

On the Cryptographic Hardness of Local Search

Nir Bitansky¹ 

Tel Aviv University, Israel
nirbitan@tau.ac.il

Idan Gerichter

Tel Aviv University, Israel
idangerichter@gmail.com

Abstract

We show new hardness results for the class of Polynomial Local Search problems (PLS):

- Hardness of PLS based on a falsifiable assumption on bilinear groups introduced by Kalai, Paneth, and Yang (STOC 2019), and the Exponential Time Hypothesis for randomized algorithms. Previous standard model constructions relied on non-falsifiable and non-standard assumptions.
- Hardness of PLS relative to random oracles. The construction is essentially different than previous constructions, and in particular is unconditionally secure. The construction also demonstrates the hardness of parallelizing local search.

The core observation behind the results is that the *unique proofs* property of incrementally-verifiable computations previously used to demonstrate hardness in PLS can be traded with a simple *incremental completeness* property.

2012 ACM Subject Classification Theory of computation → Computational complexity and cryptography

Keywords and phrases Cryptography, PLS, Lower Bounds, Incremental Computation

Digital Object Identifier 10.4230/LIPIcs.ITCS.2020.6

Funding This work was supported in part by ISF grant 18/484, and by Len Blavatnik and the Blavatnik Family Foundation.

Nir Bitansky: Supported by the Alon Young Faculty Fellowship.

1 Introduction

Local search is a well known approach for tackling optimization problems. Local search algorithms seek solutions that are *locally optimal* – they commonly try to improve on a given solution by considering small perturbations of it, called *neighbors*, and testing whether they are better according to some value function. Indeed, many algorithms use this approach with empirical success, for instance, the Simplex linear programming algorithm, the Lin-Kernighan TSP algorithm, and various machine learning algorithms [43, 38, 31]. Nevertheless, the approach has its limitations and many natural local search problems are not known to admit polynomial-time algorithms.

Aiming to characterize the computational complexity of local search, Johnson, Papadimitriou, and Yannakakis [33] introduced the class Polynomial Local Search (PLS). The class is defined by its canonical complete problem LOCAL-SEARCH (LS) [33, 29]. Here the input consists of two polynomial-size circuits: a *successor circuit* $\mathcal{S} : \{0, 1\}^n \rightarrow \{0, 1\}^n$, which given a string x outputs a string x' , and a *value circuit* \mathcal{F} , which given a string x outputs an integer value. The goal is to find a string x which is locally optimal in the sense that $\mathcal{F}(x) \geq \mathcal{F}(\mathcal{S}(x))$. Here strings correspond to solutions, with a value assigned by \mathcal{F} , and the

¹ Member of the Check Point Institute of Information Security.



successor assigned by \mathcal{S} represents “the best” neighbor. Important problems that are known to be PLS complete include finding a locally optimal max-cut [49] and finding pure Nash equilibrium in congestion games [22].

As observed in [33], PLS belongs to the class TFNP of *total search problems* in NP; namely, a local optimal solution always exists (and can be efficiently verified). As a result, it is unlikely that PLS contains NP-hard problems, as this would imply that $\text{NP} = \text{coNP}$ [33, 41]. This barrier is shared of course by other prominent subclasses of TFNP, perhaps the most famous example being the class PPA that captures the hardness of finding Nash equilibrium in bimatrix games [18, 14]. Accordingly, researchers have turned to seek for alternative evidence of hardness.

One natural place to look for such evidence is cryptography. Indeed, cryptography typically relies on problems closer to the boundary of P than to NP hard problems. For some subclasses of TFNP (such as PPP, PPA, and Ramsey) evidence of hardness has been shown based on standard cryptographic assumptions [45, 12, 32, 36, 50]. Establishing the hardness of PLS (and PPA), however, has been far more elusive. Hubáček and Yogevev [29], building on [7, 23], demonstrated a hard problem in $\text{PLS} \cap \text{PPAD}$ based on *indistinguishability obfuscation* [4], a strong cryptographic primitive, yet to be constructed under standard assumptions.² Similar hardness was then shown by Choudhuri et al. [16] assuming hardness of #SAT and the soundness of the Fiat-Shamir transform for the sumcheck protocol, which in turn can be based on *optimally-secure fully-homomorphic encryption against quasi-polynomial attackers* [13]. More recently, the same has been shown based on the hardness of iterated squaring and soundness of Fiat-Shamir for the iterated-squaring protocol of Pietrzak [46] by both Choudhuri et al. [15] and Ephraim et al. [21].

Basing the hardness of PLS (or PPA) on standard assumptions remains an open problem.

1.1 Our Results

We provide new results regarding the hardness of PLS. Our first result shows worst-case hardness of PLS based on *the KPY assumption* on bilinear groups and the randomized Exponential Time Hypothesis (ETH). The KPY assumption was introduced recently by Kalai, Paneth, and Yang [34] toward the construction of publicly-verifiable delegation schemes. The assumption is similar in spirit to previous standard assumptions that generalize Decisional Diffie Hellman (e.g., [8, 9]). It is polynomially falsifiable and holds in the generic group model. Randomized ETH postulates that solving SAT in the worst case requires exponential-time even for randomized algorithms [30].

► **Theorem 1 (Informal).** *Under the KPY assumption on bilinear groups and randomized ETH, PLS is hard in the worst case.*

The result is, in fact, more general and shows a reduction of PLS hardness to a certain type of incrementally-verifiable computation schemes (IVC) [51]. We then derive the required IVC from the work of [34]. We can also show average-case hardness at the cost of assuming superpolynomial hardness of the KPY assumption and average-case randomized ETH. See further details in the technical overview.

Our second result is a construction of PLS instances that are unconditionally hard in *the random oracle model*.

² More accurately, they show a hard problem in the class *Continuous Local Search* (CLS), which is known to be in $\text{PLS} \cap \text{PPAD}$ [19].

► **Theorem 2** (Informal). *Relative to a random oracle, LOCAL-SEARCH is unconditionally hard-on-average.*

Previously, Choudhuri et al. [16] showed hardness of $\text{CLS} \subseteq \text{PLS} \cap \text{PPAD}$ relative to a random oracle, but their result is conditioned on hardness of $\#\text{SAT}$ (the construction does not relativize with respect to this hardness and hence it cannot be derived from the random oracle itself). Indeed, our construction is quite different from theirs. Whereas their construction is based on the sum-check protocol [39] (which explains the non-black-box reliance on $\#\text{SAT}$ hardness), ours comes from recent advancements in *proofs of sequential work* motivated by blockchain applications [40, 17, 20].

The reliance on proofs of sequential work, in fact, translates to a result on the hardness of parallelizing local search. Such hardness, in the random oracle model, was recently shown for CLS (and in particular for both PLS and PPAD) by Ephraim et al. [21] further assuming that repeated squaring modulo a composite is hard to parallelize [47]. Specifically, they construct CLS instances that can be solved in some tuneable sequential polynomial time, but cannot be solved much faster by parallel algorithms. Our result gives a similar kind of hardness for PLS in the random oracle model, without relying on any unproven computational assumptions. We elaborate on this in the technical overview below.

2 Technical Overview

To motivate our approach, we start by recalling previous approaches taken toward proving PLS (or rather CLS) hardness.

2.1 Hardness via Incremental Computation

So far, the common approach toward arguing cryptographic hardness of PLS went through an intermediate problem called SINK-OF-VERIFIABLE-LINE (SVL) [1, 7]. An instance of SVL consists of a successor circuit \mathcal{S} and a verifier circuit \mathcal{V} . The successor \mathcal{S} implicitly defines a step-by-step (or, incremental) computation $s_{t+1} := \mathcal{S}(s_t)$ starting from a canonical source s_0 and ending at a sink s_T , for some superpolynomial T :

$$s_0 \xrightarrow{\mathcal{S}} \dots \xrightarrow{\mathcal{S}} s_T$$

The verifier $\mathcal{V}(s^*, t)$ is promised to accept any s^* as the t -th state if and only if it is correct, namely $s^* = \mathcal{S}^{(t)}(s_0)$. The goal is to find the sink $s_T = \mathcal{S}^{(T)}(s_0)$, given $(\mathcal{S}, \mathcal{V})$.

It is not hard to see that solving SVL can be reduced to solving a local search problem in PLS [29], and thus it is sufficient to prove the hardness of SVL.

Valiant's Incrementally-Verifiable Computation and Uniqueness. A natural way to approach SVL hardness is Valiant's notion of incrementally-verifiable computation (IVC) [51]. According to this notion, we can take a Turing machine computation given by a machine M and input x with a configuration graph:

$$M_x^0 \rightarrow \dots \rightarrow M_x^T,$$

and associate with each intermediate configuration M_x^t a short proof π_t , attesting to its correctness. The proofs are incremental – the next proof π_{t+1} can be efficiently computed from π_t .

6:4 On the Cryptographic Hardness of Local Search

At first glance, this general notion seems to directly yield SVL hardness. Indeed, we can take any polynomial-space machine M solving some hard search problem R on input x and consider the corresponding incrementally-verifiable computation:

$$(M_x^0, \pi_0) \xrightarrow{S} \dots \xrightarrow{S} (M_x^T, \pi_T) .$$

Here the successor simply advances the computation of M and incrementally computes the corresponding proofs. The SVL verifier is derived directly from the IVC verifier.

