$, $y' <_{\text{lex}} y$, we know the preimage of y'). We set the preimage of y in f' to be the output of $R^{f'', \text{SOLVE}_{\text{SIM}}(h, f', \cdot)}(1^n, y)$.

Note that the fact that y is in Y_f implies that $R^{f, \text{SOLVE}(f, \cdot)}(1^n, y) = f^{-1}(y)$, as $y \in Y_f$ are taken from the set INV_f (see line 2, Algorithm 2). As the preimage of y is computed as $R^{f'', \text{SOLVE}_{\text{SIM}}(h, f', \cdot)}(1^n, y)$, we need to show that all oracle queries in the computation are answered the same way as in the computation $R^{f, \text{SOLVE}(f, \cdot)}(1^n, y)$. Then $f'^{-1}(y) = R^{f'', \text{SOLVE}_{\text{SIM}}(h, f', \cdot)}(1^n, y) = R^{f, \text{SOLVE}(f, \cdot)}(1^n, y) = f^{-1}(y)$.

We use Claims 6 and 7 to argue that the queries were answered correctly. Note that the assumption on f' used in these claims follows from the induction hypothesis. As in the j -th iteration of the while loop f' is defined for all $x \notin X_f$ and for the $j - 1$ preimages of the lexicographically smallest $y \in Y_f$. By induction hypothesis it is defined correctly (gives the same output as f on all $x \in \text{Dom}(f')$). By Claim 6 all f -queries are answered correctly, i.e., for any

$$x' \in Q_f^{\text{dir}}(R^{f, \text{SOLVE}(f, \cdot)}(1^n, y)) \text{ we have } f''(x') = f(x').$$

To prove that SOLVE query is answered correctly observe that by f -obliviousness of the reduction the same TFNP instances are queried, i.e.,

$$Q_{\text{SOLVE}}(R^{f'', \text{SOLVE}_{\text{SIM}}(h, f', \cdot)}(1^n, y)) = Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, y)).$$

By Claim 7 the query to $\text{SOLVE}_{\text{SIM}}$ is answered identically as it would be answered by SOLVE . Thus $R^{f'', \text{SOLVE}_{\text{SIM}}(h, f', \cdot)}(1^n, y) = R^{f, \text{SOLVE}(f, \cdot)}(1^n, y)$ and the preimage of y is set correctly. \square

Finally we use Claim 8 to prove that DECODE_n correctly decodes the message of ENCODE_n and outputs the right permutation.

Proof (of Lemma 2). We assume that the first bit of the output of ENCODE_n (Algorithm 2) is 1, otherwise a full description of π has been sent to DECODE_n (Algorithm 3) and the lemma holds trivially. Let $f = h \circ \pi$ similarly as in the ENCODE_n algorithm (Algorithm 2). We have to show that DECODE_n (Algorithm 3) on line 15 holds function f' such that $f' = f$. We use Claim 8 to get that $f' \subseteq f$. Now observe that on line 15 of Algorithm 3 the function f' is defined for all inputs. Thus $f' = f$. Finally the DECODE_n algorithm returns a function $(h^{-1} \circ f) \upharpoonright \{0, 1\}^n = (h^{-1} \circ h \circ \pi) \upharpoonright \{0, 1\}^n = \pi$. \square

4.3 Encoding Is Prefix-free

To be able to directly leverage coding theory we show that our encoding is in fact prefix-free. If the set of invertible images Y_f is small the encoding starts with a “0” bit followed by a trivial (prefix-free) encoding of the given permutation. Otherwise we show that the two codewords either differ at the beginning or are of the same length. Unique decodability gives us that such codewords must differ.

Lemma 3. *Let $h \in \text{Inj}$ be any injective function and $n \in \mathbb{N}$, then the encoding given by the algorithm ENCODE_n (Algorithm 2) is prefix-free, i.e.,*

$$\forall \pi, \pi' \in \text{Inj}_n^n \text{ such that } \pi \neq \pi' : \text{ENCODE}_n^h(\pi) \text{ is not a prefix of } \text{ENCODE}_n^h(\pi').$$

We use the Claim 9 which computes the size of an encoding of a permutation π as returned by ENCODE_n (Algorithm 2).

Claim 9. *The size of the encoding computed by ENCODE_n (Algorithm 2) is*

1. $1 + \lceil \log(2^n!) \rceil$ bits in case ENCODE_n (Algorithm 2) returned $\mathcal{M} = (0, \pi)$ on line 11, and
2. $1 + n + 2 \left\lceil \log \binom{2^n}{|X_f|} \right\rceil + \lceil \log((2^n - |X_f|)!) \rceil$ bits in case ENCODE_n (Algorithm 2) returned a message $\mathcal{M} = (1, |X_f|, Y_f, X_f, \sigma)$ on line 14.

Proof. When we say that we encode an index of a combinatorial object, we mean that we enumerate all instances of this object in some natural order and then encode the index of our particular instance. Specifically for the case of k -element subsets of a set of size n , we enumerate all such subsets in the lexicographical order and then just give an index – a number between one and $\binom{n}{k}$ always using precisely $\lceil \binom{n}{k} \rceil$ bits (we use leading zeroes as a padding in case there would be less bits). Similarly, for the case of permutations of n objects, we enumerate all such permutations in lexicographical order and then give the particular index using exactly $\lceil n! \rceil$ bits (again, potentially padding the index with leading zeroes).

1. In the first case (return on line 11) the encoding consists of the single “0” bit and an index of a permutation on binary strings of length n . Thus we use $1 + \lceil \log(2^n!) \rceil$ bits.
2. In the second case (return on line 14) the encoding consists of a single “1” bit, a size of the set $X_f \subseteq \{0, 1\}^n$ (this is always encoded by n bits), two sets $X_f \subseteq \{0, 1\}^n$ and $Y_f \subseteq \text{Im}(f_n)$ both of size $|X_f|$ and thus encoded using $\lceil \log \binom{2^n}{|X_f|} \rceil$ bits each, and an index of a bijection between $\{0, 1\}^n \setminus X_f$ and $\text{Im}(f_n) \setminus Y_f$. Note that we do not have to encode the domain and the image of σ , as those can be computed (domain from the set X_f and the image from the sets Y_f and $\text{Im}(h_n) = \text{Im}(f_n)$). Thus we always use $1 + n + 2 \left\lceil \log \binom{2^n}{|X_f|} \right\rceil + \lceil \log((2^n - |X_f|)!) \rceil$ bits for the encoding. □

Now we use Claim 9 to show that the encoding is prefix-free.

Proof (of Lemma 3). Let $\pi, \pi' \in \text{Inj}_n^n$ be any permutations such that $\pi \neq \pi'$. Let $E_h(\pi) = |\text{ENCODE}_n^h(\pi)|$ and $E_h(\pi') = |\text{ENCODE}_n^h(\pi')|$. Note that by Lemma 2, $\text{ENCODE}_n^h(\pi) \neq \text{ENCODE}_n^h(\pi')$, otherwise we could not uniquely decode the permutation. First observe that if both $\text{ENCODE}_n^h(\pi), \text{ENCODE}_n^h(\pi')$ start with bit 0, then

$$E_h(\pi) = 1 + \lceil \log(2^n!) \rceil = E_h(\pi').$$

Combining the facts that $\text{ENCODE}_n^h(\pi) \neq \text{ENCODE}_n^h(\pi')$, but $E_h(\pi) = E_h(\pi')$, we get that $\text{ENCODE}_n^h(\pi)$ is not a prefix of $\text{ENCODE}_n^h(\pi')$.

We have described the encoding formally in the proof of Claim 9. It suffices to prove the lemma for the case when both encodings $\text{ENCODE}_n^h(\pi)$ and $\text{ENCODE}_n^h(\pi')$ start with bit 1. Note that the following n bits denote the size of $X_{h \circ \pi}$. Thus, if $|X_{h \circ \pi}| \neq |X_{h \circ \pi'}|$, the encodings $\text{ENCODE}_n^h(\pi)$ and $\text{ENCODE}_n^h(\pi')$ differ on the first $n + 1$ bits and $\text{ENCODE}_n^h(\pi)$ is not a prefix of $\text{ENCODE}_n^h(\pi')$. On the other hand, if $|X_{h \circ \pi}| = |X_{h \circ \pi'}| = a$ then the lengths $E_h(\pi)$ and $E_h(\pi')$ are equal:

$$E_h(\pi) = 1 + n + 2 \left\lceil \log \binom{2^n}{a} \right\rceil + \lceil \log((2^n - a)!) \rceil = E_h(\pi').$$

In combination with $\text{ENCODE}_n^h(\pi) \neq \text{ENCODE}_n^h(\pi')$, we get $\text{ENCODE}_n^h(\pi)$ is not a prefix of $\text{ENCODE}_n^h(\pi')$, which concludes the proof of the lemma. □

4.4 Bounding the Size of the Encoding

The most important part of the whole proof is of course bounding the expected size of our encoding. Claim 11 bounds the probability that a random y is in the good set (Algorithm 2, line 3), i.e., the set of challenges y where the reduction inverts and the solution returned by SOLVE does not query a preimage of the challenge itself.

