

Non-Malleable Time-Lock Puzzles and Applications

Naomi Ephraim* Cody Freitag† Ilan Komargodski‡ Rafael Pass§

Abstract

Time-lock puzzles are a mechanism for sending messages “to the future”, by allowing a sender to quickly generate a puzzle with an underlying message that remains hidden until a receiver spends a moderately large amount of time solving it.

We introduce and construct a variant of a time-lock puzzle which is *non-malleable*. A non-malleable time-lock puzzle guarantees, roughly, that it is impossible to “maul” a puzzle into one for a related message without solving it. The security of this construction relies on the existence of any (plain) time-lock puzzle and it is proven secure in the auxiliary-input random oracle model. We show that our construction satisfies bounded concurrency and prove that it is impossible to obtain full concurrency. We additionally introduce a more general non-malleability notion, termed *functional* non-malleability, which protects against tampering attacks that affect a specific function of the related messages. We show that in many (useful) cases, our construction satisfies *fully* concurrent functional non-malleability.

We use our (functional) non-malleable time-lock puzzles to give efficient multi-party protocols for desirable tasks such as coin flipping and auctions. Our protocols are (1) *fair*, meaning that no malicious party can influence the output, (2) *optimistically efficient*, meaning that if all parties are honest, then the protocol terminates within two message rounds, and (3) *publicly verifiable*, meaning that from the transcript of the protocol anyone can quickly infer the outcome, without the need to perform a long computation phase. Our protocols support an unbounded number of participants and require no adversary-independent trusted setup. Our protocol is the first protocol that satisfies all of the above properties *under any assumption*. Security is proven assuming the repeated squaring assumption and in the auxiliary-input random oracle model. Along the way, we introduce a publicly verifiable notion of time-lock puzzles which is of independent interest. This notion allows the solver of the puzzle to compute the solution together with a proof which can be quickly verified by anyone.

*Cornell Tech, nephraim@cs.cornell.edu

†Cornell Tech, cfreitag@cs.cornell.edu

‡NTT Research, ilan.komargodski@ntt-research.com

§Cornell Tech, rafael@cs.cornell.edu

Contents

1	Introduction	1
1.1	Our Result 1: Non-Malleable TLPs	2
1.2	Our Result 2: Fair Multi-Party Auctions and Coin Flipping	4
1.3	Related Work	6
2	Technical Overview	7
2.1	Non-Malleability for Time-Lock Puzzles	7
2.2	Publicly Verifiable Time-Lock Puzzles	10
2.3	The Fair Multi-Party Protocols	13
3	Preliminaries	14
3.1	Time-Lock Puzzles	14
3.2	Non-Malleable Time-Lock Puzzles	15
3.3	Time-Lock Puzzles in the Auxiliary-Input Random Oracle Model	16
4	Non-Malleable Time-Lock Puzzles	17
4.1	Functional Non-Malleable Time-Lock Puzzles	18
4.2	Non-Malleable Time-Lock Puzzles Construction	18
4.3	Proof of Concurrent Functional Non-Malleability	19
4.4	Impossibility of Fully Concurrent Non-malleability	34
5	Publicly Verifiable Non-Malleable Time-Lock Puzzles	35
5.1	Strong Trapdoor VDFs	37
5.2	One-sided PV TLPs from Strong Trapdoor VDFs	41
5.3	Non-Malleable PV TLP from One-sided PV TLPs	42
6	Applications to Multi-Party Coin-Flipping and Auctions	45
	References	49
A	One-Many Non-malleability	52
B	Proof from Section 5	54

1 Introduction

Time-lock puzzles (TLPs), introduced by Rivest, Shamir, and Wagner [RSW96], are a cryptographic mechanism for committing to a message, where a sender can (quickly) generate a puzzle with a solution that remains hidden until the receiver spends a moderately large amount of time solving it (even in the presence of parallel processors). Rivest et al. [RSW96] gave a very efficient construction of TLPs where security relies on the repeated squaring assumption. This assumption postulates, roughly, that it is impossible to significantly speed up repeated modular exponentiations in a group of unknown order. This construction and assumption have proven extremely useful in various (and sometimes unexpected) applications [BN00, LPS17, Pie19, Wes19, EFKP19, MT19, DKP20], some of which have already been implemented and deployed in existing systems.

Non-malleability. In a Man-In-the-Middle (MIM) attack, an eavesdropper tries to actively maul intermediate messages to compromise the integrity of the underlying values. To address such attacks, Dolev, Dwork and Naor [DDN91] introduced the general concept of non-malleability in the context of cryptographic commitments. Roughly speaking, non-malleable commitments are an extension of plain cryptographic commitments (that guarantee binding and hiding) with the additional property that no adversary can maul a commitment for a given value into a commitment to a “related” value. As this is a fundamental concept with many applications, there has been a tremendous amount of research on this topic [Bar02, PR05b, PR05a, LPV08, PPV08, LP09, PW10, Wee10, Goy11, LP11, GLOV12, GPR16, COSV16, COSV17, Khu17, LPS17, KS17, KK19].

Non-malleable TLPs and applications. To date, non-malleability has never been considered in the context of TLPs (or other timed primitives). Indeed, the construction of TLPs of [RSW96] *is malleable*.¹ This fact actually has negative consequences in various settings where TLPs could be useful. For instance, consider a scenario where n parties perform an auction by posting bids on a public bulletin board. To implement this fairly, a natural approach is to use a commit-and-reveal style protocol, where each party commits to its bid on the board, and once all bids are posted each party publishes its opening. Clearly, one has to use non-malleable commitments to guarantee that bids are independent (otherwise, a malicious party can potentially bid for the maximal other bid plus 1). However, non-malleability is not enough since there is a fairness issue: a malicious party may refuse to open after seeing all other bids and so other parties will never know what her bid was.

Using non-malleable TLPs to “commit” to the bids solves this problem. Indeed, the puzzle of a party who refuses to reveal its bid can be recovered after some moderately large amount of time by all honest parties. This style of protocol can also be used for fair multi-party collective coin flipping where n parties wish to agree on a common unbiased coin. There, each party encodes a random bit via a TLP and all parties will eventually agree on the parity of those bits.² This gives a highly desirable collective coin flipping protocol with an important property called *optimistic efficiency*: when all parties are honest and publish their “openings” immediately after seeing all puzzles, the protocol terminates and all parties agree on an unbiased bit. As we will see, no other known protocol for this (highly important) task has this property.

¹The puzzle of [RSW96] for a message s and difficulty T is a tuple $(g, N, T, s \oplus g^{2^T} \bmod N)$, where N is an RSA group modulus and g is a random element from \mathbb{Z}_N . The puzzle is trivially malleable since the message is one-time padded.

²For coin flipping, there is an inherent fairness issue where a malicious party can always “prematurely abort” and bias the output [Cle86]. Boneh and Naor [BN00] used timed primitives and interaction to circumvent the issue in the two-party case but we care about the multi-party case and prefer to avoid interaction as much as possible.

1.1 Our Result 1: Non-Malleable TLPs

Our first result is a construction of a non-malleable TLP. We emphasize that, as explained above, this primitive is not only natural on its own right, but also has important applications to the design of secure protocol for various basic tasks. Our construction relies on the existence of any given TLP [RSW96, BGJ⁺16] and is proven secure in the (auxiliary-input) random oracle model [Unr07]. Our construction achieves non-malleability even in a bounded *concurrent* setting [PR08], preventing any man-in-the-middle (MIM) attacker that sees an a priori bounded polynomial number of TLPs with underlying values s_1, \dots, s_n from generating puzzles with underlying solutions $\tilde{s}_1, \dots, \tilde{s}_n$ that are somehow non-trivially related to s_1, \dots, s_n .

Theorem 1.1 (Informal; See Theorem 4.2 and Corollary 4.3). *Assuming that there exists a TLP, there exists a bounded concurrent non-malleable TLP. The scheme is proven secure in the auxiliary-input random oracle model.*

We remark about two additional useful properties of our construction:

- **Depth-preserving hardness:** If the given TLP is secure against attackers of depth $T^\epsilon \cdot \text{poly}(\lambda)$, where $\epsilon > 0$ is a fixed constant and $\text{poly}(\cdot)$ is a fixed polynomial in the security parameter λ , then so is the resulting non-malleable TLP. In other words, the reduction is depth-efficient.
- **Message length proportional security loss:** If we want the resulting (one-one³) non-malleable TLP to support L bit messages, we need the given underlying TLP to be secure for bounded polynomial depth attackers of size $2^L \cdot \text{poly}(\lambda)$ for all polynomials $\text{poly}(\cdot)$ of the security parameter λ . In particular, if we need a (one-one) non-malleable TLP for messages of length up to $L = O(\log \lambda)$ bits, then we need to start off with a TLP with standard polynomial security. If we need to encode much longer messages, say of length λ bits, then we need our given TLP to be sub-exponentially secure. More generally, if we want to allow the MIM attacker to see n concurrent TLPs with messages of length L and output n TLPs, then the security loss is proportional to $2^{n \cdot L} \cdot \text{poly}(\lambda)$.

Instantiating the TLP with the construction of [RSW96], our scheme is extremely efficient: encoding a message requires a single invocation of a random oracle and few (modular) exponentiations. Additionally, our construction is very simple to describe: to generate a puzzle for a solution s with opening r , we sample a puzzle for (s, r) using randomness which itself depends (via the random oracle) on s and r . Nevertheless, the proof of security turns out to be somewhat tricky and non-trivial; see Section 2 for details. The efficiency and simplicity of our construction stand in contrast to a previous construction [DKP20] whose main feature is that it is in the plain model (without any form of setup), albeit they rely on a variety of assumptions and their construction is not practical.

We prove that our scheme is non-malleable against all polynomial-size attackers that cannot solve the puzzles (and this is inherent as the latter ones can easily maul any puzzle). We even allow the attacker’s description to depend arbitrarily on the random oracle. We formalize this notion by showing that our TLP is non-malleable in the *auxiliary-input* random oracle model, a model that was introduced by Unruh [Unr07] (see also [CDGS18]) in order to capture preprocessing attacks, where a non-uniform attacker obtains an advice string that depends arbitrarily on the random oracle. Thus, in a sense, our construction does not require any form of attacker-independent setup.

³In one-one non-malleability the MIM attacker gets to see a single TLP and is allowed to output a single TLP.

We emphasize that our protocol only achieves *bounded* concurrency, where the number of instances the attacker participates in is a priori bounded (and the scheme may depend on this bound). We show that the stronger notion of *full* concurrency, which does not place such limitations and is achievable in all other standard settings of non-malleability, is actually impossible to achieve for TLPs. Therefore, our result is best possible in this sense.

Theorem 1.2 (Informal; See Theorem 4.17). *There is no fully concurrent non-malleable TLP (even in the random oracle model).*

In a nutshell, the impossibility from Theorem 1.2 is proven by the following generic MIM attack. Given a puzzle z , if the number of “sessions” the attacker can participate in at least as large as $|z|$, she can essentially generate $|z|$ puzzles encoding *the bits of* z . Since the distinguisher in of the MIM game (which is now given those bits) can run in arbitrary polynomial time, it can simply solve the original puzzle and recover the original solution in full. This attack is circumvented in the bounded concurrency setting (Theorem 1.1) by setting the length of the puzzle to be longer than the concurrency bound.

Functional non-malleability. We note that the attack on fully concurrent non-malleable time-lock puzzles crucially relies on the fact that the *distinguisher* in the MIM game can solve the underlying puzzles. However, it is easy to see that if the distinguisher is restricted to bounded depth, this attack fails.

In light of this observation, we introduce a new definition of non-malleability that *generalizes* the standard notion considered in Theorem 1.1. We call the notion *functional* non-malleability and, as the name suggests, the security notion is parameterized by a class of functions \mathcal{F} . Denote by L the bit-length of the messages we want to support and by n the number of sessions in the input of the MIM attacker. We think of $f \in \mathcal{F}$ as some *bounded depth* function of the form $f: (\{0, 1\}^L)^n \rightarrow \{0, 1\}^m$, which is the target function of the input messages that the MIM adversary is trying to bias. Formally, we require that the distinguisher in the MIM game cannot distinguish the function f applied to the values underlying the puzzles the MIM attacker outputs. When \mathcal{F} includes all identity functions (which are bounded depth and have output length $m = n \cdot L$), functional non-malleability implies the standard definition of concurrent non-malleability (as the distinguisher just gets all the messages from the n mauled puzzles).

Naturally, it makes sense to ask what guarantees can we get if we a priori restrict f , say in its output length, without limiting the number of sessions n . This turns out to be particularly useful when the application in hand only requires non-malleability against a specific form of tampering functions (this indeed will be the case for us below). Concretely, let $\mathcal{F}_{m,d}$ to be the class of all functions whose output length is at most m bits and can be computed in fixed depth d on input length $n \cdot L$ (using the notation given above). Then, we have the following result.

Theorem 1.3 (Informal; See Theorem 4.2). *Assuming that there exists a TLP, then for every $m, d \in \text{poly}(\lambda)$ there exists a fully concurrent functional non-malleable TLP for the class functions $\mathcal{F}_{m,d}$. The scheme is proven secure in the auxiliary-input random oracle model assuming the given TLP is secure for all attackers of size at most $2^m \cdot \text{poly}(\lambda)$.*

The above construction satisfies *depth hardness preserving* property as the construction from Theorem 1.1. Further, note that as long as $m \in O(\log \lambda)$, we only require standard polynomial hardness from the given TLP.

We remark that Theorem 1.3 will turn out to be instrumental for our applications we discuss below. We also believe that the abstraction of functional non-malleability is important on its own right and view it as an independent contribution.

1.2 Our Result 2: Fair Multi-Party Auctions and Coin Flipping

As we mentioned, an appealing application of non-malleable TLPs is for tasks such as fair multi-party auctions or coin flipping. Our protocols (for both tasks) are extremely efficient and consist of just two phases: first each party “commits” to her bid/randomness using some puzzle, and then after all puzzles are made public, each party publishes its solution. If some party refuses to open their puzzle, a force-opening phase is performed.

In what follows, we focus on the task of fair multi-party coin-flipping, which is a core building blocks in recent proof-of-stake block-chain designs; see below. The application to auctions follows in a similar manner. It is convenient to consider our protocol in a setting where there is a public bulletin board. Any party can publish a puzzle to the bulletin board during the commit phase and then publish its solution after some pre-specified amount of time has elapsed. Our protocol requires no form of *adversary-independent* trusted setup (not even a common random string). We further emphasize the following highly desirable properties we achieve:

- **Optimistic Efficiency:** If all participating parties are honest, then the protocol terminates within two message rounds (without the need to wait the pre-specified amount of time for the second phase), and all parties can efficiently verify the output of the protocol.
- **Fairness:** No malicious party can bias the output of the protocol. Namely, as long as there is at least one honest participating party, the output will be a (nearly) uniformly random value.
- **Public Verifiability:** In the case that any participating party is dishonest and does not publish their solution, any party can break the puzzle in a moderate amount of time and provide a publicly verifiable proof of the solution. We even require that an honest party can prove that a published puzzle has *no* valid solution.

We note that optimistic efficiency is a typical feature of multi-party protocols using the “commit-and-reveal” paradigm, which we employ using time-lock puzzles as the analog of commitments. However, to obtain fairness and public verifiability, we make use of new additional properties of time-lock puzzles. Fairness follows from the notion of concurrent functional non-malleability, which, as described above, we achieve generically from any time-lock puzzle in the auxiliary-input random oracle model. In order to get public verifiability, we define and construct a related notion of *publicly verifiable* time-lock puzzles. We build such time-lock puzzles using a variant of verifiable delay functions (VDFs) [BBBF18, Wes19, Pie19] that can be solved *fast* using a trapdoor—the precise set of properties that we need are (non-trivially) satisfied by the known VDF construction of Pietrzak [Pie19] based on the repeated squaring assumption in the random oracle model. Put together, we achieve the following result.

Theorem 1.4 (Informal; See Theorem 6.1). *Assuming the repeated squaring assumption, there exist multi-party coin-flipping and auction protocols that satisfy optimistic efficiency, fairness, and public verifiability. The protocols support an unbounded number of participants and require no adversary-independent trusted setup. Security is proven in the auxiliary-input random oracle.*

Our protocol is the first multi-party auction protocol *under any assumption* that satisfies fairness and requires no adversary-independent setup—the construction of [BN00] relies on trusted setup and the construction of [MT19] does not satisfy fairness. For coin-flipping, a protocol due to Boneh et al. [BBBF18] (that uses VDFs and a random oracle⁴) does not satisfy optimistic efficiency: even

⁴In their protocol, each party publishes a random string r_i and then the agreed upon coin is defined by running a VDF on the seed $H(r_1 || \dots || r_n)$, where H is a random oracle.

in the fully honest execution getting the agreed upon coin flip requires a long computation. In contrast, our protocol terminates extremely fast when all parties are honest. On the other hand, our protocol requires two phases while the one of Boneh et al. [BBBF18] requires only one (if the the setup is given for free or in the random oracle model). For our protocol, we could fall back to a single phase where a single party solves all puzzles after they are published, but this is more expensive than the protocol of [BBBF18] which requires only a single party to solve one puzzle. Malavolta and Thyagarajan [MT19] address this inefficiency in the context of time-lock puzzles by constructing *homomorphic* time-lock puzzles, where many separate puzzles can be combined into a single puzzle to be solved. However, their TLP scheme is malleable and so cannot be directly used to obtain a fair protocol. We note that because our protocol satisfies public verifiability, only a single honest party needs to solve each puzzle, and this computation can easily be delegated to an external server.

We emphasize that our protocol supports an a priori unbounded number of participants. This may seem strange in light of our impossibility from Theorem 1.2. We bypass this lower bound (as mentioned above) by observing that for most natural applications (including coin-flipping and auctions), the notion of functional non-malleability from Theorem 1.3 suffices. The key insight is that we only need indistinguishability with respect to specific depth-bounded functions with short output (e.g., parity for coin-flipping, or taking the maximum for auctions). Since the output in both cases is short, we can actually support *full* concurrency which translates into having an unbounded number of participants. Additionally, we do not need complexity leveraging and can base our protocol on standard polynomial hardness.

Non-interactive coin flipping. As mentioned, our protocol has a natural non-interactive fall back to a non-interactive (single phase) protocol where all parties just publish their puzzles and then a (possibly external) evaluator solves all puzzles. Our protocol does not rely on any attacker-independent trusted setup. This is the first protocol of this kind—the only previously known protocol of [BBBF18] relies on (adversary-independent) trusted setup.

Publicly-verifiable non-malleable TLPs. As we hinted above, along the way we construct a non-malleable TLP which has an additional public verifiability property: after one party solves the puzzle, she can publish the underlying solution together with a proof which can be later used by anyone to *quickly* verify the correctness of the solution. We believe this primitive is of independent interest. We build our TLP assuming a very weak form of (partially) trusted setup. We design the setup to fit into our multi-party protocol application (Theorem 1.4) such that the parties themselves will generate this setup in the puzzle generation phase. The setup of our publicly-verifiable non-malleable TLP consists of a set of many public parameters where we only assume that at least one of them was generated honestly. We call this model the All-But-One-string (ABO-string) model.⁵

Theorem 1.5 (Informal; See Corollary 5.10). *Assuming the repeated squaring assumption, there exists a publicly verifiable non-malleable TLP in the ABO-string model. The construction is proven secure in the auxiliary-input random oracle model.*

The above construction satisfies the same *depth hardness preserving* and *message length proportional security loss* properties as the construction from Theorem 1.1.

⁵Our ABO-string model is a variant of the multi-string model of Groth and Ostrovsky [GO14], where it is assumed that a majority of the public parameters are honestly generated.

1.3 Related Work

Non-malleable TLPs. The study of non-malleability in the context of TLPs is new to this work and as we mentioned it has exciting applications. Such a puzzle was implicitly constructed in [DKP20] (presented there as a non-malleable code for bounded polynomial depth attackers). That construction is in the plain model, without any setup assumptions, however it relies on a variety of assumptions and it should be viewed as a feasibility result. Our approach, on the other hand, is more oriented towards practical constructions and therefore we are willing to assume some very weak form of setup (i.e., AI-ROM). Additionally, our construction is somewhat reminiscent to the Fujisaki-Okamoto (FO) transformation [FO13] used to generically transform any CPA-secure public-key encryption scheme into a CCA-secure one using a random oracle. Nevertheless, while conceptually the transformation is similar, since we are dealing with concurrent non-malleability in the bounded polynomial depth setting, our actual proof turns out to be quite challenging.

Fair multi-party coin flipping. A recent work of Malavolta and Thyagarajan [MT19] suggested to use plain TLPs to implement fair auctions and coin flipping protocols. They showed how efficiency in such protocols can be improved by modifying the [RSW96] repeated-squaring based TLP to be homomorphic in the sense that one can homomorphically combine many TLPs into a single one and then only solve that one. However, using plain TLPs results in a secure protocol only if one assumes that the parties act honestly (i.e., the honest-but curious model). It is possible to make this protocol maliciously secure using say concurrent non-malleable zero-knowledge proofs [BPS06, OPV10, LPTV10, LP11], proving that each party acted honestly, but this (1) makes the construction significantly less efficient, and (2) requires either trusted setup and additional hardness assumptions, or additional rounds of interaction. Our solution achieves malicious security while avoiding all of these drawbacks (but our puzzles are not homomorphic).

Let us remark that the protocols that we described guarantee fairness but not privacy. The latter, however, can be obtained in specific applications by composing our protocols with existing privacy-preserving tools such as Anonize [HMPS15].

Applications of fair coin flipping. Obtaining unbiased randomness (or a coin flip) is a natural and important building block in many systems. One well known example is online games, say card games, where participants desire to be certain that the randomness that is used in the game is unbiased.

In fact, generating unbiased bits is one of the largest bottlenecks in modern proof-of-stake crypto-currency designs [BPS16, DPS19, DGKR18]. Recall that in a proof-of-stake blockchains, the idea is, very roughly speaking, to enforce “one vote per unit of stake”. This is usually implemented by choosing random small committees at every epoch and letting that committee decide on the next block. The main question is how to obtain “pure” randomness so that the chosen committee is really “random”.

One option is to use the hash of an old-enough block as the randomness. Unfortunately, it is known that the hash of a block is not completely unbiased: an attacker can essentially fully control about logarithmically many of its bits. In existing systems, this is mitigated by “blowing up” parameters to compensate for the (small yet meaningful) advantage the attacker has, making those systems much less efficient. Using a mechanism that generates unbiased bits, we could make proof-of-stake crypto-currencies much more efficient.

Timed commitments. Boneh and Naor [BN00] introduced timed commitments, which can be viewed as a publicly verifiable and interactive TLP. They additionally require that the puzzle (which

is an interactive commitment) convinces the receiver that if they brute-force the solution, they will succeed. Because of this additional property, their commitment scheme is interactive and relies on a less standard assumption called the generalized Blum-Blum-Shub assumption. Their scheme is additionally malleable.

2 Technical Overview

In Section 2.1, we give an overview of our non-malleable time-lock puzzle construction and its proof of security. Then in Section 2.2, we overview our construction public verifiable (and non-malleable) time-lock puzzles from repeated squaring. Finally in Section 2.3, we discuss how our non-malleable and publicly verifiable time-lock puzzle construction can be used for fair multi-party coin-flipping with various desirable properties.

We start by recalling the definition of TLPs, as necessary to give an overview of our techniques. A TLP consists of two algorithms (Gen, Sol). Gen is a probabilistic procedure that takes as input an embedded solution s and a time parameter t , and outputs a puzzle z . Sol is a deterministic procedure that on input a puzzle z for time bound t , outputs a solution in depth (or parallel time) t . We note that TLPs can be thought of as a fine-grained analogue to commitments where “hardness” of the puzzle means that the puzzles are hiding against distinguishers of depth less than t . On the other hand, hiding *can be broken* in depth t (using Sol). Additionally, we require that Sol always finds the correct underlying solution s for a puzzle z . This corresponds to perfect bidding in the language of commitments.

2.1 Non-Malleability for Time-Lock Puzzles

Non-malleability, at a high level, requires that any man-in-the-middle (MIM) attacker \mathcal{A} that receives a puzzle z with solution s cannot output a different puzzle \tilde{z} for a related value \tilde{s} . Formally, we consider the (inefficient) distribution $\text{mim}_{\mathcal{A}}(s, t)$ that computes $\tilde{z} = \mathcal{A}(\text{Gen}(s, t))$ and outputs the value \tilde{s} underlying \tilde{z} . However, if $z = \tilde{z}$, then $\tilde{s} = \perp$ (since simply forwarding the commitment does not count as a valid mauling attack), and if there is no valid value underlying \tilde{z} , then \tilde{s} is also \perp . Then, non-malleability requires that for any MIM attacker with depth much less than t (so it cannot break hiding) and solution s , the distribution for value \tilde{s} given by $\text{mim}_{\mathcal{A}}(s, t)$ is statistically indistinguishable from the distribution $\text{mim}_{\mathcal{A}}(0, t)$ for an unrelated value, 0. We emphasize that indistinguishability should hold even against *unbounded* distinguishers that, in particular, can run Sol .⁶

Our construction. Our construction relies on any time-lock puzzle TLP and a common random oracle \mathcal{O} . We now describe our non-malleable TLP, which we denote nmTLP . In order to generate a puzzle for a solution s that can be broken in time t , nmTLP.Gen uses randomness r and feeds $s||r$ into the random oracle to get a string r_{tlp} . It then uses TLP.Gen to create a puzzle with difficulty t for $s||r$ using randomness r_{tlp} . That is,

$$\text{nmTLP.Gen}(s, t; r) := \text{TLP.Gen}(s||r, t; \mathcal{O}(s||r)).$$

Note that in order to solve the puzzle output by nmTLP.Gen , it suffices to just solve the puzzle generated using TLP.Gen , which takes time t . In other words, $\text{nmTLP.Sol}(z)$ simply computes $s||r = \text{TLP.Sol}(z)$ and outputs s .

⁶We could also naturally consider a computational notion of non-malleability that requires only computational indistinguishability. Still, an arbitrary polynomial time attacker could also run the Sol algorithm.

Hardness. To show the hardness of nmTLP relative to a random oracle, we rely on the hardness of TLP in the plain model, against attackers of depth much less than t . At a high level, we should that breaking the hardness of nmTLP requires either guessing the randomness r used to generate the randomness $\mathcal{O}(s||r)$ for the underlying puzzle, or directly breaking the hardness of TLP, both of which are infeasible for bounded attackers. To formalize this, we consider any depth-bounded distinguisher $\mathcal{D}^\mathcal{O}$, who receives as input a nmTLP puzzle z corresponding to solution s_0 or s_1 and distinguishes the two cases with noticeable probability. By construction, z actually corresponds to a TLP puzzle for $s_0||r_0$ or $s_1||r_1$, so we would like to use \mathcal{D} to construct a distinguisher against the hardness of TLP.

We first note that if \mathcal{D} never makes a query to \mathcal{O} containing the randomness r_b underlying z , then we can simulate \mathcal{O} by lazily sampling it in the plain model, and hence use \mathcal{D} as a distinguisher for the hardness of TLP. If \mathcal{D} does make a query containing r_b , then with overwhelming probability it must have received a puzzle corresponding to $s_b||r_b$ (since in this case, r_{1-b} is independent of \mathcal{D} and its input z). Moreover, all of its queries up until that point have uniformly random answers independent of z , so we can simulate them as well, up until receiving this query. Therefore, in both cases, we can carry out this attack in the plain model and rely on the hardness of TLP.

Non-malleability. To show non-malleability of nmTLP, we want to argue that any depth-bounded man-in-the-middle (MIM) attacker \mathcal{A} cannot maul a puzzle z for s (received on the left) to a puzzle \tilde{z} (output on the right) for a related value $\tilde{s} \neq s$. At a high level, whenever \mathcal{A} changes the underlying value s to \tilde{s} , then the output of the random oracle on \tilde{s} is now uniformly random and independent of z . Indeed, we show that for any fixed puzzle \tilde{z} and a value \tilde{s} , a randomly generated puzzle for \tilde{s} will not be equal to \tilde{z} with high probability (otherwise we show how to break the hardness of TLP). So, intuitively, the only way to generate a *valid* puzzle \tilde{z} for \tilde{s} is to “know” the underlying value \tilde{s} , but hardness intuitively implies that no depth-bounded adversary can “know” s .

We formalize this intuition by a hybrid argument to show that the MIM distribution $\tilde{s} \leftarrow \text{mim}_{\mathcal{A}}(s, t)$ is indistinguishable from $\text{mim}_{\mathcal{A}}(0, t)$. At a high level, we first replace the inefficient distribution $\text{mim}_{\mathcal{A}}(s, t)$ by a low-depth circuit \mathcal{B} . At this point, we want to use the hiding property to indistinguishably swap the puzzle to 0, so the hybrid is now unrelated to s . We describe the key ideas for these hybrids below.

For the first hybrid, the key insight is that we can compute $\text{mim}_{\mathcal{A}}(s, t)$ *in low depth* using an algorithm \mathcal{B} by simply looking at the oracle queries made by \mathcal{A} . In this sense, we are relying on the extractability property of random oracles to say that \mathcal{A} must know any valid value \tilde{s} it generates a puzzle for. Specifically, let \tilde{z} be the output of \mathcal{A} . For every query $(s_i||r_i)$ that \mathcal{A} makes to \mathcal{O} , \mathcal{B} outputs s_i if $\tilde{z} = \text{nmTLP}(s_i||r_i, t; \mathcal{O}(s_i||r_i))$. If there are no such queries, \mathcal{B} outputs \perp . \mathcal{B} requires depth comparable to the depth of \mathcal{A} since all of these checks can be done in parallel. Furthermore, the output of \mathcal{B} is indistinguishable from the true output given the above observation that \mathcal{A} cannot output a valid puzzle for a value it doesn’t query.

For the next hybrid, we would like to indistinguishably replace the underlying puzzle for s with a puzzle for 0, which would suffice to show non-malleability. Because \mathcal{B} is low-depth, it seems that we should be able to use the hiding property of nmTLP to say that the output of \mathcal{B} does not depend on the underlying value s . Specifically, we want to conclude that if the output of \mathcal{B} (who outputs many bits) is statistically far when the underlying value is s versus 0, then there exists a distinguisher (who outputs a single bit) that can distinguish puzzles for s and 0. Towards this claim, we show how to “flatten” any (possibly unbounded) distinguisher \mathcal{D} who distinguishes between the output of \mathcal{B} in the case where the underlying value is s versus 0. Specifically, we encode

the truth table of \mathcal{D} as a low-depth distinguishing circuit of size roughly $2^{|s|}$ to make this reduction go through. As a result, we need to rely on a sub-exponentially security of the underlying TLP when $|s| = \lambda$. Namely, the underlying TLP cannot be broken by sub-exponential sized circuits with depth much less than t . However, when $|s| \in O(\log \lambda)$, we only need to rely on *polynomial* security of the underlying TLP.

Impossibility of fully concurrent non-malleability. Ideally, we would like to achieve fully concurrent non-malleability, meaning that any MIM attacker that receives any polynomial n number of puzzles on the left cannot maul them to n puzzle for related values. However, we show that this is impossible to achieve.

Consider an arbitrary TLP for a polynomial time bound t . We construct a MIM attacker \mathcal{A} that receives only a single puzzle z on the left with solutions s where the length of z is L . \mathcal{A} can split z into L bits and output a puzzle on the right for each bit *of the puzzle* z . Then, the values underlying the puzzles output by \mathcal{A} when viewed together yield z , which is related to the value s ! More formally, there exists a polynomial time distinguisher that solves the puzzle z in polynomial time t and can distinguish \mathcal{A} 's output in the case when it receives a puzzle for s or an unrelated value, say 0.

This implies that for any n which is greater than the size of a puzzle, the TLP cannot be non-malleable against MIM attackers who output at most n puzzles on the right. At a high level, the impossibility follows from the fact that hardness does not hold against arbitrary polynomial-time distinguishers (which usually *is* the case for hiding of standard commitments).

Despite this impossibility, we show that we actually *can* achieve concurrent non-malleability against a *specific class of distinguishers* in the non-malleability game. We refer to this notion as concurrent *functional* non-malleability.

Achieving concurrent functional non-malleability. In many applications, we only need a form of non-malleability to hold with respect to certain classes of functions. For example, in our application to coin-flipping, we only need that a puzzle z with solution s cannot be mauled to a set of puzzles $\tilde{z}_1, \dots, \tilde{z}_n$ with underlying values $\tilde{s}_1, \dots, \tilde{s}_n$ such that $\bigoplus_{i \in [n]} \tilde{s}_i$ “depends on” s . With this in mind, we define a concurrent functional non-malleability with respect to a class of functions \mathcal{F} . We say that a TLP satisfies *functional* non-malleability for a class \mathcal{F} if the output of $f(\text{mim}_{\mathcal{A}}(s, t))$ is indistinguishable from $f(\text{mim}_{\mathcal{A}}(0, t))$ for any $f \in \mathcal{F}$, which also naturally generalizes to the concurrent setting. we note that functional non-malleability for a class \mathcal{F} actually implies standard non-malleability whenever the class \mathcal{F} contains the identity function, so functional non-malleability generalizes the standard notion of non-malleability.