The reason that the above simple transformation does not work is that *the IVC proofs may not be unique*. In particular, a correct intermediate configuration may have many different accepting proofs attesting to its correctness, whereas SVL requires that only a single string can be accepted as the t -th node (for any t). Choudhuri et al. [16] show that the SVL uniqueness requirement can be somewhat relaxed and traded with *computational uniqueness* meaning that it is computationally hard to find more than a single proof for any given statement (even if such proofs exist).

So far, however, incrementally-verifiable computation with (even computational) uniqueness has not been achieved under standard assumptions. This is, in fact, the case even for specific (hard) computations, let alone for general ones.

2.2 Incremental Completeness

We show that the hardness of PLS does not require SVL hardness nor IVC with unique proof. Rather, we observe that IVC with a conceptually simple *incremental completeness* property suffices. Then we derive such IVC for polynomially-long computations from the delegation scheme construction of Kalai, Paneth, Yang [34] and combine it with ETH to obtain PLS hardness. We next explain the notion of incremental completeness and why it suffices. In the next subsection, we explain how it is obtained.

Incremental-Complete IVC and PLS Hardness. An incremental complete IVC has the property that given *any* accepting proof π for an intermediate state M_x^t of the computation, running the prover results in an accepting proof π' for the next state M_x^{t+1} . This differs from Valiant's original formulation where completeness is only guaranteed for the prescribed proof generated by iteratively applying the prover t times.

Let us now describe how to construct hard LOCAL-SEARCH instances from incrementally-complete IVC. Once again, we start from some hard computation given by a polynomial-space machine M and instance x . Rather than considering a verifiable computation chain, we shall consider a DAG with nodes of the form (t, C, π) , where t is a timestamp, C is an alleged configuration of $M(x)$ at time t , and π is an IVC proof. We call such a node *valid* if the corresponding proof is accepting. For simplicity, we assume for now that the IVC is also perfectly sound, meaning that valid nodes are always such that C is the correct configuration M_x^t . Note that for every time t , there may very well be multiple corresponding valid nodes. We can visualize this as multiple “parallel universes”, which the verifiable computation lives in simultaneously. Incremental completeness says that in *every* such universe, as time moves forward, the computation may proceed. In particular, by induction, we always reach a *sink* of the form (T, M_x^T, π) that contains the last state of the computation.

This naturally gives an LS instance. Valid nodes (t, M_x^t, π) are given value t and are always succeeded by another valid node $(t+1, M_x^{t+1}, \pi')$ derived by advancing the computation and incrementing the IVC proof. Invalid nodes are given value -1 and are succeeded by some canonical source in the DAG $(0, M_x^0, \varepsilon)$, where ε is the empty proof, which by convention is

accepted as a proof for M_x^0 . The local maximums of the corresponding instance are exactly the sinks of the DAG, which by incremental completeness are of the form (T, M_x^T, π) , containing the last state. Therefore, finding a local-maximum reduce to solving the underlying hard computational task given by the computation $M(x)$.

Our actual construction does not assume perfect soundness of the IVC, but only computational soundness. There may exist so called *fooling nodes* (t, C^*, π^*) that pass verification although C^* is not the correct configuration M_x^t . These nodes may be local maximums and finding them may not reduce to solving the underlying problem. Nonetheless, finding fooling local maximum corresponds to breaking the soundness of the IVC, which is computationally infeasible.

2.3 Obtaining IVC with Incremental Completeness

Valiant [51] put forth a general approach toward constructing IVC by *recursive proof composition*. Oversimplifying, the basic idea is that to move from the t -th configuration $C = M_x^t$ and proof π_t to the next configuration $C' = M_x^{t+1}$ and proof π_{t+1} , the new proof π_{t+1} asserts that:

1. There exists a proof π that passes IVC verification as a valid proof that $C = M_x^t$ (this part is recursive in the sense that it relies on IVC verification for smaller times stamps),
2. C' is obtained from C by applying M_x 's transition function.

For this to be feasible, each “single-step proof” is computed using a succinct non-interactive NP proof system where verification time (and in particular, the size of the proof) is fixed and is not affected by the complexity of the NP relation. We observe that *any IVC construction following the above blueprint is incrementally complete* by definition as long as the succinct NP proof system used has perfect completeness.

Succinct Proof Systems. Succinct proofs for NP with unconditional soundness are unlikely to exist [26, 27]. Accordingly, Valiant suggested to instantiate the blueprint using computationally sound succinct arguments. This brings about some extra complications in following the blueprint. For once, proving that a previous proof *exists* is meaningless in the case of computational soundness. Rather we need the guarantee that a proof exists and can be efficiently found. Such systems are known as *succinct non-interactive arguments of knowledge* (SNARKs) and admit an efficient knowledge extractor that can efficiently extract a witness from a convincing prover.

Using such SNARKs enables the construction of IVC for arbitrary poly-time computations (or superpolynomial ones, if the SNARK is superpolynomially secure) [51, 5], and the corresponding constructions are incrementally complete by the fact that the underlying SNARKs are perfectly complete. (The actual constructions are somewhat different from the described blueprint; in particular, to deal with issues such as blowup in knowledge extraction complexity, they aggregate proofs in a tree-like fashion rather than a chain, in order to reduce the depth of the recursion. Incremental completeness holds just the same.)

Following the above, we can obtain an incrementally-complete IVC, and thus PLS hardness from SNARKs. However, the existence of SNARKs is a strong non-falsifiable assumption. It is only known under non-standard *knowledge assumptions* and its construction from standard assumptions is subject to barriers [24, 6].

Incrementally-Complete IVC from Weaker Succinct Arguments. In a recent work, Kalai, Paneth, and Yang [34] considered a relaxation of SNARKs called *quasi-arguments* that instead of standard knowledge extraction only requires a certain weaker *no-signaling extraction*

requirement. Unlike SNARKs, such arguments are not subject to any known lower bounds, and were in fact constructed from the learning with errors assumption in their privately-verifiable form [35, 44, 11, 3] and recently also in their publicly-verifiable form based on the so called KPY assumption in bilinear groups [34].

Furthermore, Kalai, Paneth, and Yang show that similarly to SNARKs such quasi-arguments can, in fact, be recursively composed toward verifying a long computation, provided that the computation is *deterministic*. Their explicit goal was not to construct IVC but rather to bootstrap quasi-arguments with a long common reference string (CRS) to succinct arguments with a short CRS. However, this is done by implicitly constructing an IVC (this connection is also apparent in previous constructions of SNARKs with a short CRS from ones with a long CRS [5]). The resulting IVC, like other IVCs based on recursive composition, is incrementally complete.

The Class of Supported Computations and the Reliance on ETH. In its native form, the IVC derived from the KPY construction supports computations of arbitrary polynomial length $T = \lambda^{O(1)}$ in the security parameter λ , with fixed prover-verifier running time (say, λ). However, as explained earlier, to get PLS hardness, we would like to apply the IVC for a hard (and thus superpolynomial) computation.

To circumvent this difficulty, we employ a fine-grained reduction. The idea is to construct instances which are tailor-made to be hard for LS algorithms of a specific polynomial running time. Fix any constant $c > 0$ and n^c -time algorithm \mathcal{A} that supposedly solve LS (here n is the size of the LS instance $(\mathcal{S}, \mathcal{F})$). Consider a search problem R and a Turing machine M that solves it in polynomial time $T(\lambda)$, while no randomized algorithm can solve it in time T^δ for constant $\delta < 1$. If we can bound the *blowup* of the IVC reduction by at most $T^{\delta/c}$, we can use \mathcal{A} to construct a randomized adversary \mathcal{A}' that solves R in time T^δ and get a contradiction.

The blowup of the IVC reduction is polynomially related to (a) the input size, (b) the space used by M , and (c) the efficiency of the IVC scheme. Accordingly, we require a computation where there is an arbitrarily large (polynomial) gap between the space used and the hardness of the underlying problem. To this end, we use the randomized ETH assumption. By appropriately choosing the size of the underlying SAT instance, we get computations that can be solved using fixed-space and some polynomial time via brute-force but not in much faster polynomial time, even by randomized algorithms. To bound the blowup due to the IVC efficiency, we follow the efficiency analysis by [34] and adapt it to the incremental setting. (In the body, we show the above in the non-uniform setting: assuming non-uniform ETH and using standard non-uniform derandomization techniques while taking special care to bound the associated blowup.)

We note that the above only gives worst-case hardness. That is, we do not show a single distribution that is hard for all polynomial-time algorithms. Assuming slight super-polynomial security of the KPY assumption and average-case randomized ETH, we can get average-case PLS hardness by essentially the same reduction (in fact, we can slightly weaken average-case ETH to a slightly subexponential time hypothesis).

2.4 Unconditional Hardness in the Random Oracle Model

Our second result, in the random model, is based on recent advancements in *proofs of sequential work* (PoSW) [40]. Roughly speaking, in a PoSW, the prover is given a statement χ and a time parameter T , and can generate a corresponding proof π by making T sequential steps. The soundness requirement is that provers that make $\ll T$ sequential steps, will fail to

generate a valid proof for χ , except with negligible probability (namely, proof computation is not parallelizable). The construction we present relies on the PoSW of Cohen and Pietrzak [17] and its extension by Döttling, Lai, and Malavolta [20] to so called *incremental proof of sequential work*; both are proven sound in the random oracle model. We now move to recalling how these systems work, and explain how we derive from them hard PLS instances relative to random oracles.

