

New Techniques in Replica Encodings with Client Setup

Rachit Garg ^{*}, George Lu ^{**}, and Brent Waters ^{***}

¹ University of Texas at Austin

² University of Texas at Austin

³ University of Texas at Austin and NTT Research

Abstract. A proof of replication system is a cryptographic primitive that allows a server (or group of servers) to prove to a client that it is dedicated to storing multiple copies or replicas of a file. Until recently, all such protocols required fine-grained timing assumptions on the amount of time it takes for a server to produce such replicas.

Damgård, Ganesh, and Orlandi (CRYPTO’ 19) [11] proposed a novel notion that we will call proof of replication with client setup. Here, a client first operates with secret coins to generate the replicas for a file. Such systems do not inherently have to require fine-grained timing assumptions. At the core of their solution to building proofs of replication with client setup is an abstraction called replica encodings. Briefly, these comprise a private coin scheme where a client algorithm given a file m can produce an encoding σ . The encodings have the property that, given any encoding σ , one can decode and retrieve the original file m . Secondly, if a server has significantly less than $n \cdot |m|$ bit of storage, it cannot reproduce n encodings. The authors give a construction of encodings from ideal permutations and trapdoor functions.

In this work, we make three central contributions.

- Our first contribution is that we discover and demonstrate that the security argument put forth by [11] is fundamentally flawed. Briefly, the security argument makes assumptions on the attacker’s storage behavior that does not capture general attacker strategies. We demonstrate this issue by constructing a trapdoor permutation which is secure assuming indistinguishability obfuscation, serves as a counterexample to their claim (for the parameterization stated).
- In our second contribution we show that the DGO construction is actually secure in the ideal permutation model (or ideal cipher model) and the random oracle (or random function) model from any trapdoor permutation when parameterized correctly. In particular, when the number of rounds in the construction is equal to $\lambda \cdot n \cdot b$ where λ is the security parameter, n is the number of replicas and b is the number of blocks. To do so we build up a proof approach from the ground up that accounts for general attacker storage behavior where we create an analysis technique that we call “sequence-then-switch”.
- Finally, we show a new construction that is provably secure in the random oracle model. Thus requiring less structure on the ideal function.

1 Introduction

In a proof of replication system [6,5], a user wants to distribute a file m and ensure that a server or group of servers will dedicate the resources to storing multiple copies or replicas of it. That is, the server should either receive or generate n replicas $\sigma_1, \dots, \sigma_n$ where the file m can be efficiently decoded from any single replica. In the original notion of proofs of replication, a server could take a file m as input and independently generate all the replicas $\sigma_1, \dots, \sigma_n$. Later it could prove possession if challenged. Since the introduction of this concept, several such solutions [25,7,17,9,16] have emerged.

However, in these solutions, there exist a tension that stems from the following attack. Consider a non-compliant server that stores just a single copy of m . When challenged to prove possession of replicas, it

^{*} Email: rachit0596@gmail.com. Supported by NSF CNS-1908611, CNS-1414082, Packard Foundation Fellowship, and Simons Investigator Award.

^{**} Email: gclu@cs.utexas.edu. Supported by NSF CNS-1908611, CNS-1414082, Packard Foundation Fellowship, and Simons Investigator Award.

^{***} Email: bwaters@cs.utexas.edu.

on the fly, generates $\sigma_1, \dots, \sigma_n$ using the legitimate generation algorithm and proceeds to prove replication using the ephemeral values as though it were storing these replicas all along.

It is easy to see that achieving meaningful security against such an attack is impossible without imposing a concrete time-bound between when a server is challenged and when it must answer. The setting of this time-bound must be coupled with an understanding of how long it takes an honest system to retrieve the replicas and produce a proof and balanced against how fast a highly provisioned server might take to produce the replicas from scratch. This balancing act creates a certain tension in that more costly replica generation will help security, but also imposes a higher burden on initiation. Moreover, other issues can arise in the context of a more extensive system. For example, if audit challenges come out at a predictable time (e.g., daily), then a cheating server could start generating the response ahead of time.

To address these issues, Damgård, Ganesh, and Orlandi [11] proposed a novel notion that we will call proof of replication with client setup. In this notion, a client that wishes to store a file m will generate replicas $\sigma_1, \dots, \sigma_n$, along with a (short) public verification key vk . The system will have the properties that (1) one can reconstruct the file from any replica along with the verification key, and (2) a server can prove possession of the replicas to any client that holds the verification key. Unlike the previous systems, proof of replication with client setup need not require fine-grained timing assumptions as a server will not be able to regenerate the replicas from only the message m and vk . Indeed the security definition says (informally) that any poly-time server that devotes significantly fewer resources than n times message length will not be able to pass the possession test.

The solution proposed in [11] combines two high-level ingredients. The first is a proof of retrievability system as proposed in prior work [29,12,8]. Roughly, if a server executes a proof of retrievability for data d with a client, this means that now, the server was capable of reconstructing d . However, a proof of retrievability in and of itself gives no guarantee about the amount of resources required to store d .

Second, the authors introduce a notion of a replica encoding. A replica encoding system consists of three algorithms: $(\text{rSetup}, \text{rEnc}, \text{rDec})$. The setup algorithm on input, a security parameter κ and the maximum number of replicas n of a scheme, outputs a public and secret key pair as $(\text{pk}, \text{sk}) \leftarrow \text{rSetup}(1^\kappa, 1^n)$. The encoding algorithm takes as input the secret key and a message m to produce an encoding as $y \leftarrow \text{rEnc}(\text{sk}, m)$. Finally, the decoding algorithm takes as input an encoding y and the public key to retrieve the message as $m \leftarrow \text{rDec}(\text{pk}, y)$ or outputs \perp to indicate failure. The algorithms are randomized, and the encoding procedure can be run multiple times to produce multiple encodings. The correctness of the scheme dictates that if one encodes any message m under a secret key and then decodes it under the corresponding public key, m will be decoded.

To capture security, we will consider a soundness game which uses a two-stage attacker $(\mathcal{A}_1, \mathcal{A}_2)$. In the first stage, \mathcal{A}_1 will be given a challenger-generated public key pk and reply with a message m . It is then given n encodings generated by the challenger as y_1, \dots, y_n . The attacker outputs a state variable state , which we will generally think of as being smaller than $|m| \cdot n$. At the second phase, the algorithm \mathcal{A}_2 is given the input state and is tasked with outputting guesses $\tilde{y}_1, \dots, \tilde{y}_n$. The security property intuitively states that if the size of the storage $|\text{state}|$ is significantly less than $v \cdot |m|$, then the number of i where $y_i = \tilde{y}_i$ will be less than v . That is, the attacker cannot do much better than simply storing a set of values y_i .

Damgård, Ganesh, and Orlandi showed how a natural compilation of existing proof of retrievability schemes along with replica encodings gave way to proofs of storage with client setup. Also, they provided a candidate construction for replica encodings from trapdoor permutations under the ideal cipher model and the random oracle model. We turn our attention to these.

The DGO construction: We now outline (a slight variant of the) construction for [11], which is given in the ideal permutation and the random oracle model. We remark that the DGO construction itself is an adaptation of one of the “hourglass” schemes of van Dijk et al. [30]. The building blocks will consist of a trapdoor permutation f, f^{-1} , along with the ideal cipher $\mathbb{T}, \mathbb{T}^{-1}$, and a random oracle \mathbb{H} . We again let κ be the security parameter and let $\lambda = \lambda(\kappa)$ be the output length of the trapdoor permutation as well as the block length of an ideal permutation $\mathbb{T} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$. For pedagogical purposes, we will assume for the sketch below that messages consist of λ bits, but in our main body, we consider the more realistic case of many block messages.

The setup algorithm simply chooses a TDP public/secret key pair as $\text{KeyGen}(1^\kappa)$ outputs (pk, sk) where KeyGen is the trapdoor permutation key generation algorithm. The public and secret key pair of the TDP serve as the keypair of the replicated encoding scheme.

The encoding algorithm $\text{rEnc}(\text{sk}, m)$ takes as input the TDP secret key and message m . It first chooses a string $\rho \xleftarrow{R} \{0, 1\}^\kappa$. It then initializes a value $Y_0 = m \oplus \text{H}(\rho)$ where H is modeled as a random oracle function. Then for $j = 1$ to r rounds it computes $Y_j = \text{f}^{-1}(\text{sk}, \text{T}(Y_{j-1}))$ where r is the number of rounds, which grows linearly with the number of replicas. The encoding is output as (Y_r, ρ) .

Finally, the decoding algorithm $\text{rDec}(\text{pk}, y = (Y_r, \rho))$ recovers a message as follows. For setting j from $r - 1$ down to 0 compute $Y_j = \text{T}^{-1}(\text{f}(\text{pk}, Y_{j+1}))$. Then output $m = Y_0 \oplus \text{H}(\rho)$.

The fact that the decoding step recovers the message follows straightforwardly from the correctness of the trapdoor permutation and ideal permutation. We also observe that it is publicly computable since it uses the public key and forward direction of the trapdoor permutation.

1.1 Our Contributions

We make three core contributions to this area:

1. Our first contribution is that we discover and demonstrate that the security argument put forth by [11] is fundamentally flawed. The security argument makes implicit assumptions about an attacker’s behavior which are not generally true. More specifically, in the security game applied to the DGO construction (in the ideal permutation and random oracle model) an attacker works in two phases. The first stage attacker \mathcal{A}_1 receives the replicas, can make several queries to the ideal permutation and then records some state `state` of limited size. This state `state` is passed to a stage two attacker \mathcal{A}_2 which can make further permutation queries and attempts to reconstruct the queries. In general a first stage attacker can apply arbitrary strategies to breaking the scheme so long as it poly-time and state `state` is sufficiently small. However, the proof argument of [11] assume that the ideal permutation queries made by the attacker will be “uniquely stored”. Roughly, they will argue that a query output bit will either be stored explicitly or not at all. This discounts the possibility of an attacker strategy such as making several oracle queries and storing the XOR of all the outputs together.

We demonstrate that the above error manifests in a false theorem statement in [11]. The authors claim that the scheme is secure for *any* trapdoor permutation (TDP) if $r = \lambda \cdot n$ rounds are applied when doing n encodings of b blocks with security parameter λ . (I.e. Claim the number of rounds does not need to scale with b .) We provide an explicit counterexample to this claim in Section 7. We give a TDP that is secure assuming indistinguishability obfuscation, but for which the scheme is attackable using these parameters. The attacker strategy actually works by XORing several query values together and is thus directly tied to the flaw in the security proof. There does not appear to be any simple “fix” to the security argument of [11] as we will see that addressing general attacker storage strategies comprises the core difficulty of proving security.

We also note that an explicit “partitioning assumption” appears in the security definition of [30] for “hourglass schemes” where the authors conjecture (but do not prove) that it seems implausible that mixing together two representations can give an advantage to an attacker. Although we do not do so formally, we believe that our counterexample can be adapted to the work of [30] as well (at least if one considered the scheme for general trapdoor permutations) and demonstrates the danger of making assumptions that restrict adversarial strategies.

2. For our second contribution we show that the DGO construction is actually secure when parameterized correctly. In particular, when the number of rounds is equal to $\lambda \cdot n \cdot b$. To do so we need to build up a proof approach from the ground up that accounts for general attacker storage behavior. We first develop an analysis technique that we call “sequence-then-switch”. We show how in this framework, we can prove security against an attacker that arbitrarily assigns state. In particular, we show how to analyze the security of a close variant of the [11] construction in the ideal permutation and random oracle model. In addition, we give an explicit construction of a trapdoor permutation using indistinguishability obfuscation which allows for an attack strategy not covered by their restricted model, showing the [11] construction as given is in fact explicitly *insecure* against general adversaries.

3. The prior construction and proof relies on the ideal permutation model. A perhaps better goal would be to have a construction secure in the random oracle or random function model as this assumes less structure on the ideal object. Typically, this is dealt with by building a random permutation from a function using a Feistel network and showing that this is “indifferentiable” in the indifferentiability framework of Mauer et al. [22]. Prior works have shown this for 14 [20] and then 8 round Feistel [10]. However, Ristenpart, Shacham, and Shrimpton [27] show that the framework does not compose for multi-round games. Since the above construction relies on a multi-round game, proof from an ideal permutation cannot be reduced to a proof to an ideal function.

We give a new construction that relies only on the ideal function model and analyze its security. Our construction uses the random function to embed a Feistel like structure into the construction. However, instead of arguing in the indifferentiability framework, we provide direct proof of security, which bypasses any composability issues. In both proofs, we allow the attacker to assign its storage arbitrarily.

1.2 Our Techniques

We begin by describing our analysis for the first construction using a TDP and ideal permutation. We focus on the construction producing many replicas on a single block, as described in the introduction for simplicity. Also, for simplicity, we consider the particular case where an attacker that asks for n replicas in the first stage and wants to produce all n of these replicas, but we significantly less than $n \cdot \lambda$ storage. In particular consider an adversary with `state` of length $n \cdot \lambda - n \cdot \omega(\log \kappa)$ bits of storage for security parameter κ and block length λ . Our central idea is to organize the proof into two parts where we first show that any storage bounded \mathcal{A}_2 must make “sequential” oracle queries on at least one replica. Then we show that on this particular replica, how one can swap out permutation output for another.

1. **Sequentiality:** In our security game, the challenger first creates n replicas of m . To create the i -th replica by choosing ρ_i randomly. It sets $Y_0^{(i)} = m \oplus H(\rho_i)$. Then for $j = 1$ to r rounds it computes $Y_j^{(i)} = f^{-1}(\text{sk}, T(Y_{j-1}^{(i)}))$. The encoding is output to \mathcal{A}_1 as $(Y_r^{(i)}, \rho_i)$ for $i \in [n]$. The attacker \mathcal{A}_1 receives the encodings, makes some more oracles queries before producing `state` of $n \cdot \lambda - n \cdot \omega(\log \kappa)$ bits and passing it to \mathcal{A}_2 .

Let’s examine the behavior of \mathcal{A}_2 whose job it is to output the encodings using the state plus oracle queries. We say that \mathcal{A}_2 “queries sequentially” on replica i if for all $j \in [0, r - 1]$ it queries the oracle T on $Y_j^{(i)}$ before it queries the oracle on $Y_{j+1}^{(i)}$. (We will think of outputting the encoding $Y_r^{(i)}$ at the end as implicitly querying on the final value.) That is for \mathcal{A}_2 to query sequentially on replica i it must both make all $r + 1$ oracle queries and make them in (relative) order. However, there could be many other queries outside the replica chain interspersed between $Y_j^{(i)}$ and $Y_{j+1}^{(i)}$.

We will first argue that except with negligible probability whenever \mathcal{A}_2 produces all the encodings, it queries sequentially on at least one replica. Observe that we cannot hope to say that it queried sequentially on all replicas as `state` could directly store several of the replica encodings, which allows the algorithm to bypass any additional queries related to that replica.

To prove this, we first define and prove a useful matching pairs lemma. Consider an algorithm \mathcal{B} that takes as input a string `advice` of length $n \cdot \lambda - n \cdot \omega(\log \kappa)$ and gets access to a string oracle access to a randomly chosen permutation $T(\cdot), T^{-1}$ of block length λ . The goal of \mathcal{B} is to provide n *distinct* pairs (x_i, y_i) such that $T(x_i) = y_i$, but *without* querying the oracle a either x_i nor y_i . Thus \mathcal{B} can make several oracle queries on many values; however, once a query is made on some x , it spoils using x as a value from one of the pairs. Note that to win in this game, \mathcal{B} needs to produce the pairs— not just distinguish them from random. Also observe that \mathcal{B} can use `advice` to help it win this game. For example, `advice` might encode the several pairs.

We prove that no attacker \mathcal{B} that makes a polynomially bounded number of queries can win in this game by a simple application of the union bound. Consider a fixed value of an `advice` string a — that is a is fixed before the permutation is chosen. We show that the probability of $\mathcal{B}(a)$ winning is at most $\frac{\text{poly}(\lambda)}{2^{n\lambda}}$. Then by the union bound the probability that there exists *any* string a which it could win with is at most $2^{n \cdot \lambda - n \cdot \omega(\log \kappa)} \cdot \frac{\text{poly}(\lambda)}{2^{n\lambda}}$ which is negligible in λ .

Now we need to show that an attacker that wins but is not sequentially querying on any replica will break our matching pairs game. We consider $(\mathcal{A}_1, \mathcal{A}_2)$ that does this. Let's think of the algorithm pair as deterministic. (If they are randomized for each security parameter, we can fix their coins that maximize success probability.) We construct an algorithm \mathcal{B} along with the process of determining an advice string that does this. Conceptually we can think of a preprocessing algorithm \mathcal{B}' that generates the advice. \mathcal{B}' will first run \mathcal{A}_1 , which makes several queries and then produce `state`. It then runs \mathcal{A}_2 on `state`. If \mathcal{A}_2 either did not produce all the replica encodings or it did sequentially query on some replica i , then abort. However, if it did not make sequential queries on all replicas, then there must be values j_1, \dots, j_n where \mathcal{A}_2 made an oracle query on $Y_{j_i}^{(i)}$ (or $f(\text{pk}, Y_{j_i}^{(i)})$), but had not yet made a query on $Y_{j_{i-1}}^{(i)}$. Let q_1, \dots, q_n be the indices of the queries (ordered chronologically) for which this occurs. Note the number of queries \mathcal{A}_2 can make is polynomial in κ , but in general, it could be much more than $r \cdot n \cdot \lambda$. The preprocessing algorithm will package its advice string as `state` along with j_1, \dots, j_n and q_1, \dots, q_n . Importantly, the size of this information is bounded by $\lg(\text{poly}(\kappa))$ for some polynomial poly since n, r , and the number of replicas is polynomially bounded. This means that if `state` is of size $n \cdot \lambda - n \cdot \omega(\log \kappa)$, then the advice string will be within $n \cdot \lambda - \omega(\log \kappa)$.