Going back to the proof of standard (non-concurrent) non-malleability for our construction nmTLP, we observe that the security we need for the underlying time-lock puzzle we use depends on $2^{|s|}$ where $|s|$ is the size of the puzzle solutions. Specifically, given any distinguisher in the non-malleability that had input of size $|s|$, we were able to construct a distinguisher for hardness of size $2^{|s|}$. In fact, this exact same proof works in the context of concurrent functional non-malleability for functions f that have *low depth* and bounded output length m . We require f to be low depth so the reduction constitutes a valid attack against hardness, and then we only require security proportional to 2^m !

We briefly discuss how our nmTLP construction works for concurrent functional non-malleability for the class \mathcal{F}_m of function with low depth and output length m . Specifically, for every m , we define a scheme nmTLP $_m$ assuming that TLP is secure against attackers of size roughly 2^m . Because TLP requires security against 2^m size attackers, our construction nmTLP $_m$ also only achieves security

against 2^m size attackers. As such, our `nmTLP.Gen` algorithm needs to use at least $m + \lambda$ bits of randomness (otherwise an attacker could cycle through all choices of randomness to break security). Recall that `nmTLPm.Gen` with randomness r outputs a puzzle using `TLP.Gen` with solution $s||r$. As a result, if we want to support solutions of size $|s|$ in `nmTLPm`, we need our underlying `TLP` to support solutions of size at least $|s| + m + \lambda$. By correctness, this implies that our schemes outputs puzzles of size roughly $O(|s| + m + \lambda)$.

Bounded concurrent non-malleability. Our construction of time-lock puzzles for concurrent functional non-malleability can also be seen as a construction for bounded concurrent (plain) non-malleability. Specifically, consider the case where the MIM attacker outputs at most n puzzles “on the right.” We can think of this as functional non-malleability where the low depth function is simply identity on $n \cdot |s|$ bits. From the above discussion, this implies a protocol assuming a `TLP` with security against size $2^{n \cdot |s|}$ attackers, with puzzles of size roughly $O(n \cdot |s| + \lambda)$.

Security in the auxiliary-input random oracle model. Finally, we remark that all of our constructions and all formal proofs in the main body of the paper are in the auxiliary-input random oracle model (AI-ROM) introduced by Unruh [Unr07]. In this model, the non-uniform attacker is allowed to depend arbitrarily on the random oracle, so there is no attacker-independent non-uniform advice. At a high level, we use the result from [Unr07] (restated in Lemma 3.6) to conclude that the view of any bounded-size MIM attacker \mathcal{A} with oracle access to \mathcal{O} (where \mathcal{A} may depend arbitrarily on \mathcal{O}) is indistinguishable the view of \mathcal{A} with access to a “lazily sampled” oracle \mathcal{P} that is fixed at a set of points F (which depend on \mathcal{A}). Formally, in the non-malleability analysis, we switch to an intermediate hybrid where the MIM attacker has access to a partially fixed, lazily sampled oracle \mathcal{P} . Then, because the MIM attacker \mathcal{A} must maul honestly generated puzzles that have high entropy, we show that it is necessary for \mathcal{A} to query the oracle \mathcal{P} outside the fixed set of points F . From this, we carefully show that a similar analysis follows as discussed above for the ROM.

2.2 Publicly Verifiable Time-Lock Puzzles

We observe that the non-malleable time-lock puzzle construction `nmTLP` we described above has a very natural—yet incomplete—public verifiability property. Solving a puzzle yields both the solution s and the randomness r use to generate that puzzle. As such, anyone who solves a valid puzzle can send the opening r to another party, and convince them that s is the unique valid solution to the puzzle. However, we emphasize that this only works for *valid* puzzles and solutions.

Consider the following problematic scenario for our `nmTLP` construction. Suppose a party “commits” to a value via a puzzle z and refuses to open the commitment. As we said before, if z is a valid puzzle, any party can solve the puzzle, get the solution s and an opening r that proves that s is the unique solution. What if the puzzle corresponds to *no solution*? We refer to this scenario by saying that the puzzle corresponds to the solution \perp . In this case (by definition), there is no solution s and opening r for any such that $z = \text{Gen}(s, t; r)$. Anyone who solve the invalid puzzle—which requires a lot of computational power—*will* be able to conclude that the puzzle is malformed, but they will not be able to convince anyone else that this is the case. Ideally, we would have a time-lock puzzle where `Sol` additionally outputs a publicly verifiable proof π that the solution it computes is correct, even if the solution may be \perp ! We refer to such a time-lock puzzle as a *publicly verifiable* time-lock puzzle. We next discuss the definition and our construction of publicly verifiable time-lock puzzles.

Defining public verifiability. More formally, a publicly verifiable time-lock puzzle consists of algorithms $(\text{Gen}, \text{Sol}, \text{Verify})$. As with normal time-lock puzzles, $\text{Gen}(s, t)$ outputs a puzzle z . $\text{Sol}(z, t)$ outputs the solution s as well as a proof π that it computed s correctly. Finally $\text{Verify}(z, (s, \pi), t)$ checks that s is indeed the correct solution for the puzzle z (corresponding to $\text{Sol}(z, t)$), using the proof π . In addition to (Gen, Sol) being a valid time-lock puzzle, we require that Sol and Verify constitute a sound non-interactive argument. In fact, we require a very strong notion of soundness. We need it to be the case that even for maliciously chosen puzzles that have no solution, the time-lock puzzle is still sound—even against the adversary that generated the malformed puzzle. In other words, we require that no attacker can compute a puzzle z , a value s' , and a proof π' such that $\text{Verify}(z, (s', \pi'), t)$ accepts yet s' is not the value s computed by $\text{Sol}(z, t)$, which may be \perp .

Ideally, we would want a publicly verifiable time-lock puzzle that requires *no setup*. We instead consider a weak form of setup which we refer to as the All-But-One-string (ABO-string) model. In this model, Sol and Verify additionally take as input a string $\text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \in (\{0, 1\}^\lambda)^n$, and we require that soundness holds as long as one of the values of crs_i is sampled uniformly (without necessarily knowing which one); this is why we refer to it as the all-but-one string model. We note that in multi-party protocols, the ABO-string model is realistic as each participant $i \in [n]$ can post a value for crs_i . Then, we require soundness to hold as long as one participant is honest, which is a reasonable assumption in this multi-party setting.

Constructing publicly verifiable time-lock puzzles. Our construction of a publicly verifiable time-lock puzzle follows the blueprint of Rivest, Shamir, and Wagner [RSW96] for constructing time-lock puzzles from repeated squaring. Namely, we use the output of a sequential function (repeated squaring in a suitable group) essentially as one-time pad to mask the value underlying the time-lock puzzle. As in [RSW96], we require that the sequential function has a trapdoor so that puzzles can be generated efficiently. Unlike [RSW96], we additionally require that the sequential function is publicly verifiable to enable public verifiability for the time-lock puzzle. Finally, we apply the non-malleability transformation described above to achieve full public verifiability. In what follows, we describe each of these steps in more detail.

For the underlying sequential function, we use what we call a *strong trapdoor verifiable delay function* (VDF). A VDF (introduced by Boneh et al. [BBBF18]) is a publicly verifiable sequential function that can be computed in time t but not much faster, even with lots of parallelism. A trapdoor VDF (formalized by Wesolowski [Wes19]) additionally has a trapdoor for quick evaluation. We require a trapdoor VDF in the ABO-string model that satisfies additional properties required by our application. While the properties we define—an achieve—are heavily tailored towards our application, we believe some of the techniques may be of independent interest. More specifically, a strong trapdoor VDF comes with a **Sample** algorithm to generate inputs for an evaluation algorithm **Eval**. We emphasize that, even in the ABO-string model, **Sample** is independent of any form of setup. Previous definitions of VDFs require the proof to be sound with probability over an *honestly sampled* input. In contrast, we require that the proof is sound for any maliciously chosen input that is in the support of the **Sample** algorithm. We note that this property is satisfied by a variant of Pietrzak’s VDF [Pie19] based on repeated squaring. At a high level, this is because Pietrzak’s VDF is sound (at least in the random oracle model) for any group of unknown order where no adversary can find a group of low order (see e.g. [BBF18] for further discussion), so by using any RSA group with no low order elements (as in [Pie19]), the proof is sound even if the group is maliciously chosen (yet still a valid RSA group), which gives the strong property we need. We note that the proof of soundness for our strong trapdoor VDF in the ABO-string + auxiliary-input random oracle model

follows by a similar argument to that of [Pie19] in the (plain) random oracle model after applying Unruh’s Lemma [Unr07] (stated in Lemma 3.6).

Next, we construct what we refer to as a *one-sided* publicly verifiable time-lock puzzle in the ABO-string model by using the strong trapdoor VDF in the RSW-style construction described above. By one-sided, we mean that completeness and soundness hold only for puzzles in the support of Gen (again, we emphasize that this is in contrast to a randomly sampled puzzle). Then, our full construction applies our non-malleability transformation to a one-sided publicly verifiable time-lock puzzle. We already argued that the non-malleability transformation provides a form of public verifiability for puzzles z in the support of Gen . Namely, anyone can prove to another party that a valid puzzle z has a solution s , but the proof may not be sound when trying to prove that a puzzle has no solution. However, we next show that if the underlying puzzle satisfies one-sided public verifiability, then the resulting (NM) PV TLP is sound for any $z \in \{0, 1\}^*$ (possibly not in the support of Gen).

Proof of full public verifiability. Let $(\text{Gen}, \text{Sol}, \text{Verify})$ be the result of applying our non-malleability transformation to a one-sided PV TLP $(\text{Gen}_{\text{tlp}}, \text{Sol}_{\text{tlp}}, \text{Verify}_{\text{tlp}})$. Consider any puzzle $z \in \{0, 1\}^*$. If z is in the support of Gen , we want to ensure that no one can prove that $s' = \perp$ is a valid solution. At the same time, if z is not in the support of Gen , we want to ensure that no one can prove that $s' \neq \perp$ is a valid solution.

When we run $\text{Sol}(z, t)$, we first run $\text{Sol}_{\text{tlp}}(z, t)$ and get a solution $s_{\text{tlp}} = \hat{s} \parallel \hat{r}$ with a proof π_{tlp} . If \hat{r} is a valid opening for the proposed solution \hat{s} , then Sol can simply output the solution $s = \hat{s}$ and the proof $\pi = \hat{r}$. If \hat{r} is not a valid opening for \hat{s} , Sol must output \perp and a proof π that this is the case. We set $\pi = (s_{\text{tlp}}, \pi_{\text{tlp}})$, which intuitively gives anyone else a way to “shortcut” the computation of Sol_{tlp} .

Now suppose that an adversary tries to falsely convince you that a puzzle z with no solution has a solution $s' \neq \perp$ using a proof $\pi' = r'$. To do so, it must be the case that r' is a valid opening for s' with respect to Gen . But if that were the case, then z *would have a solution*, in contradiction.

On the other hand, suppose that an adversary tries to falsely convince you that a puzzle z with solution s has no solution, i.e. $s' = \perp$, using a proof $\pi' = (s'_{\text{tlp}}, \pi'_{\text{tlp}})$. Since z has a solution, it means that z is in the support of Gen_{tlp} . By one-sided public verifiability, this means that π'_{tlp} is a valid proof that $s'_{\text{tlp}} = \hat{s} \parallel \hat{r}$ is the correct solution to z with respect to Gen_{tlp} . So if \hat{r} is not a valid opening for \hat{s} with respect to Gen , we know the adversary must be lying. In other words, the only way the adversary can cheat is by cheating in the underlying one-sided PV TLP on a puzzle z in the support of Gen_{tlp} .

Discussion of our non-malleable PV TLP. We note that the publicly verifiable time-lock puzzle we described above can be made to satisfy the same non-malleability guarantees as we discuss in Section 2.1 (as we construct it using the same transformation but with a specific underlying time-lock puzzle). Thus, assuming the repeated squaring assumption, we get a publicly verifiable time-lock puzzle that satisfies concurrent function non-malleability for any class of low depth functions \mathcal{F}_m with output length m . Our construction is in the ABO-string model, and we prove security in the auxiliary-input random oracle model (which is needed for soundness of the strong trapdoor VDF in the ABO-string model in addition to the non-malleability transformation). This model is reasonable for our practical applications to multi-party protocols, as we will see below.

We note that our explicit repeated squaring assumption states that repeated squaring in RSA groups for n -bit integers cannot be sped up even by size 2^m adversaries. The repeated squaring

assumption is closely related to the assumption on factoring. The current best known algorithms for factoring run in time at least $2^{n^{1/3}}$. In the case where $m \in O(\log \lambda)$, for example, we only require that polynomial-size attackers cannot speed up repeated squaring, which is a relatively mild assumption. In the case where m is larger, say $m = \lambda$, then we need to choose n to be at least λ^3 (based on known algorithms for factoring). This gives an example of the various trade-offs we get for the security and efficiency of our construction depending on the class of low depth functions \mathcal{F}^m that we want non-malleability for.

2.3 The Fair Multi-Party Protocols

We will focus on coin flipping for concreteness, and note that for auctions the ideas are similar. At a high level, the coin flipping protocol is very simple. Each party chooses a random bit and publishes a time-lock puzzle that encodes the chosen bit. After all puzzles are published, each party open their puzzle by revealing the bit that it used as well as the randomness used to generate the puzzle. Any puzzle that is not opened can be “solved” after a moderately large amount of time t . Once all puzzles have been opened, the agreed upon bit (i.e., the output of the protocol) is the XOR of all revealed bits. The above protocol template is appealing because it naturally satisfies optimistic efficiency: if all parties are honest and open their puzzles, the protocol terminates immediately. When using time-lock puzzles which are both non-malleable (as discussed in Section 2.1) and publicly verifiable (as discussed in Section 2.2), we achieve the following highly desirable properties:

- **Fairness:** No malicious party can bias the output of the protocol.

This crucially relies on non-malleability for the underlying time-lock puzzle. For a protocol with n participants, we need the time-lock puzzle to satisfy n -concurrent non-malleability. This guarantees that as long as one party is honest, the output of the protocol will be (at least statistically close to) a uniformly random bit.

- **Unbounded participants:** Anyone can participate in the protocol.

This property might come at a surprise since we show fully concurrent non-malleability is impossible to achieve. However, we emphasize that our time-lock puzzle achieves fully concurrent *functional* non-malleability for the XOR function. This allows us to deal with any a priori unbounded number of participants, which is important in many decentralized and distributed settings.

- **Public verifiability:** Only one party needs to solve each unopened puzzle, and can provide a publicly verifiable proof that it solved it correctly.

This follows immediately by the public verifiability property we achieve for the underlying time-lock puzzle. Without this property, any unopened puzzles may need to be solved by every party that want to know the output of the protocol, which is prohibitively expensive. However, public verifiability instead opens up the application to *any* party, not even involved in the protocol. Furthermore, this work can even be delegated to an external server since trust is guaranteed by the attached proof.

We note that our non-malleable and publicly verifiable time-lock puzzle is defined in the All-But-One-string (ABO-string) model, which is required for public verifiability. To implement this model, we have each participant i publish a fresh random string $\text{crs}_i \leftarrow \{0, 1\}^\lambda$ in addition to its puzzle z_i . Then, whenever some party tries to solve (or verify) a puzzle, it puts all of the random strings together as a multi-common random string $\text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n)$ from all n participants, and uses this for the publicly verifiable proof. As long as a single party is honest and publishes a

random string crs_i independent of all other participants, then the publicly verifiable proof system will be sound.

3 Preliminaries

We denote by $x \leftarrow X$ the process of sampling a value x from the distribution X . For a set \mathcal{X} , we denote by $x \leftarrow \mathcal{X}$ the process of sampling a value x from the uniform distribution on \mathcal{X} . We let $\text{Supp}(X)$ denote the support of the distribution X . For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, 2, \dots, n\}$. We use PPT as an acronym for *probabilistic polynomial time*. A function $\text{negl}: \mathbb{N} \rightarrow \mathbb{R}$ is *negligible* if it is asymptotically smaller than any inverse-polynomial function, namely, for every constant $c > 0$ there exists an integer N_c such that $\text{negl}(\lambda) \leq \lambda^{-c}$ for all $\lambda > N_c$. Two sequences of random variables $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are *computationally indistinguishable* if for any non-uniform PPT algorithm \mathcal{A} there exists a negligible function negl such that $|\Pr[\mathcal{A}(1^\lambda, X_\lambda) = 1] - \Pr[\mathcal{A}(1^\lambda, Y_\lambda) = 1]| \leq \text{negl}(\lambda)$ for all $\lambda \in \mathbb{N}$.

A non-uniform algorithm $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ is a sequence of circuits for all $\lambda \in \mathbb{N}$. We assume that \mathcal{A}_λ always implicitly receives 1^λ as its first input. We define $\text{size}(\mathcal{A}_\lambda)$ to be the size of the circuit (corresponding to total time) and $\text{depth}(\mathcal{A}_\lambda)$ to be the depth of the circuit (corresponding to parallel time). A uniform circuit family $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ is a sequence of circuits with a constant-size description. Formally, there exists a constant-size polynomial-time Turing machine M that on input 1^λ outputs \mathcal{C}_λ for all $\lambda \in \mathbb{N}$. We also consider circuit families indexed by multiple inputs, defined similarly. Let RF_λ be the set of all functions for a security parameter λ (exact parameters will be given in the construction). For a partial assignment F , we define $\text{RF}_\lambda[F]$ to be the set of functions consistent with F .

3.1 Time-Lock Puzzles

We first define the standard definition of time-lock puzzles without any additional properties.

Definition 3.1. *Let $B: \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon \in (0, 1)$. A (B, ϵ) -secure time-lock puzzle (TLP) is a tuple (Gen, Sol) with the following syntax:*

- $z \leftarrow \text{Gen}(1^\lambda, t, s)$: *A PPT algorithm that on input a security parameter $\lambda \in \mathbb{N}$, a difficulty parameter $t \in \mathbb{N}$, and a solution $s \in \{0, 1\}^\lambda$, outputs a puzzle $z \in \{0, 1\}^*$.*
- $s = \text{Sol}(1^\lambda, t, z)$: *A deterministic algorithm that on input a security parameter $\lambda \in \mathbb{N}$, a difficulty parameter $t \in \mathbb{N}$, and a puzzle $z \in \{0, 1\}^*$, outputs a solution $s \in (\{0, 1\}^\lambda \cup \{\perp\})$.*

We require (Gen, Sol) to satisfy the following properties.

- **Correctness:** *For every $\lambda, t \in \mathbb{N}$, solution $s \in \{0, 1\}^\lambda$, and $z \in \text{Supp}(\text{Gen}(1^\lambda, t, s))$, it holds that $\text{Sol}(1^\lambda, t, z) = s$.*
- **Efficiency:** *There exist a polynomial p such that for all $\lambda, t \in \mathbb{N}$, $\text{Sol}(1^\lambda, t, \cdot)$ is computable in time $t \cdot p(\lambda, \log t)$.*
- **(B, ϵ) -Hardness:** *There exist positive polynomials p, ℓ such that for all functions T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$ and every non-uniform distinguisher $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ where $\text{size}(\mathcal{A}_\lambda) \in \text{poly}(B(\lambda))$ and $\text{depth}(\mathcal{A}_\lambda) \leq (T(\lambda))^\epsilon \cdot p(\lambda)$ for all $\lambda \in \mathbb{N}$, there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$, and $s, s' \in \{0, 1\}^\lambda$,*

$$\left| \Pr[\mathcal{A}_\lambda(\text{Gen}(1^\lambda, T(\lambda), s)) = 1] - \Pr[\mathcal{A}_\lambda(\text{Gen}(1^\lambda, T(\lambda), s')) = 1] \right| \leq \text{negl}(\lambda),$$

where the probabilities are over the randomness of Gen and \mathcal{A}_λ .

In the above definition, we assume for simplicity that the solutions s are λ -bits long. We can naturally generalize this to consider the case where solutions have some specified length $L(\lambda)$.

We discuss the definition of (B, ϵ) -hardness in comparison with the definition given by Bitansky et al. [BGJ⁺16]. First, they consider only polynomial B , whereas we expand this notion to possibly allow for possibly super-polynomial functions B . Second, their definition requires that the polynomial p satisfies $p(\lambda) = 1$ for all $\lambda \in \mathbb{N}$. We instead allow for any polynomial $p(\lambda)$ multiplicative slack, which is independent of the time bound T , in the depth of \mathcal{A}_λ . We believe that the dependence on T is the more intuitive and more important quantity to consider (as $\text{poly}(\lambda)$ factors may be more dependent on the specific implementation), and furthermore, we believe our definition is easier to work with. We additionally note that if a time-lock puzzle satisfies (B, ϵ) -hardness when restricted to the polynomial $p(\lambda) = 1$ then it satisfies the full definition of (B, ϵ') -hardness for any constant $\epsilon' < \epsilon$.⁷

3.2 Non-Malleable Time-Lock Puzzles

To formalize non-malleability in the context of time-lock puzzles, we introduce a Man-In-the-Middle (MIM) adversary. Because time-lock puzzles are designed to be broken in some depth t , we restrict our MIM adversary to have at most depth $t^\epsilon \cdot \text{poly}(\lambda)$ for some $\epsilon \in (0, 1)$. Furthermore, we allow for concurrent MIM adversaries that possibly interact with many senders and receivers at the same time.

Definition 3.2 (MIM Adversaries). *Let $n_L, n_R, B, T, p: \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon \in (0, 1)$. An $(n_L, n_R, B, \epsilon, T, p)$ -Man-In-the-Middle (MIM) adversary is a non-uniform algorithm $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ satisfying $\text{depth}(\mathcal{A}_\lambda) \leq (T(\lambda))^\epsilon \cdot p(\lambda)$ and $\text{size}(\mathcal{A}_\lambda) \in \text{poly}(B(\lambda))$ for all $\lambda \in \mathbb{N}$ that receives $n_L(\lambda)$ puzzles on the left and outputs $n_R(\lambda)$ puzzles on the right.*

We next define the MIM distribution, which corresponds to the values underlying the puzzles output by the MIM adversary. To capture adversaries that simply forward one of the puzzles on the left to a receiver on the right, we set the value for any forwarded puzzle to be \perp .

Definition 3.3 (MIM Distribution). *Let $n_L, n_R, B, T, p: \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon \in (0, 1)$. Let $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ be an $(n_L, n_R, B, \epsilon, T, p)$ -MIM adversary. For any $\lambda \in \mathbb{N}$ and $\vec{s} = (s_1, \dots, s_{n_L(\lambda)}) \in (\{0, 1\}^\lambda)^{n_L(\lambda)}$, we define the distribution*

$$(\tilde{s}_1, \dots, \tilde{s}_{n_R(\lambda)}) \leftarrow \text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), \vec{s})$$

as follows. \mathcal{A}_λ receives puzzles $z_i \leftarrow \text{Gen}(1^\lambda, T(\lambda), s_i)$ for all $i \in [n_L(\lambda)]$ and outputs puzzles $(\tilde{z}_1, \dots, \tilde{z}_{n_R(\lambda)})$. Then for each $i \in [n_R(\lambda)]$, we define

$$\tilde{s}_i = \begin{cases} \perp & \text{if there exists a } j \in [n_L(\lambda)] \text{ such that } \tilde{z}_i = z_j, \\ s & \text{if there exists a unique } s \text{ such that } \tilde{z}_i = \text{Gen}(1^\lambda, T(\lambda), s; r) \text{ for some } r, \text{ or} \\ \perp & \text{otherwise.} \end{cases}$$

Intuitively, a time-lock puzzle is non-malleable if the MIM distribution of a bounded depth attacker does not depend on the solutions underlying the puzzles it receives on the left. We formalize this definition below.

⁷This follows by considering a large enough polynomial ℓ satisfying $\ell(\lambda) \geq p(\lambda)^{1/(\epsilon - \epsilon')}$ so it holds that $(T(\lambda))^{\epsilon'} \cdot p(\lambda) \leq (T(\lambda))^\epsilon$.

Definition 3.4 (Concurrent Non-malleable). Let $n_L, n_R, B, T, p: \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon \in (0, 1)$. A (B, ϵ) -secure time-lock puzzle (Gen, Sol) is (n_L, n_R) -concurrent non-malleable if there exist positive polynomials p, ℓ such that for every function T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$ and every $(n_L, n_R, B, \epsilon, T, p)$ -MIM adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, the following holds.

For any distinguisher \mathcal{D} , there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$ and $\vec{s} = (s_1, \dots, s_{n_L(\lambda)}) \in (\{0, 1\}^\lambda)^{n_L(\lambda)}$,

$$\left| \Pr \left[\mathcal{D}(\text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), \vec{s})) = 1 \right] - \Pr \left[\mathcal{D}(\text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), (0^\lambda)^{n_L(\lambda)})) = 1 \right] \right| \leq \text{negl}(\lambda).$$

We consider standard variants for the definition of non-malleable above.

Definition 3.5. Let (Gen, Sol) be a (B, ϵ) -secure time-lock puzzle. We say the following non-malleability properties are satisfied when Definition 3.4 holds against $(n_L, n_R, B, \epsilon, T, p)$ -MIM adversaries for the following settings of n_L and n_R :

- fully concurrent non-malleable if the definition holds against any $m_L, m_R \in \text{poly}(\lambda)$,
- one-many non-malleable if the definition holds for any $n_R(\lambda) \in \text{poly}(\lambda)$ and $n_L = 1$,
- n -concurrent non-malleable if the definition holds for $n_L = n_R = n$,
- one- n non-malleable for $n_L(\lambda) = 1$ and $n_R = n$,
- and simply non-malleable (not concurrent) for $n_L(\lambda) = n_R(\lambda) = 1$.

3.3 Time-Lock Puzzles in the Auxiliary-Input Random Oracle Model

In the random oracle model [BR93], security is proven in the case where all relevant parties have oracle access to a common random function. For security in the (plain) random oracle model, it is assumed that a fixed adversary is independent of the common random function and must break security with probability over the choice of the random function. Security in the auxiliary-input random oracle model, introduced by [Unr07], is proven assuming that the attacker may depend *arbitrarily* on the random oracle, as long as the attacker is of polynomial (or bounded) size.

One of the main benefits of proving security in the random oracle model is that you can argue security when the random oracle is sampled “lazily.” The following lemma, adapted from Unruh [Unr07], shows how we can use similar lazy sampling techniques in the auxiliary-input random oracle model without significant loss in security. At a high level, it says that for any computationally bounded adversary \mathcal{A} in the AI-ROM model, we can “switch” a random \mathcal{O} (that the attacker may depend arbitrarily on) with an oracle \mathcal{P} that is lazily sampled on almost all points. In other words, \mathcal{A} cannot distinguish the output of \mathcal{O} from \mathcal{P} on a *random* query.

Lemma 3.6 (Unruh [Unr07]). For any functions $p, f, \lambda \in \mathbb{N}$, and unbounded algorithm Z that on input a random oracle $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$ outputs $p(\lambda)$ -size oracle machines, there is an (inefficient) algorithm Sam that outputs a partial assignment F on $f(\lambda)$ points, such that for any $\lambda \in \mathbb{N}$ and unbounded distinguisher \mathcal{D} ,

$$\left| \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ y \leftarrow \mathcal{A}^{\mathcal{O}} \end{array} : \mathcal{D}(y) = 1 \right] - \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ F = \text{Sam}(\mathcal{A}) \\ \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ y \leftarrow \mathcal{A}^{\mathcal{P}} \end{array} : \mathcal{D}(y) = 1 \right] \right| \leq \sqrt{\frac{(p(\lambda))^2}{f(\lambda)}}.$$

We now formalize non-malleable time-lock puzzles in the auxiliary-input random oracle model. As the AI-ROM only affects the security properties against computationally bounded adversaries, the syntax, correctness, efficiency, and completeness properties are left relatively unchanged. As a result, we focus on the definitions of hardness and non-malleability. In the AI-ROM, we model computationally bounded adversaries that are allowed to depend arbitrarily on the random oracle. To formalize this, we consider inefficient algorithms Z that, for any $\lambda \in \mathbb{N}$, take as input a random oracle $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$ (of exponential size) and output circuits of bounded size. Additionally, we consider a MIM distribution for a $(n, n, B, \epsilon, T, p)$ -oracle adversary \mathcal{A} , $\text{mim}_{\mathcal{A}}^{\mathcal{O}}$, where \mathcal{A} and the distribution have oracle access to the same oracle \mathcal{O} .

Definition 3.7. *Let $B, n: \mathbb{N} \rightarrow \mathbb{N}$, and $\epsilon \in (0, 1)$ be a constant. A (B, ϵ) -secure n -concurrent non-malleable time-lock puzzle in the AI-ROM is a tuple of oracle algorithms (Gen, Sol) that satisfy correctness, efficiency, and completeness relative to any $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$, and the following:*

- **(B, ϵ) -Hardness:** *There exists positive polynomials p, ℓ such that for any function T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$ and unbounded algorithm Z that on input $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$ outputs circuits of size $\text{poly}(B(\lambda))$ and depth at most $(T(\lambda))^\epsilon \cdot p(\lambda)$, there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$ and values $s_0, s_1 \in \{0, 1\}^\lambda$,*

$$\left| \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} \leftarrow Z(\mathcal{O}) \\ z \leftarrow \text{Gen}^{\mathcal{O}}(1^\lambda, T(\lambda), s_0) \end{array} : \mathcal{A}^{\mathcal{O}}(z) = 1 \right] - \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ z \leftarrow \text{Gen}^{\mathcal{O}}(1^\lambda, T(\lambda), s_1) \end{array} : \mathcal{A}^{\mathcal{O}}(z) = 1 \right] \right| \leq \text{negl}(\lambda).$$

- **Concurrent Non-malleable:** *There exist positive polynomials p, ℓ such that for every function T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$, polynomials n_L, n_R , and unbounded algorithm Z that on input $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$ outputs an $(n_L, n_R, B, \epsilon, T, p)$ -MIM adversary, the following holds.*

For any distinguisher \mathcal{D} , there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$ and $\vec{s} = (s_1, \dots, s_{n(\lambda)}) \in (\{0, 1\}^\lambda)^{n_L(\lambda)}$,

$$\left| \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ \vec{s} \leftarrow \text{mim}_{\mathcal{A}}^{\mathcal{O}}(1^\lambda, T(\lambda), \vec{s}) \end{array} : \mathcal{D}(\vec{s}) = 1 \right] - \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} \leftarrow Z(\mathcal{O}) \\ \vec{s} \leftarrow \text{mim}_{\mathcal{A}}^{\mathcal{O}}(1^\lambda, T(\lambda), (0^\lambda)^{n(\lambda)}) \end{array} : \mathcal{D}(\vec{s}) = 1 \right] \right| \leq \text{negl}(\lambda).$$

4 Non-Malleable Time-Lock Puzzles

In Section 4.1, we define the notion of concurrent functional non-malleability for a class of functions \mathcal{F} . Then, in Section 4.2, we give a transformation from any time-lock puzzle to one that satisfies concurrent functional non-malleable for depth-bounded functions \mathcal{F} , in the auxiliary-input random oracle model. The security of our construction depends on the output length of the functions $f \in \mathcal{F}$. We discuss how this general result for concurrent functional non-malleability implies our result for bounded concurrent (standard) non-malleability. In Section 4.4, we show that no time-lock puzzle satisfies fully concurrent (standard) non-malleability.

4.1 Functional Non-Malleable Time-Lock Puzzles

We formally define concurrent *functional* non-malleability. For simplicity, we define it in the case of unbounded concurrency, but we note that it can be defined for restricted cases as in Section 3.2.