The CP PoSW. To construct a PoSW, Cohen and Pietrzak [17] suggest an elegant tree labeling procedure that is inherently sequential. They consider a binary tree whose edges are directed from the leaves toward the root. In addition, they add directed edges from every node ℓ having a sibling r on the right, to all the leaves in the tree rooted at r . We call this the CP graph (see Figure 1 in Section 6 for an illustration). The nodes of the CP graph are then given random labels as follows: the label of any given node is obtained by applying a random oracle $H_\chi = H(\chi, \cdot)$ to the labels of its incoming nodes (the oracle is seeded with χ to guarantee an independent oracle for every statement). Intuitively, the added edges enforce that in order to compute the labels of some subtree rooted at some r , it is first necessary to compute the labels of its left sibling ℓ ; in this sense, labeling must be done sequentially.

To turn this into an actual proof of sequential work π for χ , the prover publishes the label of the root of the entire tree. The verifier then responds with random challenge leaves to which the prover answers by providing all the labels in the corresponding paths toward the root along with the labels of the siblings along the path, as in standard Merkle tree authentication (indeed here the tree serves both the purpose of of a commitment and of enforcing the sequential nature). Finally, they make the proof non interactive by using the Fiat-Shamir transform. Cohen and Pietrzak prove that to successfully compute an accepting proof, a prover must sequentially call the random oracle $\approx T$ times, where T is the size of the tree.

Incremental PoSW and the DLM Construction. Cohen and Pietrzak further show that the standard streaming algorithm for Merkle commitments [42], can also be applied to their labeling procedure. That is, one can keep track of a small amount of nodes (about $\log T$), each a root of some tree in a gradually growing forest, and when needed, merging two nodes on the same level. In other words, the CP labeling can be done by an incremental computation with small space.

One thing that is of course missing toward getting PLS hardness is verifiability of the corresponding intermediate states. One Naïve idea toward such local verifiability is to apply the CP proof strategy for each subtree in the incremental labeling process. Indeed, each such subtree is, in fact, CP labeled on its own, so we can use Fiat-Shamir to compute random challenges for that particular subtree. While this is a step in the right direction it is still not enough for the goal of PLS hardness, since the intermediate proofs themselves are not computed incrementally – computing each local proof may take time proportional the size of the corresponding subtree.

Döttling, Lai, and Malavolta [20] suggested a neat idea to make the CP proofs incrementally computable.³ Roughly speaking rather than sampling fresh challenges for each subtree, they suggested to derive them at random from the previously computed challenges of this subtree. In a bit more detail, whenever merging two subtrees rooted at ℓ and r into a new

³ They were motivated by blockchains and aimed to make PoSW a process that can be advanced in a distributed fashion by different parties.

(taller) subtree, we assume by induction that each of the two already has a set of corresponding challenges S_ℓ and S_r , where say the size of each one is some parameter k . Then, for the new subtree, rooted at their parent v , we choose a new set S_v of k challenges by choosing k challenges at random from the $2k$ challenges in $S_\ell \cup S_r$ (this is again done non-interactive using Fiat-Shamir). Each of the challenges is again just a Merkle authentication path, which can be easily augmented to an authentication path for v . They prove that this choice of challenges guarantees, for every intermediate state, essentially the same soundness as the original CP construction guarantees for the final state.

Incremental Completeness and PLS Hardness. What we get from the DLM incremental PoSW (in the random oracle) is somewhat analogous to an IVC system in the standard model. However, while in the standard model we applied general IVC over some long (hard) computation, here there is no such underlying hardness. Rather, *the difficulty is only in computing the accepting proofs themselves*. Similarly to our standard model approach, here too we can address the concept of incremental completeness. Indeed, we observe that merging two accepting proofs in the DLM construction always yields an accepting proof. From here incremental completeness follows. Formally, we need to slightly augment the DLM construction to enforce consistency among different subtrees in a forest to guarantee such incremental completeness (see Section 6 for the details).

Choosing the parameters appropriately, yields problems in PLS relative to a random oracle which are exponentially hard-on-average. This translates to sub-exponential hardness of the canonical LS problem (due to polynomial blowup of the reduction). Also, by choosing the size of the tree to be a polynomial of our choice, we get search problems in PLS which are moderately hard but “depth-robust” in the sense that they cannot be solved much faster by parallel algorithms [21]. This follows from the sequential hardness of DLM.

2.5 More Related Work on Total Search Problems

The class TFNP was introduced by Megiddo and Papadimitriou [41]. They observed that TFNP is unlikely to include NP hard problems as this would imply that $\text{NP} = \text{coNP}$. Furthermore, it is strongly believed that TFNP does not have complete problems (see for instance discussion in [25]). Toward understanding of the complexity of total search problems, Papadimitriou [45] introduced several “syntactic” subclasses of TFNP that cluster problems according to the mathematical argument used to establish their totality. Recently, Goldberg and Papadimitriou [25] presented a subclass called *provable* TFNP that contains previously defined subclasses (and can be seen as generalizing them) and admits complete problems.

A long line of works have investigated the complexity of TFNP and its subclasses, searching for efficient algorithms on one hand and evidence of hardness on the other. As explained earlier in the intro, cryptography is a natural place to look for such evidence. Indeed, basing the hardness of TFNP or its subclasses, on standard cryptographic assumptions, has been successful in several cases. For example, Papadimitriou [45] observed that PPP is hard assuming one-way permutations or collision-resistant hash functions, and Jeřábek [32], building on the work of Buresh-Oppenheimer [12], demonstrated the hardness of PPA assuming factoring is hard. Hubáček, Naor and Yogev [28] showed TFNP-hardness assuming any average-case hard NP language exists (in particular one-way functions) and derandomization assumptions. Komargodski, Naor, and Yogev [37] showed that hardness of the class RAMSEY is equivalent to the existence of multi-collision resistant hash functions.

As discussed earlier, demonstrating hardness for PPAD and PLS has been more challenging and so far only achieved under non-standard cryptographic assumptions. Trying to explain our failure so far to base such hardness on standard cryptographic primitives, Rosen, Segev

and Shahaf [48] show that TFNP hard instances constructed in a black box way from random oracles (or some variants thereof) must have a nearly exponential number of solutions (which is indeed the case in our result in the random oracle model).

Finally, we note that in the smooth complexity setting, several PLS complete problems, with natural noise distributions, have been shown to be solvable in smooth polynomial-time. This include finding locally optimal Max-Cut [2] and finding pure Nash equilibrium in network coordination games [10]. These algorithms suggest an explanation of why local search may be empirically easy, while hard instances exist under reasonable assumptions and can be sampled.

Organization

In Section 3, we recall some preliminaries including the definition of PLS. In Section 4, we construct a computational reduction from a search problem having an incremental complete IVC to LS. In Section 5, we instantiate the computational reduction under the KPY assumption and ETH to demonstrate PLS hardness in the plain model. In Section 6, we argue the hardness of PLS relative to a random oracle.

3 Preliminaries

Notation. Throughout, λ will denote security parameters. When the input size is not the security parameter, we denote it by n . For $a \leq b$ integers we denote by $[a, b]$ the set $\{a, \dots, b\}$ and by $[a]$ the set $[1, a]$.

3.1 PLS Definition

The complexity class *polynomial local search* (PLS) consists of all TFNP search problems polynomial-time reducible to the LOCAL-SEARCH problem [33, 29] we denote by LS.

► **Definition 3** (LOCAL-SEARCH). *The search problem LS is the following: given two polynomial-size circuits $\mathcal{S} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $\mathcal{F} : \{0, 1\}^n \rightarrow \mathbb{N}$, find a string $v \in \{0, 1\}^n$ such that $\mathcal{F}(\mathcal{S}(v)) \leq \mathcal{F}(v)$.*

We refer to \mathcal{S} as the *successor circuit* and to \mathcal{F} as the *value circuit*. Intuitively, the circuits \mathcal{S} and \mathcal{F} could be seen as representing an implicit DAG over $\{0, 1\}^n$. In this graph, a node v is connected to $u = \mathcal{S}(v)$ if and only if $\mathcal{F}(u) > \mathcal{F}(v)$. This is a DAG since \mathcal{F} induce a topological ordering. Notice also that every sink corresponds to a local maximum with respect to \mathcal{F} . With this perspective, the totality of LS follows from the fact *every finite DAG has a sink*.

Oracle aided PLS. We denote by $\mathcal{C}^\mathcal{O} = \{\mathcal{C}_\lambda^\mathcal{O}\}$ an oracle aided circuit family. That is a circuit family with oracle gates to $\mathcal{O}_\lambda : \{0, 1\}^{r(\lambda)} \rightarrow \{0, 1\}^{\ell(\lambda)}$. The complexity class $\text{PLS}^\mathcal{O}$ relative to an oracle \mathcal{O} , is naturally defined as all TFNP $^\mathcal{O}$ search problem polynomial-time reducible to $\text{LS}^\mathcal{O}$.