Definition 4. Let (R, T_R, C, T_C, p) be any f -oblivious many-one fully black-box reduction from $iOWF$ to TFNP. By $q: \mathbb{N} \rightarrow \mathbb{N}$ we denote the function which upper bounds the number of queries which might be done indirectly, i.e.,

$$q(n) = \max_{\substack{f \in \text{Inj}, x \in \{0,1\}^n \\ y=f(x)}} \{ |Q(C^f(i, s))| \mid i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, y)), s \in \{0,1\}^*, |s| \leq p(|i|) \}.$$

Note that the queries considered above are indirect (i.e., they originated in C and we do not consider f -queries made directly from R).

Lemma 4. Let (R, T_R, C, T_C, p) be any deterministic f -oblivious fully black-box many-one reduction from $iOWF$ to TFNP. Assume $n \in \mathbb{N}$ is large enough so that $n \geq 50$ and $2q(n) + 1 \leq 2^{0.2n}$ where $q(n)$ is the maximal number of f -queries made by C on the queried instance (see Definition 4). Let the success probability of R be $\beta \geq 2^{-0.1n}$, i.e., for any $f \in \text{Inj}$ we have

$$\Pr_{x \leftarrow \{0,1\}^n} [R^{f, \text{SOLVE}}(1^n, f(x)) = x] \geq \beta \geq 2^{-0.1n}.$$

Then

$$\exists h \in \text{Inj}: \mathbb{E}_{\pi \leftarrow \text{Inj}_n^h} [|\text{ENCODE}_n^h(\pi)|] \leq \log 2^n! - \frac{8}{10} n 2^{0.1n}.$$

In Claim 10, we show that for quasi-polynomial algorithms R, C and for large enough n the bound on $q(n)$ used in the statement of Lemma 4 is without loss of generality.

Claim 10. Let (R, T_R, C, T_C, p) be any deterministic f -oblivious many-one fully black-box reduction from $iOWF$ to TFNP. Let T_C be an upper bound on the running time of the algorithm C and T_R be an upper bound on the running time of the algorithm R . We can bound $q(n)$, i.e., the maximal number of f -queries made by C on any queried instance (see Definition 4), by

$$q(n) \leq T_C(2p(T_R(101n^{\log n}))). \quad (5)$$

Moreover, if both $T_C, T_R \in O(n^{\text{polylog}(n)})$ then

$$q(n) \in o(2^{0.2n}). \quad (6)$$

Proof. Here we use the restriction on output lengths of functions from Inj (see Notation 3). Note that these restrictions are reasonable. The case where the length of output of a function depends only on the length of the input is perhaps the most natural. It also makes sense to upper bound the length of the output as for exponential stretch the reduction would just query all possible preimages and thus would be able to invert on its own. On the other hand we cover all natural injective functions – with stretch c , where c is a fixed constant.

For any $x \in \{0,1\}^n$ the reduction $R^{f, \text{SOLVE}}(1^n, f(x))$ is running on input of length at most $n + 100n^{\log n} \leq 101n^{\log n}$ (see Notation 3) and thus can query the TFNP instance i of length at most $T_R(101n^{\log n})$. Thus any considered input of C (consisting of i and $s \in \{0,1\}^*$ of length at most $p(|i|)$) is of length at most $p(T_R(101n^{\log n})) + T_R(101n^{\log n}) \leq 2p(T_R(101n^{\log n}))$. Finally, we can bound the number of queries of C by its running time by $T_C(2p(T_R(101n^{\log n})))$. This concludes the proof of inequality 5.

Note that when both T_C and T_R are in $O(n^{\text{polylog}(n)})$, then also:

$$q(n) \leq T_C(2p(T_R(101n^{\log n}))) \in O(n^{\text{polylog}(n)}) \subseteq o(2^{0.2n}),$$

which proves the inequality 6. \square

We prove the bound on the expected length of the encoding (stated in Lemma 4) with expectation taken not only over the choice of π but also over the choice of $h \leftarrow \text{Inj}$ (the expectation could be also understood to be taken over the choice of $f \leftarrow \text{Inj}$ where $f = h \circ \pi$). Then we use an averaging argument to show that there exists $h \in \text{Inj}$ for which the encoding is short enough.

Now we bound the probability that a given y is “good”. That is the reduction returns the preimage of y and, moreover, the SOLVE query is answered with a solution which is independent of the preimage of y (i.e., there is no query to the preimage of y during computation $C^f(i, s)$, where i is the queried instance and s is the returned solution). We distinguish two bad events:

1. FAIL: the reduction does not invert and
2. HIT: the SOLVE solution queries the preimage.

We bound the probability of both these events and then use them to bound the overall probability of y not being “good”.

For bounding the probability of HIT: the event that SOLVE solution queries the preimage of y , the following two bounds will come handy. We fix the queried instance i , look at only one preimage length n and distinguish two cases:

1. There are many challenges y from the image of $\{0, 1\}^n$ for which the reduction queries the instance i . In this case the bound in Claim 12 will be used to show that the solution does not “help” the reduction too much.
2. Or there exist only few challenges y of the same length n , such that i is queried. In this case we want to bound the probability that $Y_{i,n}$ (i.e., the “protected” images of f_n as in Algorithm 1) becomes “unprotected” (i.e., removed from Y_i and the “benign” solution may query the preimage of some y from $Y_{i,n}$) before the solution is returned. We bound this probability in Claim 13.

Claim 11. *Let (R, T_R, C, T_C, p) be any deterministic f -oblivious many-one fully black-box reduction from $i\text{OWF}$ to TFNP . Let $n \in \mathbb{N}$ be any natural number and let $\pi \in \text{Inj}_n^n$ be any permutation. Let $q(n)$ be the maximal number of f -queries made by C on the queried instance (see Definition 4). Then*

$$\Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [f(x) \notin G_f] \leq \frac{2q(n)}{2^{n/2}} + \Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [R^{f, \text{SOLVE}}(1^n, y) \neq f^{-1}(y)],$$

where both $f = h \circ \pi$ and G_f are as defined in the algorithm ENCODE_n (see line 2 Algorithm 2).

Claim 12 bounds the probability that SOLVE returns a solution querying the preimage of length n for some challenge y on which the reduction is running. We bound the probability for any fixed instance i . The bound is meaningful only for the case that on a given preimage length the reduction queries the instance i often (i.e., for many different preimages x , the reduction running on $f(x)$ queries the i).

In this case, we leverage the fact that there are at most $q(n)$ queries made on each solution. Thus no matter which solution is returned from the SOLVE algorithm, it queries at most $q(n)$ different preimages x . In other words the returned solution could be “useful” for the reduction for at most $q(n)$ out of many different challenges.

Claim 12. *Let (R, T_R, C, T_C, p) be any deterministic f -oblivious many-one fully black-box reduction. Let $i \in \{0, 1\}^*$ be any instance of the corresponding TFNP problem (defined by C). Let*

$$q = \max_{\substack{f \in \text{Inj}, s \in \{0,1\}^* \\ \text{s.t. } |s| \leq p(|i|)}} |Q(C^f(i, s))|.$$

Let $k \in \mathbb{N}$ and let $Y \subseteq \{0, 1\}^*$ be a set of size k and

$$\mathcal{F}_Y = \{f \mid f \in \text{Inj} \text{ and } Y \subseteq \text{Im}(f)\}.$$

Then for any $f \in \mathcal{F}_Y$ and for any $s \in \{0, 1\}^*$, such that $|s| \leq p(|i|)$:

$$\Pr_{y \leftarrow Y}[f^{-1}(y) \in Q(C^f(i, s))] \leq \frac{q}{k}.$$

Proof. The probability follows from the fact that for any $s \in \{0, 1\}^*$, $|s| \leq p(|s|)$ $C^f(i, s)$ makes at most q queries to f (i.e., $|Q(C^f(i, s))| \leq q$). Thus

$$\Pr_{y \leftarrow Y}[f^{-1}(y) \in Q(C^f(i, s))] \leq \frac{|Q(C^f(i, s))|}{|Y|} \leq \frac{q}{k}.$$

□

Now we bound the probability that for some instance i the algorithm SOLVE returns a solution which hits the preimage of some reduction's challenge $f(x)$. This bound is meaningful in the case when the reduction does not query i on many different challenges $f(x')$, where x' is of the same length as x .