Definition 4.1 (Concurrent Functional Non-malleable). *Let $B, L: \mathbb{N} \rightarrow \mathbb{N}$, $\epsilon \in (0, 1)$, and (Gen, Sol) be a (B, ϵ) -secure time-lock puzzle (Gen, Sol) for messages of length $L(\lambda)$. Let \mathcal{F} be a class of functions of the form $f: (\{0, 1\}^{L(\lambda)})^* \rightarrow \{0, 1\}^*$. We say that (Gen, Sol) is concurrent functional non-malleable for \mathcal{F} if for any function $f \in \mathcal{F}$ and polynomial n , there exist positive polynomials p, ℓ such that for every function T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$, every $(n, n, B, \epsilon, T, p)$ -MIM adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, the following holds.*

For any distinguisher \mathcal{D} , there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$ and $\vec{s} = (s_1, \dots, s_{n(\lambda)}) \in (\{0, 1\}^{L(\lambda)})^{n(\lambda)}$,

$$\left| \Pr \left[\vec{s} \leftarrow \text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), \vec{s}) : \mathcal{D}(f(\vec{s})) = 1 \right] - \Pr \left[\vec{s} \leftarrow \text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), (0^{L(\lambda)})^{n(\lambda)}) : \mathcal{D}(f(\vec{s})) = 1 \right] \right| \leq \text{negl}(\lambda).$$

We note that functional non-malleability for a class \mathcal{F} that contains the identity function id implies standard non-malleability as $\mathcal{D}(\text{id}(\vec{s})) = \mathcal{D}(\vec{s})$.

4.2 Non-Malleable Time-Lock Puzzles Construction

In this section, we give our construction of a fully concurrent functional non-malleable time-lock puzzle for functions with bounded depth and output length. We rely on the following building blocks and parameters.

- A function m denoting the output length for our function non-malleability. We require $m(\lambda) \in \lambda^{O(1)}$.
- A (B, ϵ) -secure time-lock puzzle $\text{TLP} = (\text{Gen}_{\text{tlp}}, \text{Sol}_{\text{tlp}})$. We let $\lambda_{\text{tlp}} = \lambda_{\text{tlp}}(\lambda) = m(\lambda) + \lambda$ be the bits of randomness needed for TLP on security parameter λ , and we let $\lambda + \lambda_{\text{tlp}}$ be the length of the solutions for TLP. We require that the time-lock puzzle satisfies $B(\lambda) \in 2^{m(\lambda)} \cdot \text{poly}(\lambda)$. We note that our function Gen defined below will also require λ_{tlp} bits of randomness.
- A class of functions \mathcal{F}_B^m of the form $f: (\{0, 1\}^\lambda)^* \rightarrow \{0, 1\}^{m(\lambda)}$. We assume that there exists a polynomial d such that for every polynomial n , every function $f \in \mathcal{F}_B^m$ can be computed in depth $d(\lambda, \log n(\lambda))$ and size $\text{poly}(B(\lambda))$ on inputs of length at most $\lambda \cdot n(\lambda)$.
- A random oracle $\mathcal{O} \in \text{RF}_\lambda$, where \mathcal{O} on input $(1^\lambda, t, s, r) \in \{0, 1\}^{3\lambda + \lambda_{\text{tlp}}}$ outputs a random value $r' \in \{0, 1\}^{\lambda_{\text{tlp}}}$.

Our construction $\text{nmTLP}_m = (\text{Gen}, \text{Sol})$ in the auxiliary-input random oracle model:

- $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s; r)$:
 1. Get $r' = \mathcal{O}(1^\lambda, t, s, r)$.
 2. Output $z \leftarrow \text{Gen}_{\text{tlp}}(\lambda, t, (s||r); r')$.
- $s \leftarrow \text{Sol}^{\mathcal{O}}(1^\lambda, t, z)$:
 1. Compute $s' = \text{Sol}_{\text{tlp}}(1^\lambda, t, z)$ and parse $s' = s||r$.

2. If $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s; r)$, output s .
3. If not, output \perp .

Theorem 4.2 (Fully Concurrent Functional Non-Malleable TLPs). *Let $B: \mathbb{N} \rightarrow \mathbb{N}$, $m(\lambda) \in \lambda^{O(1)}$, and $\epsilon \in (0, 1)$, where $\lambda \leq B(\lambda) \in 2^m \cdot \text{poly}(\lambda)$. Assuming TLP is a (B, ϵ) -secure time-lock puzzle, then nmTLP_m is a (B, ϵ) -secure fully concurrent functional non-malleable time-lock puzzle in the AI-ROM for the class of functions \mathcal{F}_B^m .*

We provide the proof of this theorem in Section 4.3 below. Before doing so, we observe the following corollaries to the above theorem:

- If $m(\lambda) \in O(\log(\lambda))$ then we can simply assume a polynomially-secure TLP.
- For any $m(\lambda) \in \text{poly}(\lambda)$, then our theorem follows by assuming a sub-exponentially secure TLP. Specifically, it suffices that there exists a constant $\gamma \in (0, 1)$ such that $B(\lambda) = 2^{\lambda^\gamma}$, and we can instantiate this with $\lambda_{\text{tlp}} = (\lambda \cdot m(\lambda))^{1/\gamma}$ bits of randomness.

We also observe that the above theorem can be used to get n -bounded concurrency for any polynomial n , simply by setting the output length m of the functions in \mathcal{F}_B^m to $\lambda \cdot n(\lambda)$. Specifically, let f_{id} be the identity function with input and output length $\lambda \cdot n(\lambda)$. We observe that since $f_{\text{id}} \in \mathcal{F}_B^n$, a fully concurrent functional non-malleable TLP for \mathcal{F}_B^n implies an n -concurrent non-malleable TLP, which gives the following corollary.

Corollary 4.3 (n -Concurrent Non-Malleable TLPs). *Let $B: \mathbb{N} \rightarrow \mathbb{N}$, $n(\lambda) \in \lambda^{O(1)}$, and $\epsilon \in (0, 1)$, where $\lambda \leq B(\lambda) \in 2^{\lambda \cdot n(\lambda)} \cdot \text{poly}(\lambda)$. Assuming TLP_n is a (B, ϵ) -secure time-lock puzzle, then nmTLP_n is a (B, ϵ) -secure n -concurrent non-malleable time-lock puzzle in the AI-ROM.*

4.3 Proof of Concurrent Functional Non-Malleability

We prove Theorem 4.2 by showing correctness, efficiency, hardness, and non-malleability below.

Lemma 4.4 (Correctness). *Assuming $(\text{Gen}_{\text{tlp}}, \text{Sol}_{\text{tlp}})$ satisfies correctness, then (Gen, Sol) satisfies correctness.*

Proof. Fix any $\lambda, t \in \mathbb{N}$ and $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$. We want to show that

1. For $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s; r)$ for some $s \in \{0, 1\}^\lambda$ and $r \in \{0, 1\}^{\lambda_{\text{tlp}}}$, it holds that $\text{Sol}^{\mathcal{O}}(1^\lambda, t, z) = s$, and
2. For z not in the support of $\text{Gen}(1^\lambda, t, \cdot)$, it holds that $\text{Sol}^{\mathcal{O}}(1^\lambda, t, z) = \perp$.

The first condition follows from correctness of TLP. Recall that the algorithm $\text{Sol}^{\mathcal{O}}(1^\lambda, t, z)$ runs $\widehat{s} \parallel \widehat{r} = \text{Sol}_{\text{tlp}}(1^{\lambda_{\text{tlp}}}, t, z)$ and outputs \widehat{s} if $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, \widehat{s}; \widehat{r})$. Since we assumed z was in the support of $\text{Gen}^{\mathcal{O}}$, it follows by definition of $\text{Gen}^{\mathcal{O}}$ that $z = \text{Gen}_{\text{tlp}}(1^{\lambda_{\text{tlp}}}, t, (s \parallel r); r')$ for some s, r, r' . By correctness of TLP, this implies that $\text{Sol}_{\text{tlp}}(1^{\lambda_{\text{tlp}}}, t, z) = s \parallel r$ so $s \parallel r = \widehat{s} \parallel \widehat{r}$ and Sol outputs the correct solution.

For the second condition, since z is not in the support of $\text{Gen}^{\mathcal{O}}$, then for any $s' = (s \parallel r)$ given by $\text{Sol}_{\text{tlp}}(1^\lambda, t, z)$, it cannot be the case that $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s; r)$. Therefore, $\text{Sol}^{\mathcal{O}}(1^\lambda, t, z)$ outputs \perp , as required. \square

Lemma 4.5 (Efficiency). *Assuming that $(\text{Gen}_{\text{tlp}}, \text{Sol}_{\text{tlp}})$ satisfies efficiency, then (Gen, Sol) satisfies efficiency.*

Proof. Fix any $\lambda, t \in \mathbb{N}$ and $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$. First, we show that $\text{Gen}^\mathcal{O}$ is PPT. We have that $\text{Gen}^\mathcal{O}(1^\lambda, t, \cdot)$ makes a single oracle call to \mathcal{O} , which takes time $\lambda_{\text{tlp}} = \lambda + m$ to read the output, and evaluates $\text{Gen}_{\text{tlp}}(1^\lambda, t, \cdot)$, which takes time $\text{poly}(\lambda, \log t)$. It follows that there exists a polynomial p_1 such that $\text{Gen}^\mathcal{O}(1^\lambda, t, \cdot)$ can be computed in time $p_1(\lambda, \log t)$.

For $\text{Sol}^\mathcal{O}$, recall that $\text{Sol}^\mathcal{O}(1^\lambda, t, \cdot)$ runs $\text{Sol}_{\text{tlp}}(1^\lambda, t, \cdot)$ and $\text{Gen}^\mathcal{O}(1^\lambda, t, \cdot)$. By the above, $\text{Gen}^\mathcal{O}$ can be run in time $p_1(\lambda, \log t)$. For Sol_{tlp} , recall that by the efficiency of TLP there exists a polynomial p_2 such that for all $\lambda, t \in \mathbb{N}$, $\text{Sol}_{\text{tlp}}(1^\lambda, t, \cdot)$ can be computed in time $t \cdot p_2(\lambda, \log t)$. Putting these together, we have that $\text{Sol}^\mathcal{O}(1^\lambda, t, \cdot)$ can be computed in time $p_1(\lambda, \log t) + t \cdot p_2(\lambda, \log t)$. \square

Lemma 4.6 (Hardness). *Assuming that $(\text{Gen}_{\text{tlp}}, \text{Sol}_{\text{tlp}})$ satisfies (B, ϵ) -hardness, then (Gen, Sol) satisfies (B, ϵ) -hardness in the AI-ROM.*

Proof. To show hardness, suppose for contradiction that for all positive polynomials p, ℓ , there exists a function T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$, an unbounded algorithm Z that outputs circuits of size $\text{poly}(B(\lambda))$ and depth at most $(T(\lambda))^\epsilon \cdot p(\lambda)$, and a polynomial q such that for infinitely many $\lambda \in \mathbb{N}$, there exist values $s_0, s_1 \in \{0, 1\}^\lambda$ such that

$$\left| \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{D} = Z(\mathcal{O}) \\ z \leftarrow \text{Gen}^\mathcal{O}(1^\lambda, T(\lambda), s_0) \end{array} : \mathcal{D}^\mathcal{O}(z) = 1 \right] - \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{D} = Z(\mathcal{O}) \\ z \leftarrow \text{Gen}^\mathcal{O}(1^\lambda, T(\lambda), s_1) \end{array} : \mathcal{D}^\mathcal{O}(z) = 1 \right] \right| \geq \frac{1}{q(\lambda)}. \quad (4.1)$$

As the above holds for all p, ℓ by assumption, we will give specific polynomials p and ℓ and use them to reach a contradiction. Specifically, let $p_{\text{tlp}}, \ell_{\text{tlp}}$ be the polynomials guaranteed by the hardness of TLP. We will show a contradiction to the above for $p(\lambda) = p_{\text{tlp}}(\lambda)/p'(\lambda)$ and $\ell(\lambda) = \ell_{\text{tlp}}(\lambda)$, where p' is a fixed polynomial specified in the proof of Claim 4.10.

To derive a contradiction, we will define a sequence of hybrid experiments, and we will use the fact that the statistical distance between the above probabilities is noticeable to construct an adversary that breaks the hiding of TLP. For any $\lambda \in \mathbb{N}$ and $s \in \{0, 1\}^\lambda$, we define the following sequence of hybrid experiments. Throughout these hybrids, we let $t = T(\lambda)$ and $m = m(\lambda)$.

- $\text{Hyb}_0^s(\lambda)$: This hybrid is equivalent to the terms in the probabilities above for a puzzle generated with solution $s \in \{0, 1\}^\lambda$, where $\text{Gen}^\mathcal{O}$ is written out explicitly.

$$\text{Hyb}_0^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{D} = Z(\mathcal{O}) \\ r \leftarrow \{0, 1\}^{\lambda_{\text{tlp}}}; \quad r' = \mathcal{O}(1^\lambda, t, s, r) : \mathcal{D}^\mathcal{O}(z) = 1 \\ z = \text{Gen}_{\text{tlp}}(1^\lambda, t, (s||r); r') \end{array} \right\}$$

- $\text{Hyb}_1^s(\lambda)$: Let Z' be the algorithm such that for any $\mathcal{O} \in \text{RF}_\lambda$ and $\mathcal{D} = Z(\mathcal{O})$, the algorithm $Z'(\mathcal{D})$ outputs an oracle algorithm $\mathcal{A}^\mathcal{O}$ which does the following:

1. Sample $r \leftarrow \{0, 1\}^{\lambda_{\text{tlp}}}$ and query $r' = \mathcal{O}(1^\lambda, t, s, r)$.
2. Compute $z = \text{Gen}_{\text{tlp}}(1^\lambda, t, (s||r); r')$.
3. Output $\mathcal{D}^\mathcal{O}(z)$.

This hybrid uses \mathcal{A} to determine the output of the experiment.

$$\text{Hyb}_1^s(\lambda) = \left\{ \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{D} = Z(\mathcal{O}); \quad \mathcal{A} = Z'(\mathcal{D}) : \mathcal{A}^\mathcal{O} = 1 \right\}$$

Note that $\text{size}(\mathcal{A}) \in \text{poly}(\lambda, 2^m)$, which holds by the efficiency of Gen and because $\lambda_{\text{tlp}} = m + \lambda$ and $\text{size}(\mathcal{D}) \in \text{poly}(B(\lambda))$.

- $\text{Hyb}_2^s(\lambda)$: Let $q_{\mathcal{A}}$ be the polynomial such that $\text{size}(\mathcal{A}) = q_{\mathcal{A}}(\lambda, 2^m)$. Let $f_{\text{hard}}(\lambda) = (q_{\mathcal{A}}(\lambda, 2^m) \cdot 4q(\lambda))^2$, and let Sam be the inefficient algorithm from Lemma 3.6 that on input an adversary \mathcal{A} outputs a partial assignment F on $f_{\text{hard}}(\lambda)$ points. This hybrid swaps \mathcal{O} with a random function \mathcal{P} fixed at the set of points F determined by Sam .

$$\text{Hyb}_2^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_{\lambda}; \quad \mathcal{D} = Z(\mathcal{O}); \quad \mathcal{A} = Z'(\mathcal{D}) \\ F \leftarrow \text{Sam}(\mathcal{A}); \quad \mathcal{P} \leftarrow \text{RF}_{\lambda}[F] \end{array} : \mathcal{A}^{\mathcal{P}} = 1 \right\}$$

- $\text{Hyb}_3^s(\lambda)$: This hybrid samples r' uniformly randomly from $\{0, 1\}^{\lambda_{\text{tp}}}$, rather than using \mathcal{P} to compute it. It also gives \mathcal{D} oracle access to a modified version of \mathcal{P} , where on input $(1^\lambda, t, s, r)$ it outputs r' , and agrees with \mathcal{P} on all other inputs. We denote this oracle by $\mathcal{P}[(1^\lambda, t, s, r) \rightarrow r']$. In this hybrid, we additionally write out \mathcal{A} 's actions explicitly.

$$\text{Hyb}_3^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_{\lambda}; \quad \mathcal{D} = Z(\mathcal{O}); \quad \mathcal{A} = Z'(\mathcal{D}) \\ F \leftarrow \text{Sam}(\mathcal{A}); \quad \mathcal{P} \leftarrow \text{RF}_{\lambda}[F] \\ r \leftarrow \{0, 1\}^{\lambda_{\text{tp}}}; \quad r' \leftarrow \{0, 1\}^{\lambda_{\text{tp}}} \\ z = \text{Gen}_{\text{tlp}}(1^\lambda, t, (s||r); r') \end{array} : \mathcal{D}^{\mathcal{P}[(1^\lambda, t, s, r) \rightarrow r']}(z) = 1 \right\}$$

To conclude the proof, fix any $\lambda \in \mathbb{N}$ and values $s_0, s_1 \in \{0, 1\}^\lambda$ for which Equation 4.1 holds. Since $\text{Hyb}_0^{s_b}(\lambda)$ is equal to the probability in Equation 4.1 for value s_b , we get that

$$|\text{Hyb}_0^{s_0}(\lambda) - \text{Hyb}_0^{s_1}(\lambda)| \geq \frac{1}{q(\lambda)}.$$

In the following claims, we bound the distance between each pair of consecutive hybrids. Specifically, in Claim 4.7 we show that $\Pr[\text{Hyb}_0^s(\lambda)] = \Pr[\text{Hyb}_1^s(\lambda)]$ for all $s \in \{0, 1\}^\lambda$. Then, in Claim 4.8, we bound the statistical distance between $\text{Hyb}_1^s(\lambda)$ and $\text{Hyb}_2^s(\lambda)$ by $\frac{1}{8q(\lambda)}$ for all $s \in \{0, 1\}^\lambda$. Then, we show in Claim 4.9 that there exists a negligible function negl such that $|\Pr[\text{Hyb}_2^s(\lambda)] - \Pr[\text{Hyb}_3^s(\lambda)]| \leq \text{negl}(\lambda)$ for all $s \in \{0, 1\}^\lambda$. Given these claims, it follows that the statistical distance between $\text{Hyb}_0^{s_b}(\lambda)$ and $\text{Hyb}_3^{s_b}(\lambda)$ is at most $2/(8q(\lambda)) + 2\text{negl}(\lambda)$ for $b \in \{0, 1\}$. By combining this with the above, it holds that $|\Pr[\text{Hyb}_3^{s_0}(\lambda)] - \Pr[\text{Hyb}_3^{s_1}(\lambda)]| \geq \frac{1}{2q(\lambda)}$. To conclude the proof, we show in Claim 4.10 that this implies an adversary that breaks the hardness of TLP with probability $1/2q(\lambda) - \text{negl}(\lambda) \geq 1/4q(\lambda)$, in contradiction.

Claim 4.7. For all $s \in \{0, 1\}^\lambda$, $\Pr[\text{Hyb}_0^s(\lambda)] = \Pr[\text{Hyb}_1^s(\lambda)]$.

Proof. This follows immediately from the definition of \mathcal{A} . Specifically, in $\text{Hyb}_1^s(\lambda)$, the algorithm \mathcal{A} samples r, r' , and z exactly as done in $\text{Hyb}_0^s(\lambda)$, and then outputs 1 exactly in the event that $\text{Hyb}_0^s(\lambda)$ holds. ■

Claim 4.8. For all $s \in \{0, 1\}^\lambda$,

$$|\Pr[\text{Hyb}_1^s(\lambda)] - \Pr[\text{Hyb}_2^s(\lambda)]| \geq 1/(4q(\lambda)).$$

Proof. Recall that $\text{size}(\mathcal{A}) = q_{\mathcal{A}}(\lambda, 2^m)$. We can therefore apply Lemma 3.6 (by viewing Z and Z' as a single sampling algorithm outputting \mathcal{A} based on \mathcal{O}). Therefore, the statistical distance between the output of $\mathcal{A}^{\mathcal{O}}$ and $\mathcal{A}^{\mathcal{P}}$ is at most

$$\sqrt{\frac{(q_{\mathcal{A}}(\lambda, 2^m))^2}{f_{\text{hard}}(\lambda)}} = \sqrt{\frac{(q_{\mathcal{A}}(\lambda, 2^m))^2}{(q_{\mathcal{A}}(\lambda, 2^m) \cdot 4q(\lambda))^2}} = \frac{1}{4q(\lambda)},$$

which implies the claim. ■

Claim 4.9. *There exists a negligible function negl_2 such that for all $s \in \{0, 1\}^\lambda$,*

$$|\Pr[\text{Hyb}_2^s(\lambda)] - \Pr[\text{Hyb}_3^s(\lambda)]| \leq \text{negl}_2(\lambda).$$

Proof. We start by using the definition of \mathcal{A} to rewrite $\text{Hyb}_2^s(\lambda)$ as

$$\text{Hyb}_2^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{D} = Z(\mathcal{O}); \quad \mathcal{A} = Z'(\mathcal{D}) \\ F \leftarrow \text{Sam}(\mathcal{A}); \quad \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ r \leftarrow \{0, 1\}^{\lambda_{\text{tip}}}; \quad r' = \mathcal{P}(1^\lambda, t, s, r) \\ z = \text{Gen}(1^\lambda, t, (s||r); r') \end{array} : \mathcal{D}^{\mathcal{P}}(z) = 1 \right\}.$$

The difference between this and $\text{Hyb}_3^s(\lambda)$ is that r' is sampled uniformly at random in $\text{Hyb}_3^s(\lambda)$, and \mathcal{D} has oracle access to $\mathcal{P}[(1^\lambda, t, s, r) \rightarrow r']$ rather than \mathcal{P} . Let us denote this oracle by \mathcal{P}' . We will show that these hybrids occur with the same probability, except in the case that the query to \mathcal{P} in $\text{Hyb}_2^s(\lambda)$ (which results in r') appears in F , where we recall that F is the partial assignment on $f(\lambda)$ points given by $\text{Sam}(\mathcal{A})$.

To show this, let \mathbf{E} be the event that $(1^\lambda, t, s, r) \notin F$. When \mathbf{E} holds, then the distribution of (\mathcal{P}, r, r') in $\text{Hyb}_2^s(\lambda)$ is identical to that of (\mathcal{P}', r, r') in $\text{Hyb}_3^s(\lambda)$. Namely, in both hybrids r is uniformly sampled from $\{0, 1\}^{\lambda_{\text{tip}}}$. For the oracles, in $\text{Hyb}_2^s(\lambda)$, \mathcal{P} is sampled uniformly from $\text{RF}_\lambda[F]$ and $r' = \mathcal{P}(1^\lambda, t, s, r)$, whereas in $\text{Hyb}_3^s(\lambda)$, \mathcal{P}' can be sampled according to $\text{RF}_\lambda[F]$ for all points except r , and then lazily evaluated it on r to obtain r' . Note that this relies on \mathbf{E} holding, since otherwise \mathcal{P}' on input r would be determined by F . As \mathcal{D} has oracle access to the oracle (\mathcal{P} or \mathcal{P}'), and receives z as input which is fully determined by r, r' , it follows that for all $\lambda \in \mathbb{N}$,

$$\Pr[\text{Hyb}_1^s(\lambda) \mid \mathbf{E}] = \Pr[\text{Hyb}_2^s(\lambda) \mid \mathbf{E}].$$

Furthermore, since $r \leftarrow \{0, 1\}^{\lambda_{\text{tip}}}$, the probability that \mathbf{E} fails to hold, i.e. $(1^\lambda, t, s, r) \in F$, is at most

$$\frac{f_{\text{hard}}(\lambda)}{2^{\lambda_{\text{tip}}}} = \frac{(q_{\mathcal{A}}(\lambda, 2^m) \cdot 4q(\lambda))^2}{2^{\lambda_{\text{tip}}}} = \frac{(q_{\mathcal{A}}(\lambda, 2^m) \cdot 4q(\lambda))^2}{2^{m+\lambda}} \in \frac{2^{O(m)} \cdot \text{poly}(\lambda)}{2^{m+\lambda}} = \text{negl}_2(\lambda),$$

which is negligible, where we used the fact that $\lambda_{\text{tip}} = m + \lambda$. So, it holds that for all $\lambda \in \mathbb{N}$,

$$\Pr[\neg \mathbf{E}] = \text{negl}_2(\lambda).$$

Putting these together, we conclude that for all $\lambda \in \mathbb{N}$,

$$\begin{aligned} & |\Pr[\text{Hyb}_1^s(\lambda)] - \Pr[\text{Hyb}_2^s(\lambda)]| \\ &= |\Pr[\mathbf{E}] \cdot (\Pr[\text{Hyb}_1^s(\lambda) \mid \mathbf{E}] - \Pr[\text{Hyb}_2^s(\lambda) \mid \mathbf{E}]) \\ &\quad + \Pr[\neg \mathbf{E}] \cdot (\Pr[\text{Hyb}_1^s(\lambda) \mid \neg \mathbf{E}] - \Pr[\text{Hyb}_2^s(\lambda) \mid \neg \mathbf{E}])| \\ &\leq 1 \cdot 0 + \text{negl}_2(\lambda) \cdot 1 = \text{negl}_2(\lambda), \end{aligned}$$

so the claim follows. ■

Claim 4.10. *If there exists function μ such that for infinitely many $\lambda \in \mathbb{N}$ and values $s_0, s_1 \in \{0, 1\}^\lambda$, it holds that*

$$|\Pr[\text{Hyb}_3^{s_0}(\lambda)] - \Pr[\text{Hyb}_3^{s_1}(\lambda)]| \geq \mu(\lambda),$$

then there exists an adversary \mathcal{A} that breaks the hardness of TLP with probability $\mu(\lambda) - 2\text{negl}(\lambda)$ for a negligible function negl .

Proof. We first note that by an averaging argument, the inequality in the statement of the claim implies that for infinitely many $\lambda \in \mathbb{N}$ there exists an $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$ and $F \in \text{Supp}(\text{Sam}(\mathcal{A}))$, where $\mathcal{D} = Z(\mathcal{O})$, $\mathcal{A} = Z'(\mathcal{D})$ such that

$$\left| \Pr \left[\begin{array}{l} \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ r_0, r'_0 \leftarrow \{0, 1\}^{\lambda_{\text{tp}}} \\ z_0 = \text{Gen}_{\text{tlp}}(1^\lambda, t, (s_0 \| r_0); r'_0) \end{array} : \mathcal{D}^{\mathcal{P}[(1^\lambda, t, s_0, r_0) \rightarrow r'_0]}(z_0) = 1 \right] \right. \\ \left. - \Pr \left[\begin{array}{l} \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ r_1, r'_1 \leftarrow \{0, 1\}^{\lambda_{\text{tp}}} \\ z_1 = \text{Gen}_{\text{tlp}}(1^\lambda, t, (s_1 \| r_1); r'_1) \end{array} : \mathcal{D}^{\mathcal{P}[(1^\lambda, t, s_1, r_1) \rightarrow r'_1]}(z_1) = 1 \right] \right| \geq \mu(\lambda).$$

The above implies that

$$\Pr \left[\begin{array}{l} \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ b \leftarrow \{0, 1\} \\ r_b, r'_b \leftarrow \{0, 1\}^{\lambda_{\text{tp}}} \\ z_b = \text{Gen}_{\text{tlp}}(1^\lambda, t, (s_b \| r_b); r'_b) \end{array} : \mathcal{D}^{\mathcal{P}[(1^\lambda, t, s_b, r_b) \rightarrow r'_b]}(z_b) = b \right] \geq \frac{1}{2} + \frac{\mu(\lambda)}{2}. \quad (4.2)$$

Note that \mathcal{D} 's success may depend on its ability to “hit” the randomness r_b in one of its oracle queries, since in this case, it can trivially check if z_b corresponds to a puzzle for $s_b \| r_b$. Looking ahead, we will be fixing values of r_0, r_1 and using \mathcal{D} to try to distinguish a puzzle corresponding to $s_0 \| r_0$ from a puzzle corresponding to $s_1 \| r_1$, based on which values of r_b he queries. Therefore, it will also be important to consider the event that $\mathcal{D}(z_b)$ makes a query corresponding to r_{1-b} . Next, we formally define these two events:

- Let $\text{hit} = \text{hit}(\mathcal{D}, \mathcal{P}, z)$ denote the event that \mathcal{D} on input a puzzle corresponding to $(s \| r)$ queries $(1^\lambda, t, s, r)$. Note that in the case that hit does not occur, then the answers to the oracle queries made by \mathcal{D} are distributed according to \mathcal{P} .
- Let $\text{bad} = \text{bad}(\mathcal{D}, \mathcal{P}, z_b, r_b, r_{1-b})$ denote the event that $\mathcal{D}(z_b)$ queries $(1^\lambda, t, s_{1-b}, r_{1-b})$.

We note that for a uniformly random value of r_{1-b} , it holds that

$$\Pr[\text{bad}] \in \frac{\text{poly}(B(\lambda))}{2^{|r_{1-b}|}} \in \frac{2^{O(m)} \cdot \text{poly}(\lambda)}{2^{\lambda_{\text{tp}}}} \in \frac{2^{O(m)} \cdot \text{poly}(\lambda)}{2^{m+\lambda}} \leq \text{negl}(\lambda),$$

for all $\lambda \in \mathbb{N}$, where the probability is over \mathcal{P} , r_0 , r_1 , z , b , and the randomness of \mathcal{D} . This holds because \mathcal{D} has size $\text{poly}(\lambda, m, B(\lambda)) \in \text{poly}(\lambda, 2^m)$, and r_{1-b} is independent of \mathcal{D} 's input and the answers to its queries. Combining this with Equation 4.2, we get the following, where all probabilities are over the choice of $\mathcal{P} \leftarrow \text{RF}_\lambda[F]$, $b \leftarrow \{0, 1\}$, $r_0, r_1 \leftarrow \{0, 1\}^{\lambda_{\text{tp}}}$, and $z_b \leftarrow \text{Gen}_{\text{tlp}}(1^\lambda, t, (s_b \| r_b))$:

$$\begin{aligned} \frac{1}{2} + \frac{\mu(\lambda)}{2} &\leq \Pr \left[\mathcal{D}^{\mathcal{P}[(1^\lambda, t, s_b, r_b) \rightarrow r'_b]}(z_b) = b \right] \\ &\leq \Pr \left[\mathcal{D}^{\mathcal{P}[(1^\lambda, t, s_b, r_b) \rightarrow r'_b]}(z_b) = b \wedge \overline{\text{bad}} \right] + \Pr[\text{bad}] \\ &\leq \Pr \left[\mathcal{D}^{\mathcal{P}[(1^\lambda, t, s_b, r_b) \rightarrow r'_b]}(z_b) = b \wedge \overline{\text{hit}} \wedge \overline{\text{bad}} \right] + \Pr[\text{hit} \wedge \overline{\text{bad}}] + \text{negl}(\lambda) \\ &= \Pr \left[\mathcal{D}^{\mathcal{P}}(z_b) = b \wedge \overline{\text{hit}} \wedge \overline{\text{bad}} \right] + \Pr[\text{hit} \wedge \overline{\text{bad}}] + \text{negl}(\lambda), \end{aligned}$$

where in the last line we used the fact that when $\overline{\text{hit}}$ occurs, the answers to \mathcal{D} 's queries are distributed according to \mathcal{P} .

By an averaging argument, there exist *fixed* values r_0, r_1 such that the above holds for those fixed values, namely,

$$\Pr [\mathcal{D}^{\mathcal{P}}(z_b) = b \wedge \overline{\text{hit}} \wedge \overline{\text{bad}}] + \Pr [\text{hit} \wedge \overline{\text{bad}}] + \text{negl}(\lambda) \geq \frac{1}{2} + \frac{\mu(\lambda)}{2} \quad (4.3)$$

where the probability is only over $\mathcal{P} \leftarrow \text{RF}_\lambda[F]$, $b \leftarrow \{0, 1\}$, and $z_b \leftarrow \text{Gen}_{\text{tlp}}(1^\lambda, t, (s_b || r_b))$.

We will use this to construct an adversary that breaks the hardness of TLP. The adversary \mathcal{B} that has s_0, s_1, r_0, r_1, F , and \mathcal{D} hardcoded, and receives as input a puzzle z corresponding to $s_b || r_b$ for $b \in \{0, 1\}$. It does the following:

1. Run $\mathcal{D}^{(\cdot)}(z)$. For each query q made by \mathcal{D} , if $q \in F$, give the corresponding answer; otherwise, if q has been asked previously, give the same answer as before; otherwise, give a uniformly random answer sampled from $\{0, 1\}^{\lambda_{\text{tlp}}}$.
2. If there exists unique bit b such that \mathcal{B} queries $(1^\lambda, t, s_b, r_b)$ but not $(1^\lambda, t, s_{1-b}, r_{1-b})$, output b . Otherwise, output the bit output by \mathcal{B} .

To analyze the success probability of \mathcal{B} , we can look at whether **hit** or **bad** occur. By definition of \mathcal{B} , we have the following, where all probabilities are over $\mathcal{P} \leftarrow \text{RF}_\lambda[F]$, $b \leftarrow \{0, 1\}$, and $z_b \leftarrow \text{Gen}_{\text{tlp}}(1^\lambda, t, (s_b || r_b))$:

- $\Pr [\mathcal{B}(z_b) = b \wedge \text{hit} \wedge \overline{\text{bad}}] = \Pr [\text{hit} \wedge \overline{\text{bad}}]$, since when **hit** and **bad** occur, then $\mathcal{B}(z_b)$ always outputs b .
- $\Pr [\mathcal{B}(z_b) = b \wedge \overline{\text{hit}} \wedge \overline{\text{bad}}] = \Pr [\mathcal{D}^{\mathcal{P}}(z_b) = b \wedge \overline{\text{hit}} \wedge \overline{\text{bad}}]$, since when neither **hit** nor **bad** occur, then the query answers that \mathcal{B} gives to \mathcal{D} are distributed according to \mathcal{P} , and \mathcal{B} simply outputs what \mathcal{D} outputs.
- $\Pr [\mathcal{B}(z_b) = b \wedge \text{bad}] \geq 0$.