► **Definition 4** (Oracle Aided LOCAL-SEARCH). *The search problem $\text{LS}^\mathcal{O}$ is the following: given two polynomial-size, oracle aided, circuits $\mathcal{S}^\mathcal{O} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $\mathcal{F}^\mathcal{O} : \{0, 1\}^n \rightarrow \mathbb{N}$, find a string $v \in \{0, 1\}^n$ such that $\mathcal{F}^\mathcal{O}(\mathcal{S}^\mathcal{O}(v)) \leq \mathcal{F}^\mathcal{O}(v)$.*

4 PLS Hardness from IVC

In this section, we give a computational reduction from search problems having an IVC with incremental completeness to LS. This, in turn, establishes that any (worst/average-case) hard search problem, having an IVC with incremental completeness, implies the hardness of PLS (in the worst/average-case, respectively). We start by formally defining IVC with incremental completeness.

4.1 IVC with Incremental Completeness

Conventions. Let M be a Turing machine with $T = T(\lambda)$ and $S = S(\lambda)$ bounds on its run time and space, respectively, for input of length λ . Throughout, we assume w.l.o.g that M always makes exactly T steps and the size of a configuration is exactly S . Let $x \in \{0, 1\}^\lambda$ be an input to M , we denote by $M_x^t \in \{0, 1\}^S$ the configuration in the t step of the computation $M(x)$. We denote by $M_x : \{0, 1\}^S \rightarrow \{0, 1\}^S$ the transition circuit between configurations.

► **Definition 5** (IVC with Incremental Completeness). *Let M be a Turing machine with $T = T(\lambda)$ and $S = S(\lambda)$ bounds on its run time and space, respectively, for input of length λ .*

An Incremental Verifiable Computation scheme (IVC) for M , with incremental completeness, consists of three algorithms $\text{IVC} = (\text{IVC.G}, \text{IVC.P}, \text{IVC.V})$:

- $\text{IVC.G}(x)$ is a randomized algorithm that given $x \in \{0, 1\}^\lambda$ outputs public parameters pp .
- $\text{IVC.P}(pp, t, C, \pi)$ is a deterministic algorithm that given public parameters pp , a natural number t , an arbitrary configuration C and arbitrary proof π , outputs a proof π' .
- $\text{IVC.V}(pp, t, C, \pi)$ is a deterministic algorithm that given public parameters pp , a natural number t , an arbitrary configuration C and arbitrary proof π , outputs ACC or REJ.

We make the following requirements:

1. **Incremental Completeness:** For every security parameter $\lambda \in \mathbb{N}$:

- a. For every input $x \in \{0, 1\}^\lambda$, time $t \in [0, T - 1]$ and candidate proof $\pi \in \{0, 1\}^*$:

$$\Pr \left[\begin{array}{l} \text{IVC.V}(pp, t, M_x^t, \pi) = \text{ACC} \implies \\ \text{IVC.V}(pp, t + 1, M_x^{t+1}, \pi') = \text{ACC} \end{array} \middle| \begin{array}{l} pp \leftarrow \text{IVC.G}(x) \\ \pi' = \text{IVC.P}(pp, t, M_x^t, \pi) \end{array} \right] = 1 ,$$

where M_x^t, M_x^{t+1} are the configurations in the t and $t + 1$ steps of the computation of $M(x)$, respectively.

- b. For every input $x \in \{0, 1\}^\lambda$:

$$\Pr [\text{IVC.V}(pp, 0, M_x^0, \varepsilon) = \text{ACC} \mid pp \leftarrow \text{IVC.G}(x)] = 1 ,$$

where M_x^0 is the first configuration of the computation of $M(x)$, and ε is the empty proof.

2. **Soundness:** For every efficient adversary \mathcal{A} , there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and $x \in \{0, 1\}^\lambda$:

$$\Pr \left[\begin{array}{l} C^* \neq M_x^t \\ \text{IVC.V}(pp, t, C^*, \pi^*) = \text{ACC} \end{array} \middle| \begin{array}{l} pp \leftarrow \text{IVC.G}(x) \\ (t, C^*, \pi^*) \leftarrow \mathcal{A}(x, pp) \end{array} \right] \leq \mu(\lambda) ,$$

where M_x^t is the configuration in the t step of the computation of $M(x)$.

3. **Efficiency:** The efficiency of IVC, denoted by $\mathbb{T}_{\text{IVC}}(\lambda)$, is the maximal worst-case run-time among $\text{IVC.G}, \text{IVC.P}, \text{IVC.V}$ for input having security parameter λ . We require $\mathbb{T}_{\text{IVC}}(\lambda) \leq p(\lambda)$ for a fixed polynomial p .

4.2 Computationally Sound Karp Reduction

As discussed in the introduction, since we use IVC with computational soundness, the reduction we give is also computationally sound. In what follows, we define the notion of computationally sound Karp reduction.

For $R, R' \in \text{FNP}$, a computationally sound reduction from R to R' consists of a pair of efficient algorithms $(\mathcal{X}, \mathcal{W})$. Where $\mathcal{X}(x)$ is a randomized algorithm that given an instance x of R , samples an instance $x' \leftarrow \mathcal{X}(x)$ of R' . $\mathcal{W}(w')$ is a deterministic algorithm that translates a witness w' for x' , to a candidate witness w of x . We require that its infeasible to find w' , such that $w = \mathcal{W}(w')$ is not a witness for x in R .

► **Definition 6** (Computational Karp Reduction). *For $R, R' \in \text{FNP}$, a computational Karp reduction from R to R' consists of a pair $(\mathcal{X}, \mathcal{W})$, where $\mathcal{X}(x)$ is a randomized efficient algorithm and $\mathcal{W}(w)$ is a deterministic efficient algorithm. We make the following requirement:*

- **Computational Soundness:** *For every efficient adversary \mathcal{A} , there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and $x \in \{0, 1\}^\lambda$ for which R_x is non-empty:*

$$\Pr \left[\begin{array}{l|l} w' \in R'_{x'} & x' \leftarrow \mathcal{X}(x) \\ w \notin R_x & w' \leftarrow \mathcal{A}(x') \\ & w = \mathcal{W}(w') \end{array} \right] \leq \mu(\lambda) ,$$

where R_x and $R'_{x'}$ are the set of witnesses for x in R and x' in R' respectively.

► **Definition 7** (Reduction Efficiency). *The efficiency of a computational Karp reduction $(\mathcal{X}, \mathcal{W})$, denoted $\mathbb{T}_{\text{Red}}(\lambda)$, is the maximum of $\mathbb{T}_{\mathcal{X}}(\lambda), \mathbb{T}_{\mathcal{W}}(\lambda)$ where:*

1. $\mathbb{T}_{\mathcal{X}}(\lambda)$ is the worst-case run-time of \mathcal{X} for input $x \in \{0, 1\}^\lambda$.
2. $\mathbb{T}_{\mathcal{W}}(\lambda)$ is the worst-case run-time of \mathcal{W} for input $w' \in \left\{ \begin{array}{l} x \in \{0, 1\}^\lambda \\ x' \in \text{Supp}(\mathcal{X}(x)) \\ w' \in R_{x'} \end{array} \right\}$.

Note that $\mathbb{T}_{\text{Red}}(\lambda)$ is a bound on instance size sampled using \mathcal{X} .

4.3 The Hardness Reduction

Using the notion of computational Karp reduction, our hardness result is formalized in the following theorem.

► **Theorem 8** (IVC Reduction). *Let $R \in \text{FNP}$, solvable by a polynomial-space Turing machine M . If there exists an IVC scheme with incremental completeness for M , there exists a computationally sound Karp reduction $(\mathcal{X}, \mathcal{W})$ from R to LS. The efficiency of the reduction is*

$$\mathbb{T}_{\text{Red}} = \text{poly}(\mathbb{T}_{\text{IVC}}, S, \lambda, |M|) ,$$

where $S(\lambda)$ is a bound on M space and $\mathbb{T}_{\text{IVC}}(\lambda)$ is the efficiency of the IVC. The polynomial doesn't depend on the IVC scheme, M nor R .

► **Remark 9** (Efficiency Independent of T). Jumping ahead, the fact that \mathbb{T}_{Red} does not depend directly on the time of the computation T , plays an important role in the fine-grained reduction for polynomial-time computations presented in Section 5.

Next, we describe the reduction and in the following section we analyze its security.

4.3.1 The Reduction

Let $\text{IVC} = (\text{IVC.G}, \text{IVC.P}, \text{IVC.V})$ be an incremental verifiable computation scheme with incremental completeness for M . Let $S = S(\lambda)$ be the polynomial bound on space on M when executed on $x \in \{0, 1\}^\lambda$, guaranteed by the statement. Recall that we denote by M_x^t the configuration in the t step of the computation of $M(x)$. We assume w.l.o.g that the configuration size is exactly S . Let $T = T(\lambda) = 2^{S(\lambda)} \leq 2^{\text{poly}(\lambda)}$ be an upper bound on the running time of M . We further assume w.l.o.g that M always makes exactly T steps. With the above notation, the configuration graph of M when executing x is:

$$M_x^0 \longrightarrow \cdots \longrightarrow M_x^T .$$

In the following, we describe the construction of $\mathcal{X}(x)$ for an input $x \in \{0, 1\}^\lambda$. Recall that we denote by $M_x : \{0, 1\}^S \rightarrow \{0, 1\}^S$ the transition circuit of $M(x)$, taking a configuration as input and returning the next configuration as output.