Recall that $Y_{i,n}$ is the set of challenges y of length n on which the reduction queries the instance i , i.e., this bound is used when $Y_{i,n}$ is small. Note that if the returned solution queries the preimage of any y from $Y_{i,n}$, it may “help” the reduction non-trivially (imagine $Y_{i,n}$ containing only one element - then SOLVE may potentially reveal all preimages of length n over different instances).

We want to bound the probability that SOLVE stops to “protect” $Y_{i,n}$ before returning the solution (i.e., removes $Y_{i,n}$ from Y_i). Thus we want to show that there is a solution which does not query preimage of any y from $Y_{i,n}$. The problem is that it is not sufficient to show it for a single fixed preimage length n . As even though we have such a solution, there might be smaller $Y_{i,n'}$, such that its preimages are queried on all solutions. Then SOLVE would have to stop “protecting” $Y_{i,n'}$ to be able to return a solution. The problem occurs if SOLVE stops “protecting” $Y_{i,n}$ before $Y_{i,n'}$.

We show that if we stop “protecting” the sets $Y_{i,n}$ in order given by decreasing density ($\frac{|Y_{i,n}|}{2^n}$) this is not the case. Especially we bound the probability that there exists a solution which does not query preimage neither for any y from $Y_{i,n}$ nor for any y from $Y_{i,n'}$, where the density of $Y_{i,n'}$ is smaller (i.e., $\frac{|Y_{i,n'}|}{2^{n'}} \leq \frac{|Y_{i,n}|}{2^n}$).

To bound this probability, we consider another function $g \in \text{Inj}$, which is similar to f , but we “hide” the preimages of all Y_i (the set of all “protected” images of different lengths). This means that for any $x \in f^{-1}[Y_i]$ we let $g(x) = y'$, where y' is chosen outside of the $\text{Im}(f)$. By correctness there is a solution with respect to g that cannot query the preimage of any y from Y_i (as Y_i is not in the image set). We go back to f and show that with high probability this solution remains to be a solution and that it also does not query a preimage of any $y \in Y_i$.

Claim 13. *Let (R, T_R, C, T_C, p) be any f -oblivious many-one black-box reduction. Let $i \in \{0, 1\}^*$ be any instance of the corresponding TFNP problem and*

$$q = \max_{\substack{f \in \text{Inj}, s \in \{0, 1\}^* \\ \text{s.t. } |s| \leq p(|i|)}} |Q(C^f(i, s))|.$$

Let $\delta: \mathbb{N} \rightarrow [0, 1]$ be any function, $Y \subseteq \{0, 1\}^*$ be any finite set, and $\mu \in \mathbb{T}$ be any type. Let \mathcal{F}_Y be the set defined as follows

$$\mathcal{F}_Y = \{f \in \text{Inj} \mid Y \subseteq \text{Im}(f), \tau(f) = \mu, \text{ and } \forall n \in \mathbb{N}: |Y \cap \text{Im}(f_n)| \leq \delta(n)2^n\}, \quad (7)$$

where $f_n = f \upharpoonright \{0, 1\}^n$. Let BAD_f denote the event that

$$\forall s \in \{0, 1\}^* \text{ such that } |s| \leq p(|i|) \text{ either } C^f(i, s) \neq 1 \text{ or } Y \cap f[Q(C^f(i, s))] \neq \emptyset.$$

Suppose \mathcal{F}_Y is nonempty, then

$$\Pr_{f \leftarrow \mathcal{F}_Y} [\text{BAD}_f] \leq \delta(n)q.$$

Intuitively the function δ upper-bounds the *density* of the set Y in the image of f .

Proof. We define the following set

$$\mathcal{H}_Y = \{h \in \text{Inj} \mid \tau(h) = \mu \text{ and } Y \cap \text{Im}(h) = \emptyset\}. \quad (8)$$

Intuitively the set \mathcal{H}_Y is the set of functions “avoiding” the set Y in their image and the set \mathcal{F}_Y is the set of functions “containing” the set Y in the image (where the density of Y on any image length is bounded). Let $X_{f,Y}$ be the set of preimages of Y in f (i.e., $X_{f,Y} = f^{-1}[Y]$).

Now we can bound the probability as follows (we prove the inequality below):

$$\Pr_{f \leftarrow \mathcal{F}_Y} [\text{BAD}_f] \leq \sup_{h \in \mathcal{H}_Y} \Pr_{f \leftarrow \mathcal{F}_Y} [\text{BAD}_f \mid f \upharpoonright \overline{X_{f,Y}} = h \upharpoonright \overline{X_{f,Y}}] \quad (9)$$

Let $\mathcal{H}_{Y,f}$ be defined as follows:

$$\mathcal{H}_{Y,f} = \{h \in \mathcal{H}_Y \mid f \upharpoonright \overline{X_{f,Y}} = h \upharpoonright \overline{X_{f,Y}}\}.$$

We show that the set $\mathcal{H}_{Y,f}$ is finite for any $f \in \mathcal{F}_Y$. Moreover we prove that for any $f, f' \in \mathcal{F}_Y$ the sets $\mathcal{H}_{Y,f}$ and $\mathcal{H}_{Y,f'}$ are of the same size. We compute the size of the set $\mathcal{H}_{Y,f}$ as follows:

$$|\mathcal{H}_{Y,f}| = |\{h \in \mathcal{H}_Y \mid f \upharpoonright \overline{X_{f,Y}} = h \upharpoonright \overline{X_{f,Y}}\}| = \quad (10)$$

$$= \prod_{\substack{n \in \mathbb{N}, m = \mu(n) \\ \text{s.t. } \{0,1\}^m \cap Y \neq \emptyset}} \left(\binom{2^m - 2^n}{|Y \cap \{0,1\}^m|} (|Y \cap \{0,1\}^m|!) \right). \quad (11)$$

By definition, both $f, h \in \text{Inj}$ (see equation 7 - definition of \mathcal{F}_Y and equation 8 - definition of \mathcal{H}_Y), and furthermore $\text{Im}(h) \cap Y = \emptyset$ (see equation 8 for definition of \mathcal{H}_Y). The equation 11 follows from the fact that there are $\binom{2^m - 2^n}{|Y \cap \{0,1\}^m|} (|Y \cap \{0,1\}^m|)!$ possibilities how to define $h \upharpoonright (X_f \cap \{0,1\}^n)$. Note that size of $\mathcal{H}_{Y,f}$ is finite, as Y is finite set and thus the product contains finitely many factors. Observe also that $|\mathcal{H}_{Y,f'}| = |\mathcal{H}_{Y,f}|$ as the size depends only on the type μ and set Y , but is not dependent on the exact mapping f .

It suffices to prove that for each $h \in \mathcal{H}_Y$:

$$\Pr_{f \leftarrow \mathcal{F}_Y} [\text{BAD}_f \mid f \upharpoonright \overline{X_{f,Y}} = h \upharpoonright \overline{X_{f,Y}}] \leq \delta(n)q.$$

By the correctness of the underlying TFNP problem (see Definition 2 for the exact order of the quantifiers), there exists $s \in \{0,1\}^*$ such that $|s| \leq p(i)$ and $C^h(i, s) = 1$. Let Q be the set of queries made on the solution $C^h(i, s)$, i.e., $Q = Q(C^h(i, s))$. Note that if $Q \cap X_{f,Y}$ is empty then all queries C makes on input i, s are answered equally for h as for f thus s is solution also with respect to f (i.e., $C^f(i, s) = 1$) and $Q(C^f(i, s))$ has empty intersection with $X_{f,Y}$ (as $Q(C^f(i, s)) = Q$). Thus we may rewrite the probability as follows:

$$\begin{aligned} \Pr_{f \leftarrow \mathcal{F}_Y} [\text{BAD}_f \mid f \upharpoonright \overline{X_{f,Y}} = h \upharpoonright \overline{X_{f,Y}}] &\leq \Pr_{f \leftarrow \mathcal{F}_Y} [X_{f,Y} \cap Q \neq \emptyset] \\ &\leq \sum_{n \in \mathbb{N}} \Pr_{f \leftarrow \mathcal{F}_Y} [X_{f,Y} \cap Q \cap \{0,1\}^n \neq \emptyset]. \quad (\text{by union bound}) \end{aligned}$$