Combining these with Equation 4.3, we get

$$\begin{aligned} \Pr [\mathcal{B}(z_b) = b] &\geq \Pr [\text{hit} \wedge \overline{\text{bad}}] + \Pr [\mathcal{D}^{\mathcal{P}}(z_b) = b \wedge \overline{\text{hit}} \wedge \overline{\text{bad}}] \\ &\geq \frac{1}{2} + \frac{\mu(\lambda)}{2} - \text{negl}(\lambda) \end{aligned}$$

for infinitely many $\lambda \in \mathbb{N}$, where the probabilities are over $\mathcal{P} \leftarrow \text{RF}_\lambda[F]$, $b \leftarrow \{0, 1\}$, and $z_b \leftarrow \text{Gen}_{\text{tlp}}(1^\lambda, t, (s_b || r_b))$.

To put everything together and letting $s = s_0 || r_0$ and $s' = s_1 || r_1$, the above implies that

$$\begin{aligned} &\left| \Pr [z \leftarrow \text{Gen}_{\text{tlp}}(1^\lambda, t, s) : \mathcal{B}_{r_0, r_1}(z) = 1] \right. \\ &\quad \left. - \Pr [z = \text{Gen}_{\text{tlp}}(1^\lambda, t, s') : \mathcal{B}_{r_0, r_1}(z) = 1] \right| > \mu(\lambda) - 2\text{negl}(\lambda). \end{aligned}$$

We will show that this contradicts hardness of TLP for T . Recall that $p_{\text{tlp}}, \ell_{\text{tlp}}$ are the polynomials guaranteed to exist by the hardness of TLP, and recall that we set $\ell(\lambda) = \ell_{\text{tlp}}(\lambda)$ and $p(\lambda) = p_{\text{tlp}}(\lambda)/p'(\lambda)$ for a polynomial p' specified below. To complete the proof, we need to show that $\ell_{\text{tlp}}(\lambda) \leq T(\lambda) \leq B(\lambda)$ and that the size and depth of \mathcal{B} are set such that \mathcal{B} breaks the hiding of TLP. The bounds on T are immediate as we assumed that $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ and $\ell(\lambda) = \ell_{\text{tlp}}(\lambda)$.

For the size and depth of \mathcal{B} , we have that \mathcal{B} needs to run \mathcal{D} , and needs to be able to answer \mathcal{D} 's oracle queries consistently with F and with each other. One way to implement this is for \mathcal{B} to keep

track of a set \mathcal{Q} of the queries asked so far, initialized to F . At any point during the emulation of \mathcal{B} , the set \mathcal{Q} contains at most $\text{size}(\mathcal{D}) + f_{\text{hard}}(\lambda)$ queries. For each query made by \mathcal{B} , checking if it appears in \mathcal{Q} and adding it to \mathcal{Q} if necessary can be done in size $\text{poly}(|\mathcal{Q}|)$, and so in total results in adding $\text{size}(\mathcal{D}) \cdot \text{poly}(|\mathcal{Q}|)$ to the size of \mathcal{B} to account for all the queries. For the depth, each query can be checked in $\text{poly}(\lambda, \log |\mathcal{Q}|)$ depth and added to \mathcal{Q} (while keeping \mathcal{Q} sorted) with an additional $\text{poly}(\lambda, \log |\mathcal{Q}|)$ depth. This can be done in parallel for queries made in parallel, so this results in an extra additive depth of $\text{depth}(\mathcal{D}) \cdot \text{poly}(\lambda, \log |\mathcal{Q}|)$ over the depth of \mathcal{D} .

To put everything together, recall that $f_{\text{hard}}(\lambda) = (q_{\mathcal{A}}(\lambda, 2^m) \cdot 4q(\lambda))^2$, $B(\lambda) \in 2^m \cdot \text{poly}(\lambda)$, and $\text{size}(\mathcal{D}) \in \text{poly}(B(\lambda))$. Therefore, we get that

$$\begin{aligned} \text{size}(\mathcal{B}) &\in \text{poly}(\text{size}(\mathcal{D}), f_{\text{hard}}(\lambda)) \in \text{poly}(B(\lambda), (q_{\mathcal{A}}(\lambda, 2^m) \cdot 4q(\lambda))^2) \in \text{poly}(\lambda, 2^m) \\ &\in \text{poly}(B(\lambda)) \end{aligned}$$

and

$$\begin{aligned} \text{depth}(\mathcal{B}) &\leq \text{depth}(\mathcal{D}) + \text{depth}(\mathcal{D}) \cdot \text{poly}(\lambda, \log |\mathcal{Q}|) \\ &\in \text{depth}(\mathcal{D}) \cdot \text{poly}(\lambda, \log B(\lambda), \log f_{\text{hard}}(\lambda)) \\ &\in \text{depth}(\mathcal{D}) \cdot \text{poly}(\lambda, m\lambda) \\ &\in \text{depth}(\mathcal{D}) \cdot p'(\lambda) \end{aligned}$$

for a fixed polynomial p' , depending only on m and on $\log f_{\text{hard}}(\lambda)$. Note that $\log f_{\text{hard}}(\lambda)$ is in $\text{poly}(\lambda, m, \log(q(\lambda)))$. Since q was assumed to be a polynomial, this can be upper bounded by a fixed polynomial in λ , which is independent of q , for sufficiently large λ . It follows that we can assume p' is independent of q .

Recall that $\text{depth}(\mathcal{D}) \leq (T(\lambda))^\epsilon \cdot p(\lambda)$ where we set $p(\lambda) = p_{\text{tlp}}(\lambda)/p'(\lambda)$ for p' as defined above. We can therefore upper bound $\text{depth}(\mathcal{D})$ in the above to get that

$$\begin{aligned} \text{depth}(\mathcal{B}) &\leq (T(\lambda))^\epsilon \cdot p(\lambda) \cdot p'(\lambda) \\ &= (T(\lambda))^\epsilon \cdot (p_{\text{tlp}}(\lambda)/p'(\lambda)) \cdot p'(\lambda) \\ &= (T(\lambda))^\epsilon \cdot p_{\text{tlp}}(\lambda). \end{aligned}$$

Therefore, \mathcal{B} breaks the hardness of TLP with probability $\mu(\lambda) - 2\text{negl}(\lambda)$ for infinitely many λ . ■

This completes the proof of the lemma. □

Lemma 4.11. *Assuming that (Gen, Sol) satisfies correctness and (B, ϵ) -hardness, and that $(\text{Gen}_{\text{tlp}}, \text{Sol}_{\text{tlp}})$ is a (B, ϵ) -secure time-lock puzzle, then (Gen, Sol) is fully concurrent functional non-malleable for \mathcal{F}_B^m .*

Proof. We will show that (Gen, Sol) satisfies one-many functional non-malleability for \mathcal{F}_B^m , which suffices for fully concurrent functional non-malleability for \mathcal{F}_B^m by Lemma A.1.

To show one-many functional non-malleability, suppose for contradiction that there exists an $f \in \mathcal{F}_B^m$ such that for all positive polynomials p, ℓ , there exists a function T satisfying $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$, an unbounded algorithm Z and polynomial n such that Z outputs $(1, n, B, \epsilon, T, p)$ -MIM adversaries, an unbounded distinguisher \mathcal{D} , and a polynomial q , such that for

infinitely many $\lambda \in \mathbb{N}$ and $s \in \{0, 1\}^\lambda$, it holds that

$$\left| \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ \vec{s} \leftarrow \text{mim}_{\mathcal{A}}^{\mathcal{O}}(1^\lambda, T(\lambda), s) \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right] \right. \quad (4.4)$$

$$\left. - \Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ \vec{s} \leftarrow \text{mim}_{\mathcal{A}}^{\mathcal{O}}(1^\lambda, T(\lambda), 0^\lambda) \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right] \right| > \frac{1}{q(\lambda)}. \quad (4.5)$$

Let $\ell_{\text{hard}}, p_{\text{hard}}$ be the positive polynomials from the hardness property of nmTLP_m , given by 4.6, and let $\ell_{\text{tlp}}, p_{\text{tlp}}$ be the positive polynomials from the hardness of TLP. We will derive a contradiction to the above for p, ℓ given by $p(\lambda) = p_{\text{hard}}(\lambda)/p'(\lambda)$ and $\ell(\lambda) = \ell_{\text{hard}}(\lambda) + \ell_{\text{tlp}}(\lambda) + (\ell'(\lambda))^{1/\epsilon}$, where p' is a fixed polynomial specified in the proof of Claim 4.16, and ℓ' is a fixed polynomial specified in the proof of Subclaim 4.15.

In order to show a contradiction, we will give a sequence of hybrid experiments starting with the event in the first probability above and ending with the second, and we will show that each consecutive pair either has large statistical distance, or can be used to break hardness of nmTLP_m or of TLP. Before giving the formal description of the hybrids, we give a short overview. At a high level, our strategy will be to change the way that \vec{s} is computed, so as to make it independent of s . Specifically, we start with a hybrid corresponding to the first probability above, where \vec{s} is computed using the *unbounded* sampler $\text{mim}_{\mathcal{A}}$. Next, we switch $\text{mim}_{\mathcal{A}}$ to a sampler that has *polynomial size*. This allows us to transition from the random oracle \mathcal{O} to an oracle \mathcal{P} that is lazily sampled on most points using Lemma 3.6, so that the MIM adversary \mathcal{A} only depends on a small fraction of points in the oracle (rather than the whole oracle). This enables us to then switch to sampling \vec{s} using a *depth-bounded, but sub-exponential size* sampler. Once we have done so, we can apply the hardness of our time-lock puzzle to switch the initial puzzle for s to a puzzle for 0^λ , thereby making the experiment independent of s . The hybrids are formalized next.

For any $\lambda \in \mathbb{N}$ and value $s \in \{0, 1\}^\lambda$, we consider the following sequence of hybrid experiments. Throughout these hybrids, we let $t = T(\lambda)$ and $m = m(\lambda)$.

- $\text{Hyb}_0^s(\lambda)$: This hybrid is equivalent to the terms in the probabilities above for $s \in \{0, 1\}^\lambda$.

$$\text{Hyb}_0^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{A} = Z(\mathcal{O}) \\ \vec{s} \leftarrow \text{mim}_{\mathcal{A}}^{\mathcal{O}}(1^\lambda, t, s) \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right\}$$

- $\text{Hyb}_1^s(\lambda)$: In this hybrid, we replace $\text{mim}_{\mathcal{A}}$ with a polynomial-time algorithm \mathcal{B} , which is defined based on \mathcal{A} , and sampled by an algorithm Z' . Specifically, let Z' be the algorithm such that $Z'(\mathcal{A})$ outputs the following probabilistic oracle algorithm $\mathcal{B}^{\mathcal{O}}$:

1. Sample $z \leftarrow \text{Gen}^{\mathcal{O}}(1^\lambda, t, s)$.
2. Run $\vec{z} \leftarrow \mathcal{A}^{\mathcal{O}}(z)$.
3. Do the following to compute the output vector \vec{s} . For each $i \in [n(\lambda)]$, if $\vec{z}_i = z$, set $\vec{s}_i = \perp$. Otherwise, compute $s' = \text{Sol}^{\mathcal{O}}(1^\lambda, t, \vec{z}_i)$, and set $\vec{s}_i = s'$.
4. Output \vec{s} .

Note that $\mathcal{B}^{\mathcal{O}}$ forwards all queries made by internal algorithms to its own oracle. We can now define the hybrid experiment:

$$\text{Hyb}_1^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{A} = Z(\mathcal{O}); \quad \mathcal{B} = Z'(\mathcal{A}) \\ \vec{s} \leftarrow \mathcal{B}^{\mathcal{O}} \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right\}$$

Note that $\text{size}(\mathcal{B}) \in \text{poly}(\lambda, 2^m)$, which follows by the efficiency of **Gen** and **Sol**, and because $n(\lambda) \in \text{poly}(\lambda)$ and $\text{size}(\mathcal{A}) \in \text{poly}(B)$.

- $\text{Hyb}_2^s(\lambda)$: Let $q_{\mathcal{B}}$ be the polynomial such that $\text{size}(\mathcal{B}) = q_{\mathcal{B}}(\lambda, 2^m)$. Let $f_{\text{nm}}(\lambda) = (q_{\mathcal{B}}(\lambda, 2^m) \cdot 4q(\lambda))^2$ and let **Sam** be the inefficient algorithm given in Lemma 3.6 that on input an adversary, outputs a partial assignment F on $f_{\text{nm}}(\lambda)$ points. This hybrid switches the random oracle \mathcal{O} with a random function \mathcal{P} fixed at the set of points F .

$$\text{Hyb}_1^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_{\lambda}; \quad \mathcal{A} = Z(\mathcal{O}); \quad \mathcal{B} = Z'(\mathcal{A}) \\ F \leftarrow \text{Sam}(\mathcal{B}); \quad \mathcal{P} \leftarrow \text{RF}_{\lambda}[F] \\ \vec{s} \leftarrow \mathcal{B}^{\mathcal{P}} \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right\}$$

- $\text{Hyb}_3^s(\lambda)$: In this hybrid, we change the experiment so that \vec{s} can be computed in depth less than t (given \mathcal{A} and F). Specifically, we define the distribution **bmim** (standing for “bounded MIM”) and replace \mathcal{B} with this distribution, defined as follows. Define **bmim** as follows:

$\text{bmim}_{\mathcal{A}, F}^{\mathcal{P}}(1^\lambda, t, s)$:

1. Sample $z \leftarrow \text{Gen}^{\mathcal{P}}(1^\lambda, t, s)$
2. Run $\vec{z} \leftarrow \mathcal{A}^{(\cdot)}(z)$ by forwarding all of \mathcal{A} 's queries to the oracle \mathcal{P} .
3. Let \mathcal{Q} be the set containing all oracle queries made by \mathcal{A} and all points in the partial assignment F , where $Q_j = (1^\lambda, t_j, s_j, r_j)$ for each $j \in [|\mathcal{Q}|]$.
4. Next, we form the output vector \vec{s} . For each $i \in [n(\lambda)]$, if $\vec{z}_i = z$, set $\tilde{s}_i = \perp$. Otherwise, check if there exists a query $Q_j \in \mathcal{Q}$ with $\vec{z}_i = \text{Gen}^{\mathcal{P}}(1^\lambda, t, s_j; r_j)$, and if so set $\tilde{s}_i = s_j$. Otherwise, set $\tilde{s}_i = \perp$. Note that this check can be done in parallel for each pair (i, j) .
5. Output \vec{s} .

We define this hybrid experiment from the previous one by using **bmim** instead of \mathcal{B} to sample \vec{s} , as follows.

$$\text{Hyb}_3^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_{\lambda}; \quad \mathcal{A} = Z(\mathcal{O}); \quad \mathcal{B} = Z'(\mathcal{A}) \\ F \leftarrow \text{Sam}(\mathcal{B}); \quad \mathcal{P} \leftarrow \text{RF}_{\lambda}[F] \\ \vec{s} \leftarrow \text{bmim}_{\mathcal{A}, F}^{\mathcal{P}}(1^\lambda, t, s) \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right\}$$

To complete the proof, we show that for each consecutive pair of hybrids, either the statistical distance is bounded or we can use it to break security of **TLP** or nmTLP_m . Next, we discuss how to use these claims to complete the proof, and then we give the claims.

In Claim 4.12, we show that $\Pr[\text{Hyb}_0^s(\lambda)] = \Pr[\text{Hyb}_1^s(\lambda)]$ for all $\lambda \in \mathbb{N}$ and $s \in \{0, 1\}^\lambda$. Then, in Claim 4.13 we show that there exists a negligible function negl_1 such that $|\Pr[\text{Hyb}_1^s(\lambda)] - \Pr[\text{Hyb}_2^s(\lambda)]| \leq \text{negl}_1(\lambda)$ for all $\lambda \in \mathbb{N}$ and $s \in \{0, 1\}^\lambda$. Combining these with Equation 4.4, we get that for infinitely many $\lambda \in \mathbb{N}$, there exists a value $s \in \{0, 1\}^\lambda$ such that

$$\left| \Pr[\text{Hyb}_2^s(\lambda)] - \Pr[\text{Hyb}_2^{0^\lambda}(\lambda)] \right| \geq \frac{1}{q(\lambda)} - \text{negl}_1(\lambda) \geq \frac{1}{2q(\lambda)}.$$

It follows that for infinitely many $\lambda \in \mathbb{N}$, there exists a value $s \in \{0, 1\}^\lambda$ and a pair of consecutive hybrids in the sequence $\text{Hyb}_2^s(\lambda), \text{Hyb}_3^s(\lambda), \text{Hyb}_3^{0^\lambda}(\lambda), \text{Hyb}_2^{0^\lambda}(\lambda)$ whose statistical distance is at least $1/(6q(\lambda))$. We first discuss the case where the statistical distance between $\text{Hyb}_2^s(\lambda)$ and $\text{Hyb}_3^s(\lambda)$ is

large, and then discuss the case where the distance between $\text{Hyb}_3^s(\lambda)$ and $\text{Hyb}_3^{0^\lambda}(\lambda)$ is large. The third case will follow analogously to the second one.

To show a contradiction in the first case, we show in Claim 4.14 that if for infinitely many $\lambda \in \mathbb{N}$ there exists a value $s \in \{0, 1\}^\lambda$ such that the statistical distance between $\text{Hyb}_2^s(\lambda)$ and $\text{Hyb}_3^s(\lambda)$ is at least $\mu(\lambda)$ for any function μ , then this implies an adversary that breaks the hardness of TLP with probability $\mu(\lambda)/n(\lambda)$. In this case, the statistical distance is $1/(6q(\lambda))$, which implies an adversary that breaks the hardness of TLP with probability $1/(6q(\lambda) \cdot n(\lambda))$ which is a contradiction, since n is a polynomial.

For the second case, where $\left| \Pr[\text{Hyb}_3^s(\lambda)] - \Pr[\text{Hyb}_3^{0^\lambda}(\lambda)] \right| \geq \frac{1}{6q(\lambda)}$, we show in Claim 4.16 that if for infinitely many $\lambda \in \mathbb{N}$ there exists a value $s \in \{0, 1\}^\lambda$ such that the statistical distance between these hybrids is at least $\mu(\lambda)$ for any function μ , then this implies an adversary that breaks the hardness of nmTLP_m with probability $\mu(\lambda) - \text{negl}(\lambda)$ for a negligible function negl . In our case, this results in statistical distance $1/(6q(\lambda)) - \text{negl}(\lambda) \geq 1/(12q(\lambda))$, which contradicts the hardness of nmTLP_m , and concludes the proof.

Claim 4.12. For all $s \in \{0, 1\}^\lambda$, $\Pr[\text{Hyb}_0^s(\lambda)] = \Pr[\text{Hyb}_1^s(\lambda)]$.

Proof. In both of these hybrids, the oracle \mathcal{O} and adversary \mathcal{A} are sampled equivalently. The difference between the two distributions is the following:

- In $\text{Hyb}_0^s(\lambda)$, the vector \vec{s} is sampled by $\text{mim}_{\mathcal{A}}^{\mathcal{O}}(1^\lambda, t, s)$.
- In $\text{Hyb}_1^s(\lambda)$, the vector \vec{s} is determined by letting $\mathcal{B} = Z(\mathcal{A})$ and sampling $\vec{s} \leftarrow \mathcal{B}^{\mathcal{O}}$.

We will show that these have the same distribution.

To sample \vec{s} , both \mathcal{B} and $\text{mim}_{\mathcal{A}}$ compute a puzzle z for s and run $\vec{z} \leftarrow \mathcal{A}(z)$, but then form the output vector \vec{s} differently. For ease of understanding, denote this vector as computed by $\text{mim}_{\mathcal{A}}$ in $\text{Hyb}_0^s(\lambda)$ as $\vec{s}^{(0)}$, and denote the vector as computed by \mathcal{B} in $\text{Hyb}_1^s(\lambda)$ as $\vec{s}^{(1)}$. Specifically, for each output element \tilde{s}_i , $\text{mim}_{\mathcal{A}}$ computes it as

$$\tilde{s}_i^{(0)} = \begin{cases} \perp & \text{if } \tilde{z}_i = z \\ s' & \text{if there is a unique } s' \text{ with } \tilde{z}_i = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s'; r) \text{ for some } r \\ \perp & \text{otherwise,} \end{cases}$$

while $\mathcal{B}^{\mathcal{O}}$ computes it as

$$\tilde{s}_i^{(1)} = \begin{cases} \perp & \text{if } \tilde{z}_i = z \\ s' & \text{otherwise, where } s' = \text{Sol}^{\mathcal{O}}(1^\lambda, t, \tilde{z}_i). \end{cases}$$

We will show that $\tilde{s}_i^{(0)} = \tilde{s}_i^{(1)}$. Note that if $\tilde{z}_i = z$, then $\tilde{s}_i = \perp$ in both cases, so we only need to focus on the case when $\tilde{z}_i \neq z$. To complete the proof, we will show that there exists a unique $s' \in \{0, 1\}^*$ with $\tilde{z}_i = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s'; r)$ for some r if and only if $\text{Sol}^{\mathcal{O}}(1^\lambda, t, \tilde{z}_i) \neq \perp$. This suffices for the claim since whenever this occurs, correctness of nmTLP_m implies that both $\text{mim}_{\mathcal{A}}$ and $\mathcal{B}^{\mathcal{O}}$ will set $\tilde{s}_i = s'$. Otherwise, both will set it to \perp , implying that \tilde{s}_i is identically distributed in both hybrids for all i .

For the “if” direction, suppose that there exists a unique value $s' \in \{0, 1\}^*$ with $\tilde{z}_i = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s'; r)$ for some r . Correctness of nmTLP_m directly implies that $\text{Sol}^{\mathcal{O}}(1^\lambda, t, \tilde{z}_i)$ will output s' , and hence the output of Sol is not equal to \perp .

For the “only if” direction, suppose $\text{Sol}^{\mathcal{O}}(1^\lambda, t, \tilde{z}_i) \neq \perp$. Recall that Sol solves the underlying time-lock puzzle to obtain a value $s''||r$, and then outputs s'' if $\tilde{z}_i = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s''; r)$ and \perp otherwise. As we are in the case where Sol does not output \perp , it holds that $\tilde{z}_i = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s''; r)$. Correctness of nmTLP_m implies that s'' is the unique solution to \tilde{z}_i . This completes the proof of the claim. \blacksquare

Claim 4.13. *There exists a negligible function negl_1 such that for all $s \in \{0, 1\}^\lambda$,*

$$|\Pr[\text{Hyb}_1^s(\lambda)] - \Pr[\text{Hyb}_2^s(\lambda)]| \leq \text{negl}_1(\lambda).$$

Proof. We start by noting that $\text{Hyb}_1^s(\lambda)$ can be rewritten as

$$\text{Hyb}_1^s(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{B} = Z''(\mathcal{O}) \\ \vec{s} \leftarrow \mathcal{B}^{\mathcal{O}} \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right\}$$

where $Z''(\mathcal{O})$ samples $\mathcal{A} = Z(\mathcal{O})$ and $\mathcal{B} = Z'(\mathcal{A})$. Using Z'' , we can also rewrite $\text{Hyb}_2^s(\lambda)$ as

$$\left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{B} = Z''(\mathcal{O}) \\ F \leftarrow \text{Sam}(\mathcal{B}); \quad \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ \vec{s} \leftarrow \mathcal{B}^{\mathcal{P}} \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right\}.$$

Since $|F| = f_{\text{nm}}(\lambda) = (q_{\mathcal{B}}(\lambda, 2^m) \cdot 4q(\lambda))^2$, then by Lemma 3.6 the statistical distance between these hybrids is at most

$$\sqrt{\frac{(q_{\mathcal{B}}(\lambda, 2^m))^2}{f_{\text{nm}}(\lambda)}} = \sqrt{\frac{(q_{\mathcal{B}}(\lambda, 2^m))^2}{(q_{\mathcal{B}}(\lambda, 2^m) \cdot 4q(\lambda))^2}} = \frac{1}{4q(\lambda)}$$

where we recall that $\text{size}(\mathcal{B}) = q_{\mathcal{B}}(\lambda, 2^m)$. \blacksquare

Claim 4.14. *If there exists a function μ such that for infinitely many $\lambda \in \mathbb{N}$ there exists a value $s \in \{0, 1\}^\lambda$ with*

$$|\Pr[\text{Hyb}_2^s(\lambda)] - \Pr[\text{Hyb}_3^s(\lambda)]| \geq \mu(\lambda),$$

then there exists an adversary that breaks the hardness of TLP with probability $\mu(\lambda)/n(\lambda)$ for infinitely many $\lambda \in \mathbb{N}$.

Proof. The difference between these two hybrids is that in $\text{Hyb}_2^s(\lambda)$, the values \vec{s} are sampled from $\mathcal{B}^{\mathcal{P}}$, while in $\text{Hyb}_3^s(\lambda)$ they are sampled by $\text{bmim}_{\mathcal{A}, F}^{\mathcal{P}}(1^\lambda, t, s)$. We will show that if these two distributions are far, then we can construct an adversary that breaks the hardness of TLP with probability roughly the statistical distance between the two hybrids.

The way that \vec{s} is sampled in each hybrid is as follows. Both $\mathcal{B}^{\mathcal{P}}$ and $\text{bmim}_{\mathcal{A}, F}^{\mathcal{P}}(1^\lambda, t, s)$ first sample a puzzle z for s , and then run $\vec{z} \leftarrow \mathcal{A}(z)$. They then calculate \vec{s} as follows. For ease of understanding, denote this vector as computed by \mathcal{B} in $\text{Hyb}_2^s(\lambda)$ as $\vec{s}^{(2)}$, and denote the vector as computed by bmim in $\text{Hyb}_3^s(\lambda)$ as $\vec{s}^{(3)}$. Then, we have that for each output element i , $\mathcal{B}^{\mathcal{P}}$ computes it as

$$\vec{s}_i^{(2)} = \begin{cases} \perp & \text{if } \tilde{z}_i = z \\ s' & \text{otherwise, where } s' = \text{Sol}^{\mathcal{P}}(1^\lambda, t, \tilde{z}_i) \end{cases}$$

while $\text{bmim}_{\mathcal{A},F}^{\mathcal{P}}(1^\lambda, t, s)$ computes it as

$$\tilde{s}_i^{(3)} = \begin{cases} \perp & \text{if } \tilde{z}_i = z \\ v_j & \text{if there exists a } Q_j = (1^\lambda, t, v_j, r_j) \in \mathcal{Q} \text{ with } \tilde{z}_i = \text{Gen}^{\mathcal{P}}(1^\lambda, t, v_j; r_j) \\ \perp & \text{otherwise,} \end{cases}$$

where we recall that \mathcal{Q} consists of all queries made by \mathcal{A} and all elements of the partial assignment F . Unless otherwise stated, note that all following probabilities are over $\mathcal{O} \leftarrow \text{RF}_\lambda$, $\mathcal{A} = Z(\mathcal{O})$, $\mathcal{B} = Z'(\mathcal{A})$, $F = \text{Sam}(\mathcal{B})$, $\mathcal{P} \leftarrow \text{RF}_\lambda[F]$, $\tilde{s}^{(2)} \leftarrow \mathcal{B}^{\mathcal{P}}$, $\tilde{s}^{(3)} \leftarrow \text{bmim}_{\mathcal{A},F}^{\mathcal{P}}(1^\lambda, t, s)$. By assumption, we have that

$$\begin{aligned} \mu(\lambda) &\leq |\Pr[\text{Hyb}_2^s(\lambda)] - \Pr[\text{Hyb}_3^s(\lambda)]| \leq \Pr[\tilde{s}^{(2)} \neq \tilde{s}^{(3)}] \\ &= \Pr[\exists i \in [n(\lambda)] : \tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)}] \leq \sum_{i \in [n(\lambda)]} \Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)}] \end{aligned}$$

by a union bound. Therefore, there exists some $i \in [n(\lambda)]$ such that

$$\Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)}] \geq \mu(\lambda)/n(\lambda). \quad (4.6)$$

Fix this index i . We will continue by expanding this probability.

Let \mathbf{E} be the event that there exists a $Q_j = (1^\lambda, t, s_j, r_j) \in \mathcal{Q}$ such that $\tilde{z}_i = \text{Gen}^{\mathcal{P}}(1^\lambda, t, s_j; r_j)$. We have that

$$\begin{aligned} \Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)}] &= \Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)} \wedge \mathbf{E}] + \Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)} \wedge \neg\mathbf{E}] \\ &\leq \Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)} \mid \mathbf{E}] + \Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)} \mid \neg\mathbf{E}], \end{aligned} \quad (4.7)$$

We continue by bounding each term in Equation 4.7 separately.

For the first term, we have that

$$\Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)} \mid \mathbf{E}] = 0.$$

To see this, note that conditioning on \mathbf{E} implies $\tilde{s}_i^{(3)} = s_j$, where $Q_j = (1^\lambda, t, s_j, r_j) \in \mathcal{Q}$ is the query such that $\tilde{z}_i = \text{Gen}^{\mathcal{P}}(1^\lambda, t, s_j; r_j)$. By the correctness of nmTLP_m , this implies that $\text{Sol}^{\mathcal{P}}(1^\lambda, t, \tilde{z}_i)$ outputs s_j , so $\tilde{s}_i^{(2)} = s_j$.

To bound the second term of Equation 4.7, conditioning on $\neg\mathbf{E}$ implies $\tilde{s}_i^{(3)} = \perp$, and so

$$\Pr[\tilde{s}_i^{(2)} \neq \tilde{s}_i^{(3)} \mid \neg\mathbf{E}] = \Pr[\tilde{s}_i^{(2)} \neq \perp \mid \neg\mathbf{E}].$$

Recall that $\tilde{s}_i^{(2)}$ is sampled as $\text{Sol}^{\mathcal{P}}(1^\lambda, t, \tilde{z}_i)$. By definition of Sol , it holds that its output is not \perp when $\tilde{z}_i = \text{Gen}^{\mathcal{P}}(1^\lambda, s', t; r)$ for $s' || r = \text{Sol}_{\text{tlp}}(1^\lambda, t, \tilde{z}_i)$. Therefore, the above is equal to

$$\begin{aligned} &\Pr[\text{Sol}^{\mathcal{P}}(1^\lambda, t, \tilde{z}_i) \neq \perp \mid \neg\mathbf{E}] \\ &= \Pr[\tilde{z}_i = \text{Gen}^{\mathcal{P}}(1^\lambda, s', t; r) \mid \neg\mathbf{E}] \\ &= \Pr[\tilde{z}_i = \text{Gen}_{\text{tlp}}(1^\lambda, t, (s' || r); \mathcal{P}(1^\lambda, t, s', r)) \mid \neg\mathbf{E}], \end{aligned}$$

where the above probabilities are over the choice of $\mathcal{O}, \mathcal{A}, \mathcal{B}, F, \mathcal{P}$ and the randomness used by \mathcal{A} to produce \tilde{z}_i . Since we are conditioning on $\neg E$, it follows that $\mathcal{P}(1^\lambda, t, s', r)$ is uniformly random and independent from \mathcal{A} and hence from \tilde{z}_i , so it suffices to bound the following probability, relative to any \tilde{z}_i . Therefore, by combining this with equations 4.7 and 4.6, we get the following, where the probability is only over $r' \leftarrow \{0, 1\}^{\lambda_{\text{tlp}}}$:

$$\Pr \left[r' \leftarrow \{0, 1\}^{\lambda_{\text{tlp}}} : \text{Gen}_{\text{tlp}}(1^\lambda, t, (s' || r); r') = \tilde{z}_i \right] \geq \mu(\lambda)/n(\lambda).$$

In the following sub-claim, we show that this implies we can break the hardness of TLP with the same probability, which completes the claim. Note that the following sub-claim is general, so the notation is independent from the above.

Subclaim 4.15. *If for infinitely many $\lambda \in \mathbb{N}$ there exist values $z^* \in \{0, 1\}^*$ and $s \in \{0, 1\}^{\lambda_{\text{tlp}} + \lambda}$ with*

$$\Pr \left[r \leftarrow \{0, 1\}^{\lambda_{\text{tlp}}} : \text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s; r) = z^* \right] \geq \mu(\lambda)/n(\lambda),$$

then there exists an adversary that breaks the hardness of TLP with probability $\mu(\lambda)/n(\lambda)$.