Instance translation. The algorithm $\mathcal{X}(x)$ begins by sampling $pp \leftarrow \text{IVC.G}(x)$. Given pp , we apply a deterministic procedure I to generate an LS instance $(\mathcal{S}, \mathcal{F}) = I(x, pp)$. The constructed LS instance corresponds to a graph on vertices of fixed polynomial length $\ell = \ell(\lambda) \geq \lambda$. Each vertex is of the form $(t, C, \pi) \in \{0, 1\}^\ell$ where t is parsed as an integer in $[0, T]$, C as a configuration of length S and π as a proof for the IVC padded to length ℓ if needed. We describe the successor \mathcal{S} and value circuits \mathcal{F} formally in the following.

► **Remark 10.** We assume w.l.o.g that M_x always outputs something. If the transition function of M fails to generate the next configuration given C , it will simply output M_x^0 , the initial state.

► **Remark 11.** We assume w.l.o.g that \mathcal{F} outputs values in $[-1, T]$.

In what follows, $v_0 := (0, M_x^0, \varepsilon)$ where M_x^0 is the initial state of $M(x)$ and ε is the empty proof.

Successor Circuit $\mathcal{S}(t, C, \pi)$

Hardwired: The public parameters pp , and the input x .

Algorithm:

1. If $\text{IVC.V}(pp, t, C, \pi) = \text{REJ}$, output v_0 .
2. If $t = T$, output (t, C, π) .
3. Compute $C' = M_x(C)$, $\pi' = \text{IVC.P}(pp, t, C, \pi)$. Output $(t + 1, C', \pi')$.

Value Circuit $\mathcal{F}(t, C, \pi)$

Hardwired: The public parameters pp .

Algorithm:

1. If $\text{IVC.V}(pp, t, C, \pi) = \text{REJ}$, output -1 .
2. Else, output t .

Witness translation. The algorithm $\mathcal{W}(t, C, \pi)$ simply returns the content of the output tape in the configuration C .

Efficiency. The successor circuit M_x is of size $\text{poly}(S, \lambda, |M|)$. The computations related to IVC are of size $\text{poly}(\mathbb{T}_{\text{IVC}})$ by Turing machine simulation. All other computations are polynomial in the input length ℓ . We have that $\ell = \text{poly}(S, \mathbb{T}_{\text{IVC}})$. It follows that $\mathbb{T}_{\text{Red}} = \text{poly}(\mathbb{T}_{\text{IVC}}, S, \lambda, |M|)$. All the above polynomials don't depend on M nor IVC.

By the fact M is poly-space and the efficiency requirement of IVC, that is $\mathbb{T}_{\text{IVC}} \leq \text{poly}(\lambda)$, both \mathcal{X}, \mathcal{W} are polynomial-time algorithms, as required. In turn, this further implies that \mathcal{S}, \mathcal{F} are of size $\text{poly}(\lambda)$. Since the length is a polynomial such that $\ell(\lambda) \geq \lambda$, we have that \mathcal{S}, \mathcal{F} are polynomial-size circuits.

4.4 Security Analysis

Fix $\lambda \in \mathbb{N}$ and $x \in \{0, 1\}^\lambda$. Let $(\mathcal{S}, \mathcal{F})$ be an instance in the support of $\mathcal{X}(x)$ and let pp be the public parameters hardwired in $(\mathcal{S}, \mathcal{F})$. We prove that all local maximums of $(\mathcal{S}, \mathcal{F})$ are one of two types:

- **Honest:** Nodes (T, M_x^T, π) where $\text{IVC.V}(pp, t, M_x^T, \pi) = \text{ACC}$.
- **Fooling:** Nodes (t, C^*, π^*) for which $C^* \neq M_x^t$ and $\text{IVC.V}(pp, t, C^*, \pi^*) = \text{ACC}$.

The proof of the following claim use the incremental completeness of IVC.

▷ **Claim 12.** All local maximums of $(\mathcal{S}, \mathcal{F})$ are honest-type or fooling-type.

Proof of claim. Let $v = (t, C, \pi)$ be a local maximum of $(\mathcal{S}, \mathcal{F})$, that is $\mathcal{F}(\mathcal{S}(v)) \leq \mathcal{F}(v)$. We have that $\text{IVC.V}(pp, t, C, \pi) = \text{ACC}$, as otherwise it is given value -1 and is connected to $v_0 = (0, M_x^0, \varepsilon)$ which is of value 0 . If $C \neq M_x^t$ then v is a fooling-type local maximum and if $C = M_x^t$ and $t = T$ then v is an honest-type local maximum.

We are left with the case $C = M_x^t$ and $t < T$. By construction, $\mathcal{S}(v) = (t + 1, M_x^{t+1}, \pi')$ for $\pi' = \text{IVC.P}(pp, t, M_x^t, \pi)$. By incremental completeness, $\text{IVC.V}(pp, t + 1, M_x^{t+1}, \pi') = \text{ACC}$ and thus $\mathcal{F}(\mathcal{S}(v)) = t + 1$ while $\mathcal{F}(v) = t$. Therefore it is not a local maximum, proving the claim. ◁

Note that only honest-type local maximums translate by \mathcal{W} to a valid witness for R , the underlying search problem. While there exist fooling-type local maximums, by the security of the IVC, finding them efficiently is infeasible. We proceed to prove Theorem 8.

Proof of Theorem 8. Let \mathcal{A} be an efficient adversary that attempts breaking the computational soundness of the reduction $(\mathcal{X}, \mathcal{W})$. Let $\lambda \in \mathbb{N}$ and $x \in \{0, 1\}^\lambda$ be such that R_x is non-empty. We denote by $\varepsilon = \varepsilon(x)$ the probability of success of \mathcal{A} . That is

$$\varepsilon(x) := \Pr \left[\begin{array}{l} \mathcal{F}(\mathcal{S}(w')) \leq \mathcal{F}(w') \\ w \notin R_x \end{array} \middle| \begin{array}{l} (\mathcal{S}, \mathcal{F}) \leftarrow \mathcal{X}(x) \\ w' \leftarrow \mathcal{A}(\mathcal{S}, \mathcal{F}) \\ w = \mathcal{W}(w') \end{array} \right].$$

It follows by Claim 12 that all witnesses w' for $(\mathcal{S}, \mathcal{F})$ such that $\mathcal{W}(w') = w \notin R_x$ are fooling-type local maximums. Indeed, for every honest-type local maximum, \mathcal{W} outputs the content of the output tape of M_x^T . As M solves R and R_x is non-empty, we have that the output tape of M_x^T , the last configuration, consists of w such that $w \in R_x$.

Consider $\mathcal{A}' = \{\mathcal{A}'_\lambda\}_{\lambda \in \mathbb{N}}$ that attempts breaking the soundness of the IVC defined as follows. Recall $I(x, pp)$ is the deterministic procedure used by \mathcal{X} to construct the circuits \mathcal{S} and \mathcal{F} .

$\mathcal{A}'(x, pp)$:

1. Compute $(\mathcal{S}, \mathcal{F}) = I(x, pp)$.
2. Simulate $w' \leftarrow \mathcal{A}(\mathcal{S}, \mathcal{F})$.
3. Output $w' = (t, C, \pi)$.

6:14 On the Cryptographic Hardness of Local Search

Note that, every fooling-type local maximum corresponds to (t, C^*, π^*) breaking the IVC scheme. Further notice that the distribution $(\mathcal{S}, \mathcal{F}) = I(x, pp)$ for $pp \leftarrow \text{IVC.G}(x)$ is, by definition, the distribution of $\mathcal{X}(x)$. Therefore, the success probability of $\mathcal{A}'(x, pp)$ is at least $\varepsilon(x)$. Using the soundness of IVC, there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and $x \in \{0, 1\}^\lambda$ such that R_x is non-empty we have the following:

$$\varepsilon(x) \leq \Pr \left[\begin{array}{l} C^* \neq M_x^t \\ \text{IVC.V}(pp, t, C^*, \pi^*) = \text{ACC} \end{array} \mid \begin{array}{l} pp \leftarrow \text{IVC.G}(x) \\ (t, C^*, \pi^*) \leftarrow \mathcal{A}'(x, pp) \end{array} \right] \leq \mu(\lambda) .$$

Establishing the computational soundness of the reduction. The efficiency part of the theorem has been argued in Section 4.3. \blacktriangleleft

4.5 Applying the Reduction

In this section we present results that establish that a computational Karp reduction from an underlying hard (worst-case/average-case) problem R to LS implies the hardness of LS (in worst-case/average-case respectively). The proofs for the following propositions are given in the full version of this paper.

4.5.1 Worst-Case Hardness

The existence of a computational Karp reduction from R to LS and an efficient solver for LS implies, in a straightforward way, the existence of an efficient randomized solver for R , successful with high probability. Proposition 13 extends this and show the existence of a deterministic solver for R . This is by standard non-uniform derandomization techniques. We use n to denote the size of LS instances and λ the size of R instances.