We bound the probability for a single n and the corresponding output length $m = |f(0^n)|$. Let q_n denote the number of queries of length n , i.e., $q_n = |Q \cap \{0, 1\}^n|$. We may bound the probability as follows:

$$\Pr_{f \leftarrow \mathcal{F}_Y} [X_{f,Y} \cap Q \cap \{0, 1\}^n \neq \emptyset] \leq \sup_{k \in \mathbb{N}} \Pr_{\substack{f \leftarrow \mathcal{F}_Y \\ |X_{f,Y} \cap \{0,1\}^n| = k}} [X_{f,Y} \cap Q \cap \{0, 1\}^n \neq \emptyset] \quad (12)$$

$$\leq \sup_{k \in \mathbb{N}} \frac{|Q \cap \{0, 1\}^n| \cdot \binom{2^n - 1}{k-1}}{\binom{2^n}{k}} \quad (13)$$

$$= \sup_{k \in \mathbb{N}} \frac{k}{2^n} q_n \quad (14)$$

The inequality 12 follows from the fact that condition $|X_{f,Y} \cap \{0, 1\}^n| = k$ over different k creates a partition of the probabilistic space (there are only finitely many values of k with non-zero probability which means we may use Lemma 5). The inequality 13 is computed as follows:

1. We have $\binom{2^n}{k}$ possibilities how to choose a subset of $\{0, 1\}^n$ of size k (the set $X_{f,Y} \cap \{0, 1\}^n$).
2. But there are at most $|Q \cap \{0, 1\}^n| \cdot \binom{2^n - 1}{k-1}$ ways to choose the set in such a way that it has nonempty intersection with Q . We have to pick at least one x from $Q \cap \{0, 1\}^n$ and the remaining $k - 1$ elements can be chosen arbitrarily from the remaining $2^n - 1$ elements.

Now we use the definition of \mathcal{F}_Y (see Equation 7) to bound $\frac{k}{2^n}$ and we get the following bound:

$$\Pr_{f \leftarrow \mathcal{F}_Y} [X_{f,Y} \cap Q \cap \{0, 1\}^n \neq \emptyset] \leq \sup_{k \in \mathbb{N}} \frac{k}{2^n} q_n \leq \delta(n) q_n. \quad (15)$$

Finally we summarize all the inequalities above:

$$\begin{aligned} \Pr_{f \leftarrow \mathcal{F}_Y} [\text{BAD}_f] &\leq \sup_{h \in \mathcal{H}_Y} \Pr_{f \leftarrow \mathcal{F}_Y} [\text{BAD}_f \mid f \upharpoonright \overline{X_{f,Y}} = h \upharpoonright \overline{X_{f,Y}}] && \text{(see inequality 9)} \\ &\leq \sup_{h \in \mathcal{H}_Y} \sum_{n \in \mathbb{N}} \Pr_{f \leftarrow \mathcal{F}_Y} [X_{f,Y} \cap Q \cap \{0, 1\}^n \neq \emptyset] \\ &\leq \sup_{h \in \mathcal{H}_Y} \sum_{n \in \mathbb{N}} \delta(n) q_n. \end{aligned}$$

Using the definition of $q_n = |Q \cap \{0, 1\}^n|$ we can bound $\sum_{n \in \mathbb{N}} q_n = \sum_{n \in \mathbb{N}} |Q \cap \{0, 1\}^n| = |Q| \leq q$, by assumption. This concludes the proof as the sum $\sum_{n \in \mathbb{N}} \delta(n) q_n$ can be upper bounded by $\delta(n) q$ and thus $\Pr_{f \leftarrow \mathcal{F}_Y} [\text{BAD}_f] \leq \delta(n) q$. □

We leverage the following fact from probability theory.

Lemma 5. *Let A be a random event and let $\mathcal{C} = \{C_1, C_2, \dots, C_n\}$ be a finite set of disjoint random events that cover the whole probability space. That is*

$$\forall i \neq j: \Pr[C_i \cap C_j] = 0 \text{ and } \Pr[C_1 \cup C_2 \cup \dots \cup C_n] = 1.$$

Where we define $\Pr[A \mid C] = 0$ whenever $\Pr[C] = 0$.

Then the following inequality holds:

$$\Pr[A] \leq \max_{C \in \mathcal{C}} \Pr[A \mid C].$$

Proof. Since \mathcal{C} forms a partition of the probability space we may write:

$$\begin{aligned} \Pr[A] &= \sum_{C' \in \mathcal{C}} \Pr[A | C'] \Pr[C'] \\ &\leq \sum_{C' \in \mathcal{C}} (\max_{C \in \mathcal{C}} \Pr[A | C]) \Pr[C'] \\ &= \max_{C \in \mathcal{C}} \Pr[A | C] \end{aligned}$$

□

We are ready to lower bound the probability that we can compute x from $f(x)$ using the security reduction R and f defined on values different from x , i.e., we upper bound the following probability

$$\Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [f(x) \notin G_f],$$

where $\pi \in \text{Inj}_n^n$ is a permutation and G_f is as defined in the algorithm ENCODE_n (see line 3 of Algorithm 2). Intuitively there is a trade-off based on the density: we use Claim 13 when the density is small and we use Claim 12 when the density is large enough.

Proof (of Claim 11). Recall that $y \in \text{Im}(f_n)$ is in G_f if it satisfies the following two conditions:

1. $R^{f, \text{SOLVE}}(1^n, y)$ inverts f on y , and
2. there is no indirect query for y (i.e., $f^{-1}(y) \notin Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y))$).

Thus, we bound the probability using union bound as follows:

$$\Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [f(x) \notin G_f] \leq \Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [x \notin \text{INV}_f] + \Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, f(x)))].$$

By definition of INV_f (see line 2 of Algorithm 2)

$$\Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [x \notin \text{INV}_f] = \Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [R^{f, \text{SOLVE}}(1^n, f(x)) \neq x].$$

It suffices to show that

$$\Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, f(x)))] \leq \frac{2q(n)}{2^{n/2}}. \quad (16)$$

Let us call P_n the set of possible TFNP instances that the reduction might query:

$$P_n = \{i \in \{0,1\}^* \mid \exists x \in \{0,1\}^n \exists h \in \text{Inj}, f = h \circ \pi : i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, f(x)))\}.$$

Observe that as the running time of R is limited thus the size of P_n is finite. Moreover since R is a deterministic many-one reduction the expression $i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, f(x)))$ is equivalent with $\{i\} = Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, f(x)))$, thus $\mathcal{C} = \{i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, f(x)))\}_{i \in P_n}$ defines a partition of the probability space and we may use Lemma 5 to bound:

$$\Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, f(x)))] \quad (17)$$

$$\leq \max_{i \in P_n} \Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, f(x))) \mid i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, f(x)))]. \quad (18)$$

This allows us to fix any instance $i \in \{0, 1\}^*$ and bound the probability that x was indirectly queried only for the case when TFNP instance i was queried.

We use the following notation in the rest of the proof. Let $t_{i,f}^{\text{end}} = j$ if the algorithm $\text{SOLVE}^f(i)$ returns the solution in the j -th iteration of the while loop (see line 7 of Algorithm 1). Furthermore let $t_{i,f,n}^{\text{del}}$ denote the iteration of the same while loop in which $Y_{i,n}$ is removed from Y_i (see line 14 of Algorithm 1). Note that we set $t_{i,f,n}^{\text{del}} = \infty$ in the case when $Y_{i,n}$ is never removed from Y_i .

Let $\delta: \mathbb{N} \rightarrow [0, 1]$ be a function chosen later in the proof. We can rewrite the probability as follows:

$$\Pr[x \in Q_f^{\text{indir}}(R^{f,\text{SOLVE}}(1^n, y))] = \tag{19}$$

$$= \Pr \left[x \in Q_f^{\text{indir}}(R^{f,\text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} \leq t_{i,f,n}^{\text{del}} \right] \Pr \left[t_{i,f}^{\text{end}} \leq t_{i,f,n}^{\text{del}} \right] \tag{20}$$

$$+ \Pr \left[x \in Q_f^{\text{indir}}(R^{f,\text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \right] \Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \right] \tag{21}$$

$$+ \Pr \left[x \in Q_f^{\text{indir}}(R^{f,\text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} > \delta(n) \right] \Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} > \delta(n) \right], \tag{22}$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0, 1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f,\text{SOLVE}}(1^n, y))$ holds (the TFNP instance i is fixed and we chose uniformly at random x, h such that i is queried).