Proof. We will show that there exists a depth-bounded distinguisher $\mathcal{D}' = \{\mathcal{D}'_\lambda\}_{\lambda \in \mathbb{N}}$ that breaks the hardness of the time-lock puzzle with respect to T and any inputs $s \neq s' \in \{0, 1\}^{\lambda_{\text{tlp}} + \lambda}$. Specifically, we will show that for infinitely many $\lambda \in \mathbb{N}$,

$$\left| \Pr \left[\mathcal{D}'_\lambda(\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s)) = 1 \right] - \Pr \left[\mathcal{D}'_\lambda(\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s')) = 1 \right] \right| > \mu(\lambda), \quad (4.8)$$

which contradicts the hardness of TLP. For any $\lambda \in \mathbb{N}$, we define \mathcal{D}'_λ on input z to simply output 1 if $z = z^*$ and 0 otherwise.

To show that \mathcal{D}' contradicts the hardness of TLP, we start by bounding the distinguishing probabilities. Let Λ be the infinitely large set of $\lambda \in \mathbb{N}$ such that Equation 4.8 holds. It follows that for all $\lambda \in \Lambda$,

$$\Pr \left[\mathcal{D}'_\lambda(\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s)) = 1 \right] = \Pr \left[\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s) = z^* \right] > \mu(\lambda).$$

In particular, it holds that z^* is in the support of $\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s)$ for all $\lambda \in \Lambda$. By completeness of TLP, this implies that $\text{Sol}_{\text{tlp}}(1^\lambda, T(\lambda), z^*) = s$. It follows that z^* is not in the support of $\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s')$ for any $\lambda \in \Lambda$ since otherwise it would be the case that $\text{Sol}_{\text{tlp}}(1^\lambda, T(\lambda), z^*)$ also equals s' , but $s \neq s'$. As a result, it follows that for all $\lambda \in \Lambda$ that

$$\Pr \left[\mathcal{D}'_\lambda(\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s')) = 1 \right] = \Pr \left[\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s') = z^* \right] = 0,$$

To complete the proof, we need to show that T and the size and depth of \mathcal{D}' are bounded by the necessary parameters. Recall that the hardness of TLP is given by parameters B, ϵ and polynomials $\ell_{\text{tlp}}, p_{\text{tlp}}$. We want to show that $\ell_{\text{tlp}}(\lambda) \leq T(\lambda) \leq B(\lambda)$, that $\text{size}(\mathcal{D}'_\lambda) \in \text{poly}(B(\lambda))$, and that $\text{depth}(\mathcal{D}'_\lambda) \leq (T(\lambda))^\epsilon \cdot p_{\text{tlp}}(\lambda)$. Note that the bounds on $T(\lambda)$ are immediate as we are given that $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ and we set $\ell(\lambda) \geq \ell_{\text{tlp}}(\lambda) + (\ell'(\lambda))^{1/\epsilon} \geq \ell_{\text{tlp}}(\lambda)$.

It remains to bound the size and depth of \mathcal{D}' . As z^* is in the support of $\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), \cdot)$, the efficiency of TLP implies that the size of z^* and hence $\text{size}(\mathcal{D}'_\lambda)$ can be bounded by $\text{poly}(\lambda, \log T(\lambda))$ and so is bounded by a fixed polynomial $\ell'(\lambda)$ by the above bounds on T (where ℓ' is the polynomial used to specify ℓ , and depends only on m). Note that this also bounds $\text{depth}(\mathcal{D}'_\lambda)$. It follows that

$$\text{size}(\mathcal{D}'_\lambda) \leq \ell'(\lambda) \in \text{poly}(\lambda) \in \text{poly}(B(\lambda))$$

and

$$\text{depth}(\mathcal{D}'_\lambda) \leq \ell'(\lambda) \leq (\ell(\lambda))^\epsilon \leq (T(\lambda))^\epsilon.$$

as desired, which completes the proof of the subclaim. \blacksquare

This completes the proof of Claim 4.14. \blacksquare

Claim 4.16. *If there exists a function μ such that for infinitely many $\lambda \in \mathbb{N}$ there exists a value $s \in \{0, 1\}^\lambda$ with*

$$\left| \Pr[\text{Hyb}_3^s(\lambda)] - \Pr[\text{Hyb}_3^{0^\lambda}(\lambda)] \right| \geq \mu(\lambda),$$

then there exists an adversary that breaks the hardness of nmTLP_m with probability $\mu(\lambda) - \text{negl}(\lambda)$ for a negligible function negl for infinitely many $\lambda \in \mathbb{N}$.

Proof. By an averaging argument, the inequality in the statement of the claim implies that for infinitely many $\lambda \in \mathbb{N}$, there exists an $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$, such that for $\mathcal{A} = Z(\mathcal{O})$, $\mathcal{B} = Z'(\mathcal{A})$ and $F = \text{Sam}(\mathcal{B})$, it holds that

$$\left| \Pr \left[\begin{array}{l} \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ \vec{s} \leftarrow \text{bmim}_{\mathcal{A},F}^{\mathcal{P}}(1^\lambda, t, s) \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right] - \Pr \left[\begin{array}{l} \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ \vec{s} \leftarrow \text{bmim}_{\mathcal{A},F}^{\mathcal{P}}(1^\lambda, t, 0^\lambda) \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right] \right| \geq \mu(\lambda).$$

In order to complete the proof of the claim, our goal is to use bmim and \mathcal{D} to construct an adversary against the hardness of nmTLP_m , for which we need the probability to be over the choice of a random oracle \mathcal{O}' rather than a partially fixed on \mathcal{P} . Toward that goal, for any oracle \mathcal{O}' , let $\widetilde{\text{bmim}}_{\mathcal{A},F}^{\mathcal{O}'}(1^\lambda, t, s)$ be the same as bmim , except that whenever any internal algorithm (such as \mathcal{A} or Gen) makes an oracle query Q , $\widetilde{\text{bmim}}$ first checks if $Q \in F$. If so, it returns the corresponding point as given by F , and otherwise it forwards the query to its oracle \mathcal{O}' .

We observe that sampling $\mathcal{P} \leftarrow \text{RF}_\lambda[F]$ and $\vec{s} \leftarrow \widetilde{\text{bmim}}_{\mathcal{A},F}^{\mathcal{P}}(1^\lambda, t, s)$ has the same distribution as sampling a fully-defined oracle $\mathcal{O}' \leftarrow \text{RF}_\lambda$ and then $\vec{s} \leftarrow \widetilde{\text{bmim}}_{\mathcal{A},F}^{\mathcal{O}'}(1^\lambda, t, s)$. Therefore, we have that

$$\left| \Pr \left[\begin{array}{l} \mathcal{O}' \leftarrow \text{RF}_\lambda \\ \vec{s} \leftarrow \widetilde{\text{bmim}}_{\mathcal{A},F}^{\mathcal{O}'}(1^\lambda, t, s) \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right] - \Pr \left[\begin{array}{l} \mathcal{O}' \leftarrow \text{RF}_\lambda \\ \vec{s} \leftarrow \widetilde{\text{bmim}}_{\mathcal{A},F}^{\mathcal{O}'}(1^\lambda, t, 0^\lambda) \end{array} : \mathcal{D}(f(\vec{s})) = 1 \right] \right| \geq \mu(\lambda). \quad (4.9)$$

We will use this to break the hiding of nmTLP_m by constructing an algorithm \mathcal{D}' that receives a puzzle z either to s or 0^λ , behaves exactly as $\widetilde{\text{bmim}}$ does to obtain \vec{s} , computes $f(\vec{s})$, and finally uses \mathcal{D} to distinguish between the two cases. Specifically, let $\mathcal{D}'^{\mathcal{O}'}$ be the oracle algorithm that on input a puzzle z , does the following:

1. Run $\vec{z} \leftarrow \mathcal{A}^{(\cdot)}(z)$, where for each query Q made by \mathcal{A} , if $Q \in F$ then answer with the corresponding image given by F , and otherwise forward Q to \mathcal{O}' . Let \mathcal{Q} be the set containing all queries and answers made by \mathcal{A} as well as all points in F .
2. In parallel, for each $i \in [n(\lambda)]$ and $j \in [|\mathcal{Q}|]$, compute \tilde{z}_i as done by $\widetilde{\text{bmim}}_{\mathcal{A},F}^{(\cdot)}(1^\lambda, t, \cdot)$ by checking if \tilde{z}_i is the result of generating a puzzle using $Q_j \in \mathcal{Q}$.

3. Compute $y = f(\vec{s})$.
4. Let $\text{tt}_{\mathcal{D}}$ be the circuit with width 2^m and depth $O(m)$ corresponding to the truth table of \mathcal{D} . Run $\mathcal{D}(y)$ to get $b = \mathcal{D}(y)$, and output b .⁸

To analyze the success probability of \mathcal{D}' in breaking the hardness of TLP , we observe that the only difference between running $\mathcal{D}'^{\mathcal{O}'}$ and running \mathcal{D} on the output of $\widetilde{\text{bmim}}_{\mathcal{A},F}^{\mathcal{O}'}$ is that the puzzle given to \mathcal{D}' is sampled using \mathcal{O}' , while the puzzle that bmim uses is sampled using \mathcal{O} , but replacing any queries to F with the image given in F . Specifically, let z be the puzzle given as input to \mathcal{D}' . In the first case, where z is a puzzle for s , let r be the uniformly random value such that $z = \text{Gen}^{\mathcal{O}'}(1^\lambda, t, s; r)$, so that $z = \text{Gen}_{\text{tlp}}(1^\lambda, t, (s||r); \mathcal{O}'(1^\lambda, t, s, r))$. Whenever $(1^\lambda, t, s, r) \notin F$, it follows that the output of \mathcal{D}' is distributed as in the first probability in Equation 4.9. By the same argument, when z is a puzzle for 0^λ then the output of \mathcal{D}' is distributed according to the second probability in Equation 4.9, so long as $(1^\lambda, t, 0^\lambda, d) \notin F$. As r is chosen uniformly at random from $\{0, 1\}^{\lambda_{\text{tlp}}}$, it follows that the probability that this occurs is at most

$$\frac{f_{\text{nm}}(\lambda)}{2^{\lambda_{\text{tlp}}}} = \frac{(q_{\mathcal{B}}(\lambda, 2^m) \cdot 4q(\lambda))^2}{2^{\lambda_{\text{tlp}}}} = \frac{(q_{\mathcal{B}}(\lambda, 2^m) \cdot 4q(\lambda))^2}{2^{m+\lambda}} \in \frac{\text{poly}(\lambda, 2^m)}{2^{m+\lambda}} \leq \text{negl}'(\lambda)$$

which is negligible, where we used the fact that $\lambda_{\text{tlp}} = m + \lambda$. Therefore

$$\left| \Pr \left[\begin{array}{l} \mathcal{O}' \leftarrow \text{RF}_\lambda \\ z \leftarrow \text{Gen}^{\mathcal{O}'}(1^\lambda, t, s) \end{array} : \mathcal{D}'^{\mathcal{O}'}(z) = 1 \right] \right. \\ \left. - \Pr \left[\begin{array}{l} \mathcal{O}' \leftarrow \text{RF}_\lambda \\ z \leftarrow \text{Gen}^{\mathcal{O}'}(1^\lambda, t, 0^\lambda) \end{array} : \mathcal{D}'^{\mathcal{O}'}(z) = 1 \right] \right| \geq \mu(\lambda) - \text{negl}'(\lambda).$$

It will follow that \mathcal{D}' succeeds at breaking the hardness of nmTLP_m as long as its size is bounded by $\text{poly}(B(\lambda))$ and its depth is bounded by $(T(\lambda))^\epsilon \cdot p_{\text{hard}}(\lambda)$ for some T with $\ell_{\text{hard}}(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$, where $\ell_{\text{hard}}, p_{\text{hard}}$ are the positive polynomials guaranteed by the hardness of nmTLP_m . We bound the size and depth required for each step of \mathcal{D}' next, and then bound T . For the size and depth of \mathcal{D}' , we have that:

1. The first step requires running \mathcal{A} and checking if each of its queries appears in F , and if so returning the correct answer, and so requires size $\text{poly}(|\mathcal{A}|, |F|)$ and can be done in depth $\text{depth}(\mathcal{A}) \cdot \text{poly}(\lambda, \log |F|)$ since for each of \mathcal{A} 's queries, it takes depth $\text{poly}(\lambda, \log |F|)$ to check in parallel if the query appears in F and output the answer.
2. The second step can be done with size $\text{poly}(\lambda, m, \log t, |\mathcal{Q}|)$ and depth $\text{poly}(\lambda, m, \log t)$ since it requires generating a puzzle in parallel for each $i \in [n(\lambda)]$ and $j \in [|\mathcal{Q}|]$.
3. The third step requires computing $y = f(\vec{s})$. Since the input length is $n(\lambda)$, this can be done in depth $\text{poly}(\lambda, \log n(\lambda))$ and size $\text{poly}(B(\lambda))$ by assumption on $f \in \mathcal{F}_B^m$.
4. The third step requires running $\text{tt}_{\mathcal{D}'}(y)$, which can be done with size $\text{poly}(\lambda, 2^m)$ and depth $\text{poly}(\lambda, m)$.

Putting everything together, we have that

$$\begin{aligned} \text{size}(\mathcal{D}') &\in \text{poly}(\lambda, m, n(\lambda), B(\lambda), \log t, |\mathcal{Q}|, |\mathcal{A}|, |F|, 2^m) \\ &\in \text{poly}(\lambda, m, B(\lambda), f_{\text{nm}}(\lambda), 2^m) \\ &\in \text{poly}(\lambda, m, 2^m) \\ &\in \text{poly}(B(\lambda)), \end{aligned}$$

⁸This trick, of "flattening" \mathcal{D} using its truth table, was used in this context [DKP20].

where we used the fact that $n(\lambda) \in \lambda^{O(1)}$, $B(\lambda) = 2^m \cdot \text{poly}(\lambda)$, and $|F| = f_{\text{nm}}(\lambda) = (q_B(\lambda, 2^m) \cdot 4q(\lambda))^2$. For the depth,

$$\begin{aligned} \text{depth}(\mathcal{D}') &\leq \text{depth}(\mathcal{A}) \cdot \text{poly}(\lambda, \log |F|) + \text{poly}(\lambda, m, \log n(\lambda), \log T(\lambda)) \\ &\leq \text{depth}(\mathcal{A}) \cdot \text{poly}(\lambda, m, \log n(\lambda)) \\ &\leq \text{depth}(\mathcal{A}) \cdot p'(\lambda) \end{aligned}$$

for a fixed polynomial p' (which depends only on m , $\log n$, and $\log f_{\text{hard}}(\lambda)$), where we used the fact that $m(\lambda) \in \lambda^{O(1)}$ and $T(\lambda) \leq B(\lambda)$. Note that $\log(n(\lambda))$ can be bounded by λ for sufficiently large λ . Similarly, $\log f_{\text{hard}}(\lambda)$ can be bounded by a fixed polynomial in λ which depends on $\log q(\lambda)$. As q is a polynomial, this can also be bounded by a fixed polynomial in λ independently of q , for sufficiently large λ . Moreover, we can simply have \mathcal{D}' output \perp on security parameters which are not sufficiently large. Therefore, $\text{depth}(\mathcal{D}') \leq p'(\lambda)$ for all $\lambda \in \mathbb{N}$. Recall that $\text{depth}(\mathcal{A}) \leq (T(\lambda))^\epsilon \cdot p(\lambda)$, where we set $p(\lambda) = p_{\text{hard}}(\lambda)/p'(\lambda)$. Therefore, the above is bounded by $(T(\lambda))^\epsilon \cdot p_{\text{hard}}(\lambda)$.

Finally, to bound T , we have that $T(\lambda) \geq \ell(\lambda) \geq \ell_{\text{hard}}(\lambda)$ and $T(\lambda) \leq B(\lambda)$. It follows that \mathcal{D}' breaks the hardness of nmTLP_m with probability $\mu(\lambda) - \text{negl}'(\lambda)$, which completes the claim. \blacksquare

This completes the proof of Lemma 4.11. \square

4.4 Impossibility of Fully Concurrent Non-malleability

In the following theorem, we show that there does not exist a concurrent non-malleable time-lock puzzle. Specifically, consider a time-lock puzzle with message length L that has puzzles of length at most L' . We give an explicit attack that violates n -concurrent non-malleability for $n = \lceil L'/L \rceil$. As L, L' are polynomial, this is an explicit polynomial for which non-malleability cannot hold. Furthermore, this attack works even in the random oracle model. This means that for a time-lock puzzle to satisfy n -concurrent non-malleability, the puzzles must be sufficiently long. In the context of functional non-malleability, the following theorem shows that concurrent functional non-malleability is impossible to achieve for any class \mathcal{F} that contains the identity function.

Theorem 4.17. *Let $B, L, L': \mathbb{N} \rightarrow \mathbb{N}$, and $\epsilon \in (0, 1)$ where $B(\lambda) \in 2^{\text{poly}(\lambda)}$ and $L \in \text{poly}(\lambda)$. Suppose that (Gen, Sol) is a (B, ϵ) -secure time-lock puzzle for messages of length $L(\lambda)$ with puzzles of length at most $L'(\lambda)$. Then, (Gen, Sol) does not satisfy $n(\lambda)$ -concurrent non-malleability for $n(\lambda) = \lceil L'(\lambda)/L(\lambda) \rceil$.*

Proof. We note that as n -concurrent non-malleability implies one- n non-malleability, it suffices to break one- n non-malleability.

Let p, ℓ be positive polynomials. For any function T , we define a $(1, n, B, \epsilon, T, p)$ -MIM adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ as follows. For any $\lambda \in \mathbb{N}$, let $L = L(\lambda)$, $L' = L'(\lambda)$, and $n = n(\lambda) = \lceil L'/L \rceil$. Note that, since Gen is a PPT algorithm, L' is at most $\text{poly}(\lambda, L(\lambda), \log T(\lambda))$ for any T . This implies that L' and n are both bounded by polynomials. On input a puzzle $z \in \{0, 1\}^{L'}$, \mathcal{A}_λ splits z into n parts $z_1, \dots, z_n \in \{0, 1\}^L$, padding the last part with zeroes if necessary. \mathcal{A}_λ outputs $\tilde{z}^{(1)}, \dots, \tilde{z}^{(n)}$ where $\tilde{z}^{(i)} \leftarrow \text{Gen}(1^\lambda, T(\lambda), z_i)$ for all $i \in [n]$. Note that the size (and depth) of \mathcal{A}_λ is at most $n \cdot q(\lambda, \log T(\lambda))$ for some polynomial q since Gen is a PPT algorithm. Thus, \mathcal{A} is a valid $(1, n, B, \epsilon, T, p)$ -MIM adversary for any function T such that $\ell(\lambda) + n \cdot q(\lambda, \log T(\lambda)) \leq (T(\lambda))^\epsilon \cdot p(\lambda)$. In particular, this is true for $T(\lambda) = (\ell(\lambda) + n \cdot (\log B(\lambda)) \cdot q(\lambda))^{1/\epsilon} / p(\lambda)$, which is a polynomial as ϵ is a constant and $B(\lambda) \in 2^{\text{poly}(\lambda)}$. We show for this function T , \mathcal{A} violates one- n non-malleability.

For any $\lambda \in \mathbb{N}$ and message $s \in \{0, 1\}^\lambda$, consider the following distinguisher \mathcal{D} for the MIM distribution of \mathcal{A}_λ . \mathcal{D} gets as input values $\tilde{s}^{(1)}, \dots, \tilde{s}^{(n)}$ corresponding to $(\tilde{z}^{(1)}, \dots, \tilde{z}^{(n)})$. For each

$i \in [n]$, \mathcal{D} computes $\hat{s}_i \leftarrow \text{Sol}(1^\lambda, T(\lambda), \tilde{z}^{(i)})$ and computes \hat{z} to be the first L' bits of $\hat{s}_1 \| \dots \| \hat{s}_n$. \mathcal{D} then computes $s' = \text{Sol}(1^\lambda, T(\lambda), \hat{z})$. \mathcal{D} outputs 1 if $s = s'$ and 0 otherwise. By correctness of (Gen, Sol) , it holds that \mathcal{D} outputs 1 only on input $\text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), s)$, which contradicts one- n non-malleability of (Gen, Sol) , as required. Furthermore, we note that \mathcal{D} only needs to run in depth $T(\lambda) \cdot \text{poly}(\lambda, \log T(\lambda))$. \square

5 Publicly Verifiable Non-Malleable Time-Lock Puzzles

In this section, we define and construct a publicly verifiable time-lock puzzle (PV TLP) that is additionally non-malleable as in Section 4. At a high level, a PV TLP is one where the Sol function additionally outputs a succinct proof of correctness that can be checked by a Verify function.

In order for the proof to be sound, we use a weak form of setup independent of a cheating prover’s advice. Specifically, we rely on what we call the all-but-one (ABO) string model. In this model, the Sol and Verify algorithms take as input a multi-common random string, $\text{mcrcs} \in (\{0, 1\}^\lambda)^n$ for some $n \in \mathbb{N}$, and we require soundness to hold as long as a single random string is honest. In the case that $n = 1$, this corresponds to the standard common random string model. The ABO-string model is also very related to the multi-string model of [GO14], except that the ABO-string model requires only one string—as opposed to a majority of strings—to be honestly generated.

While the ABO-string model is a weak form of setup, we prove security in the relatively strong auxiliary-input random oracle model (AI-ROM). At a high level, we do this to ensure that the puzzles generated by Gen are independent of *any* setup, while also being able to prove a meaningful notion of security in essentially the “plain” random oracle model. Furthermore, the combination is realistic for our application in Section 6.

We define a publicly verifiable time-lock puzzle in the ABO-string model as follows. We note that we don’t explicitly define the functions and properties with respect a random oracle, and defer such a treatment to the proofs of security.

Definition 5.1 (Publicly Verifiable Time-Lock Puzzle). *A tuple $(\text{Gen}, \text{Sol}, \text{Verify})$ is a (B, ϵ) -secure publicly-verifiable time-lock puzzle in the ABO-string model if $(\text{Gen}, \text{Sol}, \text{Verify})$ have the following syntax:*

- $z \leftarrow \text{Gen}(1^\lambda, t, s)$: *A PPT algorithm that on input a security parameter $\lambda \in \mathbb{N}$, a difficulty parameter $t \in \mathbb{N}$, and a solution $s \in \{0, 1\}^\lambda$, outputs a puzzle $z \in \{0, 1\}^*$, and a proof $\pi \in \{0, 1\}^*$.*
- $(s, \pi) \leftarrow \text{Sol}(1^\lambda, \text{mcrcs}, t, z)$: *A randomized algorithm that on input the security parameter $\lambda \in \mathbb{N}$, a multi-common random string $\text{mcrcs} \in (\{0, 1\}^\lambda)^*$, a difficulty parameter $t \in \mathbb{N}$, and a puzzle z , outputs a solution $s \in (\{0, 1\}^\lambda \cup \{\perp\})$ and a proof $\pi \in \{0, 1\}^*$. We denote Sol_1 and Sol_2 as the first and second outputs of Sol , respectively.*
- $b = \text{Verify}(1^\lambda, \text{mcrcs}, t, z, (s, \pi))$: *A deterministic algorithm that on input a security parameter $\lambda \in \mathbb{N}$, a multi-common random string $\text{mcrcs} \in (\{0, 1\}^\lambda)^*$, a difficulty parameter $t \in \mathbb{N}$, a puzzle z , a possible solution $s \in (\{0, 1\}^\lambda \cup \{\perp\})$, and a proof π , outputs a bit b indicating whether to accept or reject.*

We require that $(\text{Gen}, \text{Sol}, \text{Verify})$ satisfy the following properties.

- **Full correctness:** *For every $\lambda, t, n \in \mathbb{N}$, $\text{mcrcs} \in (\{0, 1\}^\lambda)^n$, and $z \in \{0, 1\}^*$, the following hold:*
 - *If $z \in \text{Supp}(\text{Gen}(1^\lambda, t, s))$ for some $s \in \{0, 1\}^\lambda$, then $\text{Sol}_1(1^\lambda, \text{mcrcs}, t, z) = s$.*

- If $z \notin \text{Supp}(\text{Gen}(1^\lambda, t, s))$ for any $s \in \{0, 1\}^\lambda$, then $\text{Sol}_1(1^\lambda, \text{mcrs}, t, z) = \perp$.
- **Efficiency:** There exist a polynomial p such that for all $\lambda, t, n \in \mathbb{N}$, $\text{mcrs} \in (\{0, 1\}^\lambda)^n$, $\text{Sol}(1^\lambda, \text{mcrs}, t, \cdot)$ is computable in time $t \cdot p(\lambda, \log t, n)$.
- **(B, ϵ) -Hardness:** The same as for time-lock puzzles as in Definition 3.1.
- **Completeness:** For any $\lambda, t, n \in \mathbb{N}$, $z \in \{0, 1\}^*$, $\text{mcrs} \in (\{0, 1\}^\lambda)^n$, it holds that

$$\text{Verify}(1^\lambda, \text{mcrs}, t, z, \text{Sol}(1^\lambda, \text{mcrs}, t, z)) = 1.$$

- **Soundness:** For all non-uniform probabilistic polynomial-time adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ and polynomials T , there exists a negligible function negl such that for all $\lambda, n \in \mathbb{N}$ and $i \in [n]$, it holds that

$$\Pr \left[\begin{array}{l} \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (z, s', \pi, \text{crs}_{-i}) \leftarrow \mathcal{A}_\lambda(1^\lambda, \text{crs}_i, T(\lambda)) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ s = \text{Sol}_1(1^\lambda, \text{mcrs}, T(\lambda), z) \end{array} : \begin{array}{l} \text{Verify}(1^\lambda, \text{mcrs}, T(\lambda), z, (s', \pi)) = 1 \\ \wedge s \neq s' \end{array} \right] \leq \text{negl}(\lambda).$$

We note that the above notion of full correctness is stronger than the standard definition of correctness for TLPs given in Definition 3.1. Also, the soundness notion is strong in the sense that the adversary can try to break soundness even for invalid puzzles.

Towards achieving this strong definition, we define a notion of a *one-sided* publicly verifiable time-lock puzzle in which correctness holds only for z in the support of Gen and soundness requires that the adversary cheats on puzzles in the support of Gen . We formalize this as follows.

Definition 5.2 (One-sided PV TLP). *A tuple $(\text{Gen}, \text{Sol}, \text{Verify})$ is a (B, ϵ) -secure one-sided publicly-verifiable time-lock puzzle in the ABO-string model if correctness holds only for $z \in \text{Supp}(\text{Gen}(1^\lambda, t, \cdot))$ and the soundness property is replaced with the following:*

- **One-sided Soundness:** For all non-uniform probabilistic polynomial-time adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ and polynomials T , there exists a negligible function negl such that for all $\lambda, n \in \mathbb{N}$ and $i \in [n]$, it holds that

$$\Pr \left[\begin{array}{l} \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (z, s', \pi, \text{crs}_{-i}) \leftarrow \mathcal{A}_\lambda(1^\lambda, \text{crs}_i, T(\lambda)) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ s = \text{Sol}_1(1^\lambda, \text{mcrs}, T(\lambda), z) \end{array} : \begin{array}{l} \text{Verify}(1^\lambda, \text{mcrs}, T(\lambda), z, (s', \pi)) = 1 \\ \wedge s \neq s' \\ \wedge z \in \text{Supp}(\text{Gen}(1^\lambda, T(\lambda), \cdot)) \end{array} \right] \leq \text{negl}(\lambda).$$

In Section 5.1, we formalize the notion of a strong trapdoor verifiable delay function (VDF), which satisfies the requirements we need for the underlying sequential function. We then give a construction based on the repeated squaring assumption. We note that it may be possible to give a construction that satisfies the necessary properties based on randomized encodings as in [BGJ⁺16], but we instead focus on a more practical construction.

In Section 5.2, we formalize the construction of a one-sided PV TLP given a strong trapdoor VDF. Finally in Section 5.3, we construct a full PV TLP by applying our non-malleability transformation of Section 4.

5.1 Strong Trapdoor VDFs

A trapdoor VDF provides a way to generate inputs to a function that take a long time to compute. At the same time, the function can be computed efficiently using a trapdoor, and even without the trapdoor, can be efficiently verified given a proof. Trapdoor VDFs were first defined by Wesolowski [Wes19] as an extension of standard VDFs [BBBF18].

We define a strong notion of a trapdoor VDF in the ABO-string model. While the formal definition is highly tailored towards our definition of PV TLPs and application to multi-party coin-flipping, we believe that some of the stronger requirements we define—and achieve—may be of independent interest. Conceptually, we require the following additional properties over previous definitions of VDFs (or trapdoor VDFs):

1. We have no setup algorithm, and instead are in the ABO-string model. In particular, this means that sampling inputs for the VDF can be done completely independently of any trusted setup.
2. We allow the sampling procedure to specify the domain \mathcal{X} of the evaluation function.
3. We require that completeness must hold for all inputs $x \in \{0, 1\}^*$ and domains $\mathcal{X} \in \{0, 1\}^*$ rather than only honestly generated inputs.
4. We require soundness to hold for any adversarially chosen—yet in the support of the Sample algorithm—input x and domain \mathcal{X} , rather than with high probability over randomly sampled x and \mathcal{X} .
5. We require a natural encoding property for elements in the domain \mathcal{X} with strings, and additionally require that \mathcal{X} defines a natural bijection that is amenable to being used as a one-time pad (like \oplus for strings or $+$ for rings, for example).

We formally define the requirements of a strong trapdoor VDF in the ABO-string model as follows.

Definition 5.3. *Let $B: \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon \in (0, 1)$. A (B, ϵ) -secure strong trapdoor verifiable delay function in the ABO-string model is a tuple $(\text{Sample}, \text{Eval}, \text{TDEval}, \text{Verify})$ with the following syntax:*

- $(x, \mathcal{X}, \text{td}) \leftarrow \text{Sample}(1^\lambda, t)$: *A PPT algorithm that on input a security parameter $\lambda \in \mathbb{N}$, a difficulty parameter $t \in \mathbb{N}$, outputs a value x in a specified domain \mathcal{X} , and a trapdoor $\text{td} \in \{0, 1\}^*$.*
- $(y, \pi) \leftarrow \text{Eval}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}))$: *An algorithm that on input a security parameter $\lambda \in \mathbb{N}$, a difficulty parameter $t \in \mathbb{N}$, and a value x in a specified domain \mathcal{X} , outputs a value $y \in \mathcal{X}$ and a proof $\pi \in \{0, 1\}^*$. We denote Eval_1 and Eval_2 as the first and second outputs of Eval, respectively, and require that Eval_1 is a deterministic function.*
- $y = \text{TDEval}(1^\lambda, t, (x, \mathcal{X}), \text{td})$: *A polynomial-time algorithm that on input a security parameter $\lambda \in \mathbb{N}$, a difficulty parameter $t \in \mathbb{N}$, a value x in a specified domain \mathcal{X} , and a trapdoor $\text{td} \in \{0, 1\}^*$, outputs a value $y \in \mathcal{X}$.*
- $b = \text{Verify}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}), (y, \pi))$: *A polynomial-time algorithm that on input a security parameter $\lambda \in \mathbb{N}$, a difficulty parameter $t \in \mathbb{N}$, values x, y in a specified domain \mathcal{X} , and a proof $\pi \in \{0, 1\}^*$, outputs a bit b indicating whether to accept or reject.*

We require that $(\text{Sample}, \text{Eval}, \text{TDEval}, \text{Verify})$ satisfy the following properties.