► **Proposition 13** (Worst-Case Hardness). *Let R be an FNP search problem having a computational Karp reduction to LS. Assume there exists an adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ of polynomial-size $s(n)$ solving LS in the worst-case. Then there exists an adversary $\mathcal{A}' = \{\mathcal{A}'_\lambda\}_{\lambda \in \mathbb{N}}$ solving R in the worst-case. The size of \mathcal{A}' is*

$$\text{size}(\mathcal{A}') = \text{poly}(s(\mathbb{T}_{\text{Red}}), \mathbb{T}_R, \lambda) ,$$

where $\mathbb{T}_{\text{Red}}(\lambda)$ is the efficiency of the reduction and $\mathbb{T}_R(\lambda)$ is the efficiency of the NP verification $R(x, y)$.

4.5.2 Average-Case Hardness

► **Proposition 14** (Average-Case Hardness). *If there exists a hard-on-average FNP problem R , with a computationally sound Karp reduction from R to LS, then LS is hard-on-average.*

Using the computational Karp reduction constructed in Section 4.3, we get the following corollary:

► **Corollary 15** (IVC Average-Case Hardness). *Let R be a hard-on-average FNP problem. If there exists an IVC scheme with incremental completeness for a polynomial-space Turing machine M solving R , then LS is hard-on-average.*

5 Instantiation under the KPY Assumption and ETH

In Section 4, we were able to prove that an IVC scheme with incremental completeness for some superpolynomially hard computation suffices to show that LS is hard. It is left to instantiate the above under the KPY assumption and construct an IVC scheme for a superpolynomially hard computation. Unfortunately, under the KPY assumption we are able to construct IVC schemes only for polynomial computations⁴. In this section, we show how to circumvent this difficulty by employing a fine-grained reduction and further assuming the non-uniform exponential time hypothesis (ETH). Finally proving the worst-case hardness of LS under the KPY assumption and non-uniform ETH.

► **Remark 16.** For ease of notation, in the following section we assume w.l.o.g $\lambda > 1$.

5.1 IVC From KPY Assumption

The proof of Theorem 18 is given in the full version of this paper.

► **Remark 17 (Dependence of Constants).** In the following, we use the constant c instead of big O notation in order to emphasize that the constant is independent of α .

► **Theorem 18 (IVC from KPY Assumption).** *Fix any constants $c, \varepsilon > 0$. Let α be a large enough constant and let M be a Turing machine of run-time $T(\lambda) \leq c\lambda^\alpha$, configuration size $S(\lambda) \leq c\lambda$ and description size $|M| \leq c \log_2 \alpha$. Under the KPY assumption, there exists an IVC scheme with incremental completeness for M having efficiency $\mathbb{T}_{\text{IVC}} = \lambda^{\varepsilon\alpha}$.*

► **Remark 19 (Non-Uniform Reduction).** The security reduction of the IVC scheme given in Theorem 18 is non-uniform and therefore we require that the IVC scheme, and accordingly the KPY assumption, to hold against non-uniform polynomial-time attackers. This is to be expected when dealing with worst-case hardness.

5.2 Fixed Space Polynomial Hardness via ETH

In this section, assuming the non-uniform ETH assumption, we show the existence of arbitrarily polynomially hard computations for circuits that can be solved by Turing machines using a fixed amount of space and slightly larger polynomial time.

► **Assumption 20 (Non-Uniform ETH).** *There exists a constant $\delta > 0$ such that no circuit family of size $O(2^{\delta n})$ solves 3SAT, where n is the length of the input 3CNF formula.*

► **Proposition 21 (Polynomial Hardness).** *Under the non-uniform ETH assumption, there exists constants $c, \delta > 0$ such that for every constant $\alpha > 1$ there exists a search problem $R_\alpha \in \text{FNP}$ satisfying the following:*

- **Algorithm:** *There exists a Turing machine M_α of run-time $T(\lambda) \leq c\lambda^\alpha$, configuration size $S(\lambda) \leq c\lambda$ and description size $|M_\alpha| \leq c \log \alpha$ solving R_α .*
- **Lower Bound:** *No circuit family of size $O(\lambda^{\delta\alpha})$ can solve R_α .*

Proof. Consider

$$R_\alpha := \left\{ (x0^{2^{|x|/\alpha}}, w) \mid (x, w) \in 3\text{SAT} \right\} .$$

⁴ This is for polynomial security of the KPY assumption. A super-polynomial security of the KPY assumption yields an IVC scheme for computations of a corresponding superpolynomial length.

6:16 On the Cryptographic Hardness of Local Search

Algorithm. The naive brute force algorithm for R_α satisfy the requirement. In the following α is not suppressed by O notation. For $x \in \{0, 1\}^n$ and input of the form $x0^{2^{n/\alpha}-n}$, brute force takes $O(2^n)$ time while the whole input is of size $\lambda = 2^{n/\alpha}$. The description size is $O(\log \alpha)$ and the memory used is $O(\lambda)$.

Lower Bound. Let $\delta > 0$ be the constant assumed to exist by non-uniform ETH. Every circuit family of size $O(\lambda^{\delta\alpha})$ solving R_α can be turned to a circuit of size $O(2^{\delta\alpha})$ solving 3SAT. Contradicting non-uniform ETH. ◀

5.3 Putting Things Together

In this section, we put things together to prove that under non-uniform ETH and the KPY assumption we have $\text{PLS} \not\subseteq \text{FP/Poly}$.

► **Theorem 22 (PLS Hardness).** *Assuming non-uniform ETH and the KPY assumption we have $\text{PLS} \not\subseteq \text{FP/Poly}$.*

Toward proving the theorem, we start by bounding the blowup associated with the IVC reduction when applied to the brute-force computation solving R_α defined by Proposition 21.

► **Lemma 23 (Bounding the Blowup).** *Fix a constant $\varepsilon > 0$. Let α be a large enough constant. Under the KPY assumption, there exists a computationally sound Karp reduction $(\mathcal{X}, \mathcal{W})$ from R_α to LS with efficiency $\mathbb{T}_{\text{Red}} \leq \lambda^{\varepsilon\alpha}$.*

Proof. Fix $\varepsilon > 0$ and let $\varepsilon' > 0$ be a constant to be chosen later. Let c be the constant guaranteed by Proposition 21. That is, for every $\alpha > 1$ there exists a Turing machine M_α of run-time $T \leq c\lambda^\alpha$, configuration size $S \leq c\lambda$ and description size $|M_\alpha| \leq c \log_2 \alpha$ solving R_α . Let α be a large enough constant such that Theorem 18 apply to M_α with constant ε' . This gives an IVC scheme with incremental completeness for M_α having efficiency $\mathbb{T}_{\text{IVC}} \leq \lambda^{\varepsilon'\alpha}$.

Applying Theorem 8 using the aforementioned IVC results in a computationally sound Karp reduction $(\mathcal{X}, \mathcal{W})$ from R_α to LS with efficiency:

$$\mathbb{T}_{\text{Red}} \leq (\mathbb{T}_{\text{IVC}} \cdot S \cdot \lambda \cdot |M_\alpha|)^\eta \leq \lambda^{\varepsilon'\alpha\eta} \cdot (c\lambda \log_2 \alpha)^{3\eta} \leq \lambda^{4\varepsilon'\alpha\eta} ,$$

where η is a constant independent of α and ε and the last inequality is true for large α . By taking $\varepsilon' < \varepsilon/4\eta$ we get $\mathbb{T}_{\text{Red}} \leq \lambda^{\varepsilon\alpha}$. This completes the proof the lemma. ◀

We are now ready to prove Theorem 22.

Proof of Theorem 22. Assume toward contradiction that under non-uniform ETH and the KPY assumption we have that $\text{PLS} \subseteq \text{FP/Poly}$. Hence, there exists a constant η and adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ of size $s(n) = n^\eta$ solving LS instances of size n .

Let $\varepsilon > 0$ be a constant to be chosen later. By bounding the blowup lemma, Lemma 23, there exists a large enough constant α such that there is a computationally sound Karp reduction $(\mathcal{X}, \mathcal{W})$ from R_α to LS having efficiency $\mathbb{T}_{\text{Red}} \leq \lambda^{\varepsilon\alpha}$. Note that the efficiency of the NP verifier for R_α is $\mathbb{T}_{R_\alpha} = O(\lambda)$ with a constant independent of α . By Proposition 13, there exists an adversary $\mathcal{A}' = \{\mathcal{A}'_\lambda\}_{\lambda \in \mathbb{N}}$, solving R_α instances of size λ of circuit size:

$$\text{size}(\mathcal{A}'_\lambda) \leq \text{poly}(s(\mathbb{T}_{\text{Red}}), \mathbb{T}_{R_\alpha}, \lambda) \leq \text{poly}(\lambda^{\varepsilon\alpha\eta}, \lambda) \underset{\alpha > 1/\varepsilon\eta}{\leq} \text{poly}(\lambda^{\varepsilon\alpha\eta}) = O(\lambda^{\varepsilon\alpha\xi}) ,$$

where ξ is a constant independent of ε and α . Let $\delta > 0$ be the constant associated with the lower bound of R_α . Recall δ is fixed and independent of α . By choosing $\varepsilon < \delta/\xi$ we have

$$\text{size}(\mathcal{A}'_\lambda) = O(\lambda^{\varepsilon\alpha\xi}) = O(\lambda^{\delta\alpha}) .$$

A contradiction to the lower bound. We conclude that $\text{PLS} \not\subseteq \text{FP/Poly}$. This completes the proof. ◀

6 Unconditional PLS Hardness in the ROM

In this section we show that relative to a random oracle, there exists exponentially average-case hard problems in PLS. We further show, relative to a random oracle, the existence of depth-robust moderately hard problems in PLS. Our result is based on the incremental proofs of sequential work (iPoSW) construction by Döttling, Lai, and Malavolta [20]. We modify the DLM construction such that an incremental completeness analogous is achieved. We proceed to reduce breaking the modified DLM scheme to LS relative to a random oracle.