We now consider algorithm \mathcal{B} , which receives the advice string. \mathcal{B} will run \mathcal{A}_2 with the following modifications. Suppose \mathcal{A}_2 makes its q -th query where $q = q_i$ for some i . This means that \mathcal{A}_2 is querying on $Y_{j_i}^{(i)}$, but had not yet made a query on $Y_{j_{i-1}}^{(i)}$. At this point \mathcal{B} determines $Y_{j_{i-1}}^{(i)}$ by querying $Y_1^{(i)} = f^{-1}(\text{sk}, \mathbb{T}(Y_0^{(i)}))$ up to $Y_{j_{i-1}}^{(i)} = f^{-1}(\text{sk}, \mathbb{T}(Y_{j_{i-2}}^{(i)}))$. It then submits $(Y_{j_{i-1}}^{(i)}, f(\text{pk}, Y_{j_i}^{(i)}))$ as one of its matching pairs noting that neither $\mathbb{T}(Y_{j_{i-1}}^{(i)})$ nor $\mathbb{T}^{-1}(f(\text{pk}, Y_{j_i}^{(i)}))$ were made before. It can also continue to run \mathcal{A}_2 without making either of these queries to the oracles since it already knows the answers to them. As this process proceeds, \mathcal{B} will eventually recover n such pairs which breaks our matching pairs lemma and arrives at a contradiction.

2. Switching:

Once sequentiality is established, we will proceed to argue that the adversary must still be sequential with good probability even when we “switch” the random oracle output of some $Y_j^{(\gamma)}$ to a random value only for \mathcal{A}_2 , allowing us to embed a trapdoor permutation challenge.

In more detail, we now consider a new switched game that is almost equivalent to the prior one. In the switched game the challenger first chooses r random values $A_{i,b} \in \{0, 1\}^\lambda$ for $j \in [1, r], b \in \{0, 1\}$ along with a bit string $x \in \{0, 1\}$. It programs the oracle T such that $Y_j^{(\gamma)} = A_{i,b}$. This game can be shown to be almost equivalent to the previous one.

Next, we consider a game where the challenger answers queries according to a string x with \mathcal{A}_1 , but switches to using a string x' (and keeps everything else the same) when responding to \mathcal{A}_2 . The challenger chooses the string x' such that the output `state` given by \mathcal{A}_1 is the same as if the queries are answered according to x' in the first phase. The attacker is considered to win only if it would produce sequential queries *both* for when x was used with \mathcal{A}_2 and when x' was used with \mathcal{A}_2 .

With high probability, such an x' will exist from the fact that $|\text{state}| \leq n \cdot \lambda - \omega(\log \kappa)$ and r is set to be $n \cdot \lambda$. We emphasize that to make this argument we do not make any further assumptions on how \mathcal{A}_1 assigns `state` other than the bound on the size. We can then use the heavy row lemma [24] to argue that if an attacker wins with probability ϵ in the previous game, it wins with probability $\approx \epsilon$ in this game. We note that the game takes exponential time to find such an x' , but this is not an issue as the closeness lemma is information-theoretic.

Finally, in order to embed a TDP challenge, we need to move to a security game that can be efficiently simulated. While it might take exponential time to find x' from x above, we observe that this is not necessary. Instead, we can embed the challenge from just knowing the shortest common prefix of x and x' . Moreover, given x , we can simply guess what the prefix is with a $\frac{1}{r}$ loss. Thus we move to a final game where the challenger simply chooses a random value j and a random permutation \mathbb{T} in the first phase and then replaces the oracle output of $Y_j^{(i)}$ with a random R in the second phase. The attacker wins if it queries $f^{-1}(\text{sk}, R)$. A simple reduction then shows that any attacker that wins in this game breaks the TDP security.

Extending to the ideal function model We can now return to our goal of building a secure construction by replacing the ideal permutation model with a random oracle model. As argued earlier, doing so is desirable as an ideal function imposes less structure and appears to be a less risky heuristic. Our solution will build upon the analysis principles established above, but proving security involves more complications.

We begin by sketching out the encoding construction. In this setting, we will have a TDP in the domain λ bits and use a random oracle T' that outputs λ bits. We will use blocks of length 2λ , and for this sketch, focus on the particular case where each replica consists of a single block message.

The setup algorithm again chooses a TDP public/secret key pair as $\mathsf{KeyGen}(1^\kappa) \rightarrow (\mathsf{pk}, \mathsf{sk})$ as before. The encoding algorithm $\mathsf{rEnc}(\mathsf{sk}, m)$ takes as input the TDP secret key and message $m \in \{0, 1\}^{2\lambda}$. It first chooses a string $\rho \xleftarrow{R} \{0, 1\}^\kappa$. It then initializes values $Y_0 = L(m \oplus \mathsf{H}(\rho))$ and $Y_1 = R(m \oplus \mathsf{H}(\rho))$ where H is a random oracle that produces an 2λ bit output and L, R are functions that take the left and right halves. Then on rounds j from 2 to r compute Y_j from Y_{j-1} and Y_{j-2} as

$$Y_j = \mathsf{f}^{-1}(\mathsf{sk}, Y_{j-2} \oplus \mathsf{T}'(Y_{j-1})).$$

The replica encoding value is 2λ bits long and consists of the last two values as $Y_{r-1} || Y_r$. The decoding algorithm rDec works backward down the Feistel structure to recover the message.

In this setting, we want to prove that in the security game, an attacker with $n \cdot 2\lambda - n \cdot \omega(\log \kappa)$ cannot produce n replica encodings. (The extra factor of two is solely due to blocks being 2λ bits here.)

Our proof will follow in the same theme of showing that there must be a form of sequential querying made on at least one replica. However, the new structure of the construction presents additional complications. For example, we could imagine an attacker \mathcal{A}_1 , which stores all the values Y_j^i for some j . This is possible since storing these only take $n\lambda$ bits, and our assumption is only that the storage is less than $2n\lambda$ bits. On the one hand, it is unclear how the attacker can leverage storing all these values because one needs consecutive values (e.g., $Y_j^{(i)}, Y_{j+1}^{(i)}$) to propagate further. And, storing n different consecutive pairs requires $2n\lambda$ bits of storage. On the other hand, the attacker can store these means at the very least we need a new notion of sequentiality.

For our new notion of sequentiality, we say that the queries to replica i meet our requirements if the longest common subsequence of the queries made and $Y_1^{(i)}, Y_2^{(i)}, \dots, Y_r^{(i)}$ is of length at least $r-3$. Intuitively, this is close to our original notion but allows for a little skipping. To prove this form of sequentiality, we invoke a random function analog of our matching pairs lemma from before. The reduction to matching pairs follows in a similar spirit to before but requires a more nuanced case analysis.

Once that is in place, our proof proceeds analogously, but again with more nuances and complications arising from the fact that we only can guarantee the weaker form of longest common subsequence.

The proposed construction is round optimal We now consider the general case of a message having b blocks and give intuition that our construction is round optimal up to constant factors. We construct a secure trapdoor permutation scheme from indistinguishability obfuscation which gives an insecure replica encoding scheme for any number of rounds $\notin \Omega(b \cdot n)$ (i.e. $\in o(b \cdot n)$). Incidentally, this also shows that the construction provided by [11], which claims to only requires $O(n)$ rounds, is insecure against general adversaries.

We provide the intuition for our construction by considering the ideal VBB notion of obfuscation. The overall idea is to construct a trapdoor permutation family where we can amortize the ‘state’ space required to invert multiple independent instances. We will consider our permutations to be on domain $\{0, 1\}^\lambda$. If we assume we have VBB obfuscation, then consider a program that takes in b many inputs $\{y_i\}_{i \in b}$ where $y_i \in \{0, 1\}^\lambda$ and an advice string also in $\{0, 1\}^\lambda$ and outputs the preimages of the messages $\{x_i = \mathsf{f}^{-1}(\mathsf{sk}, y_i)\}_{i \in b}$ iff the advice string that was input was equal to $\bigoplus_{i \in b} x_i$. The program has the secret key hardcoded and simply computes x_i and makes the check against the advice string and outputs $\{x_i\}_{i \in b}$ if the check succeeds. The VBB obfuscation of this program is then posted in the public parameters and provides a way for the adversary to compress $b \cdot \lambda$ bits to λ bits and still preserve information. Thus an adversary with outputting $r \cdot \lambda$ bits can recompute the replica from storing $o(b \cdot n\lambda)$ information. This would violate the security if we proved soundness for the same parameters as our scheme. A formal treatment is presented in Section 7.

1.3 Additional Prior Work

Proofs of Retrieability:

Proofs of retrievability guarantee to a verifier that a server is storing all of client’s data. The notion was formalized in [21], where, in an audit protocol the verifier stores a (short) verification key locally and interacts with the server to enforce accountability of the storage provider. If the server can pass an audit, then there exists an extractor algorithm, that must be able to extract the file on interaction with the server. There are different constructions for this primitive, [12,29,8]. The construction of [29] showed how to do this in the random oracle model that allow public verifiability.

Proofs of Space: Proof of space are interactive protocols between a prover (server) and a verifier (client) that guarantee that a prover has dedicated a specific amount of space. It guarantees that it would be more expensive for a dishonest prover to deviate from the honest protocol. They were introduced in [14] and have been further studied in [26,1]. Compared to a proof of replicated storage, they have an additional requirement of communication being succinct between a prover and verifier and are usually studied in the public-key setting.

Other examples of works which are different from proofs of space but enforce storage requirements similar to our soundness game on the prover are storage-enforcing commitments [19], hourglass scheme [30] and the model of computation considered by [15].

Proofs of Replicated Storage:

The formal treatment of proofs of replicated storage was given by [11,16,9]. The idea was introduced in [6,5] where they proposed Filecoin, a decentralized storage network that performs consensus using proofs of replication. Recently, [25,17,7,9,16] have given constructions for proof of replication using timing assumptions (encoding process is much slower so that a server cannot replicate data on demand). On the other hand, the scheme of [11] is not based on timing assumptions and considers the protocol with a client setup. They introduce the notion of a replica encoding that can be combined with a public verifiable proof of retrievability [29] to give a proof of replicated storage. Please see [11] for other related works such as proof of data replication.

Hourglass Scheme:

Our constructions and the construction from [11] are reminiscent of the hourglass scheme of [30]. Our construction in the ideal permutation model differs from the RSA based hourglass function of [30] in explicitly ensuring that the encoding blocks are uniformly distributed by applying a random oracle H and increasing the number of rounds suggested by their scheme. Because of our explicit encoding function, we do not need to make a partitioning assumption in our security proof. The brief analysis of their scheme gives a similar intuition to the security as used by [11] and gives a construction for the number of rounds independent of the number of blocks. But as we see in Section 7, this intuition does not hold true for general adversaries.

Technique Similarities in literature:

Some of our techniques have a flavor that appears in the the study of pebbling strategies on random oracle graphs and the memory hardness literature [15,13,3,2,4]. Pebbling strategies on random oracle graphs look at the amount of resources (the list of random-oracle calls) made by the adversary and help in proving complexity lower-bounds on the resources. Our notion of “sequentiality” is similar to the notion of a legal “ex-post-facto pebbling” on a directed acyclic graph (see [4] for details). The reductions there are proven using a core lemma which looks at a legal ex-post-facto pebbling given hints; Lemma 1 of [4,15,13] which is similar to our core lemmas for proving sequentiality Lemma 1. Interestingly, [2] considered adversaries that can store secret shares of the random oracle queries (such as a xor) and introduced the notion of an entangled pebbling game. They look at the resource of “Cumulative Memory Complexity (CMC)” and constructed an example to show that such strategies can help the adversary reduce it’s resource requirement. The followup work of [3] improved on their lower bounds results for any general adversarial strategy.

1.4 Concurrent Work

After completing our work we learned of a concurrent and independent work of Moran and Wichs [23]. They introduce a variant of replica encodings which they call incompressible encodings, and proceed to provide constructions in the random-oracle model (and the common random string model) using the Decisional Composite Residuosity or Learning with Errors assumptions. Their construction utilizes some new techniques to apply lossiness to construct said encodings. In addition, they introduce an additional “big-key” application for intrusion resilience which applies to our constructions and proofs as well.

At a very high level, our work depends on the general assumption of trapdoor permutations, whereas they use the specific number theoretic assumptions of Decisional Composite Residuosity and Learning with Errors. Comparing our construction instantiated with RSA trapdoor permutation to their DCR construction, their construction appears to be more practically efficient from a computational perspective due to the round complexity required for our construction, however, ours makes tighter use of space for small “ s ” values used in the DCR construction. An interesting future direction could be to explore concrete space and computational efficiency tradeoffs for increasing the s parameter in their DCR construction.

Similar to us, Moran and Wichs discovered foundational issues in the proof arguments of [11]. In a personal communication Wichs noted that there is a simple heuristic counterexample to the claim of [11] if one uses the heuristic of ideal obfuscation. We subsequently developed a counterexample based on the concrete assumption of indistinguishability obfuscation that we added as Section 7 of our work.

2 Preliminaries

Notation. These notations are used consistently throughout the text.

We use κ to denote the security parameter. $y \leftarrow \mathcal{B}(x)$ denotes the output of the algorithm \mathcal{B} when we run x on it. A negligible function $\text{negl}(x)$ is a function such that for every positive integer c , there exists an integer N_c such that for all $x > N_c$, $\text{negl}(x) < \frac{1}{x^c}$. $[n]$ denotes the set $\{1, 2, \dots, n\}$ and $[a, b]$ denotes the interval between a and b inclusive. $y \stackrel{R}{\leftarrow} \mathcal{D}$ implies that we are uniformly sampling y from a domain set \mathcal{D} . We say an adversary or an algorithm \mathcal{A} is probabilistic poly time (PPT) if there is a polynomial $\text{poly}(\cdot)$ such that for all κ , \mathcal{A} will halt in $\leq \text{poly}(\kappa)$ time in expectation on any input of length κ .

A trapdoor permutation is defined as a collection of three PPT algorithms $\text{KeyGen}(\cdot), f(\cdot, \cdot), f^{-1}(\cdot, \cdot)$. $i\mathcal{O}(\kappa, C)$ is an indistinguishability obfuscator for a circuit class $\{C_\kappa\}$ and security parameter κ . A puncturable pseudorandom function family (PPRF) on domain \mathcal{D}_κ and range \mathcal{R}_κ is defined using a set of algorithms $(\text{PPRF.KeyGen}, \text{PPRF.Eval}, \text{PPRF.Puncture})$. Due to space limitations, please find the complete definitions in the full version of the paper [18].

3 Defining Replica Encoding

A Replica Encoding scheme - `ReplicaEncoding` is defined as a tuple of algorithms $(\text{rSetup}, \text{rEnc}, \text{rDec})$, where `rSetup` takes in the security parameter denoted by 1^κ and the maximum number of replicas a client wishes to replicate denoted by 1^n and outputs a public key secret key pair (pk, sk) , `rEnc` is a randomized algorithm which takes a message $m \in \{0, 1\}^*$, a secret key sk and outputs a replica encoding. `rDec` is a deterministic algorithm that takes as input a public key pk , a replica encoding and outputs a message m . Formally,

$$(\text{pk}, \text{sk}) \leftarrow \text{rSetup}(1^\kappa, 1^n), y \leftarrow \text{rEnc}(\text{sk}, m), m \leftarrow \text{rDec}(\text{pk}, y).$$

Definition 1. A tuple $(\text{rSetup}, \text{rEnc}, \text{rDec})$ is a replica encoding if the following holds:

- *Correctness:* For any choice of coins of `rSetup`, the probability of incorrect decoding is

$$\forall n, m, \Pr \left[(\text{pk}, \text{sk}) \leftarrow \text{rSetup}(1^\kappa, 1^n) \right. \\ \left. \text{rDec}(\text{pk}, \text{rEnc}(\text{sk}, m)) \neq m \right] \leq \text{negl}(\kappa)$$

where the probability is over the coins of `rEnc`⁴.

⁴ There exists a generic method for converting a scheme with negligible correctness error into a perfectly correct scheme. To do so augment the `rEnc` algorithm so that it first produces the encoding. Then the new `rEnc` algorithm

- Length of the encoding scheme is denoted by a function $\text{len}(\cdot, \cdot) : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ that takes in the security parameter and the length of the message and outputs the length of the encoding, formally for any κ, m , choice of coins of rSetup ,

$$\forall \kappa, m, \Pr \left[\begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{rSetup}(1^\kappa, 1^n) \\ \text{len}(\kappa, |m|) \neq |\text{rEnc}(\text{sk}, m)| \end{array} \right] \leq \text{negl}(\kappa)$$

where the probability is over the coins of rEnc .