- **Completeness:** For every $\lambda, t, n \in \mathbb{N}$, $x, \mathcal{X} \in \{0, 1\}^*$, $\text{mcrs} \in (\{0, 1\}^\lambda)^n$, it holds that

$$\text{Verify}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}), \text{Eval}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}))) = 1.$$

- **Soundness:** For all non-uniform PPT adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ and polynomials T , there exists a negligible function negl such that for all $\lambda, n \in \mathbb{N}$, $i \in [n]$, it holds that

$$\Pr \left[\begin{array}{l} \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (x, \mathcal{X}, y', \pi, \text{crs}_{-i}) \\ \leftarrow \mathcal{A}_\lambda(1^\lambda, \text{crs}_i, T(\lambda), n) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ y = \text{Eval}_1(1^\lambda, \text{mcrs}, T(\lambda), (x, \mathcal{X})) \end{array} : \begin{array}{l} \text{Verify}(1^\lambda, \text{mcrs}, T(\lambda), (x, \mathcal{X}), (y', \pi)) = 1 \\ \wedge y \neq y' \\ \wedge (x, \mathcal{X}, *) \in \text{Supp}(\text{Sample}(1^\lambda, T(\lambda))) \end{array} \right] \leq \text{negl}(\lambda).$$

- **Trapdoor Evaluation:** For every $\lambda, t, n \in \mathbb{N}$, $(x, \mathcal{X}, \text{td}) \in \text{Supp}(\text{Sample}(1^\lambda, t))$, and $\text{mcrs} \in (\{0, 1\}^\lambda)^n$, it holds that

$$\text{Eval}_1(1^\lambda, \text{mcrs}, t, (x, \mathcal{X})) = \text{TDEval}(1^\lambda, t, (x, \mathcal{X}), \text{td}).$$

- **Honest Evaluation:** There exists a polynomial p such that for all $\lambda, t, n \in \mathbb{N}$, and $\text{mcrs} \in (\{0, 1\}^\lambda)^n$, $\text{Eval}(1^\lambda, \text{mcrs}, t, \cdot)$ is computable in time $t \cdot p(\lambda, \log t, n)$.
- **(B, ϵ) -Sequentiality:** There exist positive polynomials p, ℓ such that for all functions T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$ and every non-uniform adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ where $\text{size}(\mathcal{A}_\lambda) \in \text{poly}(B(\lambda))$ and $\text{depth}(\mathcal{A}_\lambda) \leq (T(\lambda))^\epsilon \cdot p(\lambda)$ for all $\lambda \in \mathbb{N}$, there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$, $n \in \mathbb{N}$, $\text{mcrs} \in (\{0, 1\}^\lambda)^n$,

$$\left| \Pr \left[\begin{array}{l} (x, \mathcal{X}, \text{td}) \leftarrow \text{Sample}(1^\lambda, T(\lambda)) \\ y = \text{Eval}_1(1^\lambda, \text{mcrs}, T(\lambda), (x, \mathcal{X})) \end{array} : \mathcal{A}_\lambda(x, \mathcal{X}, y) = 1 \right] \right. \\ \left. - \Pr \left[\begin{array}{l} (x, \mathcal{X}, \text{td}) \leftarrow \text{Sample}(1^\lambda, T(\lambda)) \\ r \leftarrow \mathcal{X} \end{array} : \mathcal{A}_\lambda(x, \mathcal{X}, r) = 1 \right] \right| \leq \text{negl}(\lambda).$$

- **Encoding:** For $\lambda, t \in \mathbb{N}$, and any domain \mathcal{X} output by $\text{Sample}(1^\lambda, t)$, it holds that strings in $\{0, 1\}^\lambda$ can be uniquely encoded as elements in \mathcal{X} , and elements in \mathcal{X} can be uniquely decoded to elements in $\{0, 1\}^*$. Additionally, for any element $x \in \mathcal{X}$, there is an efficiently computable bijective map $f_x: \mathcal{X} \rightarrow \mathcal{X}$, written as $f_x(y) = x \oplus y$.

In the above definition, we note that only Eval and Verify receive as input the multi-common random string mcrs , as they require it for public verifiability. In particular, TDEval can be computed without access to mcrs since we do not require it to output a proof of correctness (in fact, the trapdoor itself can be thought of as a *privately verifiable* proof). By trapdoor evaluation, this implies that the output y of the function is actually independent of mcrs .

Candidate strong trapdoor VDF. Our candidate strong trapdoor VDF is based on repeated squaring with a publicly verifiable proof of correctness from Pietrzak [Pie19]. As in [Pie19], we use the group $\mathbb{G} = \text{QR}_N^+$ of signed quadratic residues mod N , where N is the product of safe primes p, q such that $(p-1)/2$ and $(q-1)/2$ are prime. We note that QR_N^+ has size $|\text{QR}_N^+| = (p-1) \cdot (q-1)/4$ and has the property that its only subgroups are of size $(p-1)/2$ and $(q-1)/2$ which are both at least 2^λ by construction. Elements in this group can be encoded as integers in $[0, (N-1)/2]$. For any two elements $a, b \in \text{QR}_N^+$, we can define multiplication (and similarly addition) by computing

$x = a \cdot b \bmod N$ and taking the smaller of x and $N - x$, which is in $[0, (N - 1)/2]$. For further discussion of the group, we refer the reader to [Pie19].

To generate such a group, we can sample random numbers in $[2^{\lambda+1}, 2^{\lambda+2})$ until we find two safe primes p and q and output QR_N^+ where $N = p \cdot q$. As discussed in [Pie19], it is conjectured (in [VZGS13]) that for some constant c , there are $c \cdot 2^\lambda / \lambda^2$ safe λ -bit primes. Under this conjecture, it takes expected polynomial time to sample N which is a product of two safe primes. Rather than doing this, we define a group generator algorithm $\text{RSWGen}(1^\lambda)$ that searches for safe primes in $[2^{\lambda+1}, 2^{\lambda+2})$ for some *fixed* polynomial time until it fails to halt with probability at least $2^{-\lambda}$. Specifically, suppose it takes expected $p(\lambda)$ time to find two safe primes. If we run for $2 \cdot \lambda \cdot p(\lambda)$ time, we will halt with probability at least $2^{-\lambda}$. When this failure event occurs, $\text{RSWGen}(1^\lambda)$ deterministically searches for safe primes starting with $2^{\lambda+1}, 2^{\lambda+1} + 1, \dots$ until finding the first two safe primes (alternatively, we could hard code these safe primes with preprocessing). Assuming that the safe primes are evenly spread out, this will only need to search through $O(\lambda^2)$ numbers. We additionally define $\text{RSWGen}(1^\lambda)$ to output a random group element and the size of the group to be used as a trapdoor. In summary, we have the following:

- $(g, \mathbb{G}, |\mathbb{G}|) \leftarrow \text{RSWGen}(1^\lambda)$: A PPT algorithm that outputs $(g, \mathbb{G}, |\mathbb{G}|)$ where $\mathbb{G} = \text{QR}_N^+$ such that N is the product of two safe primes p and q , $g \leftarrow \mathbb{G}$ is a random element in the group, and $|\mathbb{G}| = (p - 1) \cdot (q - 1)/4$ is the size of the group. It holds that with probability at least $1 - 2^{-\lambda}$, p and q are uniformly random safe primes in $[2^{\lambda+1}, 2^{\lambda+2})$.

We next define formalize the repeated squaring assumption we make for RSWGen . We note that this assumption for QR_N^+ is implied by the more standard repeated squaring assumption for \mathbb{Z}_N^* with at most a factor of 8 loss in advantage, as discussed in [Pie19].

Assumption 5.4 (Repeated Squaring Assumption). *Let $B: \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon \in (0, 1)$. We say that the (B, ϵ) -repeated squaring assumption for RSWGen holds if there exist functions ℓ, p such that for all functions T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$ and every non-uniform distinguisher $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ where $\text{size}(\mathcal{A}_\lambda) \in \text{poly}(B(\lambda))$ and $\text{depth}(\mathcal{A}_\lambda) \leq (T(\lambda))^\epsilon \cdot p(\lambda)$ for all $\lambda \in \mathbb{N}$, there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$,*

$$\left| \Pr \left[\begin{array}{l} (g, \mathbb{G}, |\mathbb{G}|) \leftarrow \text{RSWGen}(1^\lambda) \\ y = g^{2^{T(\lambda)}} \end{array} : \mathcal{A}_\lambda(g, \mathbb{G}, y) = 1 \right] \right. \\ \left. - \Pr \left[\begin{array}{l} (g, \mathbb{G}, |\mathbb{G}|) \leftarrow \text{RSWGen}(1^\lambda) \\ r \leftarrow \mathbb{G} \end{array} : \mathcal{A}_\lambda(g, \mathbb{G}, r) = 1 \right] \right| \leq \text{negl}(\lambda).$$

Proof of repeated squaring. We next discuss the publicly verifiable proof of repeated squaring for the group QR_N^+ given in Pietrzak [Pie19]. We use a slightly modified version of this protocol, which is in the ABO-string model and where we explicitly assume access to a random oracle $\mathcal{O}: \{0, 1\}^* \rightarrow [2^\lambda]$.

We first describe the interactive protocol of [Pie19]. In order to prove the claim that $x^{2^T} = y$ in QR_N^+ ⁹, the interactive protocol does the following. If $T = 1$, the verifier directly checks that $x^2 = y$. Otherwise, the prover sends to the verifier the value $\mu = x^{2^{T/2}}$ and the verifier replies with $r \leftarrow [2^\lambda]$. The prover and verifier recursively engage in a protocol to prove that $(x^r \mu)^{T/2} = \mu^r y$. This interactive proof consists of $2 \cdot \log T$ rounds of communication, where the prover sends a single group element and the verifier responds with a λ -bit challenge. The proof has soundness error at most $3 \cdot \log T / 2^\lambda$ even against unbounded cheating provers. In particular, [Pie19] shows that at

⁹we assume for simplicity that T is a power of 2, which can be easily dealt with as in [Pie19] if this is not the case

there are at most 3 “bad challenges” that the verifier might send in each round of the protocol, in the sense that a bad challenge might cause the verifier to wrongly accept false statements.

In order to make the above protocol non-interactive via the Fiat-Shamir heuristic, the prover uses the random oracle \mathcal{O} on the transcript so far to generate the random challenges of the verifier itself. The verifier then accepts the proof if all challenges are consistent with H and the interactive verifier would have accepted. We emphasize that in our version of the protocol, the oracle \mathcal{O} is indexed with mcrs , which is assumed to have at least λ bits of entropy *independent* of any possible adversary. Following the above idea, we formalize the algorithms RSWProve and RSWVerify that we use:

- $\pi = \text{RSWProve}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y)$:
 1. Let $x_0 = g, y_0 = y$.
 2. For $i = 1, \dots, \log t$,
 - (a) Compute $\pi_i = x_{i-1}^{2^{t/2^i}}$.
 - (b) Let $r_i = \mathcal{O}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y, \pi_1, \dots, \pi_i)$.
 - (c) Compute $x_i = x_{i-1}^{r_i}$ and $y_i = \pi_i^{r_i} \cdot y_{i-1}$.
 3. Output $\pi = (\pi_1, \dots, \pi_{\log t})$.
- $b = \text{RSWVerify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y, \pi)$:
 1. Let $x_0 = g, y_0 = y$.
 2. For $i = 1, \dots, \log t$,
 - (a) Let $r_i = \mathcal{O}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y, \pi_1, \dots, \pi_i)$.
 - (b) Compute $x_i = x_{i-1}^{r_i}$ and $y_i = \pi_i^{r_i} \cdot y_{i-1}$.
 3. Output 1 if and only if $x_{\log t}^2 = y_{\log t}$.

We note that RSWProve can be computed in time $t \cdot \text{poly}(\lambda, \log t, n)$ and RSWVerify can be computed in time $\text{poly}(\lambda, \log t, n)$. We note that [Pie19] shows how to compute RSWProve in time $t + \sqrt{t}$ when using \sqrt{t} memory from computing $y = g^{2^t}$ (ignoring $\text{poly}(\lambda, n)$ factors).

In the following lemmas, we show that these algorithms satisfy the following completeness and soundness properties.

Lemma 5.5 (Completeness). *For any $\lambda, t, n \in \mathbb{N}$, \mathbb{G} which is a valid representation of QR_N^+ for some N , $g \in \mathbb{G}$, and $\text{mcrs} \in (\{0, 1\}^\lambda)^n$, let $y = g^{2^t}$ and $\pi = \text{RSWProve}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y)$. Then, it holds that*

$$\text{RSWVerify}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y, \pi) = 1.$$

Completeness follows immediately from [Pie19].

Lemma 5.6 (Soundness). *For any polynomial T and unbounded algorithm Z that on input $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$ outputs polynomial-size circuits, there exists a negligible function negl such that for all $\lambda, n \in \mathbb{N}$, $i \in [n]$, it holds that*

$$\Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (g, \mathbb{G}, y', \pi, \text{crs}_{-i}) \\ \quad \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda, \text{crs}_i, T(\lambda), n) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \end{array} \quad : \quad \begin{array}{l} \text{RSWVerify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, T(\lambda), (g, \mathbb{G}), y', \pi) = 1 \\ \wedge y' \neq g^{2^{T(\lambda)}} \\ \wedge (g, \mathbb{G}, *) \in \text{Supp}(\text{RSWGen}(1^\lambda, T(\lambda))) \end{array} \right] \leq \text{negl}(\lambda).$$

At a high level, the proof follows by a simple application of Lemma 3.6 of Unruh [Unr07] followed by an analysis of Pietrzak’s VDF [Pie19] in the (plain) random oracle model. We give the proof in Appendix B.

Strong trapdoor VDF construction. We are now ready to state our strong trapdoor VDF construction $\text{VDF} = (\text{Sample}_{\text{vdf}}, \text{Eval}_{\text{vdf}}, \text{TDEval}_{\text{vdf}}, \text{Verify}_{\text{vdf}})$ in the ABO-string model. We give Eval_{vdf} and $\text{Verify}_{\text{vdf}}$ access to an oracle function \mathcal{O} . We note that we can efficiently check if \mathbb{G} is a valid representation of QR_N^+ for some N , and we can efficiently check membership in $\mathbb{G} = \text{QR}_N^+$. In particular, we say that (g, \mathbb{G}) are valid if \mathbb{G} can be parsed as a valid representation of QR_N^+ and $g \in \mathbb{G}$.

- $(g, \mathbb{G}, |\mathbb{G}|) \leftarrow \text{Sample}_{\text{vdf}}(1^\lambda, t)$:
 1. Output $(g, \mathbb{G}, |\mathbb{G}|) \leftarrow \text{RSWGen}(1^\lambda)$.
- $(y, \pi) \leftarrow \text{Eval}_{\text{vdf}}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}))$:
 1. If (g, \mathbb{G}) are invalid, output $y = \pi = \perp$.
 2. Otherwise, compute $y = g^{2^t}$ in \mathbb{G} and $\pi = \text{RSWProve}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y)$.
- $y = \text{TDEval}_{\text{vdf}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), |\mathbb{G}|)$:
 1. If (g, \mathbb{G}) are invalid, output $y = \perp$.
 2. Otherwise, output $y = g^{(2^t \bmod |\mathbb{G}|)}$.
- $b = \text{Verify}_{\text{vdf}}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), (y, \pi))$:
 1. If (g, \mathbb{G}) are invalid, output 1 if and only if $y = \pi = \perp$.
 2. Otherwise, output 1 if and only if $\text{RSWVerify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y, \pi) = 1$.

Theorem 5.7. *Assuming the (B, ϵ) -repeated squaring assumption holds for RSWGen , then there exists a (B, ϵ) -secure strong trapdoor VDF in the ABO-string model. Soundness holds in the auxiliary-input random oracle model.*

Each of the required properties follow almost directly from the definition of a strong trapdoor VDF. We provide the full proof for completeness in Appendix B.

5.2 One-sided PV TLPs from Strong Trapdoor VDFs

Let $\text{VDF} = (\text{Sample}_{\text{vdf}}, \text{Eval}_{\text{vdf}}, \text{TDEval}_{\text{vdf}}, \text{Verify}_{\text{vdf}})$ be any strong trapdoor VDF. Our construction is given by the tuple $\text{TLP} = (\text{Gen}_{\text{tlp}}, \text{Sol}_{\text{tlp}}, \text{Verify}_{\text{tlp}})$, where Sol_{tlp} and $\text{Verify}_{\text{tlp}}$ have oracle access to a function \mathcal{O} , defined as follows.

- $z \leftarrow \text{Gen}_{\text{tlp}}(1^\lambda, t, s)$:
 1. Compute $(x, \mathcal{X}, \text{td}) \leftarrow \text{Sample}_{\text{vdf}}(1^\lambda, t)$ and $y = \text{TDEval}_{\text{vdf}}(1^\lambda, t, (x, \mathcal{X}), \text{td})$.
 2. Encode s as an element $x_s \in \mathcal{X}$ and compute $x_c = x_s \oplus y$.
 3. Output $z = (x, \mathcal{X}, c)$.
- $(s, \pi) \leftarrow \text{Sol}_{\text{tlp}}(1^\lambda, \text{mcrs}, t, z)$:

1. Parse z as (x, \mathcal{X}, c) . If z cannot be parsed this way, output (\perp, \perp) .
 2. Compute $(y, \pi_{\text{vdf}}) \leftarrow \text{Eval}_{\text{vdf}}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}))$.
 3. Let $x_s = y \oplus c$, and let s be the string encoding of x_s .
 4. Output $(s, (y, \pi_{\text{vdf}}))$.
- $b = \text{Verify}_{\text{tlp}}(1^\lambda, \text{mcrs}, t, z, (s, \pi))$:
 1. Parse z as (x, \mathcal{X}, c) and π as (y, π_{vdf}) . If z cannot be parsed this way, output 1 if and only if $s = \pi = \perp$.
 2. Otherwise, output 1 if and only if $\text{Verify}_{\text{vdf}}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}), (y, \pi_{\text{vdf}})) = 1$ and $c \oplus y = x_s$ where x_s is the encoding of s in \mathcal{X} .

In the following theorem, we show that TLP is one-sided publicly verifiable time-lock puzzle in the ABO-string model, assuming that VDF is *any* strong trapdoor VDF. However, we emphasize that for our explicit construction of VDF, soundness holds in the auxiliary-input random oracle model, so we achieve the same soundness for our explicit TLP construction.

Theorem 5.8. *Suppose there exists a (B, ϵ) -secure strong trapdoor VDF. Then, there exists a (B, ϵ) -secure one-sided publicly verifiable time-lock puzzle.*

We provide a full proof of this theorem in Appendix B.

5.3 Non-Malleable PV TLP from One-sided PV TLPs

By applying a similar transformation as in Section 4 to any one-sided PV TLP, we achieve a publicly verifiable time-lock puzzle (with full correctness and soundness) that is additionally non-malleable. For completeness, we restate the transformation with the syntax of *publicly verifiable* time-lock puzzles, and note that we assume the same parameters as in Section 4. Let m be a polynomial, and let $\text{TLP}_{\text{os}}^m = (\text{Gen}_{\text{os}}, \text{Sol}_{\text{os}}, \text{Verify}_{\text{os}})$ be a one-sided PV TLP. We construct a non-malleable PV TLP $(\text{Gen}, \text{Sol}, \text{Verify})$ in the ABO-string model, where all algorithms have oracle access to a function \mathcal{O} .

- $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s; r)$:
 1. Get $r_{\text{os}} = \mathcal{O}(1^\lambda, t, s, r)$.
 2. Output $z \leftarrow \text{Gen}_{\text{os}}(1^{\lambda_{\text{os}}}, t, (s||r); r_{\text{os}})$.
- $(s, \pi) \leftarrow \text{Sol}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z)$:
 1. Compute $(s_{\text{os}}, \pi_{\text{os}}) = \text{Sol}_{\text{os}}(1^{\lambda_{\text{tp}}}, \text{mcrs}, t, z)$ and parse $s_{\text{os}} = s||r$.
 2. If $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s; r)$, output (s, r) .
 3. Otherwise, output $(\perp, (s_{\text{os}}, \pi_{\text{os}}))$.
- $b = \text{Verify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z, (s, \pi))$:
 1. If $s \neq \perp$, parse $\pi = r$. Output 1 if and only if $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s; r)$.
 2. If $s = \perp$, parse $\pi = (s_{\text{os}}, \pi_{\text{os}})$ and $s_{\text{os}} = s||r$. Output 1 if and only if $z \neq \text{Gen}^{\mathcal{O}}(1^\lambda, t, s; r)$ and $\text{Verify}_{\text{os}}(1^{\lambda_{\text{os}}}, \text{mcrs}, t, z, (s_{\text{os}}, \pi_{\text{os}})) = 1$.

Recall the class of depth-bounded functions \mathcal{F}_B^m with $m(\lambda)$ -bit outputs defined in Section 4. We get the following theorem.

Theorem 5.9 (Non-malleable PV TLP from one-sided PV TLP). *Let $B: \mathbb{N} \rightarrow \mathbb{N}$, $m(\lambda) \in \lambda^{O(1)}$, and $\epsilon \in (0, 1)$, where $\lambda \leq B(\lambda) \in 2^{m(\lambda)} \cdot \text{poly}(\lambda)$. Assuming the existence of a (B, ϵ) -secure one-sided publicly verifiable time-lock puzzle TLP_{os} in the ABO-string model, then there exists a (B, ϵ) -secure publicly verifiable time-lock puzzle in the ABO-string model. Soundness holds in the auxiliary-input random oracle model. Furthermore, the construction satisfies fully concurrent functional non-malleability for the class of functions \mathcal{F}_B^m .*

Before proving the above theorem, we state the following corollary by combining Theorems 5.7, 5.8, and 5.9. We emphasize that the resulting construction is proven secure in the auxiliary-input random oracle model.

Corollary 5.10 (NM PV TLP from Repeated Squaring). *Let $B: \mathbb{N} \rightarrow \mathbb{N}$, $m(\lambda) \in \lambda^{O(1)}$, and $\epsilon \in (0, 1)$, where $\lambda \leq B(\lambda) \in 2^{m(\lambda)} \cdot \text{poly}(\lambda)$. Assuming the (B, ϵ) -repeated squaring assumption holds for RSWGen , then there exists a (B, ϵ) -secure publicly verifiable time-lock puzzle in the ABO-string model. Soundness holds in the auxiliary-input random oracle model. Furthermore, the construction satisfies concurrent functional non-malleability for the class of functions \mathcal{F}_B^m .*

The proof of Theorem 5.9 follows by considering the construction $(\text{Gen}, \text{Sol}, \text{Verify})$ assuming $(\text{Gen}_{\text{os}}, \text{Sol}_{\text{os}}, \text{Verify}_{\text{os}})$ is a (B, ϵ) -secure one-sided publicly verifiable time-lock puzzle. The proofs of efficiency, hardness, and concurrent functional non-malleability follow immediately from Theorem 4.2. We provide proofs of full correctness and completeness in Appendix B. We proceed to prove that the construction satisfies the full notion of soundness in the auxiliary-input random oracle model assuming the underlying construction satisfies only one-sided soundness.

Lemma 5.11 (Soundness). *Assuming $(\text{Gen}_{\text{os}}, \text{Sol}_{\text{os}}, \text{Verify}_{\text{os}})$ is a (B, ϵ) -secure one-sided publicly verifiable time-lock puzzle, then the construction $(\text{Gen}, \text{Sol}, \text{Verify})$ satisfies the soundness property in the auxiliary-input random oracle model for publicly verifiable time-lock puzzles in the ABO-string model.*

Proof. Suppose by way of contradiction that $(\text{Gen}, \text{Sol}, \text{Verify})$ does not satisfy soundness. Namely, there exists a polynomial T , unbounded algorithm Z that outputs polynomial-size circuits, a polynomial q , and integers $n \in \mathbb{N}$, $i \in [n]$ such that for infinitely many $\lambda \in \mathbb{N}$, it holds that

$$\Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (z, s', \pi', \text{crs}_{-i}) \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda, \text{crs}_i, T(\lambda)) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ s = \text{Sol}_1^{\mathcal{O}}(1^\lambda, \text{mcrs}, T(\lambda), z) \end{array} : \begin{array}{l} \text{Verify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, T(\lambda), z, (s', \pi')) = 1 \\ \wedge s' \neq s \end{array} \right] > 1/q(\lambda).$$

Fix any λ for which this holds. Throughout the proof, we let $t = T(\lambda)$. Additionally, it will be helpful to define $(s, \pi) = \text{Sol}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z)$ to be the output of Sol for the values given by \mathcal{A} . Furthermore, by completeness, we know that $\text{Verify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z, (s, \pi)) = 1$.

We consider three possible events that may occur when \mathcal{A} succeeds in the above experiment. Either (1) $s \neq \perp$ and $s' \neq \perp$, (2) $s = \perp$ and $s' \neq \perp$, or (3) $s \neq \perp$ and $s' = \perp$. As $s \neq s'$ when \mathcal{A} succeeds, this covers all of the possible cases. We show that (1) and (2) cannot occur, and then proceed to bound the probability that (3) occurs.

For (1), suppose that it is the case that neither s nor s' are equal to \perp . We know that $\text{Verify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z, (s, \pi)) = 1$ and $\text{Verify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z, (s', \pi')) = 1$. This implies that z is both in the support of $\text{Gen}^{\mathcal{O}}(1^\lambda, t, s)$ and the support of $\text{Gen}^{\mathcal{O}}(1^\lambda, t, s')$, but this contradicts correctness since it means that $\text{Sol}_1^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z)$ must be equal to both s and s' .

For (2), suppose that $s = \perp$ and $s' \neq \perp$. Because $\text{Verify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z, (s', \pi')) = 1$ and $s' \neq \perp$, this means that $z = \text{Gen}^{\mathcal{O}}(1^\lambda, t, s'; r')$ where $r' = \pi'$. Then by correctness, it holds that $\text{Sol}_1^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, z) = s'$, which contradicts the facts that $s' \neq s$.

As a result, we know that (3) occurs with probability at least $1/q(\lambda)$, meaning that when \mathcal{A} wins, $s \neq \perp$ and $s' = \perp$. We show that this can be used to break the one-sided soundness of $(\text{Gen}_{\text{os}}, \text{Sol}_{\text{os}}, \text{Verify}_{\text{os}})$.

We define a non-uniform algorithm $\mathcal{B} = \{\mathcal{B}_\lambda\}_{\lambda \in \mathbb{N}}$ as follows. In order for \mathcal{B}_λ to use \mathcal{A} in the reduction, \mathcal{B}_λ needs to simulate the oracle calls that \mathcal{A} makes to \mathcal{O} . To do so, we need to make use of Lemma 3.6 of Unruh [Unr07]. Specifically, let $p(\lambda)$ be an upper bound on the size of the adversaries that Z outputs, which is polynomial by assumption. For the polynomial function $f(\lambda) = 4 \cdot (p(\lambda))^2 \cdot (q(\lambda))^2$, we consider the inefficient algorithm Sam that on input \mathcal{A} outputs a partial assignment F of size $f(\lambda)$. By Lemma 3.6, it holds that the output distribution of \mathcal{A} with access to $\mathcal{P} \leftarrow \text{RF}_\lambda[F]$ has statistical distance at most $\sqrt{p(\lambda)^2/f(\lambda)} = 1/(2 \cdot q(\lambda))$ from before. Thus, it holds that

$$\Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{A} = Z(\mathcal{O}) \\ F = \text{Sam}(\mathcal{A}); \quad \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (z, s', \pi', \text{crs}_{-i}) \leftarrow \mathcal{A}^{\mathcal{P}}(1^\lambda, \text{crs}_i, T(\lambda)) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ s = \text{Sol}_1^{\mathcal{P}}(1^\lambda, \text{mcrs}, T(\lambda), z) \end{array} : \begin{array}{l} \text{Verify}^{\mathcal{P}}(1^\lambda, \text{mcrs}, T(\lambda), z, (s', \pi')) = 1 \\ \wedge s' \neq s \end{array} \right] > 1/(2 \cdot q(\lambda)).$$

Now, for each $\lambda \in \mathbb{N}$, Z wins with at least $1/(2 \cdot q(\lambda))$ probability over a random $\mathcal{O} \leftarrow \text{RF}_\lambda$. By a simple averaging argument, there must exist a fixed oracle \mathcal{O} such that Z wins with at least $1/(2 \cdot q(\lambda))$ probability on that oracle. Let \mathcal{A} be the output of Z on such an \mathcal{O} . We give \mathcal{B}_λ the description of \mathcal{A} hardcoded as non-uniform advice as well as the set of points F . As $|F|$ and \mathcal{A} are polynomial-size, it holds that \mathcal{B}_λ is also be polynomial-size.

We are now ready to define the behavior of \mathcal{B}_λ . On input $(1^\lambda, \text{crs}_i, T(\lambda))$, \mathcal{B}_λ computes $(z, s', \pi', \text{crs}_{-i}) \leftarrow \mathcal{A}^{\mathcal{P}}(1^\lambda, \text{crs}_i, T(\lambda))$. Whenever \mathcal{A} queries the oracle \mathcal{P} , \mathcal{B}_λ responds using F if possible, and otherwise responds with a uniformly random value (responding consistently if the same query is made multiple times). \mathcal{B}_λ then parses π' as $(s_{\text{os}}, \pi_{\text{os}})$ and outputs $(z, s_{\text{os}}, \pi_{\text{os}}, \text{crs}_i)$ if possible.

For correctness, we have already shown that when \mathcal{A} wins, it must be the case that $s' = \perp$ and $s \neq \perp$. Since $s' = \perp$ and $\text{Verify}^{\mathcal{P}}(1^\lambda, \text{mcrs}, T(\lambda), z, (s', \pi')) = 1$, it holds that $\pi' = (s'_{\text{os}}, \pi'_{\text{os}})$ and $\text{Verify}_{\text{os}}(1^{\lambda_{\text{os}}}, \text{mcrs}, t, z, (s'_{\text{os}}, \pi'_{\text{os}})) = 1$. However, because $s \neq \perp$ and $\text{Verify}^{\mathcal{P}}(1^\lambda, \text{mcrs}, T(\lambda), z, (s, \pi)) = 1$, it must be the case that $z = \text{Gen}^{\mathcal{P}}(1^\lambda, t, s; r)$ for $\pi = r$. This implies that $z \in \text{Supp}(\text{Gen}_{\text{os}}(1^{\lambda_{\text{os}}}, t, s_{\text{os}}))$ where $s_{\text{os}} = s||r$. By correctness and completeness of the underlying TLP, $(s_{\text{os}}, \pi_{\text{os}}) = \text{Sol}_{\text{os}}(1^{\lambda_{\text{os}}}, \text{mcrs}, t, z)$. However, this implies that \mathcal{B}_λ wins whenever \mathcal{A} wins as $s'_{\text{os}} \neq s_{\text{os}}$, $\text{Verify}_{\text{os}}(1^{\lambda_{\text{os}}}, \text{mcrs}, T(\lambda), z, (s'_{\text{os}}, \pi'_{\text{os}})) = 1$, and $z \in \text{Supp}(\text{Gen}_{\text{os}}(1^{\lambda_{\text{os}}}, t, \cdot))$. More precisely, \mathcal{B}_λ satisfies the following for infinitely many $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (z, s'_{\text{os}}, \pi'_{\text{os}}, \text{crs}_{-i}) \\ \leftarrow \mathcal{B}_\lambda(1^\lambda, \text{crs}_i, T(\lambda)) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ s_{\text{os}} = \text{Sol}_{\text{os},1}(1^{\lambda_{\text{os}}}, \text{mcrs}, T(\lambda), z) \end{array} : \begin{array}{l} \text{Verify}_{\text{os}}(1^\lambda, \text{mcrs}, T(\lambda), z, (s'_{\text{os}}, \pi'_{\text{os}})) = 1 \\ \wedge s'_{\text{os}} \neq s_{\text{os}} \\ \wedge z \in \text{Supp}(\text{Gen}_{\text{os}}(1^{\lambda_{\text{os}}}, t, \cdot)) \end{array} \right] > 1/(2 \cdot q(\lambda)),$$

which contradicts one-sided soundness of the underlying TLP. \square

6 Applications to Multi-Party Coin-Flipping and Auctions

In this section, we discuss our fair multi-party protocols. We focus on the case of multi-party coin-flipping and address auctions in Remark 1 below.

Our multi-party coin-flipping protocol is based on a publicly verifiable time-lock puzzle that satisfies concurrent functional non-malleability for the XOR function f_{\oplus} . Specifically, in order to produce L bits of randomness, we need concurrent function non-malleability for the function $f_{\oplus}: (\{0, 1\}^L)^* \rightarrow \{0, 1\}^L$ that on input (r_1, \dots, r_n) outputs $\bigoplus_{r_i \neq \perp} r_i$.

We describe our protocol in a public bulletin board model, where any party may “publish” a message that all other parties will see within some fixed time. Our protocol consists four phases: a commit phase, open phase, force open phase, and output phase. The commit and open phases consist of a single synchronous round of communication where all participating parties publish a message on the bulletin board. The force open phase can be computed by any party, and only needs to be computed by a single (honest) party. Once all puzzles have been opened (or force opened), any party can run the output phase to get the output of the protocol. When we refer to an *honest* participant, we mean a party that runs the protocol as specified, independent of all other participants.

For any $B, L: \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon \in (0, 1)$, let $(\text{Gen}, \text{Sol}, \text{Verify})$ be a (B, ϵ) -secure publicly verifiable time-lock puzzle with message length $L(\lambda)$ that satisfies concurrent functional non-malleability for the function f_{\oplus} (which has output length $L(\lambda)$). The protocol takes as common input a security parameter λ and a time bound $t = T(\lambda)$ where T is any sufficiently large function smaller than B .