6.1 Graph Related Definitions

We recall some graph related definitions used in the DLM scheme.

Notation. Throughout, we consider directed acyclic graphs (DAGs) $G = (V, E)$ where $V \subseteq \{0, 1\}^{\leq d}$ is the vertex set and we refer to the parameter d as the depth of G . A vertex $v \in V$ is a parent of $u \in V$ if $(v, u) \in E$. The set of parents of u is denoted by $\text{parents}(u)$. A vertex v is an ancestor of u if there is a directed path from v to u in G . The set of ancestors of u is denoted by $\text{ancestors}(u)$. Any vertex $v \in V \cap \{0, 1\}^n$ is called a leaf. Denote by $\text{leaves}(u)$ the set of all leaf nodes that are ancestors of u .

6.1.1 CP Graphs

We recall the definition of CP graphs (see Figure 1), a family of directed acyclic graphs defined first in the PoSW construction of Cohen and Pietrzak [17]. We denote the aforementioned graph family by $CP = \{CP_d\}_{d \in \mathbb{N}}$.




Figure 1 The CP_4 graph; red edges corresponds to E'' .

► **Definition 24** (CP Graphs, [17]). Let $B_d = (V, E')$ be the complete binary tree of depth d with edges pointing from the leaves to the root. The graph $CP_d = (V, E)$ is constructed from B_d with some added edges E'' . Where E'' contains (v, u) for any leaf $u \in \{0, 1\}^d$ and any v which is a left sibling to a node on the path from u to the root. Formally, $E = E' \cup E''$ where:

$$E'' := \{(v, u) \mid u \in \{0, 1\}^d, u = a||1||a', v = a||0, \text{ for some } a, a' \in \{0, 1\}^{\leq d}\} .$$

The following fact on non-leaves in CP graphs is used in the construction and analysis.

► **Fact 25** (Left and Right Parents). Let $CP_d = (V, E)$ and let $v \in V$ be a non-leaf node which is not the root. The node v has exactly two parents ℓ and r . We have the following:

$$\text{leaves}(v) = \text{leaves}(\ell) \cup \text{leaves}(r) .$$

6.1.2 Graph Labeling

► **Definition 26** (Partial Vertex Labeling). For a graph $G = (V, E)$, a partial vertex labeling is a function $\mathcal{L} : V \rightarrow \{0, 1\}^* \cup \{\perp\}$. The label for node $u \in V$ is denoted by $\mathcal{L}[u]$.

► **Definition 27** (Valid DAG Labeling). Let $G = (V, E)$ be a DAG with $V \subseteq \{0, 1\}^*$ and let $f : \{0, 1\}^* \rightarrow \{0, 1\}^w$ be a labeling function. The labeling of G with respect to f is defined recursively, for every $v \in V$:


$$\mathcal{L}[v] := f(v, \mathcal{L}[p_1], \dots, \mathcal{L}[p_r]) ,$$

where $(p_1, \dots, p_r) = \text{parents}(v)$, ordered in lexicographic order.

Note that the above labeling is well defined since G is acyclic and can be computed with $|V|$ queries to f in a topological order.

In the following lemma we recall the CP labeling procedure given in [17]. See Figure 2 for an illustration of the procedure.

► **Lemma 28** (CP Labeling Procedure, Lemma 3 in [17]). The labeling of CP_d with respect to $f : \{0, 1\}^* \rightarrow \{0, 1\}^w$ can be computed in a topological ordering using $O(w \cdot d)$ bits of memory.



■ **Figure 2** The CP labeling procedure for CP_2 ; black nodes represent labels the algorithm keep track of; denoted \mathcal{U}_t when computing the t label.

We require additional properties from the labeling procedure for the modified DLM scheme. For that purpose, we give a description of the procedure below and prove additional lemmas regarding it.

Labeling Procedure Description. The algorithm is recursive. The base case CP_1 is easy. For $d > 1$, let v be the root of CP_d and let ℓ, r be the left and right parents of v , respectively. The subtree rooted by ℓ is isomorphic to CP_{d-1} and hence its labels can be computed recursively. Keep track of the last label $\mathcal{L}[\ell]$ and proceed to the subtree rooted by r . Apart from edges coming from ℓ to the leaves, it is isomorphic to CP_{d-1} . One can use $\mathcal{L}[\ell]$, which is in memory, and recursively compute the right subtree. Keep track of the last label $\mathcal{L}[r]$. Compute the root label $\mathcal{L}[v] = f(v, \mathcal{L}[\ell], \mathcal{L}[r])$ and forget both $\mathcal{L}[\ell]$ and $\mathcal{L}[r]$. Output the root label $\mathcal{L}[v]$.

Incremental Labeling. We consider an incremental version of the above procedure. Fix $d \in \mathbb{N}$ and let v_1, v_2, \dots, v_T be the topological ordering in which labels are being outputted by the CP labeling procedure with depth d . Let $\mathcal{U}_t \subseteq V$ be the set of nodes the procedure keep track of after outputting v_t . For the incremental version, we are interested in algorithms that efficiently compute \mathcal{U}_t and v_t given t .

► **Lemma 29** (Computing \mathcal{U}_t Efficiently). *For CP_d , computing \mathcal{U}_t can be done in $\text{poly}(d)$ time.*

Proof. The algorithm is recursive and follows the same post-order traversal as the labeling procedure. For CP_d , let v be the root node and let ℓ, r be its left and right parents respectively. If $t = 2^{d+1} - 1$ return the root node, $\{v\}$. If $t \leq 2^d - 1$, proceed to recursively compute \mathcal{U}_t on the left subtree rooted by ℓ , which is isomorphic to CP_{d-1} . Output the resulting set. Else $t > 2^d - 1$, recursively compute $\mathcal{U}_{t-(2^d-1)}$ on the subtree rooted by r while ignoring the extra edges from ℓ . This graph is also isomorphic to CP_{d-1} . Append ℓ to the resulting set and output.

The algorithm correctness follows by induction. For efficiency, notice that the depth of the recursion is d and every level takes $\text{poly}(d)$ time. ◀

► **Lemma 30** (Computing v_t Efficiently). *For CP_d , computing v_t can be done in $\text{poly}(d)$ time.*

Proof. Notice that by construction of the CP labeling procedure, v_1, \dots, v_T is a post-order traversal on the binary tree B_d which is the base of CP_d . Finding the t -th node in a post-order traversal can be done by a simple recursion with efficiency $\text{poly}(d)$. ◀

► **Lemma 31** (One Vertex at a Time). *For every $t \in [2, T]$ we have $\mathcal{U}_t = \{v_t\} \uplus (\mathcal{U}_t \cap \mathcal{U}_{t-1})$.*

Proof. The theorem statement is equivalent to proving $\mathcal{U}_t \setminus \mathcal{U}_{t-1} = \{v_t\}$. By definition we have that $v_t \in \mathcal{U}_t$. Apart from v_t , the procedure only forget nodes, hence $\mathcal{U}_t \subseteq \{v_t\} \cup \mathcal{U}_{t-1}$. It is left to show that $v_t \notin \mathcal{U}_{t-1}$. Indeed, using the aforementioned, by induction $\mathcal{U}_{t-1} \subseteq \{v_1, \dots, v_{t-1}\}$, in particular $v_t \notin \mathcal{U}_{t-1}$. ◀

6.2 Construction

In this section, we present a modified version of the incremental proof of sequential work (iPoSW) due to [20]. The main difference, apart from syntax, is that we require stricter conditions when verifying nodes.

Parameters. Let λ be a computational security parameter. Let $T(\lambda) = T = 2^{\lambda+1} - 1$ and let $s(\lambda) = s = \Theta(\lambda^3)$ that is a power of 2. Let $\mathsf{H} : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ and $\mathsf{H}' : \{0, 1\}^* \rightarrow \{0, 1\}^{3s}$ be independently chosen random oracle ensembles depending on λ .

► **Remark 32** (Random Coins Count). The use of $3s$ random coins allows to sample a statistically close to uniform subset. See [20] for details.

Syntax. The scheme consists of two algorithms, $\text{DLM.P}^{\mathsf{H}, \mathsf{H}'}(\chi, t, \pi)$, $\text{DLM.V}^{\mathsf{H}, \mathsf{H}'}(\chi, t, \pi)$ with the following syntax.

1. $\text{DLM.P}^{\mathsf{H}, \mathsf{H}'}(\chi, t, \pi)$ is a deterministic algorithm with oracle access to H, H' . Given a statement χ , a natural number t and a proof π , it outputs a proof π' .
2. $\text{DLM.V}^{\mathsf{H}, \mathsf{H}'}(\chi, t, \pi)$ is a deterministic algorithm with oracle access to H, H' . Given a statement χ , a natural number t and a proof π , it outputs ACC or REJ.