- **s-Sound:** Consider the game $\text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n)$ between an adversary pair $(\mathcal{A}_1, \mathcal{A}_2)$ and a challenger defined in Figure 1. A replica encoding scheme is **s-sound** ($\mathfrak{s} : \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]$), if for any probabilistic poly-time adversaries $(\mathcal{A}_1, \mathcal{A}_2)$, for all $n \in \mathbb{N}$, there exists a function negl such that the following holds.

$$\Pr \left[\begin{array}{l} (v, \text{state}, m) \leftarrow \text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n), \text{ s.t.} \\ |\text{state}| < v \cdot \mathfrak{s}(\kappa, |m|) \cdot \text{len}(\kappa, |m|) \end{array} \right] \leq \text{negl}(\kappa).$$

where the probability is over the coins with the challenger and the two adversaries $\mathcal{A}_1, \mathcal{A}_2$.

Game $\text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n)$

- **Setup:** The challenger (denoted by \mathcal{C}) runs $(\text{pk}, \text{sk}) \leftarrow \text{rSetup}(1^\kappa, 1^n)$ and sends pk to \mathcal{A}_1 . It keeps the secret key sk for itself.
- **File Challenge:** The adversary \mathcal{A}_1 , on input $(1^\kappa, \text{pk})$, chooses a file $m \in \{0, 1\}^*$. It sends m to \mathcal{C} .
- The challenger outputs n encodings of m by calling rEnc n times.

$$\forall i \in [n], y^{(i)} \leftarrow \text{rEnc}(\text{sk}, m)$$

and returns $y^{(1)}, \dots, y^{(n)}$ to \mathcal{A}_1 .

- **State Sharing:** \mathcal{A}_1 outputs $\text{state} \leftarrow \mathcal{A}_1(1^\kappa, \text{pk}, y^{(1)}, \dots, y^{(n)})$ and sends state , the number of replicas 1^n and message m to \mathcal{A}_2 .
- **Guess:** \mathcal{A}_2 on receiving state state , outputs the replica guess to \mathcal{C} ,

$$(\tilde{y}^{(1)}, \dots, \tilde{y}^{(n)}) \leftarrow \mathcal{A}_2(1^\kappa, 1^n, \text{pk}, m, \text{state})$$

- **Verify:** Let $v_i = 1$ if $\tilde{y}^{(i)} = y^{(i)}$ and 0 otherwise. Output $(v = \sum_{i=1}^n v_i, \text{state}, m)$.

Fig. 1: The soundness game for the replica encoding scheme.

A remark on the efficiency. We remark that there can exist trivial constructions of replica encoding by simply concatenating a string m with the randomness ρ i.e. let $\text{rEnc}(\text{sk}, m) = m||\rho$. These schemes are secure for $\mathfrak{s} \in \frac{|\rho|}{|m|+|\rho|} - \frac{\omega(\log \kappa)}{|m|+|\rho|}$. If we consider long ρ , we can construct a sound replica encoding scheme for arbitrary $\mathfrak{s}(\kappa, |m|)$. As a specific example, imagine $\text{rEnc}(\text{sk}, m) = m||\rho$ where $\rho \xleftarrow{R} \{0, 1\}^{99|m|}$. This scheme is trivially correct as m is output in the clear and $\text{len}(\kappa, |m|) = 100|m|$. For all functions \mathfrak{s} such that $\mathfrak{s}(\kappa, |m|) \in \frac{99}{100} - \frac{\omega(\log \kappa)}{100|m|}$, the proposed scheme is **s sound**. Intuitively, for each encoding \mathcal{A}_2 has $99|m| - \omega(\log \kappa)$ information in state and is supposed to output $99|m|$ random bits. Even if they randomly guess the remaining bits the probability of success will be negligible in κ . For this reason we are interested in schemes that do better than the soundness efficiency tradeoffs of this trivial solution.

Definitions in prior work. The formal definition of proof of replica encoding was given by Damgård et al [11]. The soundness game can also be defined from the proof of space literature where the input message to be stored is generated through a private key setup (not revealed to the prover and the verifier) and the time bound for the prover is polynomial. We simply clean up the definitions proposed by [11] and highlight a few differences.

run the deterministic rDec algorithm on the encoding to check that the message was recovered. If not, output the message in the clear and a flag bit indicating that the message is output in plain instead of the encoding. This adds a negligible hit in the security as opposed to the correctness.

The earlier soundness definition is stated in terms of a constant c ,

$$\Pr [|\text{state}| < c \cdot v \cdot \text{len}(\kappa, |m|) \mid (v, \text{state}, m) \leftarrow \text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n)] \leq \text{negl}(\kappa).$$

We make 2 changes to this definition. First, rather than using a constant c , our soundness is stated in terms of a function $\mathfrak{s}(\kappa, |m|)$. This change is purely for increasing the flexibility of the definition, as we can always take $\mathfrak{s}(\kappa, |m|)$ to be a constant function. We make this change to highlight our theorem statement \mathfrak{s} soundness for a larger class of functions. Next, we consider the total probability of the adversary winning the soundness game with any number v replicas rather than the conditional probability per fixed v value. In the original definition, the security can trivially be broken.

To see the issue let's consider an attack algorithm that tries to guess the secret information used by \mathcal{C} when constructing the challenge (e.g. it tries to guess the TDP secret key and the randomness used during the encryption algorithms). If its guess is correct, it can recover the replica encodings by running rEnc in the forward direction and outputs the n replicas; otherwise, it simply gives up and outputs all 0's. Clearly such an adversary should not be viewed as successful since it only succeeds a negligible fraction of the time. However, if its guess is correct (which happens only with negligible probability) it wins the game with $v = n$ and no state bits used. Otherwise, if the guess is incorrect even for some encoding then $v < n$. Even though the winning probability of winning is negligible, when conditioned on $v = n$, this adversary succeeds with probability 1.

Tweaking their definition to include v, state as output of the game and not conditioning on events where the correct replica is output solves the issue.

Other minor differences between our definitions include a rSetup algorithm that sets up the parameters for the scheme. We do this to formalize the alignment and use the KeyGen environment of the underlying trapdoor permutation. The formal definition of replica encoding in DGO includes an efficiency condition defined as exactly $|m| + O(\kappa)$. We do not restrict the efficiency in the formal definition in our work and state it as a desired property that should be required for a practical replica encoding scheme.

4 Lemmas on Random Functions and Permutations

This section contains useful information theoretic lemmas on analyzing random permutations. The first is a result on the hardness of outputting relations on the required ideal primitive given limited advice and restricted behavior. We will use this later in showing adversaries capable in distinguishing between certain games will be able to do the following with noticeable probability. Due to space constraints, we defer the proofs of these lemmas to the full version [18]. Additionally, in the full version we discuss and prove the random function analogues of these lemmas.

Lemma 1. *Let $\mathbb{T}, \mathbb{T}^{-1} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ be oracles to a random permutation and its inverse. Consider any computationally unbounded adversary \mathcal{B} that makes polynomially bounded (in λ) queries to $\mathbb{T}, \mathbb{T}^{-1}$ on input a bounded advice and outputs n pairs (x_i, y_i) without querying them explicitly. If advice is bounded by $n \cdot \lambda - \omega(\log \lambda)$ bits where n is polynomial in λ , the probability that it succeeds is negligible in λ .*

More formally, let the inputs and outputs by \mathcal{B} to oracle \mathcal{O} be denoted by lists $\mathfrak{s}_{\mathcal{B}}^{\mathcal{O}}, \mathfrak{S}_{\mathcal{B}}^{\mathcal{O}}$ respectively. Then,

$$\Pr \left[\begin{array}{l} \exists \text{ advice} \in \{0, 1\}^* \text{ s.t. } |\text{advice}| \leq n \cdot \lambda - \omega(\log \lambda), \\ \{(x_i, y_i)\}_{i=1}^n \leftarrow \mathcal{B}^{\mathbb{T}(\cdot), \mathbb{T}^{-1}(\cdot)}(\text{advice}) \text{ where} \\ \forall i \neq j \in [n], x_i \neq x_j, \mathbb{T}(x_i) = y_i, x_i \notin \mathfrak{s}_{\mathcal{B}}^{\mathbb{T}} \text{ and } y_i \notin \mathfrak{s}_{\mathcal{B}}^{\mathbb{T}^{-1}} \end{array} \right] \leq \text{negl}(\lambda),$$

the probability is over the choice of the permutation \mathbb{T} .

Definition 2. *Let π be a permutation or permutation oracle with domain \mathcal{D} , and let $x_1, x_2 \in \mathcal{D}$. We define the notation $\pi' = \pi[\text{swap}(x_1, x_2)]$ to imply π' to be same as π but swapped on points $x_1, \pi^{-1}(x_2)$. Concretely,*

$$\pi'(x) = \begin{cases} x_2 & x = x_1 \\ \pi(x_1) & x = \pi^{-1}(x_2) \\ \pi(x) & \text{otherwise.} \end{cases}$$

Lemma 2. Let $S_{\mathcal{D}}$ denote the symmetric group on \mathcal{D} . Let $x, r \xleftarrow{R} \mathcal{D}$, and $\pi \xleftarrow{R} S_{\mathcal{D}}$. Then $(x, \pi[\text{swap}(x, r)])$ is uniform on $\mathcal{D} \times S_{\mathcal{D}}$ - i.e. x is independent of $\pi[\text{swap}(x, r)]$.

Definition 3. Multiple invocations of the swap notation are defined as following $\pi[\text{swap}(x_1, y_1), \dots, \text{swap}(x_k, y_k)]$:

- Let $\pi_0 = \pi$.
- Iterate from $i = 1$ to k ,
 - Perform i^{th} swap, $\pi_i = \pi_{i-1}[\text{swap}(x_i, y_i)]$.
- Output π_k .

Lemma 3. Let $\{r_0, r_1, \dots, r_k\} \xleftarrow{R} \mathcal{D}$ and π be a random permutation. Let $S_{\mathcal{D}}$ be the set of all permutations. Let τ be a fixed permutation on \mathcal{D} . Then

$$(r_k, \pi[\text{swap}(r_0, \tau(r_1)), \dots, \text{swap}(r_{k-1}, \tau(r_k))])$$

is uniform on $\mathcal{D} \times S_{\mathcal{D}}$.

We introduce another useful result on the probability of finding collisions on a deterministic function h .

Lemma 4. Let $\mathcal{D}(\kappa), \mathcal{R}(\kappa)$ represent domain, range respectively dependent on the security parameter. Let h be any deterministic function that maps values in domain $\mathcal{D}(\kappa)$ to range $\mathcal{R}(\kappa)$. Then,

$$\Pr_a[\exists b \neq a \in \mathcal{D}(\kappa), h(a) = h(b)] \geq \frac{|\mathcal{D}(\kappa)| - |\mathcal{R}(\kappa)| + 1}{|\mathcal{D}(\kappa)|}.$$

5 Replica Encoding in the Ideal Permutation Model.

We now give the construction and proof of our replica encoding scheme from trapdoor permutations in the ideal permutation model and the random oracle model. As stated in the introduction, the construction itself is a close variant of [11]. However, our proof will introduce new analysis techniques that account for an attacker that stores state in an arbitrary manner.

Let κ denote the security parameter. Let $\lambda(\kappa)$ (denoted by λ) be a function polynomial in κ and represents block length in our construction. We use a trapdoor permutation ($\text{KeyGen}, f(\cdot, \cdot), f^{-1}(\cdot, \cdot)$) where the domain for the family of trapdoor functions is $\mathcal{D}_{\text{pk}} = \{0, 1\}^\lambda$ where KeyGen is setup with security parameter κ . Let T, T^{-1} be random permutation oracles on the same domain $\{0, 1\}^\lambda$ and H be a random oracle on the range $\{0, 1\}^\lambda$.

5.1 Construction

Let $r(\kappa, n, |m|)$ (denoted by r) be the number of rounds in our scheme. For our construction, it depends on the security parameter, maximum number of replicas chosen during setup and the message length.

rSetup($1^\kappa, 1^n$):

Run $\text{KeyGen}(1^\kappa) \rightarrow (\text{pk}, \text{sk})$. Output $(\text{pk}' = (\text{pk}, n), \text{sk}' = (\text{sk}, n))$.

rEnc(sk', m):

- Parse $\text{sk}' = (\text{sk}, n)$.
- Choose a string $\rho \xleftarrow{R} \{0, 1\}^\kappa$.
- Divide m into b blocks of length λ i.e. $m = m_1 || m_2 || \dots || m_b, b = \lceil |m|/\lambda \rceil$.
- Set $r = n \cdot b \cdot \lambda$.
- Compute $\forall t \in [b]$,

$$Y_{t,0} = m_t \oplus H(\rho || t).$$

- For rounds j from 1 to r , compute:

$$Y_{t,j} = f^{-1}(\text{sk}, \mathbb{T}(Y_{t,j-1})).$$

- Let $y_r = Y_{1,r} || \dots || Y_{b,r}$ and output (y_r, ρ) .

rDec(pk', y):

- Parse $\text{pk}' = (\text{pk}, n)$.
- Parse y as (y_r, ρ) . Parse y_r as $Y_{1,r} || \dots || Y_{b,r}$, where $b = \lceil |y_r|/\lambda \rceil$ and $r = n \cdot b \cdot \lambda$.
- For rounds j from $r - 1$ to 0:
 - Compute $\forall t \in [b]$,

$$Y_{t,j} = \mathbb{T}^{-1}(f(\text{pk}, Y_{t,j+1})).$$

- $\forall t \in [b]$ compute,

$$m_t = Y_{t,0} \oplus \mathbb{H}(\rho || t)$$

Output $m = m_1 || \dots || m_b$.

The encoding length for our scheme is $\text{len}(\kappa, |m|) = |m| + O(\kappa)$.⁵

5.2 Security of Replica Encoding Scheme

Theorem 1. *Assuming $(\text{KeyGen}(1^\kappa), f(\cdot, \cdot), f^{-1}(\cdot, \cdot))$ is a secure trapdoor permutation and $\mathbb{T}, \mathbb{T}^{-1}$ are oracles to a random permutation on domain and range $\{0, 1\}^\lambda$ and \mathbb{H} is a random oracle on the same range. Then our construction for ReplicaEncoding described above is s -sound according to Definition 1 for all $\kappa, n \in \mathbb{N}$ and $s \in 1 - \frac{\omega(\log \kappa)}{\lambda}$.*

Sequence of Games Our proof proceeds via a sequence of games as described below. We assume that adversaries have their randomness non-uniformly fixed in each game to maximize their success. The changes in each game in comparison to the previous one are indicated with red. Details of the previous game are copied without explicit rewriting.

Game 0: This is the original $\text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n)$ security game where we record the queries made by the adversaries in lists. We also assume that any list is ordered and stores distinct elements. More concretely, when in Phase 1 a query x is made on \mathcal{O} , \mathcal{C} checks if $x \notin \mathbf{u}^\mathcal{O}$ and updates the list $\mathbf{u}^\mathcal{O}$ if the condition is true. It performs this operation of maintaining the list for each Phase and oracle separately. Denote $\mathbf{q}_1^\mathcal{O}, \mathbf{q}_2^\mathcal{O}, \mathbf{q}_3^\mathcal{O}$ as the functions that take in the security parameter and output the total distinct queries made by the adversaries to oracle \mathcal{O} during the three phases respectively. Note that the functionality of the oracles is still the same, we just record queries.

- **Setup:** The challenger (denoted by \mathcal{C}) runs $(\text{pk}', \text{sk}') \leftarrow \text{rSetup}(1^\kappa, 1^n)$ and sends public key pk' to \mathcal{A}_1 . It keeps the secret key sk' for itself.
- **Phase 1:** The adversary \mathcal{A}_1 issues queries on $\mathbb{T}, \mathbb{T}^{-1}, \mathbb{H}$, \mathcal{C} responds the query back to \mathcal{A}_1 . **Let the queries on oracle \mathcal{O} be denoted by an ordered and distinct list $\mathbf{u}^\mathcal{O} = (u_1^\mathcal{O}, \dots, u_{q_1^\mathcal{O}}^\mathcal{O})$ and their outputs be denoted by an ordered and distinct list $\mathbf{U}^\mathcal{O} = (U_1^\mathcal{O}, \dots, U_{q_1^\mathcal{O}}^\mathcal{O})$.**
- **File Challenge:** $m \in \{0, 1\}^* \leftarrow \mathcal{A}_1^{\mathbb{H}(\cdot), \mathbb{T}(\cdot), \mathbb{T}^{-1}(\cdot)}(1^\kappa, \text{pk}')$. It sends m to \mathcal{C} who parses pk' as (pk, n) ; sk' as (sk, n) and does the following:
 - Divide m into b blocks of length λ i.e. $m = m_1 || m_2 || \dots || m_b$, $b = \lceil |m|/\lambda \rceil$.
 - For $i \in [n]$,
 - * Choose a string $\rho_i \xleftarrow{R} \{0, 1\}^\kappa$.