Observe that,

$$\Pr \left[x \in Q_f^{\text{indir}}(R^{f,\text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} \leq t_{i,f,n}^{\text{del}} \right] = 0$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0, 1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f,\text{SOLVE}}(1^n, y))$ holds. This holds because the returned solution does not contain preimage of any element from Y_i by definition and $y \in Y_i$ during the iteration $t_{i,f}^{\text{end}}$.

Using this and bounding some of the probabilities by 1 we get the following bound:

$$\Pr[x \in Q_f^{\text{indir}}(R^{f,\text{SOLVE}}(1^n, y))] \leq \Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \right] + \Pr \left[x \in Q_f^{\text{indir}}(R^{f,\text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} > \delta(n) \right],$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0, 1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f,\text{SOLVE}}(1^n, y))$ holds.

We bound the second probability using Claim 12 for instance i , $k = |Y_{i,n}| \geq \delta(n)2^n$, the set $Y = Y_{i,n}$ and $q(n)$ as the upper bound on the number of f -queries made by C . This bounds the probability by

$$\Pr \left[x \in Q_f^{\text{indir}}(R^{f,\text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} > \delta(n) \right] \leq \frac{q(n)}{k} \leq \frac{q(n)}{\delta(n)2^n},$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0, 1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f,\text{SOLVE}}(1^n, y))$ holds.

We can use Claim 13 to bound the first probability. We actually bound the probability that $Y_{i,n}$ is not deleted before solve returns supposing that $Y_{i,n}$ has small density, i.e.,

$$\Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \right] \leq \Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \mid \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \right],$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0, 1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f,\text{SOLVE}}(1^n, y))$ holds.

We want to argue that $Y_{i,n}$ is not deleted before SOLVE returns. There might be preimage lengths n' such that whenever $Y_{i,n}$ is “protected” $Y_{i,n'}$ is also “protected” (in particular all n' satisfying $|Y_{i,n'}|/2^{n'} \leq |Y_{i,n}|/2^n$).

Thus we need to use Claim 13 for the set Y which is a union of $Y_{i,n'}$ with small density (i.e., small $|Y_{i,n'}|/2^{n'}$). As the density is computed with respect to the length of the preimage (not the length of the image) it is practical to consider probability just over functions of the same type. For these reasons we upper bound the probability as follows:

$$\Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \mid \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \right] \leq \sup_{\mu \in \mathbb{T}} \Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \mid \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \text{ and } \tau(h) = \mu \right],$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0,1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f,\text{SOLVE}}(1^n, y))$ holds. Note that the set of all types is not even countable so we cannot use Lemma 5 directly. However for a fixed n we only care about type up to some input length (given by the running time of R). Thus we could define a partition of the probability space to finitely many events. However this notation is a bit cleaner.

To bound the right hand side we fix some type μ and continue just with functions of the type μ .

We need to determine the set Y we use in the Claim 13. Let Y_i^{del} be the set Y_i as defined in the algorithm SOLVE in the iteration in which $Y_{i,n}$ is deleted. We want $Y \subseteq \{0,1\}^*$ to match the set Y_i^{del} . More precisely we want Y such that it can be partitioned into sets $Y_{i,n'}$ where $n' \in \mathbb{N}$ in such a way that

1. $Y = \bigcup_{n' \in \mathbb{N}} Y_{i,n'}$,
2. $\forall n' \in \mathbb{N} \forall y \in Y_{i,n'}: |y| = \mu(n)$, and
3. $\forall n' \in \mathbb{N}: |Y_{i,n'}| \leq \delta(n)/2^{n'}$.

We can bound the probability as follows

$$\Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \mid \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \text{ and } \tau(h) = \mu \right] \leq \max_Y \Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \mid \frac{|Y_{i,n}|}{2^n} \leq \delta(n), \tau(h) = \mu, \text{ and } Y_i^{\text{del}} = Y \right],$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0,1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f,\text{SOLVE}}(1^n, y))$ holds and the supremum is over all sets Y satisfying the conditions above.

The first two conditions are satisfied trivially by setting each $Y_{i,n'}$ to the corresponding set in SOLVE. The third condition follows from the fact that we are in the iteration in which $Y_{i,n}$ should be removed from Y_i . We are removing the set $Y_{i,n}$ with the highest density and $|Y_{i,n}|/2^n \leq \delta(n)$ thus all other sets $Y_{i,n'} \subseteq Y_i^{\text{del}}$ have density at most $\delta(n)$.

Now we fix any set Y satisfying the conditions above and we upper bound the probability using Claim 13 for instance i , $q(n)$ being an upper bound on the number of f -queries made by C , $\delta(n)$ being the density of challenges y and our fixed set Y .

Recall the definition of \mathcal{F}_Y from Claim 13:

$$\mathcal{F}_Y = \{f \in \text{Inj} \mid Y \subseteq \text{Im}(f) \text{ and } \forall n \in \mathbb{N}: |Y \cap \text{Im}(f_n)| \leq \delta(n)2^n\}, \quad (23)$$

Observe that \mathcal{F}_Y is not empty. Claim 13 gives us that

$$\Pr_{f \leftarrow \mathcal{F}_Y} [\forall s \in \{0,1\}^* \text{ such that } |s| \leq p(|i|): \text{ either } C^f(i, s) \neq 1 \text{ or } Y \cap f[Q(C^f(i, s))] \neq \emptyset] \leq \delta(n)q(n).$$

Observe that by the definition of $Y = Y_i^{\text{del}}$ and \mathcal{F}_Y (see Claim 13), this bounds the probability that there was a suitable solution for returning before $Y_{i,n}$ was deleted from Y_i , i.e., it bounds the probability $t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}}$ and thus we get

$$\Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \mid \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \text{ and } \tau(h) = \mu \text{ and } Y_i^{\text{del}} = Y \right] \leq \delta(n)q(n),$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0,1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f,\text{SOLVE}}(1^n, y))$ holds.

Summarizing the above inequalities we get

$$\begin{aligned}
& \Pr[x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y))] \\
& \leq \Pr \left[t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} \leq \delta(n) \right] + \Pr \left[x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} > t_{i,f,n}^{\text{del}} \text{ and } \frac{|Y_{i,n}|}{2^n} > \delta(n) \right] \\
& \leq \frac{q(n)}{\delta(n)2^n} + \delta(n)q(n),
\end{aligned}$$

where the probability is over the choice of $h \leftarrow \text{Inj}$, $x \leftarrow \{0, 1\}^n$ such that when we set $f = h \circ \pi$, $y = f(x)$ then $i \in Q_{\text{SOLVE}}(R^{f, \text{SOLVE}}(1^n, y))$ holds.

Finally we set $\delta(n) = \frac{1}{2^{n/2}}$ to get

$$\Pr_{\substack{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj} \\ f = h \circ \pi}} [x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, f(x)))] \leq \frac{q(n)}{\delta(n)2^n} + \delta(n)q(n) \leq \frac{2q(n)}{2^{n/2}}.$$

This concludes the proof of Claim 11 as we proved that

$$\begin{aligned}
\Pr_{\substack{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj} \\ f = h \circ \pi}} [f(x) \notin G_f] & \leq \Pr[x \notin \text{INV}_f] + \Pr[x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, f(x)))] \\
& \leq \Pr_{x \leftarrow \{0,1\}^n, h \leftarrow \text{Inj}} [R^{f, \text{SOLVE}}(1^n, y) \neq f^{-1}(y)] + \frac{2q(n)}{2^{n/2}}.
\end{aligned}$$

□

Finally, we upper bound the expected size of encoding produced by ENCODE_n (see Algorithm 2) stated in Lemma 4.