- **Commit phase:** Each participant i samples $s_i \leftarrow \{0, 1\}^{L(\lambda)}$ and $r_i, \text{crs}_i \leftarrow \{0, 1\}^\lambda$, computes $z_i \leftarrow \text{Gen}(1^\lambda, t, s_i; r_i)$, and publishes z_i and crs_i . Let $\text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n)$. We assume that only unique puzzle are published to the bulletin board.
- **Open phase:** Each participant i that published in the commit phase publishes the solution s_i and with an opening r_i .
- **Force open phase:** For each puzzle z_j , if either (a) there is no published solution s_j and opening r_j or (b) if $z_j \neq \text{Gen}(1^\lambda, t, s_j; r_j)$, compute and publish $(s_j, \pi_j) \leftarrow \text{Sol}(1^\lambda, \text{mcrs}, t, z_j)$ (where s_j might be \perp).
- **Output phase:** If for every puzzle z_j and solution s_j , either (a) there is a published opening r_j such that $z_j = \text{Gen}(1^\lambda, t, s_j; r_j)$ or (b) a published proof π_j such that $\text{Verify}(1^\lambda, \text{mcrs}, t, z_j, (s_j, \pi_j)) = 1$, then output $s = \bigoplus_{s_j \neq \perp} s_j$.

We note that the protocol above does not assume an a priori bound on the number of participants. Furthermore, there is no external setup needed by the protocol. All participants, however, do publish a random string $\text{crs}_i \leftarrow \{0, 1\}^\lambda$ to implement the ABO-string model for $(\text{Gen}, \text{Sol}, \text{Verify})$.

Theorem 6.1. *Let $B: \mathbb{N} \rightarrow \mathbb{N}$, $L(\lambda) \in \lambda^{O(1)}$, and $\epsilon \in (0, 1)$ where $\lambda \leq B(\lambda) \in 2^{L(\lambda)} \cdot \text{poly}(\lambda)$. Assume the existence of a (B, ϵ) -secure publicly verifiable time-lock puzzle for $L(\lambda)$ bit messages in the ABO-string model that satisfies concurrent function non-malleability for f_{\oplus} with $L(\lambda)$ bit output. Then, there exists a multi-party coin-flipping protocol that satisfies optimistic efficiency, fairness, and public verifiability. The protocol supports an unbounded number of participants and requires no adversary-independent trusted setup.*

We note the following result by plugging Corollary 5.10 into Theorem 6.1.

Corollary 6.2. *Let $B, L: \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon \in (0, 1)$ where $\lambda \leq B(\lambda) \in 2^{L(\lambda)} \cdot \text{poly}(\lambda)$. Assuming the (B, ϵ) -repeated squaring assumption for RSWGen, there exists a multi-party coin-flipping protocol that satisfies optimistic efficiency, fairness, and public verifiability. The protocol supports an unbounded number of participants and requires no adversary-independent trusted setup. Security is proven in the auxiliary-input random oracle.*

We remark how we can adapt our protocol to deal with auctions.

Remark 1 (Multi-Party Auctions). *For our application to auctions, we consider a standard second-price, sealed-bid auction, in which the auctioned item is assigned to the highest bidder who pays the second highest bid for the item. We assume some form of authenticated channels so we can know the bidders' identities in order to distribute the auctioned items. We leave these as external implementation details for the protocol. The main protocol proceeds as follows.*

In the commit phase, each participant computes a time-lock puzzle to their bid. The open and force open phases are identical to the case of coin-flipping. Then in the output phase, we need to determine the identity of the highest bidder and the value of the second highest bid.

The function that computes the output consists of finding the top two values in a set. This can be computed in low depth (doing a tree of comparisons in parallel) and has output length $\log n + \log M$ where n is the number of participants and M is a bound on the largest valid bid. Thus, using our publicly verifiable time-lock puzzle that satisfies concurrent functional non-malleability for this function, the resulting protocol is secure assuming $n \cdot M \cdot \text{poly}(\lambda)$ security for the repeated squaring assumption. Assuming n and M are polynomially bounded, we only need polynomial security assumptions.

In the remainder of this section, we discuss the various properties satisfied by our construction of Theorem 6.1.

Complexity. We discuss the efficiency of each phase in the above protocol. As mentioned above, the commit and open phase consist of a single round of synchronous communication. For the commit phase, each participating party requires at most $\text{poly}(\lambda, \log t)$ local computation time by the efficiency of Gen, and we require that all messages be posted within a specified time at most t^ϵ . After this specified time, all participating parties can publish their solutions and openings as part of the open phase, and if needed, unopened puzzles can be force opened. By the efficiency of Sol, force open requires time $t \cdot \text{poly}(\lambda, \log t)$ per unopened puzzle and can be computed by a single party. The output phase can be computed by anyone and requires at most $n \cdot \text{poly}(\lambda, \log t)$ time by the efficiency of Gen and Verify, where n is the number of puzzles submitted during the commit phase.

Optimistic efficiency. If all parties that publish a puzzle in the commit phase also publish a valid solution and opening in the open phase, then the output phase can be run immediately without running the force open phase.

Public verifiability. Let z_j be any puzzle which was not opened, or was opened incorrectly during the open phase. Suppose an honest party runs the force open phase for puzzle z_j and publishes $(s_j, \pi_j) \leftarrow \text{Sol}(1^\lambda, \text{mcrs}, t, z_j)$. By completeness of (Gen, Sol, Verify), it follows that $\text{Verify}(1^\lambda, \text{mcrs}, t, z_j, (s_j, \pi_j)) = 1$ for any $z_j \in \{0, 1\}^*$, so that check in the output phase will pass. Thus, if a single honest party runs the entire force open phase, then any party can run the output phase and all checks will pass.

Fairness. We formalize and prove fairness in the following lemma. At a high level, we show that as long as there is a single honest participant (who only needs to publish an honestly sampled puzzle independent of all other participants), the output of the protocol is statistically close to a uniformly random distribution over $L(\lambda)$ bits.

Lemma 6.3 (Fairness). *For any distinguisher \mathcal{D} , there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$, the following holds. Suppose that at most $n(\lambda) \in \text{poly}(\lambda)$ participants for the commit phase, and at least one honest party runs the commit phase. Let s be the output of the open phase at the end of the protocol for security parameter λ , and $r \leftarrow \{0, 1\}^{L(\lambda)}$. It holds that*

$$|\Pr[\mathcal{D}(s) = 1] - \Pr[\mathcal{D}(r) = 1]| \leq \text{negl}(\lambda).$$

Proof. Suppose by way of contradiction that there exists a polynomial q such that for infinitely many $\lambda \in \mathbb{N}$,

$$|\Pr[\mathcal{D}(s) = 1] - \Pr[\mathcal{D}(r) = 1]| > 1/q(\lambda).$$

Fix any $\lambda \in \mathbb{N}$ for which the above holds, and let $n = n(\lambda)$, $L = L(\lambda)$, and $t = T(\lambda)$. Let z_1, \dots, z_n be the (unique) puzzles published in the commit phase and let $\text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n)$. Without loss of generality, suppose that participant “1” is honest, so z_1 and crs_1 are generated honestly and published to the bulletin board. More specifically, $s_1 \leftarrow \{0, 1\}^L$, $r_1, \text{crs}_1 \leftarrow \{0, 1\}^\lambda$, $z_1 \leftarrow \text{Gen}(1^\lambda, t, s_1; r_1)$, and only z_1 and crs_1 are published before the open phase begins. Note that since z_1 is generated honestly and independently of all other parties, it will be unique with all but negligible probability, so will be successfully published to the bulletin board.

Let $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ be a non-uniform algorithm where \mathcal{A}_λ consists of the actions of all parties in the protocol other than participant 1 that outputs puzzles z_2, \dots, z_n . By definition of the protocol, \mathcal{A}_λ has polynomial size and depth bounded by $(T(\lambda))^\epsilon$. Thus, \mathcal{A} is a valid $(1, n, B, \epsilon, T, 1)$ -MIM adversary, and we can consider the corresponding distribution $\text{mim}_{\mathcal{A}}(1^\lambda, t, v)$ for any $v \in \{0, 1\}^L$.

Let $\mathcal{B} = \{\mathcal{B}_\lambda\}_{\lambda \in \mathbb{N}}$ be a non-uniform algorithm where \mathcal{B}_λ receives as input $(1^\lambda, \text{crs}_1, t)$, where crs_1 is the random string output by the honest party, and then simulates the actions of all parties in the protocol and eventually computes the output s . In particular, we assume that \mathcal{B}_λ publishes all of the relevant values on the public bulletin board.

We consider the following hybrid distributions:

$$\begin{aligned} \text{Hyb}_0(\lambda) &= s, \\ \text{Hyb}_1(\lambda) &= \mathcal{B}_\lambda(1^\lambda, \text{crs}_1, t), \\ \text{Hyb}_2(\lambda) &= f_\oplus(s_1, \text{mim}_{\mathcal{A}}(1^\lambda, t, s_1)), \\ \text{Hyb}_3(\lambda) &= f_\oplus(s_1, \text{mim}_{\mathcal{A}}(1^\lambda, t, 0^L)), \\ \text{Hyb}_4(\lambda) &= r. \end{aligned}$$

By assumption, we have that

$$|\Pr[\mathcal{D}(\text{Hyb}_0(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_4(\lambda)) = 1]| > 1/q(\lambda).$$

Towards a contradiction, we will show that for each $i \in \{1, 2, 3, 4\}$ that $|\Pr[\mathcal{D}(\text{Hyb}_{i-1}(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_i(\lambda)) = 1]| \leq 1/(4 \cdot q(\lambda))$.

- $|\Pr[\mathcal{D}(\text{Hyb}_0(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_1(\lambda)) = 1]| \leq 1/(4 \cdot q(\lambda))$

By definition of \mathcal{B} , $\text{Hyb}_0(\lambda)$ and $\text{Hyb}_1(\lambda)$ are identically distributed, so it holds that

$$|\Pr[\mathcal{D}(\text{Hyb}_0(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_1(\lambda)) = 1]| = 0.$$

- $|\Pr[\mathcal{D}(\text{Hyb}_1(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_2(\lambda)) = 1]| \leq 1/(4 \cdot q(\lambda))$

Assume by way of contradiction that $|\Pr[\mathcal{D}(\text{Hyb}_1(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_2(\lambda)) = 1]| > 1/(4 \cdot q(\lambda))$. Under this assumption, we show how to break soundness of $(\text{Gen}, \text{Sol}, \text{Verify})$.

Consider a non-uniform algorithm $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ defined as follows. \mathcal{C}_λ on input $(1^\lambda, \text{crs}_1, t)$ simulates $\mathcal{B}_\lambda(1^\lambda, \text{crs}_1, t)$. For $i \in [2, n]$, let crs_i be the random string posted, and s'_i be the solution published on the bulletin board for puzzle z_i with valid opening r'_i or proof π'_i . If z_i is equal to a puzzle z_j for $j < i$, set $s_i = \perp$, and otherwise let $(s_i, \pi_i) = \text{Sol}(1^\lambda, \text{mcrs}, t, z_i)$. If there exists an $i \in [n]$ such that $s_i \neq s'_i$, \mathcal{C}_λ outputs $(z_i, s'_i, \pi'_i, \text{crs}_{-1})$, and outputs \perp otherwise. Since there are a polynomial n number of participants, each of which is polynomial time, \mathcal{B}_λ and hence \mathcal{C}_λ runs in polynomial time. We proceed to argue that \mathcal{C}_λ breaks soundness with $1/(4 \cdot q(\lambda))$ probability.

Recall that s_i is the solution given by Sol . By full correctness of $(\text{Gen}, \text{Sol}, \text{Verify})$, for every $i \in [n]$, s_i is either the unique value such that $z_i \in \text{Supp}(\text{Gen}(1^\lambda, t, s_i))$ or s_i is \perp . By definition of $\text{mim}_\mathcal{A}$, it follows that (s_2, \dots, s_n) is the output of $\text{mim}_\mathcal{A}(1^\lambda, t, s_1)$ (as we can assume without loss of generality that duplicate puzzles are ignored by \mathcal{B}_λ). Thus, in the event where $s_i = s'_i$ for all $i \in [n]$, it follows that $\text{Hyb}_1(\lambda) = \text{Hyb}_2(\lambda)$. However, we assumed that with probability $\text{Hyb}_1(\lambda) \neq \text{Hyb}_2(\lambda)$ with at least $1/(4 \cdot q(\lambda))$ probability, so it follows that there exists an $i \in [n]$ such that $s_i \neq s'_i$ with the same probability.

Whenever $s_i \neq s'_i$, we claim that \mathcal{C}_λ breaks soundness. We have that \mathcal{B}_λ outputs $f_\oplus(s'_1, \dots, s'_n)$. We know that the checks in the output phase pass for s'_i , so either (a) $z_i = \text{Gen}(1^\lambda, t, s'_i; r'_i)$ or (b) $\text{Verify}(1^\lambda, \text{mcrs}, t, z_i, (s'_i, \pi'_i)) = 1$. We have that (a) is impossible since if $z_i \in \text{Supp}(\text{Gen}(1^\lambda, t, s'_i))$, this means that $s_i = s'_i$ by correctness. Thus, it must be the case that (b) holds with at least $1/(4 \cdot q(\lambda))$ probability. This implies that for infinitely many $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} \text{crs}_1 \leftarrow \{0, 1\}^\lambda \\ (z, s'_i, \pi'_i, \text{crs}_{-1}) \leftarrow \mathcal{C}_\lambda(1^\lambda, \text{crs}_1, t) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ s_i = \text{Sol}_1(1^\lambda, \text{mcrs}, t, z_i) \end{array} : \begin{array}{l} \text{Verify}(1^\lambda, \text{mcrs}, t, z_i, (s'_i, \pi'_i)) = 1 \\ \wedge s_i \neq s'_i \end{array} \right] > 1/(4 \cdot q(\lambda)),$$

which contradicts soundness.

- $|\Pr[\mathcal{D}(\text{Hyb}_2(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_3(\lambda)) = 1]| \leq 1/(4 \cdot q(\lambda))$

Assume by way of contradiction that $|\Pr[\mathcal{D}(\text{Hyb}_2(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_3(\lambda)) = 1]| > 1/(4 \cdot q(\lambda))$. Under this assumption, we show how to break concurrent functional non-malleability of $(\text{Gen}, \text{Sol}, \text{Verify})$ for the function f_\oplus .

We already mentioned that \mathcal{A} is a valid $(1, n, B, \epsilon, T, 1)$ -MIM adversary. Thus, it suffices to come up with a distinguisher \mathcal{D}_{nm} for non-malleability that can distinguish $f_\oplus(\text{mim}_\mathcal{A}(1^\lambda, t, v))$ for $v = s_1$ or $v = 0^L$.

\mathcal{D}_{nm} on input $f_\oplus(\text{mim}_\mathcal{A}(1^\lambda, t, v))$ has s_1 hardcoded and simply outputs $\mathcal{D}(s_1 \oplus f_\oplus(\text{mim}_\mathcal{A}(1^\lambda, t, v)))$. By definition of f_\oplus , it holds that

$$s_1 \oplus f_\oplus(\text{mim}_\mathcal{A}(1^\lambda, t, v)) = f_\oplus(s_1, \text{mim}_\mathcal{A}(1^\lambda, t, v)).$$

Thus, \mathcal{D}_{nm} distinguishes $v = s_1$ and $v = 0^L$ with the same probability as \mathcal{D} , so it holds that

$$|\Pr[\mathcal{D}_{\text{nm}}(f_\oplus(\text{mim}_\mathcal{A}(1^\lambda, t, s_1))) = 1] - \Pr[\mathcal{D}_{\text{nm}}(f_\oplus(\text{mim}_\mathcal{A}(1^\lambda, t, 0^L))) = 1]| > 1/(4 \cdot q(\lambda)),$$

in contradiction.

- $|\Pr[\mathcal{D}(\text{Hyb}_3(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_4(\lambda)) = 1]| \leq 1/(4 \cdot q(\lambda))$

By assumption that participant 1 is honest, it holds that $\text{mim}_{\mathcal{A}}(1^\lambda, t, 0^L)$ is independent of s_1 . Since $s_1 \leftarrow \{0, 1\}^\lambda$, this implies that $f_{\oplus}(s_1, \text{mim}_{\mathcal{A}}(1^\lambda, t, 0^L))$ is uniformly random. Thus, it is identically distributed to $r \leftarrow \{0, 1\}^\lambda$. It follows that

$$|\Pr[\mathcal{D}(\text{Hyb}_3(\lambda)) = 1] - \Pr[\mathcal{D}(\text{Hyb}_4(\lambda)) = 1]| = 0.$$

This completes the proof of the lemma. □

Acknowledgements

This work was supported in part by NSF Award SATC-1704788, NSF Award RI-1703846, AFOSR Award FA9550-18-1-0267, and by NSF Award DGE-1650441. This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2019-19-020700006. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [Bar02] Boaz Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In *43rd Symposium on Foundations of Computer Science FOCS*, pages 345–355, 2002.
- [BBBF18] Dan Boneh, Joseph Boneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *Advances in Cryptology - CRYPTO*, pages 757–788, 2018.
- [BBF18] Dan Boneh, Benedikt Bünz, and Ben Fisch. A survey of two verifiable delay functions. *IACR Cryptol. ePrint Arch.*, 2018:712, 2018.
- [BGJ⁺16] Nir Bitansky, Shafi Goldwasser, Abhishek Jain, Omer Paneth, Vinod Vaikuntanathan, and Brent Waters. Time-lock puzzles from randomized encodings. In *ITCS*, 2016.
- [BN00] Dan Boneh and Moni Naor. Timed commitments. In *Advances in Cryptology - CRYPTO*, pages 236–254, 2000.
- [BPS06] Boaz Barak, Manoj Prabhakaran, and Amit Sahai. Concurrent non-malleable zero knowledge. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 345–354, 2006.
- [BPS16] Iddo Bentov, Rafael Pass, and Elaine Shi. Snow white: Provably secure proofs of stake. *IACR Cryptology ePrint Archive*, 2016:919, 2016.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM Conference on Computer and Communications Security*, pages 62–73. ACM, 1993.

- [CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. Random oracles and non-uniformity. In *Advances in Cryptology - EUROCRYPT*, pages 227–258, 2018.
- [Cle86] Richard Cleve. Limits on the security of coin flips when half the processors are faulty (extended abstract). In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, STOC*, pages 364–369, 1986.
- [COSV16] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Concurrent non-malleable commitments (and more) in 3 rounds. In *Advances in Cryptology - CRYPTO*, pages 270–299, 2016.
- [COSV17] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Four-round concurrent non-malleable commitments from one-way functions. In *Advances in Cryptology - CRYPTO*, pages 127–157, 2017.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, STOC*, pages 542–552, 1991.
- [DGKR18] Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Advances in Cryptology - EUROCRYPT*, pages 66–98, 2018.
- [DKP20] Dana Dachman-Soled, Ilan Komargodski, and Rafael Pass. Non-malleable codes for bounded polynomial depth tampering, 2020.
- [DPS19] Phil Daian, Rafael Pass, and Elaine Shi. Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake. In *Financial Cryptography and Data Security - 23rd International Conference, FC*, pages 23–41, 2019.
- [EFKP19] Naomi Ephraim, Cody Freitag, Ilan Komargodski, and Rafael Pass. Continuous verifiable delay functions. *IACR Cryptology ePrint Archive*, 2019:619, 2019.
- [FO13] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *J. Cryptology*, 26(1):80–101, 2013.
- [GLOV12] Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 51–60, 2012.
- [GO14] Jens Groth and Rafail Ostrovsky. Cryptography in the multi-string model. *J. Cryptology*, 27(3):506–543, 2014.
- [Goy11] Vipul Goyal. Constant round non-malleable protocols using one way functions. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC*, pages 695–704, 2011.
- [GPR16] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, pages 1128–1141, 2016.
- [HMPS15] Susan Hohenberger, Steven Myers, Rafael Pass, and Abhi Shelat. An overview of ANONIZE: A large-scale anonymous survey system. *IEEE Security & Privacy*, 13(2):22–29, 2015.

- [Khu17] Dakshita Khurana. Round optimal concurrent non-malleability from polynomial hardness. In *Theory of Cryptography - 15th International Conference, TCC*, pages 139–171, 2017.
- [KK19] Yael Tauman Kalai and Dakshita Khurana. Non-interactive non-malleability from quantum supremacy. In *Advances in Cryptology - CRYPTO*, pages 552–582, 2019.
- [KS17] Dakshita Khurana and Amit Sahai. How to achieve non-malleability in one or two rounds. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 564–575, 2017.
- [LP09] Huijia Lin and Rafael Pass. Non-malleability amplification. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC*, pages 189–198, 2009.
- [LP11] Huijia Lin and Rafael Pass. Constant-round non-malleable commitments from any one-way function. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC*, pages 705–714, 2011.
- [LPS17] Huijia Lin, Rafael Pass, and Pratik Soni. Two-round and non-interactive concurrent non-malleable commitments from time-lock puzzles. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 576–587. IEEE Computer Society, 2017.
- [LPTV10] Huijia Lin, Rafael Pass, Wei-Lung Dustin Tseng, and Muthuramakrishnan Venkatasubramanian. Concurrent non-malleable zero knowledge proofs. In *Advances in Cryptology - CRYPTO*, 2010.
- [LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkatasubramanian. Concurrent non-malleable commitments from any one-way function. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC*, pages 571–588, 2008.
- [MT19] Giulio Malavolta and Sri Aravinda Krishnan Thyagarajan. Homomorphic time-lock puzzles and applications. In *Advances in Cryptology - CRYPTO*, volume 11692 of *Lecture Notes in Computer Science*, pages 620–649. Springer, 2019.
- [OPV10] Rafail Ostrovsky, Omkant Pandey, and Ivan Visconti. Efficiency preserving transformations for concurrent non-malleable zero knowledge. In *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC*, pages 535–552, 2010.
- [Pie19] Krzysztof Pietrzak. Simple verifiable delay functions. In *10th Innovations in Theoretical Computer Science Conference, ITCS*, pages 60:1–60:15, 2019.
- [PPV08] Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive one-way functions and applications. In *Advances in Cryptology - CRYPTO*, pages 57–74, 2008.
- [PR05a] Rafael Pass and Alon Rosen. Concurrent non-malleable commitments. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 563–572, 2005.
- [PR05b] Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, STOC*, pages 533–542, 2005.

- [PR08] Rafael Pass and Alon Rosen. Concurrent nonmalleable commitments. *SIAM J. Comput.*, 37(6):1891–1925, 2008.
- [PW10] Rafael Pass and Hoeteck Wee. Constant-round non-malleable commitments from sub-exponential one-way functions. In *Advances in Cryptology - EUROCRYPT*, 2010.
- [RSW96] R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto, 1996. Technical report, Massachusetts Institute of Technology, Cambridge, MA, USA.
- [Unr07] Dominique Unruh. Random oracles and auxiliary input. In *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, pages 205–223, 2007.
- [VZGS13] Joachim Von Zur Gathen and Igor E Shparlinski. Generating safe primes. *Journal of Mathematical Cryptology*, 7(4):333–365, 2013.
- [Wee10] Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *FOCS*, 2010.
- [Wes19] Benjamin Wesolowski. Efficient verifiable delay functions. In *Advances in Cryptology - EUROCRYPT*, pages 379–407, 2019.

A One-Many Non-malleability

We show that one- m non-malleable time-lock puzzles are in fact m -concurrent non-malleable. The structure of the proof follows from [LPV08], who show the claim for standard non-malleable puzzles.

Lemma A.1 (One-many to concurrent). *Let $B(\lambda) \geq \lambda$ and $\epsilon \in (0, 1)$. If (Gen, Sol) is a (B, ϵ) -secure time-lock puzzle that is one-many functional non-malleable for a class of functions \mathcal{F} , then it is also fully concurrent function non-malleable for \mathcal{F} .*

Proof. Let (Gen, Sol) be a (B, ϵ) -secure one-many functional non-malleable time-lock puzzle for a class of functions \mathcal{F} .

To show the claim, suppose for contradiction that the time-lock puzzle is not fully-concurrent non-malleable, meaning that there exists a function $f \in \mathcal{F}$ such that for every positive polynomials p, ℓ , there exists a function T with $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$ and polynomial n such that the following holds. Fix a probabilistic polynomial-time $(n, n, B, \epsilon, T, p)$ -MIM adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, and suppose there exists a distinguisher $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$, a polynomial q , and a vector $\vec{s} = (s_1, \dots, s_{n(\lambda)})$ such that for infinitely many $\lambda \in \mathbb{N}$ it holds that

$$\left| \Pr \left[\mathcal{D}_\lambda(f(\text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), \vec{s})) = 1) \right] - \Pr \left[\mathcal{D}_\lambda(f(\text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), (0^\lambda)^{m(\lambda)})) = 1) \right] \right| \geq \frac{1}{q(\lambda)}.$$

Looking ahead, let $\ell_{\text{tlp}}, p_{\text{tlp}}$ be the positive polynomials associated with the functional non-malleability of (Gen, Sol) for f . We will show that the above implies a $(1, n, B, \epsilon, T, p_{\text{tlp}})$ -MIM adversary \mathcal{B} such that for infinitely many $\lambda \in \mathbb{N}$, there exists a value $s \in \{0, 1\}^\lambda$ such that \mathcal{D}_λ can be used to distinguish the output of $f(\text{mim}_{\mathcal{B}}(1^\lambda, T(\lambda), s))$ from $f(\text{mim}_{\mathcal{B}}(1^\lambda, T(\lambda), 0^\lambda))$, which will contradict the one-many functional non-malleability of the time-lock puzzle.

To so do, fix any λ for which \mathcal{D}_λ succeeds at distinguishing above, and for $i \in \{0, \dots, n(\lambda)\}$ define $\vec{v}^{(i)}$ to be the vector where the first i entries are (s_1, \dots, s_i) and the remaining $n(\lambda) - i$ entries

are set to 0^λ . As $\vec{v}^{(n(\lambda))} = \vec{s}$ and $\vec{v}^{(0)} = (0^\lambda)^{n(\lambda)}$, it follows by a hybrid argument that there exists an $i \in [n(\lambda)]$ such that

$$\begin{aligned} & \left| \Pr \left[\mathcal{D}_\lambda(f(\text{mim}_\mathcal{A}(1^\lambda, T(\lambda), \vec{v}^{(i-1)})) = 1) \right] \right. \\ & \left. - \Pr \left[\mathcal{D}_\lambda(f(\text{mim}_\mathcal{A}(1^\lambda, T(\lambda), \vec{v}^{(i)})) = 1) \right] \right| \geq \frac{1}{n(\lambda) \cdot p(\lambda)} \end{aligned} \quad (\text{A.1})$$

We can now define the $(1, n, B, \epsilon, T, p_{\text{tlp}})$ -MIM adversary \mathcal{B}_λ , which we will use to contradict the one-many non-malleability of the time-lock puzzle relative to the value s_i . The adversary \mathcal{B}_λ has $\vec{v}^{(i)}$ hardcoded, and receives as input a puzzle z^* either to s_i or 0^λ . It then does the following:

1. Sample puzzles $z_j \leftarrow \text{Gen}(1^\lambda, T(\lambda), \vec{v}_j^{(i)})$ in parallel for $j \in [n(\lambda)] \setminus \{i\}$.
2. Let $\vec{z} = (z_1, \dots, z_{i-1}, z^*, z_{i+1}, \dots, z_{n(\lambda)})$, and run $\tilde{\vec{z}} \leftarrow \mathcal{A}_\lambda(\vec{z})$.
3. In parallel, check if any of the puzzles in $\tilde{\vec{z}}$ are equal to any puzzles in \vec{z} . If so, replace the matching puzzles with \perp , and output the resulting vector.

We will show that \mathcal{B}_λ succeeds at breaking the one- n functional non-malleability of the time-lock puzzle, which suffices to break one-many functional non-malleability. We first analyze its success probability, and then its efficiency. To analyze its success probability, we claim that when \mathcal{B}_λ receives a puzzle for s_i , then

$$\text{mim}_\mathcal{B}(1^\lambda, T(\lambda), s_i) \equiv \text{mim}_\mathcal{A}(1^\lambda, T(\lambda), \vec{v}^{(i)})$$

and when \mathcal{B}_λ receives a puzzle for 0 , it holds that

$$\text{mim}_\mathcal{B}(1^\lambda, T(\lambda), 0^\lambda) \equiv \text{mim}_\mathcal{A}(1^\lambda, T(\lambda), \vec{v}^{(i-1)}).$$

To see this, recall that for any vector \vec{s} , the distribution $\text{mim}_\mathcal{A}(1^\lambda, T(\lambda), \vec{s})$ is given by (a) sampling a vector of puzzles for \vec{s} , (b) running \mathcal{A}_λ on that vector to obtain its output vector, (c) replacing any puzzle that appear in both vectors with \perp , and (d) finding the unique solutions to the resulting puzzles (or \perp). Next, we compare this to the output of $\text{mim}_\mathcal{B}$, both on input 0 and s_i :

- The distribution $\text{mim}_\mathcal{B}$ first samples the puzzle z^* , either to 0^λ or to s_i , and proceeds to run $\mathcal{B}_\lambda(z^*)$, who computes the vector \vec{z} . When \mathcal{B}_λ receives a puzzle for s_i , then it sets \vec{z} to be a puzzle vector to $\vec{v}^{(i)}$, and when \mathcal{B}_λ receives a puzzle for 0^λ , it sets the puzzle vector to correspond to $\vec{v}^{(i-1)}$. Therefore, \vec{z} as computed by $\text{mim}_\mathcal{B}(1^\lambda, T(\lambda), s_i)$ has the same distribution as the input to \mathcal{A} computed by $\text{mim}_\mathcal{A}(1^\lambda, T(\lambda), \vec{v}^{(i)})$, and \vec{z} as computed by $\text{mim}_\mathcal{B}(1^\lambda, T(\lambda), 0^\lambda)$ has the same distribution as the input to \mathcal{A} computed by $\text{mim}_\mathcal{A}(1^\lambda, T(\lambda), \vec{v}^{(i-1)})$.
- Next, \mathcal{B}_λ computes $\tilde{\vec{z}} \leftarrow \mathcal{A}_\lambda(\vec{z})$, exactly as done in step (b) of $\text{mim}_\mathcal{A}$.
- \mathcal{B}_λ then replaces any puzzles in $\tilde{\vec{z}}$ that have been copied from \vec{z} with \perp , exactly as in step (c) of $\text{mim}_\mathcal{A}$. Note that at this point \mathcal{B}_λ outputs $\tilde{\vec{z}}$, so $\text{mim}_\mathcal{B}$ is done running \mathcal{B}_λ .
- Finally, $\text{mim}_\mathcal{B}$ compares the puzzles in $\tilde{\vec{z}}$ with z^* and replaces any that have been copied, but since z^* is part of \vec{z} , this has already been done by \mathcal{B}_λ and so doesn't change $\tilde{\vec{z}}$. Lastly, $\text{mim}_\mathcal{B}$ uniquely solves these puzzles, exactly as in step (d) of $\text{mim}_\mathcal{A}$.

We conclude that $\text{mim}_{\mathcal{B}}(1^\lambda, T(\lambda), s_i) \equiv \text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), \vec{v}^{(i)})$ and $\text{mim}_{\mathcal{B}}(1^\lambda, T(\lambda), 0^\lambda) \equiv \text{mim}_{\mathcal{A}}(1^\lambda, T(\lambda), \vec{v}^{(i-1)})$. When coupling this with Equation A.1, it implies that for infinitely many $\lambda \in \mathbb{N}$, the distinguisher \mathcal{D}_λ can distinguish $f(\text{mim}_{\mathcal{B}}(1^\lambda, T(\lambda), s_i))$ and $f(\text{mim}_{\mathcal{B}}(1^\lambda, T(\lambda), 0^\lambda))$ with inverse polynomial probability.