⁵ Upto additional rounding factors.

* Compute $\forall t \in [b]$,

$$Y_{t,0}^{(i)} = m_t \oplus H(\rho_i || t).$$

* For rounds j from 1 to r and $\forall t \in [b]$,

· Compute $Y_{t,j}^{(i)}$ from $Y_{t,j-1}^{(i)}$ as

$$Y_{t,j}^{(i)} = f^{-1}(\text{sk}, T(Y_{t,j-1}^{(i)})).$$

* Let $y_r^{(i)} = Y_{1,r}^{(i)} || \dots || Y_{b,r}^{(i)}$ and set $y^{(i)} = (y_r^{(i)}, \rho_i)$.

\mathcal{C} returns $y^{(1)}, y^{(2)}, \dots, y^{(n)}$ to \mathcal{A}_1 .

– **Phase 2:** \mathcal{A}_1 issues additional queries on T, T^{-1}, H , \mathcal{C} responds the query back to \mathcal{A}_1 . Let the queries on oracle \mathcal{O} be denoted by an ordered and distinct list $\mathbf{v}^\mathcal{O} = (v_1^\mathcal{O}, \dots, v_{q_2^\mathcal{O}}^\mathcal{O})$ and their outputs be denoted by an ordered and distinct list $\mathbf{V}^\mathcal{O} = (V_1^\mathcal{O}, \dots, V_{q_2^\mathcal{O}}^\mathcal{O})$.

– **State Sharing:** \mathcal{A}_1 outputs state $\text{state} \leftarrow \mathcal{A}_1^{H(\cdot), T(\cdot), T^{-1}(\cdot)}(1^\kappa, \text{pk}', y)$ and sends state to \mathcal{A}_2 .

– **Phase 3:** The adversary \mathcal{A}_2 queries on T, T^{-1}, H , \mathcal{C} responds the query back to \mathcal{A}_2 . Let the queries on oracle \mathcal{O} be denoted by an ordered and distinct list $\mathbf{w}^\mathcal{O} = (w_1^\mathcal{O}, \dots, w_{q_3^\mathcal{O}}^\mathcal{O})$ and their outputs be denoted by an ordered and distinct list $\mathbf{W}^\mathcal{O} = (W_1^\mathcal{O}, \dots, W_{q_3^\mathcal{O}}^\mathcal{O})$.

– **Guess:** \mathcal{A}_2 outputs the replica guesses to \mathcal{C} .

$$\{\tilde{y}^{(i)}\} \leftarrow \mathcal{A}_2(1^\kappa, \text{pk}', m, \text{state}).$$

– **Verify:** Let $v_i = 1$ if $\tilde{y}^{(i)} = y^{(i)}$ and 0 otherwise. Adversary wins if $|\text{state}| < \sum v_i \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Game 1 : In this game we remove the sk and rely on the public key with an additional reprogramming step at oracle H . This helps us further down the road in showing a reduction to the security of the trapdoor permutation.

– **Setup:** The challenger (denoted by \mathcal{C}) runs $(\text{pk}', \text{sk}') \leftarrow \text{rSetup}(1^\kappa, 1^n)$ and sends public key pk' to \mathcal{A}_1 . It keeps the secret key sk' for itself. **Set flag = 0.**

– **Phase 1:** ...

– **File Challenge:** $m \in \{0, 1\}^* \leftarrow \mathcal{A}_1^{H(\cdot), T(\cdot), T^{-1}(\cdot)}(1^\kappa, \text{pk}')$. It sends m to \mathcal{C} who parses pk' as (pk, n) ; sk' as (sk, n) and does the following:

• Divide m into b blocks of length λ i.e. $m = m_1 || m_2 || \dots || m_b$, $b = \lceil |m|/\lambda \rceil$.

• For $i \in [n]$,

* Choose a string $\rho_i \xleftarrow{R} \{0, 1\}^\kappa$.

Prequery Check H If $\exists t \in [b] : \rho_i || t \in \mathbf{u}^H$, set **flag = 1**.

* Sample $\{Y_{t,r}^{(i)}\}_{t \in [b]} \xleftarrow{R} \{0, 1\}^\lambda$

* For rounds j from r to 1 and $\forall t \in [b]$,

· Compute $Y_{t,j-1}^{(i)}$ from $Y_{t,j}^{(i)}$ as

$$Y_{t,j-1}^{(i)} = T^{-1}(f(\text{pk}, Y_{t,j}^{(i)})).$$

* For each block $\forall t \in [b]$, reprogram H

$$H(\rho_i || t) = m_t \oplus Y_{t,0}^{(i)}$$

* Let $y_r^{(i)} = Y_{1,r}^{(i)} || \dots || Y_{b,r}^{(i)}$ and set $y^{(i)} = (y_r^{(i)}, \rho_i)$.

\mathcal{C} returns $y^{(1)}, y^{(2)}, \dots, y^{(n)}$ to \mathcal{A}_1 .

– **Phase 2, State Sharing, Phase 3, Guess:** ...

– **Verify:** Let $v_i = 1$ if $\tilde{y}^{(i)} = y^{(i)}$ and 0 otherwise. Adversary wins if **flag = 0** and $|\text{state}| < \sum v_i \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Game 2: In this game an adversary wins if they query on the oracle rather than outputting the replica. This helps us ease the notation by only focussing at the oracle query lists.

– **Setup, Phase 1, File Challenge, Phase 2, State Sharing, Phase 3:** ...

– **Guess:** ...

\mathcal{C} adds the guess to \mathcal{A}_2 's lists of queries to T in Phase 3, i.e. $\forall i \in [n]$, let $\tilde{y}^{(i)} = (\tilde{Y}_{0,r}^{(i)} || \dots || \tilde{Y}_{b,r}^{(i)}, \tilde{\rho}_i)$. $\forall t \in [b]$ add $\tilde{Y}_{t,r}^{(i)}$ to list of queries to T by \mathcal{A}_2 in Phase 3.

- **Verify:** Let $v_i = 1$ if $\forall t \in [b]$, T is queried on $Y_{t,r}^{(i)}$ and 0 otherwise. Adversary wins if $\text{flag} = 0$ and $|\text{state}| < \sum v_i \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Game 3: In this game, we look at the queries made by the adversary and require that it traverses atleast one block in some replica sequentially.

- **Setup, Phase 1, File Challenge, Phase 2, State Sharing, Phase 3, Guess:**
- **Sequentiality:**

We consider going through \mathcal{A}'_2 's ordered list of queries to T and T^{-1} . If $\forall i \in [n] \forall t \in [b]$, there is a point in time such that some $Y_{t,j+1}^{(i)}$ was queried on T or $f(\text{pk}, Y_{t,j+1}^{(i)})$ was queried on T^{-1} when \mathcal{A}_2 has not made a query to T for $Y_{t,j}^{(i)}$, then set $\text{flag} = 1$.

- **Verify:** Let $v_i = 1$ if $\forall t \in [b]$, T is queried on $Y_{t,r}^{(i)}$ and 0 otherwise. Adversary wins if $\text{flag} = 0$ and $|\text{state}| < \sum v_i \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Game 4: In this game, we guess the block which the adversary traversed sequentially. We concentrate on one randomly chosen block and replica and the adversary wins if it outputs the correct encoding for this block. We lose a multiplicative factor of $b \cdot n$ in the reduction due to this change.

- **Setup:** The challenger (denoted by \mathcal{C}) runs $(\text{pk}', \text{sk}') \leftarrow \text{rSetup}(1^\kappa, 1^n)$ and sends public key pk' to \mathcal{A}_1 . It keeps the secret key sk' for itself. Set $\text{flag} = 0$.
- **Choose a random $\beta \in [b]$ and $\gamma \in [n]$.**
- **Phase 1, File Challenge, Phase 2, State Sharing, Phase 3, Guess:**
- **Sequentiality:**

We consider going through \mathcal{A}'_2 's list of queries to T and T^{-1} . If there is a point in time such that some $Y_{\beta,j+1}^{(\gamma)}$ was queried on T or $f(\text{pk}, Y_{\beta,j+1}^{(\gamma)})$ was queried on T^{-1} when \mathcal{A}_2 has not made a query to T for $Y_{\beta,j}^{(\gamma)}$, then set $\text{flag} = 1$.

- **Verify:** ~~Let $v_i = 1$ if $\forall t \in [b]$, T is queried on $Y_{t,r}^{(i)}$ and 0 otherwise.~~ Adversary wins if T is queried on $Y_{\beta,r}^{(\gamma)}$, $\text{flag} = 0$, and $|\text{state}| < n \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Game 5: In this game, we reprogram the oracles H, T to have a permutation which we can analyze cleanly. The primary idea behind this game is that there will exist two sequences of values on the chosen block and replica for which any adversary \mathcal{A}_1 produces the same state. These possibilities for a “switch” are set up in this game. H is programmed to output $Y_{\beta,0}^{(\gamma)}$ and for $i \in [r]$, the values $A_{i,0}, A_{i,1}$ have a choice to be mapped to either of the two $A_{i+1,0}, A_{i+1,1}$ depending on the sampled index x . The collision check makes sure that the reprogramming preserves the permutation property of T and the prequery check is done to make sure that none of the values were queried in the oracle lists in the previous phase. The oracle T_x is then reprogrammed according to the swap operation defined in [Definition 3](#) where for $i \in [r]$, x_i is now mapped to $f(\text{pk}, x_{i+1})$ where x_i is used to indicate the notation for $A_{i,x[i]}$.

- **Setup, Phase 1:**
- **Sampling a new permutation:**

- Sample, $Y_{\beta,0}^{(\gamma)}, A_{1,0}, \dots, A_{r,0}, A_{1,1}, \dots, A_{r,1} \xleftarrow{R} \{0, 1\}^\lambda$.

Let $\mathcal{Z}_1 = \{Y_{\beta,0}^{(\gamma)}, A_{1,0}, \dots, A_{r,0}, A_{1,1}, \dots, A_{r,1}\}$.

Let $\mathcal{Z}_2 = \{f(\text{pk}, A_{1,0}), \dots, f(\text{pk}, A_{r,0}), f(\text{pk}, A_{1,1}), \dots, f(\text{pk}, A_{r,1})\}$.

Collision Check: If $|\mathcal{Z}_1| \neq 2r + 1$, set $\text{flag} = 1$.

Prequery Check T : If $(\mathcal{Z}_1 \cup \mathcal{Z}_2) \cap (\mathbf{u}^\top \cup \mathbf{u}^{\top^{-1}} \cup \mathbf{U}^\top \cup \mathbf{U}^{\top^{-1}}) \neq \emptyset$, set $\text{flag} = 1$.

- Sample a random setting $x \xleftarrow{R} \{0, 1\}^r$. Let $x[k]$ denote the k^{th} bit of x . We will write x_j to refer to $A_{j,x[j-1]}$ and denote $A_{j,\bar{x}[j-1]}$ with \bar{x}_j . Set x_0 to denote $Y_{\beta,0}^{(\gamma)}$.
- Define T_x using swap ([Definition 3](#)):

$$\mathsf{T}_x = \mathsf{T}[\text{swap}(x_0, f(\text{pk}, x_1)), \dots, \text{swap}(x_{r-1}, f(\text{pk}, x_r))].$$

- Let T_x^{-1} be the inverse of T_x .

- **File Challenge:** $m \in \{0, 1\}^* \leftarrow \mathcal{A}_1^{\mathsf{H}(\cdot), \mathsf{T}(\cdot), \mathsf{T}^{-1}(\cdot)}(1^\kappa, \text{pk}')$. It sends m to \mathcal{C} who parses pk' as (pk, n) ; sk' as (sk, n) and does the following:

- Divide m into b blocks of length λ i.e. $m = m_1 || m_2 || \dots || m_b$, $b = \lceil |m|/\lambda \rceil$.
- For $i \in [n]$,

* Choose a string $\rho_i \xleftarrow{R} \{0, 1\}^\kappa$.

Prequery Check H If $\exists t \in [b] : \rho_i || t \in \mathbf{u}^H$, set **flag** = 1.

* Sample $\{Y_{t,r}^{(i)}\}_{t \in [b]} \xleftarrow{R} \{0, 1\}^\lambda$

* For rounds j from r to 1 and $\forall t \in [b]$, **continue** if $t \neq \beta$ or $i \neq \gamma$,

· Compute $Y_{t,j-1}^{(i)}$ from $Y_{t,j}^{(i)}$ as

$$Y_{t,j-1}^{(i)} = T^{-1}(f(\text{pk}, Y_{t,j}^{(i)})).$$

* For each block $\forall t \in [b]$, reprogram **H**

$$H(\rho_i || t) = m_t \oplus Y_{t,0}^{(i)}.$$

* Let $y_r^{(i)} = Y_{1,r}^{(i)} || \dots || Y_{b,r}^{(i)}$ and set $y^{(i)} = (y_r^{(i)}, \rho_i)$.

\mathcal{C} returns $y^{(1)}, y^{(2)}, \dots, y^{(n)}$ to \mathcal{A}_1 .

– **Phase 2:** Use T_x, T_x^{-1} to answer queries for T, T^{-1} respectively.

– **State Sharing:**

– **Phase 3:** Use T_x, T_x^{-1} to answer queries for T, T^{-1} respectively.

– **Guess, Sequentiality:**

– **Verify:** Adversary wins if T is queried on $Y_{\beta,r}^{(\gamma)}$, **flag** = 0, and $|\text{state}| < n \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Game 6: In this game, \mathcal{C} has unbounded computation time and calls $\mathcal{A}_1, \mathcal{A}_2$ exponentially many times to find a collision to **state** through the procedure **search**. The setting y' for which the procedure **search** outputs a collision in **state** is stored in a set which is outputted at the end of the procedure. **search**($1^\kappa, y, \text{state}; \zeta$) takes input $y \in \{0, 1\}^r$, **state** and runs algorithms $\mathcal{A}_1, \mathcal{A}_2$ on Game 5. Let ζ be the randomness used by the procedure and denotes all the random coins (except those used to sample x) used by \mathcal{C} . The procedure is described in [Figure 2](#).

search($1^\kappa, y, \text{state}; \zeta$)

Inputs: Security parameter - 1^κ
Oracle Settings on T - $y \in \{0, 1\}^r$
State - **state**
Randomness used in the game - ζ

Output: Set containing all oracle settings with collision in **state** - \mathcal{S}

– Set $\mathcal{S} = \emptyset$.

– $\forall y' \neq y \in \{0, 1\}^r$,

- Run $\mathcal{A}_1, \mathcal{A}_2$ on Game 5 with randomness defined by ζ and using y' instead of x in the game.
- Let **state'** be the state shared between $\mathcal{A}_1, \mathcal{A}_2$.
- If **state'** = **state** and \mathcal{A}_2 wins Game 5, then $\mathcal{S} = \mathcal{S} \cup \{y'\}$.

– Output \mathcal{S} .

Fig. 2: Routine search

– **Setup, Phase 1, Sampling a New Permutation, File Challenge, Phase 2, State Sharing:**

– **Running search:** Let ζ be all the random coins (except those used to sample x) used by \mathcal{C} . Let $\mathcal{S} \leftarrow \text{search}(1^\kappa, x, \text{state}; \zeta)$.

If $\mathcal{S} = \emptyset$ set **flag** = 1 and $x' = x$, otherwise sample $x' \xleftarrow{R} \mathcal{S}$.

– **Setting switched oracle:**

• Let $x'[k]$ denote the k^{th} bit of x' . We will write x'_j to refer to $A_{j,x'[j-1]}$ and denote $A_{j,x'[j-1]}$ with \bar{x}'_j . Set x'_0 to denote $Y_{\beta,0}^{(\gamma)}$.

• Define $T_{x'}$ to be:

$$T_{x'} = T[\text{swap}(x'_0, f(\text{pk}, x'_1)), \dots, \text{swap}(x'_{r-1}, f(\text{pk}, x'_r))].$$

• Let $T_{x'}^{-1}$ be the inverse of $T_{x'}$.

- **Phase 3:** Use $T_{x'}, T_{x'}^{-1}$ to answer queries for T, T^{-1} respectively.
- **Guess:**
- **Sequentiality:**
If $\exists j \in [0, r] : (x'_{j+1})$ was queried on T or $f(\text{pk}, x'_{j+1})$ was queried on T^{-1} while T had not been queried on x'_j , set $\text{flag} = 1$.
- **Verify:** Adversary wins if T is queried on x'_r , $\text{flag} = 0$, and $|\text{state}| < n \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Game 7: In this game we modify the verification step for which an adversary can win this game. We increase it's winning probability so that the adversary can win if it doesn't query the full sequence, but queries at the point where the sequences x, x' diverge. Notice that we define another oracle $T_{x'}^\delta$ here that doesn't reprogram the complete sequence. This change is statistically indistinguishable to the adversary.