Proof (of Lemma 4). We denote $E_h(\pi)$ the size of encoding, i.e., $E_h(\pi) = |\text{ENCODE}_n^h(\pi)|$ and set $f = h \circ \pi$. First we bound the size of encoding for permutation π in the case that ENCODE_n returned $\mathcal{M} = (1, |X_f|, Y_f, X_f, \sigma)$, i.e., $|X_f| \geq 2^{0.6n}$. We denote the size of X_f by A and set $N = 2^n$. As stated in Claim 9 the size of encoding from algorithm ENCODE_n in case $A \geq 2^{0.6n}$ is

$$\begin{aligned}
\mathbb{E}_{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj}} [E_h(\pi) \mid A \geq 2^{0.6n}] & = 1 + n + 2 \left[\log \left(\binom{N}{A} \right) \right] + \lceil \log((N - A)!) \rceil \\
& \leq 4 + n + \log \left(\binom{N}{A}^2 \cdot (N - A)! \right) \\
& = 4 + n + \log(N!) + \log \left(\binom{N}{A} \frac{1}{A!} \right).
\end{aligned}$$

We bound the last term as follows:

$$\begin{aligned}
\log \left(\binom{N}{A} \frac{1}{A!} \right) & \leq \log \left(\left(\frac{eN}{A} \right)^A \left(\frac{e}{A} \right)^A \right) \\
& = A(n + 2 \log e - 2 \log A).
\end{aligned}$$

Recall that by assumption $2^{0.6n} \leq A$ on the other hand $A = |X_f| \leq 2^n$ as $X_f \subseteq \{0, 1\}^n$. Now we differentiate by A and get

$$(A(n + 2 \log e - 2 \log A))' = n + 2 \log e - 2 \log A - \frac{2}{\ln 2}.$$

The derivative is negative and the function is decreasing for $A \in [2^{0.6n}, 2^n]$ and $n \geq 50$. We can bound

$$\begin{aligned} \log \left(\binom{N}{A} \frac{1}{A!} \right) &\leq A(n + 2 \log e - 2 \log A) \\ &\leq 2^{0.6n} (n + 2 \log e - 2 \log 2^{0.6n}) && \text{(decreasing for } A \in [2^{0.6n}, 2^n]) \\ &\leq 2^{0.6n} (-0.1n). && \text{(using } n \geq 50) \end{aligned}$$

If we combine the above bounds, we get the following bound on the size of the encoding for case $A \geq 2^{0.6n}$:

$$\begin{aligned} \mathbb{E}_{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj}} [E_h(\pi) \mid A \geq 2^{0.6n}] &= 4 + n + \log(N!) + \log \left(\binom{N}{A} \frac{1}{A!} \right) \leq \\ &\leq 2n + \log(N!) + 2^{0.6n} (-0.1n) \leq \\ &\leq \log(N!) - n2^{0.2n}. \end{aligned}$$

Now we bound the probability that $|X_f| \geq 2^{0.6n}$. By assumption of the lemma we can bound $(2q(n) + 1) \leq 2^{0.2n}$ and we get

$$\begin{aligned} \Pr_{\substack{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj} \\ f = h \circ \pi}} [|X_f| \geq 2^{0.6n}] &\geq \Pr[|G_f| \geq (2q(n) + 1)2^{0.6n}] \\ &= \Pr[|\overline{G}_f| \leq 2^n - (2q(n) + 1)2^{0.6n}] \\ &= \Pr[|\overline{G}_f| \leq 2^n - 2^{0.8n}]. && (2q(n) + 1 \leq 2^{0.2n}) \end{aligned}$$

By Claim 11 we can bound the expected size of \overline{G}_f as follows:

$$\begin{aligned} \mathbb{E}_{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj}, f = h \circ \pi} [|\overline{G}_f|] &\leq \left(\frac{2q(n)}{2^{n/2}} + (1 - \beta(n)) \right) 2^n \\ &\leq \left(\frac{2q(n)}{2^{n/2}} + (1 - 2^{-0.1n}) \right) 2^n && \text{(as } \beta(n) \geq 2^{-0.1n}) \\ &\leq 2 \cdot 2^{0.7n} + 2^n - 2^{0.9n} && \text{(as } 2q(n) \leq 2^{0.2n}) \\ &\leq 2^n - \frac{9}{10} 2^{0.9n}. && \text{(as } n \geq 50) \end{aligned}$$

We can use the Markov inequality to bound

$$\Pr_{\substack{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj} \\ f = h \circ \pi}} [|\overline{G}_f| > 2^n - 2^{0.8n}] \leq \frac{2^n - \frac{9}{10} 2^{0.9n}}{2^n - 2^{0.8n}}.$$

Combining these inequalities we get:

$$\begin{aligned} \Pr_{\substack{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj} \\ f = h \circ \pi}} [|X_f| \geq 2^{0.6n}] &\geq \Pr[|\overline{G}_f| \leq 2^n - 2^{0.8n}] \\ &= 1 - \Pr[|\overline{G}_f| > 2^n - 2^{0.8n}] \\ &\geq 1 - \frac{2^n - \frac{9}{10} 2^{0.9n}}{2^n - 2^{0.8n}} \\ &\geq \frac{8}{10} 2^{-0.1n}. && (n \geq 50) \end{aligned}$$

We are ready to bound the size of the encoding as follows:

$$\begin{aligned}
\mathbb{E}_{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj}}[E_h(\pi)] &= \Pr_{\pi, h} [|X_{\pi \circ h}| < 2^{0.6n}] (1 + \log N!) + \Pr_{\pi, h} [|X_{\pi \circ h}| \geq 2^{0.6n}] \mathbb{E}_{\pi, h}[E_h(\pi) \mid |X_{\pi \circ h}| \geq 2^{0.6n}] \\
&\leq \Pr_{\pi, h} [|X_{\pi \circ h}| < 2^{0.6n}] (1 + \log N!) + \Pr_{\pi, h} [|X_{\pi \circ h}| \geq 2^{0.6n}] (\log(N!) - n2^{0.2n}) \\
&\leq 1 + \log(N!) - \Pr_{\pi, h} [|X_{\pi \circ h}| \geq 2^{0.6n}] n2^{0.2n} \\
&\leq 1 + \log(N!) - \left(\frac{8}{10} 2^{-0.1n}\right) n2^{0.2n} \\
&\leq \log(N!) - \frac{8}{10} n2^{0.1n}.
\end{aligned}$$

Finally by averaging argument there exists $h \in \text{Inj}$ such that $\mathbb{E}_{\pi \leftarrow \text{Inj}_n^n}[E_h(\pi)] \leq \log(N!) - \frac{8}{10} n2^{0.1n}$. \square

5 Extensions

5.1 Non-Adaptive Reductions

It is possible to extend our proof from Section 3.2 to rule out even *non-adaptive security reductions* which submit multiple queries to the oracle SOLVE in parallel, though still f -obliviously, as defined in Definition 3.

Notice that the algorithms SOLVE, ENCODE $_n$, DECODE $_n$, and SOLVESIM are well defined even for non-adaptive reductions and we can use them without any change. Our analysis showing that SOLVE always returns a solution (Lemma 1), that DECODE $_n$ correctly decodes the permutation (Lemma 2), and that the encoding is prefix-free (Lemma 3) remain unchanged as well. The statement of Lemma 4 also holds without a change.

The only change in our proof would be in the proof of Claim 11 in Equation 17, where we are upper bounding the probability of an indirect hit. To deal with non-adaptive reductions, we would have to partition the probability space by sets of queried TFNP instances and then use union bound to bound the probability that there would be a queried instance causing an indirect hit. Thus the maximum in Equation 18 will be multiplied by $T_R(n + 100n^{\log(n)})$ (the upper bound on the number of SOLVE queries).

Accordingly, the bound on the indirect hit (Equation 16) changes to

$$\Pr_{\substack{x \in \{0,1\}^n, h \in \text{Inj} \\ f = h \circ \pi}} [x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, f(x)))] \leq \frac{T_R(n + 100n^{\log(n)}) \cdot 2q(n)}{2^{n/2}}. \quad (24)$$

This is still $\mathcal{O}(n^{\text{polylog}(n)})$. This propagates in the proof of Lemma 4 when bounding

$$\mathbb{E}_{\substack{\pi \leftarrow \text{Inj}_n^n, h \leftarrow \text{Inj} \\ f = h \circ \pi}} [|\overline{G_f}|]$$

in the same way as we bound the term by $2^{0.2n}$. Therefore, we can prove the following strengthening of Theorem 1.

Theorem 14. *There is no fully black-box construction (R, T_R, C, T_C, p) of a hard TFNP problem from injective one-way functions with deterministic f -oblivious non-adaptive reduction with success probability at least $2^{-0.1n}$ such that both running times $T_R, T_C \in O(n^{\text{polylog}(n)})$.*

5.2 Randomized Reductions

In this section, we describe how to generalize our proof to handle fully black-box constructions of hard TFNP problems from iOWF with *randomized security reductions*.

Algorithm 5: The oracle SOLVE (for a randomized security reduction).