We next analyze the efficiency of \mathcal{B}_λ . Recall that \mathcal{B}_λ samples $n(\lambda) - 1$ puzzles to form the vector \vec{z} , runs $\mathcal{A}_\lambda(\vec{z})$, and then compares the resulting puzzles given by \mathcal{A}_λ to those in \vec{z} . As $\text{Gen}(1^\lambda, T(\lambda), \cdot)$ runs in $\text{poly}(\lambda, \log T(\lambda))$ time and $\text{size}(\mathcal{A}_\lambda) \in \text{poly}(B(\lambda))$, it follows that

$$\text{size}(\mathcal{B}_\lambda) \in \text{poly}(\lambda, n(\lambda), \log T(\lambda), B(\lambda)) \in \text{poly}(B(\lambda)).$$

To analyze its depth, since the puzzles are sampled in parallel, this can be done in fixed polynomial depth $d(\lambda, \log T(\lambda))$ (independent of n) by the efficiency of Gen . Running \mathcal{A}_λ requires at most $(T(\lambda))^\epsilon \cdot p(\lambda)$ depth since it is an $(n, n, B, \epsilon, T, p)$ -MIM adversary. Finally, comparing the resulting puzzles to the original also requires depth $d(\lambda, \log T(\lambda))$, by comparing each of the $(n(\lambda))^2$ pairs in parallel, since d bounds the length of each puzzle. Putting everything together, we have that

$$\begin{aligned} \text{depth}(\mathcal{B}_\lambda) &= 2d(\lambda, \log T(\lambda)) + (T(\lambda))^\epsilon \cdot p(\lambda) \\ &\in (T(\lambda))^\epsilon \cdot p'(\lambda) \cdot p(\lambda), \end{aligned}$$

for a polynomial p' depending only on $\log B(\lambda)$ and d (which are fixed in advance). Since the above holds for any p , we can choose p related to p_{tlp} . Specifically, set $p(\lambda) = p_{\text{tlp}}(\lambda)/p'(\lambda)$. Then, $\text{depth}(\mathcal{B}_\lambda) \in (T(\lambda))^\epsilon \cdot p(\lambda)$, as desired. Lastly, we need to show that $\ell_{\text{tlp}}(\lambda) \leq T(\lambda) \leq B(\lambda)$. We are given that $\ell(\lambda) \leq T(\lambda) \leq B(\lambda)$, and the analysis holds for all positive polynomials ℓ . Therefore, setting $\ell(\lambda) = \ell_{\text{tlp}}(\lambda)$ gives the desired bound on T , which gives the contradiction. \square

B Proof from Section 5

In this section, we state and prove the theorems for our construction of publicly-verifiable, non-malleable time-lock puzzles of Section 5.

Lemma B.1 (Restatement of Lemma 5.6). *For any polynomial T and unbounded algorithm Z that on input $\mathcal{O} \in \text{Supp}(\text{RF}_\lambda)$ outputs polynomial-size circuits, there exists a negligible function negl such that for all $\lambda, n \in \mathbb{N}$, $i \in [n]$, it holds that*

$$\Pr \left[\begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda \\ \mathcal{A} = Z(\mathcal{O}) \\ \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (g, \mathbb{G}, y', \pi, \text{crs}_{-i}) \\ \quad \leftarrow \mathcal{A}^\mathcal{O}(1^\lambda, \text{crs}_i, T(\lambda), n) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \end{array} : \begin{array}{l} \text{RSWVerify}^\mathcal{O}(1^\lambda, \text{mcrs}, T(\lambda), (g, \mathbb{G}), y', \pi) = 1 \\ \wedge y' \neq g^{2^{T(\lambda)}} \\ \wedge (g, \mathbb{G}, *) \in \text{Supp}(\text{RSWGen}(1^\lambda, T(\lambda))) \end{array} \right] \leq \text{negl}(\lambda).$$

Proof. Let C be a checking algorithm that has hardcoded $\lambda, T(\lambda), n, i$ and on input $(g, \mathbb{G}, y', \pi, \text{mcrs})$ outputs 1 if and only if the event in the experiment of the lemma statement holds. Then, we can rephrase what we want to show as $\Pr[\text{Hyb}_0(\lambda)] \leq \text{negl}(\lambda)$ where $\text{Hyb}_0(\lambda)$ is defined as follows:

$$\text{Hyb}_0(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{A} = Z(\mathcal{O}) \\ \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (g, y', \pi, \text{crs}_{-i}) \leftarrow \mathcal{A}^\mathcal{O}(1^\lambda, \text{crs}_i, T(\lambda), \mathbb{G}, n) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \end{array} : C(g, \mathbb{G}, y', \pi, \text{mcrs}) = 1 \right\}.$$

Let $f(\lambda) = 2^{\lambda/2}$ and let Sam be the inefficient algorithm given in Lemma 3.6 that on input an adversary, outputs a partial assignment F on $f(\lambda)$ points. Consider the hybrid distribution $\text{Hyb}_1(\lambda)$, defined as follows:

$$\text{Hyb}_1(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{A} = Z(\mathcal{O}) \\ F \leftarrow \text{Sam}(\mathcal{A}); \quad \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (g, y', \pi, \text{crs}_{-i}) \leftarrow \mathcal{A}^{\mathcal{P}}(1^\lambda, \text{crs}_i, T(\lambda), \mathbb{G}, n) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \end{array} : C(g, \mathbb{G}, y', \pi, \text{mcrs}) = 1 \right\}.$$

By Lemma 3.6, it follows that there exists a negligible function $\text{negl}_1(\lambda) \in \sqrt{\text{poly}(\lambda)/2^{\lambda/2}}$ such that $|\Pr[\text{Hyb}_1(\lambda)] - \Pr[\text{Hyb}_0(\lambda)]| \leq \text{negl}_1(\lambda)$.

Next, we define an inefficient checking algorithm C' that additionally receives F as input and outputs 1 if and only if $C(g, \mathbb{G}, y', \pi, \text{mcrs}) = 1$ and F doesn't contain an assignment for any inputs that contain crs_i . We next consider the hybrid $\text{Hyb}_2(\lambda)$ defined as follows:

$$\text{Hyb}_2(\lambda) = \left\{ \begin{array}{l} \mathcal{O} \leftarrow \text{RF}_\lambda; \quad \mathcal{A} = Z(\mathcal{O}) \\ F \leftarrow \text{Sam}(\mathcal{A}); \quad \mathcal{P} \leftarrow \text{RF}_\lambda[F] \\ \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (g, y', \pi, \text{crs}_{-i}) \leftarrow \mathcal{A}^{\mathcal{P}}(1^\lambda, \text{crs}_i, T(\lambda), \mathbb{G}, n) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \end{array} : C'(g, \mathbb{G}, y', \pi, \text{mcrs}, F) = 1 \right\}.$$

Since $|F| = 2^{\lambda/2}$ and crs_i is uniform over 2^λ different values, this implies that $|\Pr[\text{Hyb}_2(\lambda)] - \Pr[\text{Hyb}_1(\lambda)]| \leq \text{negl}_2(\lambda)$ for negligible function $\text{negl}_2(\lambda) = 2^{\lambda/2}/2^\lambda = 2^{-\lambda/2}$.

$\Pr[\text{Hyb}_2(\lambda)]$ is bounded by the probability that a non-uniform PPT algorithm can break the soundness of Pietrzak's VDF in the plain random oracle model for any group QR_N^+ where N is the product of two safe primes (which is the case since \mathbb{G} is in the support of RSWGen). By the analysis in [Pie19], it follows that any adversary that makes at most Q queries to the random oracle can break soundness with probability at most $Q \cdot (3/2^\lambda) \leq \text{negl}_3(\lambda)$, which is negligible when Q is bounded by a polynomial.

Finally, we conclude that there exists a negligible function $\text{negl}(\lambda) = \text{negl}_1(\lambda) + \text{negl}_2(\lambda) + \text{negl}_3(\lambda)$ such that $\Pr[\text{Hyb}_0(\lambda)] \leq \text{negl}(\lambda)$, as required. \square

Theorem B.2 (Restatement of Theorem 5.7). *Assuming the (B, ϵ) -repeated squaring assumption holds for RSWGen , then there exists a (B, ϵ) -secure strong trapdoor VDF in the ABO-string model. Soundness holds in the auxiliary-input random oracle model.*

Proof. Let RSWGen , RSWProve , RSWVerify be defined as above. The theorem follows by considering $\text{VDF} = (\text{Sample}_{\text{vdf}}, \text{Eval}_{\text{vdf}}, \text{TDEval}_{\text{vdf}}, \text{Verify}_{\text{vdf}})$. In what follows, we separately argue completeness, soundness, trapdoor evaluation, honest evaluation, and sequentiality.

Completeness. Let $\lambda, t, n \in \mathbb{N}$, $g, \mathbb{G} \in \{0, 1\}^*$, $\text{mcrs} \in (\{0, 1\}^\lambda)^n$. If (g, \mathbb{G}) is invalid, then Eval_{vdf} outputs $y = \pi = \perp$. In this case, $\text{Verify}_{\text{vdf}}$ outputs 1 as required. Otherwise, it must be the case that $g \in \mathbb{G}$ and $\mathbb{G} = \text{QR}_N^+$ for some N . In this case, Eval_{vdf} computes $y = g^{2^t}$ and $\pi = \text{RSWProve}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y)$. $\text{Verify}_{\text{vdf}}$ outputs 1 if $\text{RSWVerify}^{\mathcal{O}}(1^\lambda, \text{mcrs}, t, (g, \mathbb{G}), y, \pi) = 1$, which holds by Lemma 5.5.

Soundness. Let $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ be a non-uniform PPT adversary and T be a polynomial. We note that we only require soundness to hold for valid (g, \mathbb{G}) which are in the support of $\text{Sample}_{\text{vdf}}(1^\lambda, T(\lambda))$. By definition of $\text{Sample}_{\text{vdf}}$, this implies that (g, \mathbb{G}) are in the support of $\text{RSWGen}(1^\lambda, T(\lambda))$. It follows by definition of Eval_{vdf} and $\text{Verify}_{\text{vdf}}$ that if \mathcal{A} violates soundness of VDF, then it also violates soundness for the proof of repeated squaring. By Lemma 5.6, this can happen with at most negligible probability in the auxiliary-input random oracle model.

Trapdoor evaluation. Let $\lambda, t, n \in \mathbb{N}$, $(g, \mathbb{G}, |\mathbb{G}|) \in \text{Supp}(\text{Sample}_{\text{vdf}}(1^\lambda, t))$, and $\text{mcrs} \in (\{0, 1\}^\lambda)^n$. Since (g, \mathbb{G}) are in the support of $\text{Sample}_{\text{vdf}}$, they must be valid. In this case, $\text{TDEval}_{\text{vdf}}$ outputs $y = g^{2^t \bmod |\mathbb{G}|}$ and Eval_{vdf} outputs $y' = g^{2^t}$. Since $|\mathbb{G}|$ is the order of the group, $y = y'$ by definition.

Honest evaluation. Computing $y = g^{2^t}$ takes time $t \cdot \text{poly}(\lambda)$ to compute t sequential squares in \mathbb{G} . As discussed above, RSWProve takes time $t \cdot \text{poly}(\lambda, \log t, n)$ to compute. It follows that there exists a polynomial p such that for all $\lambda, t, n \in \mathbb{N}$, and $\text{mcrs} \in (\{0, 1\}^\lambda)^n$, $\text{Eval}_{\text{vdf}}(1^\lambda, \text{mcrs}, t, \cdot)$ can be computed in time $t \cdot p(\lambda, \log t, n)$, as required.

Sequentiality. (B, ϵ) -Sequentiality of VDF follows immediately from the (B, ϵ) -repeated squaring assumption for RSWGen stated in Assumption 5.4. Namely, $\text{Sample}_{\text{vdf}}$ simply outputs (g, \mathbb{G}) given by RSWGen , and Eval_1 simply computes $g^{2^{T(\lambda)}}$. So sequentiality follows syntactically.

Encoding. The group $\mathbb{G} = \text{QR}_N^+$ is a ring. The group can be encoded by the integer N represented by a string in $\{0, 1\}^*$. Elements can be encoded by the string representation of integers in $[0, (N - 1)/2]$. The natural bijective function $f_x(y)$ is simply addition $x + y$ in this group. \square

Theorem B.3 (Restatement of Theorem 5.8). *Suppose there exists a (B, ϵ) -secure strong trapdoor VDF. Then, there exists a (B, ϵ) -secure publicly verifiable time-lock puzzle with honest soundness.*

Proof. Let $\text{VDF} = (\text{Gen}_{\text{vdf}}, \text{Eval}_{\text{vdf}}, \text{TDEval}_{\text{vdf}}, \text{Verify}_{\text{vdf}})$ be any one-time trapdoor VDF. The theorem follows by considering the construction $\text{TLP} = (\text{Gen}_{\text{tlp}}, \text{Sol}_{\text{tlp}}, \text{Verify}_{\text{tlp}})$ as defined above. In what follows, we separately argue correctness, efficiency, hardness, completeness, and soundness for this transformation.

Correctness. Let $\lambda, t \in \mathbb{N}$, $s \in \{0, 1\}^\lambda$, $z = (x, \mathcal{X}, c) \in \text{Supp}(\text{Gen}_{\text{tlp}}(1^\lambda, t, s))$, and $s' = \text{Sol}_{\text{tlp}, 1}(1^\lambda, \text{mcrs}, t, z)$ for any $n \in \mathbb{N}$ and $\text{mcrs} \in (\{0, 1\}^\lambda)^n$. We need to show that $s' = s$. By definition of Gen_{tlp} , it holds that $(x, \mathcal{X}, \text{td}) \in \text{Supp}(\text{Gen}_{\text{vdf}}(1^\lambda, t))$, $y = \text{TDEval}_{\text{vdf}}(1^\lambda, t, (x, \mathcal{X}), \text{td})$, and $c = y \oplus x_s$ where x_s is the encoding of s . In particular, this implies that $x_s = y \oplus c$ since \oplus is a bijection over \mathcal{X} by the encoding property of VDF. By definition of Sol_{tlp} , it also holds that $y' = \text{Eval}_{\text{vdf}, 1}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}))$ and $x_{s'} = y' \oplus c$ where $x_{s'}$ is the encoding of s' . By the trapdoor evaluation property of VDF, it holds that $y' = y$, so $x_{s'} = x_s$. Since the string encodings are unique, this implies that $s' = s$, as required.

Efficiency. We note that by honest evaluation property of VDF, there exists a polynomial p_{vdf} such that $\text{Eval}_{\text{vdf}}(1^\lambda, \text{mcrs}, t, \cdot)$ is computable in time $t \cdot p_{\text{vdf}}(\lambda, \log t, n)$. Computing $y \oplus c$ and encoding s as x_s takes fixed polynomial time $\alpha(\lambda)$ independent of t . This implies that there exists a polynomial $p_{\text{tlp}}(\lambda, \log t) \geq p_{\text{vdf}}(\lambda, \log t, n) + \alpha(\lambda)/t$ such that $\text{Sol}_{\text{tlp}}(1^\lambda, \text{mcrs}, t, \cdot)$ can be computed in time $t \cdot p_{\text{vdf}}(\lambda, \log t, n) + \alpha(\lambda) \leq t \cdot p_{\text{tlp}}(\lambda, \log t, n)$, as required.

Hardness. We show that (B, ϵ) -hardness of TLP follows from the (B, ϵ) -sequentiality of VDF. Suppose by way of contradiction that (B, ϵ) -hardness of TLP does not hold. Specifically, for any positive polynomial ℓ_{tlp} and p_{tlp} there exists a function T with $\ell_{\text{tlp}}(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$, a non-uniform adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ where $\text{size}(\mathcal{A}_\lambda) \in \text{poly}(B(\lambda))$ and $\text{depth}(\mathcal{A}_\lambda) \leq (T(\lambda))^\epsilon \cdot p_{\text{tlp}}(\lambda)$ for all $\lambda \in \mathbb{N}$, a polynomial q , and strings $s, s' \in \{0, 1\}^\lambda$ such that for infinitely many λ , it holds that

$$\left| \Pr \left[\mathcal{A}_\lambda(\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s)) = 1 \right] - \Pr \left[\mathcal{A}_\lambda(\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), s')) = 1 \right] \right| > 1/q(\lambda),$$

where the probabilities are over the randomness of Gen_{tlp} and \mathcal{A}_λ .

For any value $v \in \{0, 1\}^\lambda$, let $\text{Hyb}_v(\lambda)$ be the distribution $\text{Gen}_{\text{tlp}}(1^\lambda, T(\lambda), v)$. Let $\text{Hyb}_v^{\text{rand}}(\lambda)$ be the distribution that outputs $z = (x, \mathcal{X}, c)$ where $(x, \mathcal{X}, \text{td}) \leftarrow \text{Gen}_{\text{vdf}}(1^\lambda, T(\lambda))$ and $c = r \oplus x_s$ where x_s is the encoding of s and $r \leftarrow \mathcal{X}$ is uniformly sampled. By assumption, it holds that

$$|\Pr[\mathcal{A}_\lambda(\text{Hyb}_s(\lambda)) = 1] - \Pr[\mathcal{A}_\lambda(\text{Hyb}_{s'}(\lambda)) = 1]| > 1/q(\lambda).$$

We next note that $\text{Hyb}_s^{\text{rand}}(\lambda)$ and $\text{Hyb}_{s'}^{\text{rand}}(\lambda)$ are identically distributed. This is because random element $r \leftarrow \mathcal{X}$ is uniformly random and, by the encoding property of VDF, \oplus is a bijective function on \mathcal{X} . It follows that c is uniformly distributed over \mathcal{X} in both $\text{Hyb}_s^{\text{rand}}(\lambda)$ and $\text{Hyb}_{s'}^{\text{rand}}(\lambda)$.

As a result, it must be the case that

$$|\Pr[\mathcal{A}_\lambda(\text{Hyb}_s(\lambda)) = 1] - \Pr[\mathcal{A}_\lambda(\text{Hyb}_s^{\text{rand}}(\lambda)) = 1]| > 1/(2 \cdot q(\lambda))$$

for either s or s' . Assume that this holds for s without loss of generality. We show that this breaks the (B, ϵ) -sequentiality of VDF.

We first define a reduction from adversaries against hardness of TLP to adversaries against sequentiality of VDF. Given any adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ for TLP and string $s \in \{0, 1\}^\lambda$, we construct an adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}_{\lambda \in \mathbb{N}}$ against the sequentiality of the one-time trapdoor VDF as follows. \mathcal{B}_λ on input (x, \mathcal{X}, a) computes $c = a \oplus x_s$ where x_s is the encoding of s and outputs $\mathcal{A}_\lambda(x, \mathcal{X}, c)$. Whenever a is equal to $y = \text{TDEval}_{\text{vdf}}(1^\lambda, t, x)$, we note that the distribution output by \mathcal{B}_λ is identically distributed to $\mathcal{A}_\lambda(\text{Hyb}_s(\lambda))$. Whenever a is a random element $r \leftarrow \mathcal{X}$, it holds that the output of \mathcal{B}_λ is identically distributed to $\mathcal{A}_\lambda(\text{Hyb}_s^{\text{rand}}(\lambda))$. Thus, for any adversary \mathcal{A} and string s , it holds that

$$|\Pr[\mathcal{B}_\lambda(x, y) = 1] - \Pr[\mathcal{B}_\lambda(x, r) = 1]| = \left| \Pr[\mathcal{A}_\lambda(\text{Hyb}_s(\lambda)) = 1] - \Pr[\mathcal{A}_\lambda(\text{Hyb}_s^{\text{rand}}(\lambda)) = 1] \right|$$

We next argue that this reduction can be used to break the (B, ϵ) -sequentiality of VDF. Let ℓ_{vdf} and p_{vdf} be any positive polynomials, and let $\alpha(\lambda)$ be the polynomial function representing the time to do an encoding and compute \oplus for any domain \mathcal{X} specified by VDF. Consider the function ℓ_{tlp} defined as

$$\ell_{\text{tlp}}(\lambda) = \ell_{\text{vdf}}(\lambda) + \left(\frac{2 \cdot \alpha(\lambda)}{p_{\text{vdf}}(\lambda)} \right)^{1/\epsilon}$$

and then p_{tlp} is defined as

$$p_{\text{tlp}}(\lambda) = p_{\text{vdf}}(\lambda) - \frac{\alpha(\lambda)}{(\ell_{\text{tlp}}(\lambda))^\epsilon}.$$

Note that ℓ_{tlp} and p_{tlp} are polynomials as ϵ is a constant and ℓ_{vdf} , p_{vdf} , and α are all polynomials. Furthermore, p_{tlp} is positive as p_{vdf} is positive and

$$\frac{\alpha(\lambda)}{(\ell_{\text{tlp}}(\lambda))^\epsilon} \leq \frac{p_{\text{vdf}}(\lambda)}{2}$$

by construction. Thus, by assumption, there exists a function T with $\ell_{\text{tlp}}(\lambda) \leq T(\lambda) \leq B(\lambda)$, an adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ with $\text{size}(\mathcal{A}_\lambda) \in \text{poly}(B(\lambda))$ and $\text{depth}(\mathcal{A}_\lambda) \leq (T(\lambda))^\epsilon \cdot p_{\text{tlp}}(\lambda)$, a polynomial q , and a string s satisfying

$$\left| \Pr[\mathcal{A}_\lambda(\text{Hyb}_s(\lambda)) = 1] - \Pr[\mathcal{A}_\lambda(\text{Hyb}_s^{\text{rand}}(\lambda))] \right| > 1/(2 \cdot q(\lambda))$$

for infinitely many $\lambda \in \mathbb{N}$, as defined above.

For the same function T , consider an adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}_{\lambda \in \mathbb{N}}$ based on \mathcal{A} and s given by the reduction above. First note that $\ell_{\text{tlp}}(\lambda) \geq \ell_{\text{vdf}}(\lambda)$, so it holds that T satisfies $\ell_{\text{vdf}}(\lambda) \leq T(\lambda) \leq B(\lambda)$ for all $\lambda \in \mathbb{N}$. Next, we note that $\text{size}(\mathcal{B}_\lambda) \leq \text{size}(\mathcal{A}_\lambda) + \alpha(\lambda) \in \text{poly}(B(\lambda))$. Also, $\text{depth}(\mathcal{B}_\lambda) \leq \text{depth}(\mathcal{A}_\lambda) + \alpha(\lambda) \leq (T(\lambda))^\epsilon \cdot p_{\text{tlp}}(\lambda) + \alpha(\lambda)$. Given that $p_{\text{tlp}}(\lambda) = p_{\text{vdf}}(\lambda) - \alpha(\lambda)/(\ell_{\text{tlp}}(\lambda))^\epsilon$ and $T(\lambda) \geq \ell_{\text{tlp}}(\lambda)$, it holds that

$$\alpha(\lambda) = (\ell_{\text{tlp}}(\lambda))^\epsilon \cdot (p_{\text{vdf}}(\lambda) - p_{\text{tlp}}(\lambda)) \leq (T(\lambda))^\epsilon \cdot (p_{\text{vdf}}(\lambda) - p_{\text{tlp}}(\lambda)).$$

This implies that

$$\begin{aligned} \text{depth}(\mathcal{B}_\lambda) &\leq (T(\lambda))^\epsilon \cdot p_{\text{tlp}}(\lambda) + \alpha(\lambda) \\ &\leq (T(\lambda))^\epsilon \cdot p_{\text{tlp}}(\lambda) + (T(\lambda))^\epsilon \cdot (p_{\text{vdf}}(\lambda) - p_{\text{tlp}}(\lambda)) \\ &= (T(\lambda))^\epsilon \cdot p_{\text{vdf}}(\lambda), \end{aligned}$$

as required. Finally, we recall that, for $x \leftarrow \text{Gen}_{\text{vdf}}(1^\lambda, t)$, $y = \text{TDEval}_{\text{vdf}}(1^\lambda, t, x)$ and $r \leftarrow \{0, 1\}^\lambda$,

$$\begin{aligned} &|\Pr[\mathcal{B}_\lambda(x, y) = 1] - \Pr[\mathcal{B}_\lambda(x, r) = 1]| \\ &= \left| \Pr[\mathcal{A}_\lambda(\text{Hyb}_s(\lambda)) = 1] - \Pr[\mathcal{A}_\lambda(\text{Hyb}_s^{\text{rand}}(\lambda)) = 1] \right| \\ &> 1/(2 \cdot q(\lambda)), \end{aligned}$$

for infinitely many $\lambda \in \mathbb{N}$ by assumption, in contradiction.

Completeness. Let $\lambda, t, n \in \mathbb{N}$, $z \in \{0, 1\}^*$, and $\text{mcrs} \in (\{0, 1\}^\lambda)^n$. If z cannot be parsed as (x, \mathcal{X}, c) , Sol_{tlp} outputs $(s, \pi) = (\perp, \perp)$. As $\text{Verify}_{\text{tlp}}$ can also check if z is parsed this way, it outputs 1 in this case, as required. Otherwise, z can be parsed as (x, \mathcal{X}, c) . Let $(y, \pi_{\text{vdf}}) = \text{Eval}_{\text{vdf}}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}))$ and $x_s = y \oplus c$ with string encoding s as computed by Sol_{tlp} . As $x, \mathcal{X} \in \{0, 1\}^*$, it follows by completeness of VDF that $\text{Verify}_{\text{vdf}}(1^\lambda, \text{mcrs}, t, (x, \mathcal{X}), (y, \pi_{\text{vdf}})) = 1$. Since $x_s = y \oplus c$ by definition, it follows that $\text{Verify}_{\text{tlp}}$ outputs 1, as required.

Honest soundness. Suppose there is a non-uniform PPT adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ that breaks soundness of TLP . Specifically, suppose there exist polynomials T, q , integers $n \in \mathbb{N}$ and $i \in [n]$ such that for infinitely many $\lambda \in \mathbb{N}$, it holds that

$$\Pr \left[\begin{array}{l} \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (z, s', \pi, \text{crs}_{-i}) \leftarrow \mathcal{A}_\lambda(1^\lambda, \text{crs}_i, T(\lambda), n) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ s = \text{Sol}_1(1^\lambda, \text{mcrs}, T(\lambda), z) \end{array} : \begin{array}{l} \text{Verify}(1^\lambda, \text{mcrs}, T(\lambda), z, (s', \pi)) = 1 \\ \wedge s \neq s' \\ \wedge z \in \text{Supp}(\text{Gen}(1^\lambda, T(\lambda), \cdot)) \end{array} \right] > 1/q(\lambda).$$

We construct an adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}_{\lambda \in \mathbb{N}}$ such that for the polynomial T , integers $n \in \mathbb{N}$ and $i \in [n]$, \mathcal{B} breaks that soundness of VDF with probability $1/q(\lambda)$ for infinitely many $\lambda \in \mathbb{N}$. \mathcal{B}_λ on input $(1^\lambda, \text{crs}_i, T(\lambda), n)$ computes $(z, s', \pi, \text{crs}_{-i}) = \mathcal{A}_\lambda(1^\lambda, T(\lambda), (x, \mathcal{X}, c))$, parses π as (y', π_{vdf}) , and

outputs $(x, \mathcal{X}, y', \pi_{\text{vdf}}, \text{crs}_{-i})$. First, we note that since \mathcal{A}_λ runs in polynomial time, so does \mathcal{B}_λ . We next analyze the success probability of \mathcal{B} for security parameter λ . First we note that when \mathcal{A}_λ succeeds, the puzzle $z = (x, \mathcal{X}, c)$ output by \mathcal{A}_λ is in the support of $\text{Gen}(1^\lambda, T(\lambda), \cdot)$, so it also holds that $(x, \mathcal{X}, *) \in \text{Supp}(\text{Sample}_{\text{vdf}}(1^\lambda, T(\lambda)))$. Next, whenever $\text{Verify}_{\text{tlp}}(1^\lambda, \text{mcrs}, T(\lambda), (x, \mathcal{X}, c), (s', \pi))$ outputs 1, it holds that $x_{s'} = c \oplus y'$ and $\text{Verify}_{\text{vdf}}(1^\lambda, \text{mcrs}, T(\lambda), (x, \mathcal{X}), (y', \pi_{\text{vdf}})) = 1$. By completeness of TLP, it holds that $x_s = c \oplus y$. However, since $s \neq s'$, it holds that $x_s \neq x_{s'}$ so $y \neq y'$. It follows that for the polynomial T , integers $n \in \mathbb{N}$ and $i \in [n]$, for infinitely many $\lambda \in \mathbb{N}$, it holds that

$$\Pr \left[\begin{array}{l} \text{crs}_i \leftarrow \{0, 1\}^\lambda \\ (x, \mathcal{X}, y', \pi_{\text{vdf}}, \text{crs}_{-i}) \\ \quad \leftarrow \mathcal{B}_\lambda(1^\lambda, \text{crs}_i, T(\lambda), n) \\ \text{mcrs} = (\text{crs}_1, \dots, \text{crs}_n) \\ y = \text{Eval}_{\text{vdf}, 1}(1^\lambda, \text{mcrs}, T(\lambda), (x, \mathcal{X})) \end{array} : \begin{array}{l} \text{Verify}_{\text{vdf}}(1^\lambda, \text{mcrs}, T(\lambda), (x, \mathcal{X}), (y', \pi_{\text{vdf}})) = 1 \\ \wedge y \neq y' \\ \wedge (x, \mathcal{X}, *) \in \text{Supp}(\text{Sample}_{\text{vdf}}(1^\lambda, T(\lambda))) \end{array} \right] > 1/q(\lambda),$$

in contradiction. \square

Lemma B.4 (Correctness and completeness proofs). *Let $(\text{Gen}, \text{Sol}, \text{Verify})$ be the construction of a publicly verifiable time-lock puzzle from any one-sided publicly verifiable time-lock puzzle from Section 5.3. $(\text{Gen}, \text{Sol}, \text{Verify})$ satisfy full correctness and completeness.*

Proof. We separately prove full correctness and completeness.

Full correctness. Let $\lambda, t \in \mathbb{N}$, $z \in \{0, 1\}^*$, and $s' = \text{Sol}_{\text{tlp}, 1}(1^\lambda, \text{mcrs}, t, z)$ for any $n \in \mathbb{N}$ and $\text{mcrs} \in (\{0, 1\}^\lambda)^n$.

In the case that $z \in \text{Supp}(\text{Gen}^\mathcal{O}(1^\lambda, t, s))$ for some $s \in \{0, 1\}^\lambda$, we need to show that $s' = s$. Let r be a value such that $z = \text{Gen}^\mathcal{O}(1^\lambda, t, s; r)$. First note that $z \in \text{Supp}(\text{Gen}_{\text{os}}(1^{\lambda_{\text{tlp}}}, t, (s||r)))$. Next, Sol computes $(\hat{s}_{\text{os}}, \pi_{\text{os}}) = \text{Sol}_{\text{os}}(1^{\lambda_{\text{tlp}}}, \text{mcrs}, t, z)$ and parses $\hat{s}_{\text{os}} = \hat{s}||\hat{r}$. By correctness of Sol_{os} , it holds that $\hat{s} = s$ and $\hat{r} = r$. Then, by assumption, it holds that $z = \text{Gen}^\mathcal{O}(1^\lambda, t, s; r)$, so Sol outputs $s' = s$, as required.

In the case that $z \notin \text{Supp}(\text{Gen}^\mathcal{O}(1^\lambda, t, s))$ for any $s \in \{0, 1\}^\lambda$, we need to show that $s' = \perp$. Note that Sol only outputs a value $s' \neq \perp$ if $z = \text{Gen}^\mathcal{O}(1^{\lambda_{\text{tlp}}}, t, s; r)$ for some s . By definition, this can only be the case for $z \in \text{Supp}(\text{Gen}(1^\lambda, t, s))$ for some s . Thus, s' must be equal to \perp in this case.

Completeness. Let $\lambda, t, n \in \mathbb{N}$, $z \in \{0, 1\}^*$, $\text{mcrs} \in (\{0, 1\}^\lambda)^n$. Let $(s, \pi) = \text{Sol}^\mathcal{O}(1^\lambda, \text{mcrs}, t, z)$. We need to show that $\text{Verify}^\mathcal{O}(1^\lambda, \text{mcrs}, t, z, (s, \pi)) = 1$.

First, we consider the case where $z \in \text{Supp}(\text{Gen}^\mathcal{O}(1^\lambda, t, s))$ for some $s \in \{0, 1\}^\lambda$. This means that $z = \text{Gen}^\mathcal{O}(1^\lambda, t, s; r)$ for some $r \in \{0, 1\}^*$. In the proof of correctness above, we showed that $\text{Sol}^\mathcal{O}(1^\lambda, \text{mcrs}, t, z)$ output (s, r) in this case. Since $s \neq \perp$, Verify outputs 1 if and only if $z = \text{Gen}^\mathcal{O}(1^\lambda, t, s; r)$, which it does by assumption.

Next, we consider the case where $z \notin \text{Supp}(\text{Gen}^\mathcal{O}(1^\lambda, t, s))$ for any $s \in \{0, 1\}^\lambda$. In this case, we argued in correctness that $s = \perp$, so it must be the case that $\pi = (s_{\text{os}}, \pi_{\text{os}}) = \text{Sol}_{\text{os}}(1^{\lambda_{\text{os}}}, \text{mcrs}, t, z)$. By completeness of the underlying TLP, we know that $\text{Verify}_{\text{os}}(1^{\lambda_{\text{os}}}, \text{mcrs}, t, z, (s_{\text{os}}, \pi_{\text{os}})) = 1$. Thus, Verify outputs 1 as long as $z \neq \text{Gen}^\mathcal{O}(1^\lambda, t, s; r)$, but this must be the case as $z \notin \text{Supp}(\text{Gen}^\mathcal{O}(1^\lambda, t, \cdot))$ by assumption. \square