- **Setup, Phase 1, Sampling a New Permutation, File Challenge, Phase 2, State Sharing, Running search:**
- **Setting switched oracle:**
 - Let $x'[k]$ denote the k^{th} bit of x' . We will write x'_j to refer to $A_{j, x'_{[j-1]}}$ and denote $A_{j, x'_{[j-1]}}$ with \bar{x}'_j . Set x'_0 to denote $Y_{\beta, 0}^{(\gamma)}$.
 - Let δ be the first index for which $x_\delta \neq x'_\delta$.
 - Define $T_{x'}^\delta$ to be:

$$\begin{aligned} T_{x'}^\delta &= T[\text{swap}(x'_0, f(\text{pk}, x'_1)), \dots, \text{swap}(x'_{\delta-1}, f(\text{pk}, x'_\delta))] \\ &= T[\text{swap}(x_0, f(\text{pk}, x_1)), \dots, \text{swap}(x_{\delta-1}, f(\text{pk}, \bar{x}_\delta))]. \end{aligned}$$

- Let $T_{x'}^{-1}$ be the inverse of $T_{x'}$.
- **Phase 3, Guess:**
- **Sequentiality:**
- **Verify:** Adversary wins if T is queried on \bar{x}_δ , $\text{flag} = 0$ and $|\text{state}| < n \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Game 8: In this game we observe that \mathcal{C} need not be unbounded computation time and only needs to the guess the first prefix at which x, x' differ to successfully output one sequential query.

- **Setup, Phase 1, Sampling a New Permutation, File Challenge, Phase 2, State Sharing:**
- **Running search:**
- **Setting switched oracle:**
 - Let $x'[k]$ denote the k^{th} bit of x' . We will write x'_j to refer to $A_{j, x'_{[j-1]}}$ and denote $A_{j, x'_{[j-1]}}$ with \bar{x}'_j . Set x'_0 to denote $Y_{\beta, 0}^{(\gamma)}$.
 - Let $\delta \xleftarrow{R} [r]$
 - Define $T_{x'}^\delta$ to be:

$$T_{x'}^\delta = T[\text{swap}(x_0, f(\text{pk}, x_1)), \dots, \text{swap}(x_{\delta-1}, f(\text{pk}, \bar{x}_\delta))].$$

- Let $T_{x'}^{-1}$ be the inverse of $T_{x'}$.
- **Phase 3, Guess:**
- **Verify:** Adversary wins if T is queried on \bar{x}_δ and $\text{flag} = 0$.

Indistinguishability of Games Let $F_i(\kappa)$ (denoted by F_i) be the probability that the adversaries win at the end of Game i .

Game 0

Proof. Game 0 is a restatement of the original $\text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n)$ game with two differences, (i) the syntactical change to note down queries to each oracle, (ii) expands on rEnc . Both the syntactical changes do not change the functioning of the game. \square

Game 1

Lemma 5. $\Pr[F_1] \geq \Pr[F_0] - \text{negl}(\kappa)$.

Proof. Game 1 differs from Game 0 in how queries are answered to the adversaries and the possibility of flag being set. Let $\Pr[F_0] = \epsilon$ be the probability of adversary winning in Game 0. Let E^H be the event flag is set due to **Prequery Check H**. We can upper bound the difference in probability of F_1 and F_0 by the sum of (i) the probability that E^H occurs, (ii) and the statistical difference of the output of \mathcal{C} from the alternate method of encoding generation.

Claim. $\Pr[E^H] = \text{negl}(\kappa)$.

Proof. Since each ρ_i is generated uniformly on $\{0, 1\}^\kappa$ and independent of \mathcal{A}_1 , we can bound the probability that any fixed query is equal to a particular ρ_i as $k \in [q_1^H]$ $i \in [n]$ $\Pr[u_k^H = \rho_i] = 2^{-\kappa}$. From this, we can union bound the

$$\Pr[\exists k \in [q_1^H] \ i \in [n] \ t \in [b] : u_k^H = \rho_i || t] \leq q_1^H \cdot n \cdot b \cdot 2^{-\kappa}$$

Since $q_1^H, n, b \in \text{poly}(\kappa)$, this is negligible □

Claim. The distribution of $T \times \{Y_{t,r}^{(i)}\}$ in Game 0 is statistically close to uniform.

Proof. First, observe that with probability $1 - \text{negl}(\kappa)$, $\forall i \in [n], t \in [b]$ $\rho_i || t$ is not queried H on by \mathcal{A}_1 before submitting m . This is apparent as each ρ_i is uniformly random on a domain of size 2^κ and \mathcal{A}_1 can make at most $\text{poly}(\kappa)$ queries. Since $Y_{t,0}^{(i)} = H(\rho_i || t) \oplus m_t$, and m is independent of $H(\rho_i || t)$, we can say $\{Y_{t,0}^{(i)}\}$ is uniform and independent of at least a $1 - \text{negl}(\kappa)$ fraction of T , and once an T is fixed, this defines a bijective relation from $Y_{t,0}^{(i)}$ to $Y_{t,r}^{(i)}$, so the latter is also uniform and independent of at least a $1 - \text{negl}(\kappa)$ fraction of T . Bounding the statistical distance with $\text{negl}(\kappa)$. □

Since in Game 1, it is apparent that $T \times \{Y_{t,r}^{(i)}\}$ is uniform by the fact that $\{Y_{t,r}^{(i)}\}$ are generated independent of T , [Claim 5.2](#) bounds the statistical distance between the responses of \mathcal{C} between Game 0 and 1 with $\text{negl}(\kappa)$. Combined with the previous claim, we can conclude the total difference between Game 0 and 1 is negligible. □

Game 2

Claim. $\Pr[F_2] \geq \Pr[F_1]$.

Proof. Game 2 differs from Game 1 in the winning condition. Let $\Pr[F_1] = \epsilon$ be the probability of adversary winning in winning Game 1. Then the adversary is sure to win Game 2 as \mathcal{C} records the output in the query. □

Game 3

Lemma 6. $\Pr[F_3] \geq \Pr[F_2] - \text{negl}(\kappa)$

Proof. Observe that these games only differ when flag is set to 1 and adversary ends up winning Game 2. Let ϵ be the probability that $(\mathcal{A}_1, \mathcal{A}_2)$ wins in such a manner. We will refer to this as winning non-sequentially.

Claim. \mathcal{A}_1 will only query on $Y_{t,j}^{(i)}$ or $f(\text{pk}, Y_{t,j}^{(i)})$ in **Phase 1** with negligible probability.

$$\Pr[\exists(i, j, t) : Y_{t,j}^{(i)} \in \mathbf{u}^T \vee f(\text{pk}, Y_{t,j}^{(i)}) \in \mathbf{u}^{T^{-1}}] = \text{negl}(\kappa)$$

Proof. Once we fix \mathbb{T} , any $Y_{t,j}^{(i)}$ and $f(\mathbf{pk}, Y_{t,j}^{(i)})$ have a bijective relation onto $Y_{t,r}^{(i)}$, which is uniform and independently generated, so we can union bound the probability

$$\begin{aligned} & \Pr[\exists(i, j, t) : Y_{t,j}^{(i)} \in \mathbf{u}^{\mathbb{T}} \vee f(\mathbf{pk}, Y_{t,j}^{(i)}) \in \mathbf{u}^{\mathbb{T}^{-1}}] \\ & \leq \sum_{(i,j,t) \in [n] \times [r] \times [b]} \left(\sum_{q=0}^{\mathbf{q}_1^{\mathbb{T}}} \Pr[Y_{t,j}^{(i)} = \mathbf{u}_q^{\mathbb{T}}] + \sum_{q=0}^{\mathbf{q}_1^{\mathbb{T}^{-1}}} \Pr[f(\mathbf{pk}, Y_{t,j}^{(i)}) = \mathbf{u}_q^{\mathbb{T}^{-1}}] \right) \\ & = n \cdot b \cdot r \left(\mathbf{q}_1^{\mathbb{T}} \cdot 2^{-\lambda} + \mathbf{q}_1^{\mathbb{T}^{-1}} \cdot 2^{-\lambda} \right) \end{aligned}$$

Since $b, r, n, \mathbf{q}_1^{\mathbb{T}}, \mathbf{q}_1^{\mathbb{T}^{-1}}$ are all $\text{poly}(\kappa)$, this is negligible. \square

Claim. With all but negligible probability, the $\{Y_{t,j}^{(i)}\}$ are unique.

$$\Pr[\exists(i_1, j_1, t_1) \neq (i_2, j_2, t_2) : Y_{t_1, j_1}^{(i_1)} = Y_{t_2, j_2}^{(i_2)}] \leq \text{negl}(\kappa)$$

Proof. We do a case by case analysis,

- Assume $(i_1, t_1) = (i_2, t_2)$ - where both values are on the same block, but different rounds. Lets consider $j_1 \in [r]$. Let j_2 be the smallest index greater than j_1 such that $Y_{t_1, j_1}^{(i_1)} = Y_{t_2, j_2}^{(i_2)}$. Now consider the set of all permutations such that $\mathbb{T}(Y_{t_1, j}^{(i_1)}) = f(\mathbf{pk}, Y_{t_1, j+1}^{(i_1)})$ for $j \in [j_1, j_2 - 2]$. This set fixes the permutation on $j_2 - j_1 - 1$ points. Thus the probability over the remaining permutations that $\mathbb{T}(Y_{t_1, j_2-1}^{(i_1)}) = Y_{t_2, j_1}^{(i_2)}$ is $1/(2^\lambda - j_2 + j_1 + 1)$. Union bounding over j_1, j_2, i_1, t_1 , which are all $\text{poly}(\kappa)$ the probability is still negligible.
- Assume $(i_1, t_1) \neq (i_2, t_2)$ - where the two values are on different blocks. Lets consider $j_1, j_2 \in [r]$. Observe that if $\max(j_1, j_2) < r$, $Y_{t_1, j_1}^{(i_1)} = Y_{t_2, j_2}^{(i_2)} \Rightarrow Y_{t_1, j_1+1}^{(i_1)} = Y_{t_2, j_2+1}^{(i_2)}$. Without loss of generality, assume $j_1 \leq j_2$, this implies $Y_{t_1, j_1-j_2+r}^{(i_1)} = Y_{t_2, r}^{(i_2)}$. However, since $Y_{t_2, r}^{(i_2)}$ was independently randomly chosen, the probability it is equal to any $Y_{t_1, j}^{(i_1)}$ occurs with probability $\leq \frac{r}{2^\lambda}$, which is negligible. We can union bound over t_1, i_1, t_2, i_2 , which are all $\text{poly}(\kappa)$, so the probability is still negligible. \square

Now consider the following computationally unbounded algorithm \mathcal{B}' with access to oracle $\mathbb{T}, \mathbb{T}^{-1}$. This algorithm will translate a non-sequential \mathcal{A}_2 into a reduction to the game outlined in [Lemma 1](#) by working backwards from m and using its unbounded computation to invert the trapdoor permutation and recover the output of a non-sequential random oracle query without ever querying on its input.

Reduction $\mathcal{B}'^{\mathbb{T}(\cdot), \mathbb{T}^{-1}(\cdot)}$ (advice):

Goal: Produce Input Output oracle pairs without explicitly querying the oracle.

– **Setup:**

- Sample a random function $H(\cdot)$ and use it to answer oracle queries made by $\mathcal{A}_1, \mathcal{A}_2$.
- Perform, **Setup** and **Phase 1** as in Game 3.
- Receive $m \leftarrow \mathcal{A}_1^{\mathbb{H}(\cdot), \mathbb{T}(\cdot), \mathbb{T}^{-1}(\cdot)}(1^\kappa, \mathbf{pk}')$ after **Phase 1**. Parse \mathbf{pk}' as (\mathbf{pk}, n) .
- Choose a set of random $\{\rho_i\}_{i \in [n]}$ and compute $\{Y_{t,0}^{(i)} = H(\rho_i || t) \oplus m_t\}_{t \in [b], i \in [n]}$.
- Parse advice as $(\text{state}, \mathbf{Q} = \{(q_k, i_k, t_k, j_k)\}_k)$ where k is polynomial in security parameter and q_k represents query that will be made by algorithm \mathcal{A}_2 , i_k, t_k, j_k represent the replica, block and round number respectively.

– **Simulate:**

- Run $\mathcal{A}_2^{\mathsf{H}(\cdot), \mathsf{T}(\cdot), \mathsf{T}^{-1}(\cdot)}(1^\kappa, \mathsf{pk}', m, \mathsf{state})$, interacting with the random oracle queries it makes to $\mathsf{T}, \mathsf{T}^{-1}$. Let $\mathbf{w}^{\mathsf{T}, \mathsf{T}^{-1}}$ denote the ordered and distinct list of queries \mathcal{A}_2 makes to T or T^{-1} , and let $w_q^{\mathsf{T}, \mathsf{T}^{-1}}$ refer to the q^{th} element in this list. We perform the operations below while running \mathcal{A}_2 .

- $\forall k$, let x_k be $w_{q_k}^{\mathsf{T}, \mathsf{T}^{-1}}$. Compute $Y_{t_k, j_k}^{(i_k)}$ from $Y_{t_k, 0}^{(i_k)}$.
 - * For j from 1 to j_k ,

$$Y_{t_k, j}^{(i_k)} = \mathsf{f}^{-1}(\mathsf{sk}, \mathsf{T}(Y_{t_k, j-1}^{(i_k)})).$$

Let $y_k = x_k$ if the q_k^{th} query was made to T^{-1} and $y_k = \mathsf{f}(\mathsf{pk}, x_k)$ if it were made to T .

- Any time \mathcal{A}_2 attempts to query T on $Y_{t_k, j_k}^{(i_k)}$ or T^{-1} on y_k , \mathcal{B}' doesn't query the true oracle and instead returns y_k or $Y_{t_k, j_k}^{(i_k)}$ to \mathcal{A}_2 respectively.
- For any other queries \mathcal{A}_2 makes to an oracle, \mathcal{B}' simply queries the appropriate oracle, returns the query results and completes execution of \mathcal{A}_2 .

- **Return:** Output the pairs $\{(Y_{t_k, j_k}^{(i_k)}, \mathsf{f}(\mathsf{pk}, x_k))\}_k$.

Claim. Suppose $Y_{t, j+1}^{(i)}$ is queried on T or $\mathsf{f}(\mathsf{pk}, Y_{t, j+1}^{(i)})$ is queried on T^{-1} by \mathcal{A}_2 when $Y_{t, j}^{(i)}$ has not yet been queried on T . Then \exists hint q' such that \mathcal{B}' on input $\mathsf{advice} = (\mathsf{state}, \mathsf{Q})$ where $(q', i, t, j) \in \mathsf{Q}$ and \mathcal{B}' outputs $(Y_{t, j}^{(i)}, \mathsf{f}(\mathsf{pk}, Y_{t, j+1}^{(i)}))$ and in **Simulate** phase never queries T on $Y_{t, j}^{(i)}$ or T^{-1} on $\mathsf{f}(\mathsf{pk}, Y_{t, j+1}^{(i)})$ and only queries T on $Y_{t, j'}^{(i)}$ for $j' < j$ (in addition to \mathcal{A}_2 's queries).

Proof. Let q' be the smallest index in $\mathbf{w}^{\mathsf{T}, \mathsf{T}^{-1}}$ which represent a query of $Y_{t, j+1}^{(i)}$ to T or $\mathsf{f}(\mathsf{pk}, Y_{t, j}^{(i)})$ to T^{-1} by \mathcal{A}_2 . \mathcal{B}' from our construction can from now on answer $\mathsf{T}(Y_{t, j}^{(i)})$ and $\mathsf{T}^{-1}(\mathsf{f}(\mathsf{pk}, Y_{t, j}^{(i)}))$ without querying the true oracle in **Simulate** phase. Since by assumption, $Y_{t, j}^{(i)}$ has not yet been queried on T by \mathcal{A}_2 yet, and q' was minimum, so T^{-1} was not queried on $\mathsf{f}(\mathsf{pk}, Y_{t, j}^{(i)})$ before either. \square

Claim. If \mathcal{A}_2 wins nonsequentially with probability ϵ ,

$$\Pr \left[\begin{array}{l} \exists \mathsf{advice} \in \{0, 1\}^* \text{ s.t. } |\mathsf{advice}| \leq n' \cdot \lambda - \omega(\log \lambda), \\ \{(x_i, y_i)\}_{i=1}^{n'} \leftarrow \mathcal{B}'^{\mathsf{T}(\cdot), \mathsf{T}^{-1}(\cdot)}(\mathsf{advice}) \text{ where} \\ \forall i \neq j \in [n'], x_i \neq x_j, \mathsf{T}(x_i) = y_i, x_i \notin \mathsf{s}_{\mathcal{B}'}^{\mathsf{T}} \text{ and } y_i \notin \mathsf{s}_{\mathcal{B}'}^{\mathsf{T}^{-1}} \end{array} \right] \geq \epsilon - \mathsf{negl}(\kappa),$$

Proof. We can take advice to be the state produced by $\mathcal{A}_1^{\mathsf{H}(\cdot), \mathsf{T}, \mathsf{T}^{-1}}(1^\kappa, \mathsf{pk})$ and $(q, i, t, j) \in \mathsf{Q}$ to be a j such that $Y_{t, j+1}^{(i)}$ was queried before $Y_{t, j}^{(i)}$ for every $i \in [n], t \in [b]$.