Hardwired : a fully black-box construction (R, T_R, C, T_C, p) of a hard TFNP problem from iOWF
Oracle access: an injective function $f = \{f_n\}_{n \in \mathbb{N}} \in \text{Inj}$
Input : an instance $i \in \{0, 1\}^*$
Output : a solution $s \in \{0, 1\}^*$ such that $C^f(i, s) = 1$

- 1 Compute $Y_i = \bigcup_{n=1}^{T_C(|i|+p(|i|))} \{y \in \text{Im}(f_n) \mid \Pr_{r \leftarrow \{0,1\}^*} [i \in Q_{\text{SOLVE}}(R^f, \text{SOLVE}(1^n, y; r))] > 0\}$
- 2 **for** $y \in Y_i$ **do**
- 3 Let $n = |f^{-1}(y)|$
- 4 Set // compute probability reduction runs on y conditioned on i has been queried
- 5
$$p_i(y) = \frac{\Pr_{r \leftarrow \{0,1\}^*} [i \in Q_{\text{SOLVE}}(R^f, \text{SOLVE}(1^n, y; r))]}{\Pr_{r \leftarrow \{0,1\}^*, y' \leftarrow Y_i, |y'|=|y|} [i \in Q_{\text{SOLVE}}(R^f, \text{SOLVE}(1^n, y'; r))]}$$
- 6 **end**
- 7 Compute $S_{i,f} = \{s \in \{0, 1\}^* \mid |s| \leq p(|i|) \text{ and } C^f(i, s) = 1\}$
- 8 **while** *True* **do**
- 9 $B_{i,f} = \{s \in S_{i,f} \mid f[Q(C^f(i, s))] \cap Y_i = \emptyset\}$
- 10 **if** $B_{i,f} \neq \emptyset$ **then**
- 11 | **return** lexicographically smallest $s \in B_{i,f}$
- 12 **end**
- 13 Choose $y \in Y_i$ such that $p_i(y)$ is minimized.
- 14 Set $Y_i = Y_i \setminus \{y\}$
- 15 **end**

First, we need to change our algorithm SOLVE. One could imagine that the construction has some instance i created for a concrete challenge y , on which R queries i with high probability. But R might also query the instance i for many other challenges y' (on each of them with small probability) to hide the real challenge y . Thus we need to take the probability of querying the instance i into account.

Our new SOLVE works as shown in Algorithm 5. Note that it is an extension of the previous algorithm (i.e., if the security reduction ignores its random bits, SOLVE removes strings from Y_i in the same order but Algorithm 5 removes challenges one by one like instead of in “batches” $Y_{i,n}$ like Algorithm 1). The proof that the new SOLVE returns a solution is fairly similar to the proof of Lemma 1. Note that due to f -obliviousness we are still able to simulate this algorithm (all added computation depends only on the $\text{Im}(f)$ and is independent of the concrete mapping f).

We also need to change the algorithms ENCODE_n and DECODE_n as these algorithms need to share the randomness used for the reduction. Thus we equip them with a set of strings r_y for $y \in \text{Im}(h_n)$, where every r_y will be used for running the security reduction on challenge y . Note that we can bound the length of each r_y by $T_R(n + 100n^{\log n})$. This is because we can bound the input on which R is run by $n + 100n^{\log n}$ and thus bound the running time of R by $T_R(n + 100n^{\log n})$. If r_y is longer the reduction is not able to use it.

We want encoding to be uniquely decodable and prefix-free. We show this not just for any choice of $h \in \text{Inj}$ but also for any choice of the random bits r_y , where $y \in \text{Im}(h_n)$. The proofs follows the same blueprint as the proofs of Lemmas 2 and 3.

The main difference is in the part when we are bounding the size of the encoding (Section 4.4). We need to change Claims 12 and 13 to reflect the fact that we are removing challenges y from the set Y_i based on the probabilities $p_i(y)$.

Both statements of Claims 12 and 13 use a finite set of challenges Y . In our new version of these claims we partition the set $Y = \bigcup_{n \in \mathbb{N}} Y_n$ where each $Y_n \subseteq \{0, 1\}^{\mu(n)}$ (where μ is a fixed function type) and for each Y_n we have a distribution $\mathcal{D}_{Y,n}$.

Here we describe an alternative to Claim 12. Instead of the set Y from the claim we consider Y with the distributions $\mathcal{D}_{Y,n}$, note that for this we need to consider just functions of some fixed type μ (thus \mathcal{F}_Y is restricted to functions of type μ). We take $Y' = \{y \in Y \mid \Pr_{z \leftarrow \mathcal{D}_Y} [z = y] \leq \gamma (|f^{-1}(y)|)\}$. By union bound (k -

th query was on a preimage of some $y \in Y$) we can show that for any $f \in \text{Inj}$ and any $s \in \{0, 1\}^*$, $|s| \leq p(|i|)$:

$$\forall n \in \mathbb{N}: \Pr_{y \leftarrow \mathcal{D}_{Y,n}} [y \in Y' \text{ and } f^{-1}(y) \in Q(C^f(i, s))] \leq q\gamma(n).$$

In the alternative to Claim 13 we also use any finite set Y . Now we assume that for each $n \in \mathbb{N}$ and for each $y \in Y_n$ $\Pr_{z \leftarrow \mathcal{D}_{Y,n}} [z = y] \geq \gamma(n)$. Recall the definition of \mathcal{F}_y (Claim 13):

$$\mathcal{F}_Y = \{f \in \text{Inj} \mid Y \subseteq \text{Im}(f), \tau(f) = \mu, \text{ and } \forall n \in \mathbb{N}: |Y \cap \text{Im}(f_n)| \leq \delta(n)2^n\}. \quad (25)$$

Observe that for each n we have $|Y \cap \text{Im}(f_n)| = |Y_n| \leq 1/\gamma(n)$ and thus we may use Claim 13 for δ defined as $\forall n \in \mathbb{N}: \delta(n) = \frac{1}{\gamma(n)2^n}$ to get that

$$\Pr_{f \leftarrow \mathcal{F}_Y} [\forall s \in \{0, 1\}^* \text{ such that } |s| \leq p(|i|) \text{ either } C^f(i, s) \neq 1 \text{ or } Y \cap f[Q(C^f(i, s))] \neq \emptyset] \leq \frac{q}{\gamma(n)2^n}.$$

To put things together like in Claim 11 we again follow the flow of the original proof. We need to redefine the deletion time $t_{i,f,n}^{\text{del}}$ when the set $Y_{i,n}$ was deleted from Y_i . We instead work with the time $t_{i,f,y}^{\text{del}}$ when the single challenge y was deleted for each $y \in Y_i$.

The bound Equation 19 is substituted by the following:

$$\begin{aligned} \Pr[x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y))] &= \\ &= \Pr[x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} \leq t_{i,f,y}^{\text{del}}] \Pr[t_{i,f}^{\text{end}} \leq t_{i,f,y}^{\text{del}}] \\ &+ \Pr[x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} > t_{i,f,y}^{\text{del}} \text{ and } p_i(y) \geq \gamma(n)] \Pr[t_{i,f}^{\text{end}} > t_{i,f,y}^{\text{del}} \text{ and } p_i(y) \geq \gamma(n)] \\ &+ \Pr[x \in Q_f^{\text{indir}}(R^{f, \text{SOLVE}}(1^n, y)) \mid t_{i,f}^{\text{end}} > t_{i,f,y}^{\text{del}} \text{ and } p_i(y) < \gamma(n)] \Pr[t_{i,f}^{\text{end}} > t_{i,f,y}^{\text{del}} \text{ and } p_i(y) < \gamma(n)], \end{aligned}$$

Now similarly as in the previous proof we show that the first summand is zero, the second can be bounded using our alternative of Claim 13 and the third by the alternative of Claim 12 (where we additionally need to fix the type of the function similarly as we did when we were using Claim 13 in the original proof).

After the averaging argument which chooses the function h we average again to get suitable randomness for the security reduction (both are independent of π).

This can be combined with the previous changes to get the full strength theorem:

Theorem 15. *There is no fully black-box construction (R, T_R, C, T_C, p) of a hard TFNP problem from injective one-way functions with randomized f -oblivious non-adaptive reduction with success probability at least $2^{-0.1n}$ such that both running times $T_R, T_C \in O(n^{\text{polylog}(n)})$.*

6 Conclusions

In this work, we have shown that there are intrinsic barriers preventing *simple* fully black-box constructions of hard TFNP problems from injective one-way functions. The main technical contribution of our work is the technique of designing a ‘‘TFNP-breaker’’ oracle SOLVE which depends on the reduction.

The natural direction towards extending our results would be attempting to lift the restriction to f -oblivious and non-adaptive reductions. One reason for why this might be challenging is that a black-box separation of TFNP and injective one-way functions would subsume the separation of collision-resistant hash functions and one-way functions [Sim98], for which, despite the recent progress, there are only non-trivial proofs [MM11, HHR15, BD19].

Acknowledgements

We wish to thank the anonymous reviewers for their comments which helped us to clarify the presentations of our results.

References

- AKV04. Tim Abbot, Daniel Kane, and Paul Valiant. On algorithms for Nash equilibria. Unpublished manuscript, 2004. <http://web.mit.edu/tabbott/Public/final.pdf>.
- Bae14. Paul Baecker. Simon’s circuit. *IACR Cryptol. ePrint Arch.*, 2014:476, 2014.
- BCE⁺98. Paul Beame, Stephen A. Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. The relative complexity of NP search problems. *J. Comput. Syst. Sci.*, 57(1):3–19, 1998.
- BD19. Nir Bitansky and Akshay Degwekar. On the complexity of collision resistant hash functions: New and old black-box separations. In *TCC (1)*, volume 11891 of *Lecture Notes in Computer Science*, pages 422–450. Springer, 2019.
- BDV17. Nir Bitansky, Akshay Degwekar, and Vinod Vaikuntanathan. Structure vs. hardness through the obfuscation lens. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, pages 696–723, 2017.
- BG20. Nir Bitansky and Idan Gerichter. On the cryptographic hardness of local search. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, pages 6:1–6:29, 2020.
- BKSY11. Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhimovich. Limits on the power of zero-knowledge proofs in cryptographic constructions. In *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 559–578. Springer, 2011.
- BPR15. Nir Bitansky, Omer Paneth, and Alon Rosen. On the cryptographic hardness of finding a Nash equilibrium. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1480–1498, 2015.
- Bur06. Joshua Buresh-Oppenheimer. On the TFNP complexity of factoring. Unpublished, <http://www.cs.toronto.edu/~bureshop/factor.pdf>, 2006.
- CDT09. Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3), 2009.
- CHK⁺19a. Arka Rai Choudhuri, Pavel Hubáček, Chethan Kamath, Krzysztof Pietrzak, Alon Rosen, and Guy N. Rothblum. Finding a Nash equilibrium is no easier than breaking Fiat-Shamir. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 1103–1114. ACM, 2019.
- CHK⁺19b. Arka Rai Choudhuri, Pavel Hubáček, Chethan Kamath, Krzysztof Pietrzak, Alon Rosen, and Guy N. Rothblum. PPAD-hardness via iterated squaring modulo a composite. *IACR Cryptology ePrint Archive*, 2019:667, 2019.
- CT12. Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- DGP09. Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
- DP11. Constantinos Daskalakis and Christos H. Papadimitriou. Continuous local search. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 790–804, 2011.
- EFKP20. Naomi Ephraim, Cody Freitag, Ilan Komargodski, and Rafael Pass. Continuous verifiable delay functions. In *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part III*, volume 12107 of *Lecture Notes in Computer Science*, pages 125–154, 2020.
- FG18. Aris Filos-Ratsikas and Paul W. Goldberg. Consensus halving is PPA-complete. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 51–64, 2018.
- FG19. Aris Filos-Ratsikas and Paul W. Goldberg. The complexity of splitting necklaces and bisecting ham sandwiches. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 638–649, 2019.
- Fis12. Marc Fischlin. Black-box reductions and separations in cryptography. In *AFRICACRYPT*, volume 7374 of *Lecture Notes in Computer Science*, pages 413–422. Springer, 2012.
- GH19. Paul W. Goldberg and Alexandros Hollender. The hairy ball problem is PPAD-complete. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, pages 65:1–65:14, 2019.
- GL89. Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, pages 25–32, 1989.

- GMR01. Yael Gertner, Tal Malkin, and Omer Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 126–135. IEEE Computer Society, 2001.
- GPS16. Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. Revisiting the cryptographic hardness of finding a Nash equilibrium. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, pages 579–604, 2016.
- GT00. Rosario Gennaro and Luca Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *FOCS*, pages 305–313. IEEE Computer Society, 2000.
- HHRS15. Iftach Haitner, Jonathan J. Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols - tight lower bounds on the round and communication complexities of statistically hiding commitments. *SIAM J. Comput.*, 44(1):193–242, 2015.
- HCY17. Pavel Hubáček, Moni Naor, and Eylon Yogev. The journey from NP to TFNP hardness. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, pages 60:1–60:21, 2017.
- HR04. Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In *CRYPTO*, volume 3152 of *Lecture Notes in Computer Science*, pages 92–105. Springer, 2004.
- HY17. Pavel Hubáček and Eylon Yogev. Hardness of continuous local search: Query complexity and cryptographic lower bounds. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1352–1371, 2017.
- IR89. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, pages 44–61. ACM, 1989.
- Jeř16. Emil Jeřábek. Integer factoring and modular square roots. *J. Comput. Syst. Sci.*, 82(2):380–394, 2016.
- JK20. Ruta Jawale and Dakshita Khurana. Lossy correlation intractability and PPAD hardness from sub-exponential LWE. *IACR Cryptol. ePrint Arch.*, 2020:911, 2020.
- JKKZ20. Ruta Jawale, Yael Tauman Kalai, Dakshita Khurana, and Rachel Zhang. SNARGs for bounded depth computations and PPAD hardness from sub-exponential LWE. *IACR Cryptol. ePrint Arch.*, 2020:980, 2020.
- JPY88. David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37(1):79–100, 1988.
- KPR⁺13. Shiva Kintali, Laura J. Poplawski, Rajmohan Rajaraman, Ravi Sundaram, and Shang-Hua Teng. Reducibility among fractional stability problems. *SIAM J. Comput.*, 42(6):2063–2113, 2013.
- KPY20. Yael Tauman Kalai, Omer Paneth, and Lisa Yang. Delegation with updatable unambiguous proofs and PPAD-hardness. In *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 652–673. Springer, 2020.
- KS17. Ilan Komargodski and Gil Segev. From Minicrypt to Obfustopia via private-key functional encryption. In *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part I*, pages 122–151, 2017.
- KZ20. Yael Tauman Kalai and Rachel Zhang. SNARGs for bounded depth computations from sub-exponential LWE. *IACR Cryptol. ePrint Arch.*, 2020:860, 2020.
- LV20. Alex Lombardi and Vinod Vaikuntanathan. Fiat-Shamir for repeated squaring with applications to PPAD-hardness and VDFs. In *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 632–651. Springer, 2020.
- MM11. Takahiro Matsuda and Kanta Matsuura. On black-box separations among injective one-way functions. In *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings*, volume 6597 of *Lecture Notes in Computer Science*, pages 597–614. Springer, 2011.
- MP91. Nimrod Megiddo and Christos H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theor. Comput. Sci.*, 81(2):317–324, 1991.
- Pap94. Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.

- PV10. Rafael Pass and Muthuramakrishnan Venkitasubramaniam. Private coins versus public coins in zero-knowledge proof systems. In *TCC*, volume 5978 of *Lecture Notes in Computer Science*, pages 588–605. Springer, 2010.
- PV19. Rafael Pass and Muthuramakrishnan Venkitasubramaniam. A round-collapse theorem for computationally-sound protocols; or, TFNP is hard (on average) in Pessiland. *CoRR*, abs/1906.10837, 2019.
- RS10. Alon Rosen and Gil Segev. Chosen-ciphertext security via correlated products. *SIAM J. Comput.*, 39(7):3058–3088, 2010.
- RSS17. Alon Rosen, Gil Segev, and Ido Shahaf. Can PPAD hardness be based on standard cryptographic assumptions? In *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, volume 10678 of *Lecture Notes in Computer Science*, pages 747–776. Springer, 2017.
- RTV04. Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings*, pages 1–20, 2004.
- Rud88. Steven Rudich. *Limits on the Provable Consequences of One-way Functions*. PhD thesis, EECS Department, University of California, Berkeley, Dec 1988.
- Sim98. Daniel R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *Advances in Cryptology - EUROCRYPT '98, International Conference on the Theory and Application of Cryptographic Techniques, Espoo, Finland, May 31 - June 4, 1998, Proceeding*, volume 1403 of *Lecture Notes in Computer Science*, pages 334–345. Springer, 1998.
- Wee07. Hoeteck Wee. Lower bounds for non-interactive zero-knowledge. In *Theory of Cryptography, 4th Theory of Cryptography Conference, TCC 2007, Amsterdam, The Netherlands, February 21-24, 2007, Proceedings*, volume 4392 of *Lecture Notes in Computer Science*, pages 103–117. Springer, 2007.