By [Claim 5.2](#), such a q exists for each (i, t, j) , and \mathcal{B}' has outputted $n \cdot b$ pairs $(Y_{t, j}^{(i)}, \mathsf{f}(\mathsf{pk}, Y_{t, j+1}^{(i)}))$ without querying on $\mathsf{T}, \mathsf{T}^{-1}$ in the **Simulate** phase. By [Claim 5.2](#) and [Claim 5.2](#) that these pairs (x_i, y_i) are distinct, and that \mathcal{B} will not have queried any x_i or y_i in the **Setup** phase with all but negligible probability.

Since (q, i, t, j) are all $\leq \mathsf{poly}(\kappa)$ by the running time of \mathcal{A}_2 , Q only needs $b \cdot n \cdot O(\log \kappa)$ bits. As $s \in 1 - \frac{\omega(\log \kappa)}{\lambda}$, we get,

$$\begin{aligned} |\mathsf{state}| &\leq b \cdot n \cdot \lambda - b \cdot n \cdot \omega(\log(\kappa)) \Rightarrow |\mathsf{advice}| \leq b \cdot n \cdot \lambda - b \cdot n \cdot \omega(\log(\kappa)) + b \cdot n \cdot O(\log \kappa) \\ &\leq b \cdot n \cdot \lambda - \omega(\log(\kappa)). \end{aligned}$$

This proves the claim for $n' = b \cdot n$ which is polynomial in κ and hence polynomial in λ . \square

By [Lemma 1](#), $\epsilon - \mathsf{negl}(\kappa) \in \mathsf{negl}(\lambda) = \mathsf{negl}(\kappa) \Rightarrow \epsilon \leq \mathsf{negl}(\kappa)$. \square

Game 4

Claim. $\Pr[\mathbf{F}_4] \geq \frac{\Pr[\mathbf{F}_3]}{bn}$.

Proof. Game 4 differs from Game 3 in the winning condition. Let $\Pr[\mathbf{F}_3] = \epsilon$ be the probability of adversary winning in winning Game 3. Let this adversary be \mathcal{A}_2 . From the sequentiality condition we have that \mathcal{A}_2 is sequential on at least one block and replica. The probability that this guess was made correctly in Game 4 is $\geq \frac{1}{bn}$. This adversary thus wins Game 4 with probability $\geq \frac{\epsilon}{bn}$. \square

Game 5

Lemma 7. $\Pr[\mathbf{F}_5] \geq \Pr[\mathbf{F}_4] - \text{negl}(\kappa)$.

Proof. Game 5 differs from Game 4 in how queries are answered to the adversaries and the possibility of flag being set. Let $\Pr[\mathbf{F}_4] = \epsilon$ be the probability of adversary winning in winning Game 4. Let \mathbf{E}^\top be the event that it is set due to **Prequery Check T**, and let \mathbf{E}^\times be the probability that flag is due to **Collision Check**. We can bound the probability that the adversary wins Game 5 by the sum of (i) the probability that \mathbf{E}^\times occurs, (ii) the probability \mathbf{E}^\top occurs, and (iii) the statistical difference of the output of \mathcal{C} from using $\mathbf{T}, \mathbf{T}_\times$.

Claim. $\Pr[\mathbf{E}^\times] = \text{negl}(\kappa)$.

We note that since \mathcal{Z}_1 are uniform and independently random, we can bound the probability that $z_a, z_b \in \mathcal{Z}_1$ collide for any fixed $a \neq b$ is $\frac{1}{2^\lambda}$, so we can union bound the probability that

$$\Pr[\exists a \neq b : z_a = z_b] \leq \sum_{a=0}^{2r+1} \sum_{b=a+1}^{2r+1} \Pr[z_a = z_b] = \binom{2r+1}{2} \frac{1}{2^\lambda} = \text{negl}(\kappa)$$

Claim. $\Pr[\mathbf{E}^\top] = \text{negl}(\kappa)$

Proof. We note that since all $4r+1$ elements of $\mathcal{Z}_1 \cup \mathcal{Z}_2$ are uniform and independent of \mathbf{T} , we can bound the probability that some $z \in \mathcal{Z}_1 \cup \mathcal{Z}_2$ is equal to some $z' \in (\mathbf{u}^\top \cup \mathbf{u}^{\top^{-1}} \cup \mathbf{U}^\top \cup \mathbf{U}^{\top^{-1}})$ with $\frac{1}{2^\lambda}$. Thus, we can union bound

$$\Pr[(\mathcal{Z}_1 \cup \mathcal{Z}_2) \cap (\mathbf{u}^\top \cup \mathbf{u}^{\top^{-1}} \cup \mathbf{U}^\top \cup \mathbf{U}^{\top^{-1}}) \neq \emptyset] \leq (4r+1) \cdot (2q_1^\top + 2q_1^{\top^{-1}}) \cdot 2^{-\lambda}$$

Since $q_1^\top, q_1^{\top^{-1}}, r$ are all $\text{poly}(\kappa)$, this is $\text{negl}(\kappa)$. \square

For (iii), applying [Lemma 3](#) with $\pi = \mathbf{T}$, $\{r_0, \dots, r_k\} = \{x_0, \dots, x_r\}$, and $\tau = f(\text{pk}, \cdot)$ tells us that $(x_r, \mathbf{T}[\text{swap}(x_0, f(\text{pk}, x_1)), \dots, \text{swap}(x_{r-1}, f(\text{pk}, x_r))])$ is uniformly random. Since the only place \mathcal{C} 's responses differ are answers to \mathbf{T}_\times and returning the encoding $Y_{\beta, r}^{(\gamma)} = x_r$. We recall that $(Y_{\beta, r}^{(\gamma)}, \mathbf{T})$ is uniform in Game 4 by construction, so these distributions are identical.

Since (i), (ii), (iii) are all negligible, the adversary thus wins Game 5 with probability $\epsilon - \text{negl}(\kappa)$. \square

Game 6

Let $\Pr[\mathbf{F}_5] = \epsilon$ be the probability of adversary winning in winning Game 5.

Lemma 8. $\Pr[\mathbf{F}_6] \geq \left(\frac{\epsilon}{2}\right) \left(\frac{\epsilon 2^{r-1} - 2^{|\text{state}|} + 1}{\epsilon 2^{r-1}}\right)$.

Here we apply the bound on state size to argue that x' which ensure that \mathcal{A} wins must be fairly frequent.

Proof. Note that by construction, \mathcal{A} will win Game 6 as long as an appropriate x' exists. We will lower bound said probability below. Let us consider the random coins used in Game 5. The randomness for \mathcal{C} is over the choice of permutation picked \mathbb{T} and $\gamma, \beta, \mathcal{Z}, x, \delta, \rho'$ where ρ' denotes the randomness over ρ_i sampled for each replica and the coins used by the $r\text{Setup}$. Let us denote $\eta = (\gamma, \beta, \mathbb{T}, \mathcal{Z}, \rho')$ for ease of notation. Let $\mathcal{A}_1, \mathcal{A}_2$ denote the adversaries that solves Game 5 with ϵ probability.

Define the set $\mathcal{N} = \{(\eta, x) | F_5(\eta, x)\}$ where $F_5(\eta, x)$ denotes the event that adversaries win game 5 with given parameters. Thus $\Pr_{\eta, x}[(\eta, x) \in \mathcal{N}] = \epsilon$.

Let us define a heavy set, $\mathcal{H} = \{(\eta, x) | \Pr_{\eta, x'}[(\eta, x') \in \mathcal{N}] \geq \frac{\epsilon}{2}\}$. Then from the heavy row lemma of [24], $\Pr_{\eta, x}[(\eta, x) \in \mathcal{H} | (\eta, x) \in \mathcal{N}] \geq 1/2$ i.e. probability that our winning instance lies on a heavy set is atleast half.

Note that for \mathcal{A}_2 to successfully solve Game 6, we must see the following events,

1. Adversary solves the Game 5 i.e. $(\eta, x) \in \mathcal{N}$, and let this be denoted by event E_1 .

$$\Pr_{\eta, x}[E_1] \geq \epsilon.$$

2. Let the event that $(\eta, x) \in \mathcal{H}$ lies on a heavy set be denoted by E_2 . By heavy row lemma, we have that,

$$\begin{aligned} \Pr_{\eta, x}[E_2 | E_1] &\geq 1/2 \\ \Rightarrow \Pr_{\eta, x}[E_2] &\geq \frac{\epsilon}{2}. \end{aligned}$$

3. Let the event that there exists a collision of x consistent with state that is also successful on Game 5 be denoted by E_3 .

Claim. $\Pr_{\eta, x}[E_3 | E_2] \geq \left(\frac{\epsilon 2^{r-1} - 2^{|\text{state}|} + 1}{\epsilon 2^{r-1}} \right)$.

Proof. Given that $(\eta, x) \in \mathcal{H}$ i.e. it lies on a heavy set. Define a possible set of switches by $\mathcal{Q} = \{x' | (\eta, x') \in \mathcal{N}\}$. From the definition of \mathcal{H} , clearly $|\mathcal{Q}| \geq \frac{\epsilon 2^r}{2}$.

From Lemma 4 (Section 4), define h_η as a function from $\mathcal{Q} \rightarrow \{0, 1\}^{|\text{state}|}$ where $h_\eta(a) = \text{state}_a$ where state_a denotes the state shared by \mathcal{A}_1 to \mathcal{A}_2 when playing Game 5 with parameters η and x set to a . Since \mathcal{A}_1 's coins are deterministically fixed, h_η is a deterministic function.

The lemma implies the statement that,

$$\Pr_x[\exists x' \neq x \in \mathcal{Q}, h_\eta(x) = h_\eta(x')] \geq \frac{|\mathcal{D}| - |\mathcal{R}| + 1}{|\mathcal{D}|}.$$

This proves the claim. □

$$\begin{aligned} \Pr_{\eta, x}(\mathcal{A}_2 \text{ wins Game 6}) &\geq \Pr[E_2 \cap E_3] \\ &\geq (\Pr[E_2]) (\Pr[E_3 | E_2]) \\ &\geq \left(\frac{\epsilon}{2}\right) \left(\frac{\epsilon 2^{r-1} - 2^{|\text{state}|} + 1}{\epsilon 2^{r-1}}\right) \end{aligned}$$

□

We note that if F_5 is non negligible then so is F_6 - i.e. $\epsilon \geq \frac{1}{\text{poly}(\kappa)}$, then

$$\epsilon 2^{r-1} \geq 2^{r-O(\log \kappa)} \in \omega(2^{|\text{state}|} + 1) \Rightarrow \frac{\epsilon 2^{r-1} - 2^{|\text{state}|} + 1}{\epsilon 2^{r-1}} \in (1 - \text{negl}(\kappa)).$$

Game 7

Lemma 9. $\Pr[F_7] \geq \Pr[F_6] - \text{negl}(\kappa)$.

Proof. Game 7 is played on a different permutation than Game 6, with a different winning condition. By our definition of $\text{swap}(\cdot, \cdot)$, we can observe that these permutations differ when querying $T_{x'_j}^\delta$ on x'_j or $T^{-1}(f(\text{pk}, x'_j))$ for $j > \delta$ and when querying $T_{x'_j}^\delta$ on $f(\text{pk}, x'_j)$ or $T(f(\text{pk}, x'_j))$ again for $j > \delta$.

Claim. Let $a, b, c, d \xleftarrow{R} \mathcal{D}$ and $T \xleftarrow{R} S_{\mathcal{D}}$. Then the distributions of $(a, b, T[\text{swap}(a, b)])$ and $(a, b, T[\text{swap}(a, c), \text{swap}(d, b), \text{swap}(a, b)])$ are identical.

Proof. By Lemma 2, $(a, T[\text{swap}(a, c)])$ is identically distributed to (a, T) , so since b is independent of a , we can apply Lemma 2 again to say $(a, b, T[\text{swap}(a, c), \text{swap}(b, d)])$ is identically distributed to (a, b, T) , after the result follows from applying $\text{swap}(a, b)$ to both distributions. \square

By Claim 5.2, we know the distribution of

$$(\{x'_j\}, T[\dots, \text{swap}(x'_j, f(\text{pk}, x'_{j+1})), \dots])$$

is identical to that of

$$(\{x'_j\}, T[\dots, \text{swap}(x'_j, c), \text{swap}(d, f(\text{pk}, x'_{j+1})), \text{swap}(x'_j, f(\text{pk}, x'_{j+1})), \dots])$$

for any $j \in [r]$, but observe in the latter case, d and c are the preimage and image of $f(\text{pk}, x'_{j+1})$ and x'_j respectively before $\text{swap}(x'_j, f(\text{pk}, x'_{j+1}))$, and are independently random, which tells us that $T^{-1}(f(\text{pk}, x'_j))$ and $T(f(\text{pk}, x'_j))$ are random independent of $\{x'_j, T_{x'_j}\}$, so the probability that some $z \in w^T \cup w^{T^{-1}}$ is equal a one of those is $2^{-\lambda}$, which allows us to union bound the total probability

$$\Pr[(T^{-1}(\{f(\text{pk}, x_j)\}) \cup \{T(f(\text{pk}, x'_j)\})] \cap (w^T \cup w^{T^{-1}}) \neq \emptyset] \leq 2r \cdot (q_3^T + q_3^{T^{-1}}) \cdot 2^{-\lambda}$$

Since $q_3^T, q_3^{T^{-1}}, r$ are all $\text{poly}(\kappa)$, this is $\text{negl}(\kappa)$. In addition \mathcal{A} wins Game 6, by sequentiality, \mathcal{A}_2 must have queried on x'_δ before querying on any x'_j or $f(\text{pk}, x'_j)$ for $j > \delta$, so Game 6 and Game 7 have identical queries with all but negligible probability before x'_δ is queried on, at which point \mathcal{A} wins Game 7. \square

Game 8

Claim. $\Pr[F_8] \geq \frac{\Pr[F_7]}{r}$.

Proof. Our main observation here is that \mathcal{C} does not need to guess the switch completely, \mathcal{C} guesses δ . The probability that this index was correctly guessed in Game 8 is exactly equal to $\frac{1}{r}$. This adversary thus queries x'_δ and hence wins Game 8 with probability $\geq \frac{\Pr[F_7]}{r}$. \square

Claim. $\Pr[F_8] = \text{negl}(\kappa)$

Proof. We will show through a reduction that an adversary able to win Game 8 with probability ϵ will also be able to invert our trapdoor permutation with the same probability.

Let $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\kappa), y \xleftarrow{R} \mathcal{D}_{\text{pk}} = \{0, 1\}^\lambda, z = f(\text{pk}, y)$,

Reduction $\mathcal{B}(\text{pk}, z)$

Goal: Output z such that, $f(\text{pk}, y) = z$.

- **Setup:** The challenger (denoted by \mathcal{C}) runs $(\text{pk}', \text{sk}') \leftarrow \text{rSetup}(1^\kappa, 1^n)$ and sends public key pk' to \mathcal{A}_1 . ~~It keeps the secret key sk' for itself.~~
- **Phase 1, Sampling a New Permutation, File Challenge, Phase 2, State Sharing:** . . .

- **Setting switched oracle:**
 - \dots
 - Let $\delta \xleftarrow{R} [r]$
 - Define T_x using swap ([Definition 3](#)):

$$T_x^\delta = T[\text{swap}(x_0, f(\text{pk}, x_1)), \dots, \text{swap}(x_{\delta-1}, z)]$$

- Let $T_{x'}^{-1}$ be the inverse of $T_{x'}$.
- **Phase 3, Guess:** \dots
- ~~**Verify:** Adversary wins if T is queried on \bar{x}_δ and $\text{flag} = 0$.~~
- **Embed:** Output embedding w_i^T , where $w_i^T \in \mathbf{w}^T$, s.t. $f(\text{pk}, w_i^T) = z$.

Since \bar{x}_δ is uniformly distributed and not used before here, $f(\text{pk}, \bar{x}_\delta)$ is uniform and independent, and indistinguishable from a uniformly random z . Note that since we no longer check for flag , it is no longer necessary to perform collision and prequery checks, the only place where x'_δ is used instead of $f(\text{pk}, x'_\delta)$. We can simulate the random permutation oracles and regular random oracle by picking uniformly random values on the appropriate domain whenever \mathcal{A} makes a new query and maintaining a map of queries already made. Thus, if \mathcal{A}_2 won game 8, they must have queried

$$\bar{x}_\delta = f^{-1}(\text{sk}, f(\text{pk}, \bar{x}_\delta)) = f^{-1}(\text{sk}, z)$$

Which, from our definition of a secure trapdoor permutation, can only happen with negligible probability, \square

To complete our proof of [Theorem 1](#), from our sequence of games, we conclude that $\Pr[F_0] = \text{negl}(\kappa)$, fulfilling our soundness definition.

6 Replica Encodings in the Random Function Model

We now turn toward building Replica Encodings from trapdoor permutations in the ideal function model. Our construction will embed a Feistel like structure into the replica encoding construction. We will directly prove security of this construction. Our construction makes use of the KeyGen, f , and f^{-1} defined for a trapdoor permutation on domain $\{0, 1\}^\lambda$ and a random function T' on the same domain. Let H be a random oracle on the range $\{0, 1\}^\lambda$.

Define functions $L, R : \{0, 1\}^* \rightarrow \{0, 1\}^*$ on even length inputs as follows. If $x = y||z$, where $x, y, z \in \{0, 1\}^*$, $|y| = |z|$, then the function $L(\cdot)$ denotes the left half of x i.e. $L(x) = y$ and the function $R(\cdot)$ denotes the right half of x i.e. $R(x) = z$.

6.1 Construction

$r\text{Setup}(1^\kappa, 1^n)$:

Run $\text{KeyGen}(1^\kappa) \rightarrow (\text{pk}, \text{sk})$. Output $(\text{pk}' = (\text{pk}, n), \text{sk}' = (\text{sk}, n))$.

$r\text{Enc}(\text{sk}', m)$:

- Parse $\text{sk}' = (\text{sk}, n)$.
- Choose a string $\rho \xleftarrow{R} \{0, 1\}^\kappa$.
- Divide m into b blocks of length 2λ i.e. $m = m_1||m_2||\dots||m_b$, $b = \lceil |m|/2\lambda \rceil$.
- Set $r = n \cdot b \cdot \lambda$.
- Compute $\forall t \in [b]$,

$$Y_{t,0} = L(m_t \oplus H(\rho||t)).$$

$$Y_{t,1} = R(m_t \oplus H(\rho||t)).$$

- For rounds j from 2 to r compute:
 - Compute $Y_{t,j}$ from $Y_{t,j-1}$ and $Y_{t,j-2}$ as

$$Y_{t,j} = f^{-1}(\text{sk}, Y_{t,j-2} \oplus T'(Y_{t,j-1}))$$

- Let $Z_t = Y_{t,r-1} || Y_{t,r}$
- Let $y_r = Z_1 || \dots || Z_b$ and output (y_r, ρ) .

rDec(pk', y):

- Parse $\text{pk}' = (\text{pk}, n)$.
- Parse y as (y_r, ρ) . Parse y_r as $Z_1 || Z_2 || \dots || Z_b$, where $b = \lceil |y_r| / 2\lambda \rceil$ and $r = n \cdot b \cdot \lambda$.
- For each $Z_t = Y_{t,r-1} || Y_{t,r}$ and for rounds j from $r-2$ to 0 compute:
 - Compute $Y_{t,j}$ from $Y_{t,j+1}$ and $Y_{t,j+2}$ as

$$Y_{t,j} = f(\text{pk}, Y_{t,j+2}) \oplus \Gamma'(Y_{t,j+1})$$

- $\forall t \in [b]$ compute,

$$m_t = Y_{t,0} || Y_{t,1} \oplus H(\rho || t)$$

Output $m = m_1 || \dots || m_b$.

The encoding length for our scheme is $\text{len}(\kappa, |m|) = |m| + O(\kappa)$.⁶

6.2 Proof of Security

Theorem 2. *Assuming $(\text{KeyGen}(1^\kappa), f(\cdot, \cdot), f^{-1}(\cdot, \cdot))$ is a secure trapdoor permutation on domain and range $\{0, 1\}^\lambda$ and Γ' is an oracle to a random function on the same domain and range, and H is a random oracle with range $\{0, 1\}^{2\lambda}$. Then our construction for ReplicaEncoding described above is s -sound according to Definition 1 for all $\kappa, n \in \mathbb{N}$ and $s \in 1 - \frac{\omega(\log \kappa)}{2\lambda}$.*

Due to space constraints, we defer the sequence of games and the proof of this theorem to the full version of our paper, [18].

7 Counterexample for Round Function Independent of Blocks

We gave intuition in Section 1.2 using VBB obfuscation that our construction is round optimal up to constant factors i.e. is insecure for any number of rounds $\in o(b \cdot n)$. Below we formalize the notion by giving a construction from $i\mathcal{O}$ that captures this intuition formally and constructs a scheme in the end that breaks soundness security.

We assume the existence of a trapdoor permutation $(\text{KeyGen}, f(\cdot, \cdot), f^{-1}(\cdot, \cdot))$ with domain $\{0, 1\}^\lambda$ for $\lambda \in \omega(\kappa)$ ⁷, a puncturable PRF family $(\text{PPRF.KeyGen}, \text{PPRF.Eval}, \text{PPRF.Puncture})$, indistinguishability obfuscation $i\mathcal{O}$ for all polynomial sized circuits.

7.1 Construction

Let $(\text{KeyGen}, f(\cdot, \cdot), f^{-1}(\cdot, \cdot))$ be a trapdoor permutation on $\{0, 1\}^\kappa$, where KeyGen uses some $r(\kappa)$ bits of randomness. Let $(\text{PPRF.KeyGen}, \text{PPRF.Eval}, \text{PPRF.Puncture})$ be a puncturable PRF on domain $\{0, 1\}^\kappa$ and range $\{0, 1\}^{r(\kappa)}$. We will construct a trapdoor permutation $(\text{keygen}'_n, f'_n(\cdot, \cdot), f'^{-1}_n(\cdot, \cdot))$ on domain $\{0, 1\}^{2\kappa}$ parameterized by a quantity $n \in \text{poly}(\kappa)$.

$\text{keygen}'_n(1^\kappa)$

1. Sample $K \leftarrow \text{PPRF.KeyGen}(1^\kappa)$
2. Let $\mathcal{O}\text{Program } f = i\mathcal{O}(\kappa, \text{Program } f)$ and $\mathcal{O}\text{Program } f^{-1} = i\mathcal{O}(\kappa, \text{Program } f^{-1})$.
3. Output $(\text{pk} = (\mathcal{O}\text{Program } f, \mathcal{O}\text{Program } f^{-1}), \text{sk} = K)$

$f'_n(\text{pk} = (\mathcal{O}\text{Program } f, \mathcal{O}\text{Program } f^{-1}), (z, x))$

1. Let $y \leftarrow \mathcal{O}\text{Program } f(z, x)$ and output (z, y) .

$f'^{-1}_n(\text{sk} = K, (z, y))$

1. Let $r \leftarrow \text{PPRF.Eval}(K, z)$.
2. Let $(\text{pk}_0, \text{sk}_0) \leftarrow \text{KeyGen}(1^\kappa; r)$.
3. Output $(z, f^{-1}(\text{sk}_0, y))$.

⁶ Upto additional rounding factors.

⁷ note it suffices to have some trapdoor permutation with domain $\lambda \in \omega(\kappa^\epsilon)$ for $\epsilon > 0$, and we can generically transform this by taking said TDP on security parameter $\kappa' = \kappa^{1/\epsilon}$

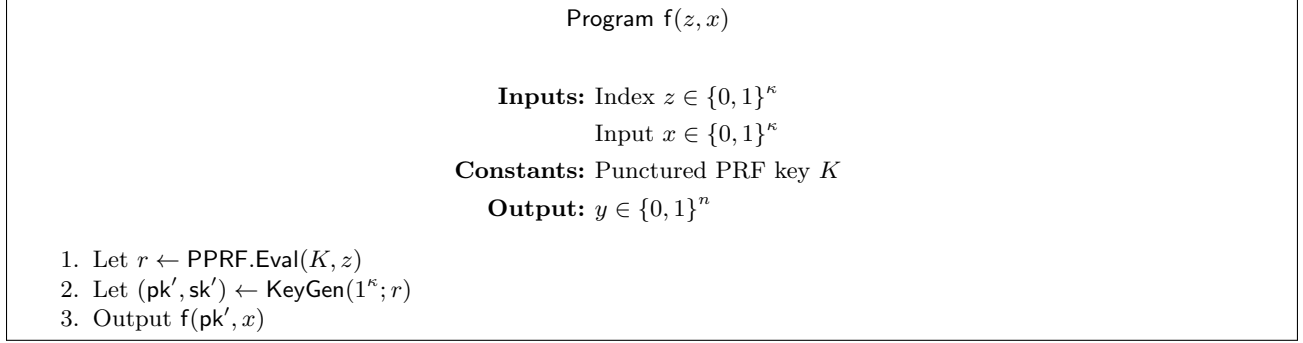


Fig. 3: Routine Program f

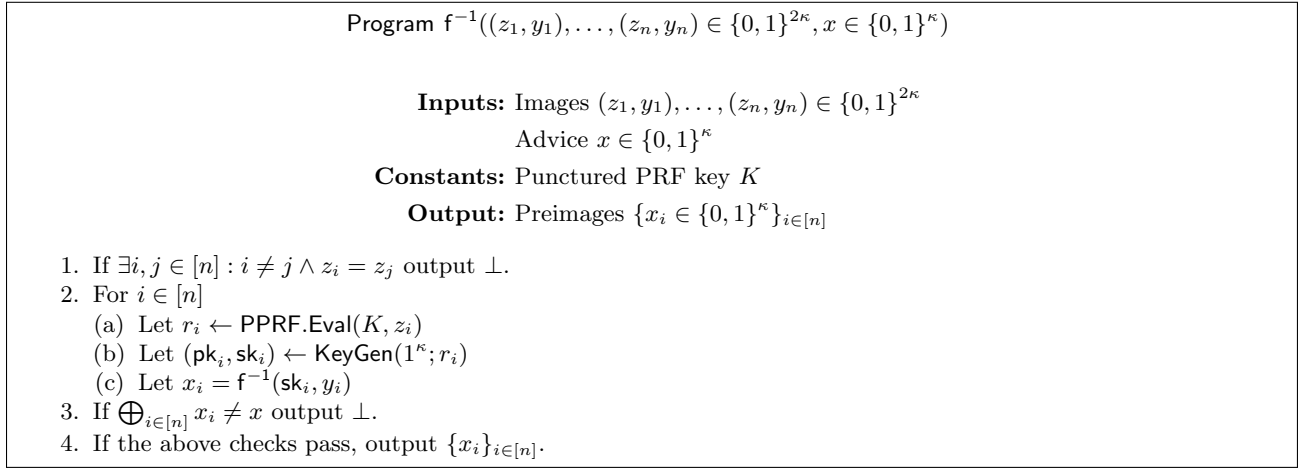


Fig. 4: Routine Program f^{-1}

7.2 Proofs

Efficiency

Claim. $(\text{keygen}'_n, f'_n(\cdot, \cdot), f_n^{-1}(\cdot, \cdot))$ are polynomial time algorithms

Proof. keygen'_n simply calls $i\mathcal{O}$ twice. The programs Program f and Program f^{-1} simply call the underlying PRF and trapdoor primitives at most $n \in \text{poly}(\kappa)$ times. By the efficiency of the underlying PRF and trapdoor permutation, Program f and Program f^{-1} are poly-sized circuits and $i\mathcal{O}$ runs in poly time, thus keygen'_n runs in polynomial time.

f'_n simply evaluates a polynomial sized circuit, which is polynomial time.

f_n^{-1} does a single call to PPRF.Eval , KeyGen , $f^{-1}(\cdot, \cdot)$, which are all polynomial time algorithms by definition. \square

Correctness

Claim. The correctness of $i\mathcal{O}$ and correctness of $f^{-1}(\text{sk}', \cdot)$ computing inverse of $f(\text{pk}', \cdot)$ implies $f_n^{-1}(\text{sk}, \cdot)$ computes the inverse of $f'_n(\text{pk}, \cdot)$, i.e.

$$\forall \kappa, (\text{pk}, \text{sk}) \leftarrow \text{keygen}'_n(1^\kappa), \forall x' \in \{0, 1\}^{2\kappa}, f_n^{-1}(\text{sk}, f'_n(\text{pk}, x')) = x'.$$

Proof. Let $x' = (z, x) \in \{0, 1\}^{2\kappa}$ be an arbitrary input to \cdot . Recall that f'_n simply runs $\mathcal{O}\text{Program } f$ on (z, x) and is same as the result of outputting Program f on (z, x) from correctness of $i\mathcal{O}$. The output produced is $(z, f(\text{pk}', x))$ where $(\text{pk}', \text{sk}') = \text{KeyGen}(1^\kappa, F(K, z))$. f_n^{-1} when run on $(z, f(\text{pk}', x))$ produces the same (pk', sk') pair as $\mathcal{O}\text{Program } f$. Since,

$$\forall \kappa, (\text{pk}', \text{sk}') \leftarrow \text{KeyGen}(1^\kappa), \forall x \in \{0, 1\}^\kappa, f^{-1}(\text{sk}', f(\text{pk}', x)) = x.$$

$f_n'^{-1}$ returns $(z, f^{-1}(\text{sk}', f(\text{pk}', x))) = (z, x)$. □

Security

Theorem 3. *Assuming $(\text{KeyGen}, f(\cdot, \cdot), f^{-1}(\cdot, \cdot))$ is a secure one way permutation, indistinguishability of $i\mathcal{O}$ and a puncturable PRF family $(\text{PPRF.KeyGen}, \text{PPRF.Eval}, \text{PPRF.Puncture})$ secure, $(\text{keygen}'_n, f'_n, f_n'^{-1})$ is a secure one way permutation - i.e., that for all PPT algorithms \mathcal{A}*

$$\Pr \left[\begin{array}{c} f'_n(\text{pk}, (z_0, x_0)) = (z, y) \text{ s.t.} \\ (\text{pk}, \text{sk}) \leftarrow \text{keygen}'_n(1^\kappa), (z_0, x_0) \xleftarrow{R} \{0, 1\}^{2\kappa}, (z, y) = f'_n(\text{pk}, (z_0, x_0)), (z', x') \leftarrow \mathcal{A}(\text{pk}, (z, y)) \end{array} \right] \leq \text{negl}(\kappa),$$

over the random coins of keygen'_n and sampling of (z_0, x_0) .

We will show this via a sequence of games, where the view of the adversary between successive games is indistinguishable. Due to space constraints, the formal proof is deferred to the full version [18]. The proof follows the punctured programming technique of [28].

Game 0

This is the original security game.

1. Challenger samples a random $(z_0, x_0) \xleftarrow{R} \{0, 1\}^{2\kappa}$
 - (a) Sample $K \leftarrow \text{PPRF.KeyGen}(1^\kappa)$
 - (b) Let $\mathcal{O}\text{Program } f = i\mathcal{O}(\kappa, \text{Program } f)$ and $\mathcal{O}\text{Program } f^{-1} = i\mathcal{O}(\kappa, \text{Program } f^{-1})$.⁸
 - (c) Output $(\text{pk} = (\mathcal{O}\text{Program } f, \mathcal{O}\text{Program } f^{-1}), \text{sk} = K)$
2. Challenger runs $f'_n(\text{pk}, (z_0, x_0))$
 - (a) Let $y_0 \leftarrow \mathcal{O}\text{Program } f(z_0, x_0)$.
3. Adversary is given $(\text{pk}, (z_0, y_0))$.
4. Adversary outputs (z', x')
5. If $f'_n(\text{pk}, (z', x')) = (z_0, y_0)$ then output 1 else output 0

Game 1

In this game, we change the way $\text{Program } f$ is programmed.

Program $f^*(z, x)$

Inputs: Index $z \in \{0, 1\}^\kappa$
Input $x \in \{0, 1\}^\kappa$

Constants: Punctured PRF key $K(\{z_0\})$
Public Key pk_0

Output: $y \in \{0, 1\}^n$

1. If $z = z_0$, output $f(\text{pk}_0, x)$.
2. Let $r \leftarrow \text{PPRF.Eval}(K(\{z_0\}), z)$
3. Let $(\text{pk}', \text{sk}') \leftarrow \text{KeyGen}(1^\kappa; r)$
4. Output $f(\text{pk}', x)$

Fig. 5: Routine $\text{Program } f^*$

1. Challenger samples a random $(z_0, x_0) \xleftarrow{R} \{0, 1\}^{2\kappa}$
 - (a) Sample $K \leftarrow \text{PPRF.KeyGen}(1^\kappa)$
 - (b) Compute $(\text{pk}_0, \text{sk}_0) \leftarrow \text{KeyGen}(1^\kappa, F(K, z_0))$

⁸ The security parameter in the input to $i\mathcal{O}$ algorithm is the smallest λ for which $\text{Program } f, \text{Program } f^*$ are in \mathcal{C}_λ which will be polynomial in κ as the circuits are polynomial. We denote this by κ here for notation clarity.

- (c) Compute punctured key $K(\{z_0\})$
 - (d) Let $\mathcal{O}\text{Program } f^* = i\mathcal{O}(\kappa, \text{Program } f^*)$ and $\mathcal{O}\text{Program } f^{-1} = i\mathcal{O}(\kappa, \text{Program } f^{-1})$.
 - (e) Output $(\text{pk} = (\mathcal{O}\text{Program } f^*, \mathcal{O}\text{Program } f^{-1}), \text{sk} = K)$
2. Challenger runs $f'_n(\text{pk}, (z_0, x_0))$
 - (a) Let $y_0 \leftarrow \mathcal{O}\text{Program } f^*(z_0, x_0)$.
 3. Adversary is given $(\text{pk}, (z_0, y_0))$.
 4. Adversary outputs (z', x')
 5. If $f'_n(\text{pk}, (z', x')) = (z_0, y_0)$ then output 1 else output 0

Game 2

In this game, we change the way $\text{Program } f^{-1}$ is programmed.

Program $f^{-1*}((z_1, y_1), \dots, (z_n, y_n) \in \{0, 1\}^{2\kappa}, x \in \{0, 1\}^\kappa)$

Inputs: Images $(z_1, y_1), \dots, (z_n, y_n) \in \{0, 1\}^{2\kappa}$
 Advice $x \in \{0, 1\}^\kappa$

Constants: Punctured PRF key $K(\{z_0\})$
 Public Key pk_0

Output: Preimages $\{x_i \in \{0, 1\}^\kappa\}_{i \in [n]}$

1. If $\exists i, j \in [n] : i \neq j \wedge z_i = z_j$ output \perp .
2. For $i \in [n]$ when $z_i \neq z_0$
 - (a) Let $r_i \leftarrow \text{PPRF.Eval}(K(\{z_0\}), z_i)$
 - (b) Let $(\text{pk}_i, \text{sk}_i) \leftarrow \text{KeyGen}(1^\kappa; r_i)$
 - (c) Let $x_i = f^{-1}(\text{sk}_i, y_i)$
3. If $\exists i' : z_{i'} = z_0$
 - (a) Let $x_{i'} = x \oplus (\bigoplus_{i \in [n] \setminus \{i'\}} x_i)$.
 - (b) If $f(\text{pk}_0, x_{i'}) \neq y_{i'}$, output \perp .
4. If $\bigoplus_{i \in [n]} x_i \neq x$ output \perp .
5. If the above checks pass, output $\{x_i\}_{i \in [n]}$.

Fig. 6: Routine $\text{Program } f^{-1*}$

1. Challenger samples a random $(z_0, x_0) \xleftarrow{R} \{0, 1\}^{2\kappa}$
 - (a) Sample $K \leftarrow \text{PPRF.KeyGen}(1^\kappa)$
 - (b) Compute $(\text{pk}_0, \text{sk}_0) \leftarrow \text{KeyGen}(1^\kappa, F(K, z_0))$
 - (c) Compute punctured key $K(\{z_0\})$
 - (d) Let $\mathcal{O}\text{Program } f^* = i\mathcal{O}(\kappa, \text{Program } f^*)$ and $\mathcal{O}\text{Program } f^{-1*} = i\mathcal{O}(\kappa, \text{Program } f^{-1*})$.
 - (e) Output $(\text{pk} = (\mathcal{O}\text{Program } f^*, \mathcal{O}\text{Program } f^{-1*}), \text{sk} = K)$
2. Challenger runs $f'_n(\text{pk}, (z_0, x_0))$
 - (a) Let $y_0 \leftarrow \mathcal{O}\text{Program } f^*(z_0, x_0)$.
3. Adversary is given $(\text{pk}, (z_0, y_0))$.
4. Adversary outputs (z', x')
5. If $f'_n(\text{pk}, (z', x')) = (z_0, y_0)$ then output 1 else output 0

Game 3

In this game, we compute $(\text{pk}_0, \text{sk}_0) \leftarrow \text{KeyGen}(1^\kappa, r_0)$ using true randomness r_0 .

1. Challenger samples a random $(z_0, x_0) \xleftarrow{R} \{0, 1\}^{2\kappa}$
 - (a) Sample $K \leftarrow \text{PPRF.KeyGen}(1^\kappa)$
 - (b) Sample $r_0 \xleftarrow{R} \{0, 1\}^{r(\kappa)}$

- (c) Compute $(pk_0, sk_0) \leftarrow \text{KeyGen}(1^\kappa, r_0)$
- (d) Compute punctured key $K(\{z_0\})$
- (e) Let $\mathcal{O}\text{Program } f^* = i\mathcal{O}(\kappa, \text{Program } f^*)$ and $\mathcal{O}\text{Program } f^{-1*} = i\mathcal{O}(\kappa, \text{Program } f^{-1*})$.
- (f) Output $(pk = (\mathcal{O}\text{Program } f^*, \mathcal{O}\text{Program } f^{-1*}), sk = K)$
2. Challenger runs $f'_n(pk, (z_0, x_0))$
 - (a) Let $y_0 \leftarrow \mathcal{O}\text{Program } f^*(z_0, x_0)$.
3. Adversary is given $(pk, (z_0, y_0))$.
4. Adversary outputs (z', x')
5. If $f'_n(pk, (z', x')) = (z_0, y_0)$ then output 1 else output 0

7.3 Attack on Replica Encoding Scheme

First, we restate the security game in the context of the above TDP. We consider a variation of our construction in Section 5 with $r \in o(b \cdot n)$ instantiated with the above trapdoor permutation and present a concrete attack adversary which breaks the s -soundness of our replica encoding scheme for any constant $s \in (0, 1)$. We remark that this attack also applies to the construction in Section 6.

Game 0: $\text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n)$

- **Setup:** The challenger (denoted by \mathcal{C}) runs $(pk', sk') \leftarrow r\text{Setup}(1^\kappa, 1^n)$ and sends public key $pk' = (C'_0, C'_1, n)$ to \mathcal{A}_1 . It keeps the secret key $sk' = (K, n)$ for itself.
- **Phase 1:** The adversary \mathcal{A}_1 issues queries on $\mathbb{T}, \mathbb{T}^{-1}, \mathbb{H}$, \mathcal{C} responds the query back to \mathcal{A}_1 .
- **File Challenge:** $m \in \{0, 1\}^*$ $\leftarrow \mathcal{A}_1^{\mathbb{H}(\cdot), \mathbb{T}(\cdot), \mathbb{T}^{-1}(\cdot)}(1^\kappa, pk')$. It sends m to \mathcal{C} who parses pk' as (C'_0, C'_1, n) ; sk' as (K, n) and does the following:
 - Divide m into b blocks of length λ i.e. $m = m_1 || m_2 || \dots || m_b$, $b = \lceil |m|/\lambda \rceil$.
 - For $i \in [n]$,
 - * Choose a string $\rho_i \xleftarrow{R} \{0, 1\}^\kappa$.
 - * Compute $\forall t \in [b]$,

$$Y_{t,0}^{(i)} = m_t \oplus \mathbb{H}(\rho_i || t).$$

- * For rounds j from 1 to r and $\forall t \in [b]$,
 - Let $z_{t,j}^{(i)}, x_{t,j}^{(i)} \in \{0, 1\}^{\lambda/2}$
 - Let $z_{t,j}^{(i)} || x_{t,j}^{(i)} = \mathbb{T}(Y_{t,j}^{(i)})$
 - Compute $Y_{t,j}^{(i)}$ from $Y_{t,j}^{(i)}$ as

$$(pk_{t,j}^{(i)}, sk_{t,j}^{(i)}) = \text{KeyGen}(1^\kappa; \text{PPRF.Eval}(K, z_{t,j}^{(i)})).$$

$$Y_{t,j}^{(i)} = z_{t,j}^{(i)} || f^{-1}(sk_{t,j}^{(i)}, x_{t,j}^{(i)})$$

- * Let $y_r^{(i)} = Y_{1,r}^{(i)} || \dots || Y_{b,r}^{(i)}$ and set $y^{(i)} = (y_r^{(i)}, \rho_i)$.

\mathcal{C} returns $y^{(1)}, y^{(2)}, \dots, y^{(n)}$ to \mathcal{A}_1 .

- **Phase 2:** \mathcal{A}_1 issues additional queries on $\mathbb{T}, \mathbb{T}^{-1}, \mathbb{H}$, \mathcal{C} responds the query back to \mathcal{A}_1 .
- **State Sharing:** \mathcal{A}_1 outputs state $\text{state} \leftarrow \mathcal{A}_1^{\mathbb{H}(\cdot), \mathbb{T}(\cdot), \mathbb{T}^{-1}(\cdot)}(1^\kappa, pk', y)$ and sends state to \mathcal{A}_2 .
- **Phase 3:** The adversary \mathcal{A}_2 queries on $\mathbb{T}, \mathbb{T}^{-1}, \mathbb{H}$, \mathcal{C} responds the query back to \mathcal{A}_2 .
- **Guess:** \mathcal{A}_2 outputs the replica guesses to \mathcal{C} .

$$\{\tilde{y}^{(i)}\} \leftarrow \mathcal{A}_2(1^\kappa, pk', m, \text{state}).$$

- **Verify:** Let $v_i = 1$ if $\tilde{y}^{(i)} = y^{(i)}$ and 0 otherwise. Adversary wins if $|\text{state}| < \sum v_i \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$.

Now below, we present our construction of adversaries $\mathcal{A}_1, \mathcal{A}_2$.

$\mathcal{A}_1(1^\kappa, pk' = (pk, n))$

- Choose any message $m \in \{0, 1\}^{b \cdot \lambda}$ where $b \geq 1$.
- Send m to challenger.
- Receive $\{y^{(i)} = \{Y_{t,r}^{(i)}\}_{t \in [b]}, \rho_i\}_{i \in [n]}$.
- For each $j \in [r]$, set $x_j = \bigoplus_{t \in [b], i \in [n]} Y_{t,j}^{(i)}$.

- Send $\{x_j\}, \{\rho_i\}$ as state.
- $\mathcal{A}_2(1^\kappa, \mathbf{pk}' = (C'_0, C'_1, n), m, (\{x_j\}, \rho_i))$
 - Divide m into b blocks of length λ , $m = m_1 || \dots || m_b$
 - For $i \in [r], t \in [b]$
 - Compute $Y_{t,0}^{(i)} = H(\rho_i || t) \oplus m_t$
 - For $j \in [r]$
 - Set $\{Y_{t,j}^{(i)}\}_{i \in [r], t \in [b]} = C'_1(\{T(Y_{t,j-1}^{(i)})\}, x_j)$
 - For $i \in [r]$
 - Let $y_r^{(i)} = Y_{1,r}^{(i)} || \dots || Y_{b,r}^{(i)}$ and output $(y_r^{(i)}, \rho_i)$

Lemma 10. $(\mathcal{A}_1, \mathcal{A}_2)$ use $\lambda \cdot o(b \cdot n) + n \cdot o(\lambda)$ space.

Proof. We observe that the state output is $\{x_j\}, \{\rho_i\}$, which use $r \cdot \lambda$, and $\kappa \cdot n$ space respectively. We use the fact that $r \in o(b \cdot n)$ and $\lambda \in \omega(\kappa)$ to give us our result. \square

Lemma 11. $\forall n \geq 1$, there exists a negligible function negl such that the probability that $\sum_i v_i = n$ in the verification stage of $\text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n)$ for adversaries $(\mathcal{A}_1, \mathcal{A}_2)$ is $1 - \text{negl}(\kappa)$.

Proof. We recall that C'_1 is simply an obfuscation of program $\text{Program } f^{-1}$. Thus, as long as the collection $\{z_{t,j}^{(i)}\}_{t \in [b], i \in [n]}$ is unique for every $j \in [r]$ and x_j is the \oplus of $\{x_{t,j}^{(i)}\}_{t \in [b], i \in [n]}$, then C'_1 will successfully invert. By the fact that H is a random oracle, and that T and f are permutations, we use the fact that a uniform random variable under a permutation is uniformly random to get that $\{z_{t,j}^{(i)}\}_{t \in [b], i \in [n]}$ is a uniform and independently random set. Thus, we can union bound the probability that any of them collide with $\binom{b \cdot n}{2} \cdot 2^{-\lambda/2} \in \text{negl}(\kappa)$. In addition, we know that x_j is the aforementioned value by construction. Thus, we can inductively reason that our adversary computes $\{Y_{t,j}^{(i)}\}_{i \in [r], t \in [b]}$ correctly, and thus can recover the original encodings. \square

Lemma 12. When instantiated with a round function $r \in o(b \cdot n)$, the construction in Section 5 is not s -sound for any constant functions $s(\kappa, |m|) = c \in (0, 1)$.

Proof. Recall the definition of s -soundness as

$$\Pr \left[\begin{array}{l} (v, \text{state}, m) \leftarrow \text{Sound}_{\mathcal{A}_1, \mathcal{A}_2}(\kappa, n), s.t. \\ |\text{state}| < v \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|) \end{array} \right] \leq \text{negl}(\kappa).$$

We know by Lemma 11 that $v = n$ with all but negligible probability, and from Lemma 10 that $|\text{state}| \in \lambda \cdot o(b \cdot n) + n \cdot o(\lambda)$. We recall that $\text{len}(\kappa, |m|) = |m| + O(\kappa) > \lambda \cdot b$. From this, we can conclude that for sufficiently large $\kappa, |m|$, we know

$$\lambda \cdot o(b \cdot n) + n \cdot o(\lambda) < n \cdot c \cdot \lambda \cdot b < v \cdot s(\kappa, |m|) \cdot \text{len}(\kappa, |m|)$$

with all but negligible probability, and so this scheme is not s -sound. \square

References

1. Hamza Abusalah, Joël Alwen, Bram Cohen, Danylo Khilko, Krzysztof Pietrzak, and Leonid Reyzin. Beyond hellman’s time-memory trade-offs with applications to proofs of space. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 357–379. Springer, 2017.
2. Joël Alwen, Binyi Chen, Chethan Kamath, Vladimir Kolmogorov, Krzysztof Pietrzak, and Stefano Tessaro. On the complexity of script and proofs of space in the parallel random oracle model. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 358–387. Springer, 2016.
3. Joël Alwen, Binyi Chen, Krzysztof Pietrzak, Leonid Reyzin, and Stefano Tessaro. Script is maximally memory-hard. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 33–62. Springer, 2017.
4. Joël Alwen and Vladimir Serbinenko. High parallel complexity graphs and memory-hard functions. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 595–603, 2015.

5. J Benet and N Greco. Filecoin: A decentralized storage network. *Protoc. Labs*, 2018.
6. Juan Benet, David Dalrymple, and Nicola Greco. Proof of replication. *Protoc. Labs*, 2017.
7. Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 757–788. Springer, 2018.
8. Kevin D Bowers, Ari Juels, and Alina Oprea. Proofs of retrievability: Theory and implementation. In *Proceedings of the 2009 ACM workshop on Cloud computing security*, pages 43–54, 2009.
9. Ethan Cecchetti, Ben Fisch, Ian Miers, and Ari Juels. Pies: Public incompressible encodings for decentralized storage. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 1351–1367, 2019.
10. Yuanxi Dai and John Steinberger. Indifferentiability of 8-round feistel networks. In *Annual International Cryptology Conference*, pages 95–120. Springer, 2016.
11. Ivan Damgård, Chaya Ganesh, and Claudio Orlandi. Proofs of replicated storage without timing assumptions. In *Advances in Cryptology - CRYPTO 2019*, 2019.
12. Yevgeniy Dodis, Salil P. Vadhan, and Daniel Wichs. Proofs of retrievability via hardness amplification. In Omer Reingold, editor, *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, volume 5444 of *Lecture Notes in Computer Science*, pages 109–127. Springer, 2009.
13. Cynthia Dwork, Moni Naor, and Hoeteck Wee. Pebbling and proofs of work. In *Annual International Cryptology Conference*, pages 37–54. Springer, 2005.
14. Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. In *Annual Cryptology Conference*, pages 585–605. Springer, 2015.
15. Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. One-time computable self-erasing functions. In *Theory of Cryptography Conference*, pages 125–143. Springer, 2011.
16. Ben Fisch. Tight proofs of space and replication. Cryptology ePrint Archive, Report 2018/702, 2018. <https://eprint.iacr.org/2018/702>.
17. Ben Fisch, Joseph Bonneau, Juan Benet, and Nicola Greco. Proofs of replication using depth robust graphs. *Blockchain Protocol Analysis and Security Engineering*, 2018, 2018.
18. Rachit Garg, George Lu, and Brent Waters. New techniques in replica encodings with client setup. 2020.
19. Philippe Golle, Stanislaw Jarecki, and Ilya Mironov. Cryptographic primitives enforcing communication and storage complexity. In *International Conference on Financial Cryptography*, pages 120–135. Springer, 2002.
20. Thomas Holenstein, Robin Künzler, and Stefano Tessaro. The equivalence of the random oracle model and the ideal cipher model, revisited. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 89–98. ACM, 2011.
21. Ari Juels and Burton S Kaliski Jr. Pors: Proofs of retrievability for large files. In *Proceedings of the 14th ACM conference on Computer and communications security*, pages 584–597, 2007.
22. Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings*, volume 2951 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 2004.
23. Tal Moran and Daniel Wichs. Incompressible encodings. Manuscript, 2020.
24. Kazuo Ohta and Tatsuaki Okamoto. On concrete security treatment of signatures derived from identification. In *Annual International Cryptology Conference*, pages 354–369. Springer, 1998.
25. Krzysztof Pietrzak. Proofs of catalytic space. Cryptology ePrint Archive, Report 2018/194, 2018. <https://eprint.iacr.org/2018/194>.
26. Ling Ren and Srinivas Devadas. Proof of space from stacked expanders. In *Theory of Cryptography Conference*, pages 262–285. Springer, 2016.
27. Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indifferentiability framework. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 487–506. Springer, 2011.
28. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 475–484, 2014.
29. Hovav Shacham and Brent Waters. Compact proofs of retrievability. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 90–107. Springer, 2008.
30. Marten Van Dijk, Ari Juels, Alina Oprea, Ronald L Rivest, Emil Stefanov, and Nikos Triandopoulos. Hourglass schemes: how to prove that cloud files are encrypted. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 265–280, 2012.