6:20 On the Cryptographic Hardness of Local Search

Proof Structure. The proofs π are of the form $\pi = (\mathcal{L}, (u_i, S_{u_i}, \mathcal{P}_{u_i})_{i=1}^m)$ where:

- \mathcal{L} is a partial vertex labeling for CP_λ .
- u_i is a node of CP_λ .
- S_{u_i} is a set of challenge leaves, indexed by the node u_i .
- \mathcal{P}_{u_i} is a set of candidate authentication paths, indexed by the node u_i .
- $\text{Path} \in \mathcal{P}_{u_i}$ is an ordered list of the form $\text{Path} = (w_j)_j$ where each w_j is a node of CP_λ .

► **Remark 33 (Labeling Data Structure).** The partial vertex labeling \mathcal{L} is given by a dictionary data structure. Where keys are nodes v and the value associated is the corresponding label $\mathcal{L}[v]$. A node which is not in the dictionary is given the default value \perp . Let n be the amount of nodes v for which $\mathcal{L}[v] \neq \perp$. We use a dictionary implementation where size and retrieval time are $\text{poly}(n)$.

6.2.1 Algorithms Description

Conventions. A random oracle salted with the statement $\chi \in \{0, 1\}^*$ is denoted by $H_\chi := H(\chi, \cdot)$. Let v_1, v_2, \dots, v_T be the topological order given by Lemma 28 for nodes of CP_λ . Recall that \mathcal{U}_t is the list of nodes the CP labeling procedure keeps track of when computing the label for v_t . We extend the definition such that $\mathcal{U}_0 = \emptyset$.

DLM.P^{H, H'}(χ, t, π)

Input:

1. String $\chi \in \{0, 1\}^\lambda$.
2. Time $t \in [T]$.
3. Candidate proof $\pi = (\mathcal{L}, (u_i, S_{u_i}, \mathcal{P}_{u_i})_{i=1}^m)$.

Algorithm:

1. If $\{u_1, \dots, u_m\} \neq \mathcal{U}_{t-1}$ output ε .
2. Set $\mathcal{L}[v_t] = H_\chi(v_t, \mathcal{L}[p_1], \dots, \mathcal{L}[p_r])$ where $(p_1, \dots, p_r) = \text{parents}(v_t)$ in lexicographic order.
3. If v_t is a leaf, let $S_{v_t} = \{v_t\}$ and $\mathcal{P}_{v_t} = \{(v_t)\}$.
4. Otherwise, v_t has two parents, $\ell, r \in \mathcal{U}_{t-1}$.
 - a. Sample a random subset S_{v_t} of size $\min(s, |\text{leaves}(v_t)|)$ from $S_\ell \cup S_r$ using $\text{rand} \leftarrow H'_\chi(v_t, \mathcal{L}[v_t])$ as random coins.
 - b. Let $\mathcal{P}_{v_t} = \emptyset$. For every $\text{Path} \in \mathcal{P}_\ell \cup \mathcal{P}_r$ starting from a leaf in S_{v_t} , extend with v_t and append to \mathcal{P}_{v_t} .

5. Remove labels from \mathcal{L} for all nodes but $\left\{ v \left| \begin{array}{l} u \in \mathcal{U}_t \\ \text{Path} \in \mathcal{P}_u \\ w \in \text{Path} \\ v = w \text{ or } v \in \text{parents}(w) \end{array} \right. \right\}$.

6. Output $\pi' = (\mathcal{L}, (u, S_u, \mathcal{P}_u)_{u \in \mathcal{U}_t})$.

Note that in the above we use the efficient algorithms due to computing \mathcal{U}_t efficiently Lemma 29 and computing v_t efficiently Lemma 30 to compute $\mathcal{U}_{t-1}, \mathcal{U}_t$ and v_t . We also use one vertex at a time Lemma 31 in Line 5 and Line 6.

$$\text{DLM.V}^{\text{H},\text{H}'}(\chi, t, \pi)$$
Input:

1. String $\chi \in \{0, 1\}^\lambda$.
2. Time $t \in [T]$.
3. Candidate proof $\pi = (\mathcal{L}, (u_i, S_{u_i}, \mathcal{P}_{u_i})_{i=1}^m)$.

Algorithm:

1. Verify that $\{u_1, \dots, u_m\} = \mathcal{U}_{t-1}$ and distinct. Otherwise, output REJ.

2. Verify that the nodes having label in \mathcal{L} are $\left\{ v \left| \begin{array}{l} u \in \mathcal{U}_{t-1} \\ \text{Path} \in \mathcal{P}_u \\ w \in \text{Path} \\ v = w \text{ or } v \in \text{parents}(w) \end{array} \right. \right\}$.

Otherwise output REJ.

3. For every $i = 1 \dots m$:
 - a. Validate $S_{u_i} \subseteq \text{leaves}(u_i)$ and is of size $\min(s, |\text{leaves}(u_i)|)$. Otherwise output REJ.
 - b. Check that every path in \mathcal{P}_{u_i} is a valid path from a leaf in S_{u_i} to u_i . Check that every leaf in S_{u_i} has exactly one corresponding path in \mathcal{P}_{u_i} . Otherwise output REJ.
 - c. Let $(\text{Path}_1, \dots, \text{Path}_k) = \mathcal{P}_{u_i}$ ordered in lexicographic order of corresponding source leaves.
 - d. For $j = 1, \dots, k$:
 - i. If $\text{VerifyMerklePath}^{\text{H}}(\text{Path}_j, \mathcal{L}, \chi)$ rejects, output REJ.
 - ii. If $\text{VerifyRandChoice}^{\text{H}'}(\text{Path}_j, \mathcal{L}, j, \chi)$ rejects, output REJ.
4. Output ACC.

In the above description we used the subroutines VerifyAuthPath and $\text{VerifyRandChoice}^{\text{H}'}$ that are described in the following.

Merkle Authentication Path Verification $\text{VerifyAuthPath}^{\text{H}}(\text{Path}, \mathcal{L}, \chi)$. Let $\text{Path} = (w_j)_j$ be a candidate authentication path. For every node w_j with $(p_1, \dots, p_r) = \text{parents}(w_j)$ ordered in lexicographic order, verify

$$\mathcal{L}[w_j] = \text{H}_\chi(w_j, \mathcal{L}[p_1], \dots, \mathcal{L}[p_r]) .$$

If all checks pass, output ACC, otherwise output REJ.

Random Choice Verification $\text{VerifyRandChoice}^{\text{H}'}(\text{Path}, \mathcal{L}, \text{ind}, \chi)$. We give a sketch of the verification, for further details see [20]. Let $\text{Path} = w_1 \rightarrow \dots \rightarrow w_r$ be a path where $w_1 \in \text{leaves}(w_r)$. In order to verify the random choice done by H'_χ , we recreate the choice going from the top down. We are given ind , the index of Path in \mathcal{P} ordered by source leaves. This is also the index of the leaf that was picked in the last random choice.

Note that since s is a power of 2, the random choice procedure we use in the prover either choose an s subset from $2s$ set or take the whole set. Starting from w_r , use H'_χ to generate the random coins rand . Using rand pick an s subset from the set $[2s]$. Let i be the ind -th element in the resulting subset. If $i \leq s$, we know that w_{r-1} should be a left parent to w_r . Otherwise $i > s$ which tells us that w_{r-1} should be a right parent. We continue to w_{r-1} with $\text{ind} = i$ if $i \leq s$ and $\text{ind} = i - s$ otherwise. We continue with this procedure until we get to a node with less than s leaves. If all verifications of the path passed, output ACC and otherwise output REJ.

Efficiency. All the computations done above are polynomial in the input length that is of size $\text{poly}(\lambda, |\pi|)$. Notice that due to the restriction on the size of \mathcal{L} (Line 2 of DLM.V), the size of all accepted proofs is $\text{poly}(\lambda)$. Therefore, for all accepted proofs the runtime is $\text{poly}(\lambda)$ for a fixed polynomial. Hence, we may modify DLM.V to reject and DLM.P to output ε once this $\text{poly}(\lambda)$ time limit exceeds. This way, the efficiency of DLM.V and DLM.P is $\text{poly}(\lambda)$ for every input.

► **Remark 34 (Proofs are Not Unique).** Proofs in the DLM scheme are not computationally unique. To see this, consider an adversary that attempts to find a collision for accepting proofs of time $t = \Theta(s^2) = \Theta(\lambda^6)$ that is of the form $2^{k+1} - 1$. The first proof is generated honestly. For the second proof, select randomly $u \in \text{leaves}(v_t)$. Compute the proof honestly up until the label for u is computed. Alter the label for u . Continue as the honestly generated proof. With overwhelming probability, the two proofs generated are different. The second proof is accepting if there is no authentication path starting with u . Since t is of the form $2^{k+1} - 1$, there are s authentication paths starting with leaves in $\text{leaves}(v_t)$. Also, all the leaves have the same probability to be chosen. There are 2^k leaves, therefore the probability u is selected is $s/2^k = \Omega(1/\lambda^3)$. The described adversary is efficient and successful with noticeable probability.

6.2.2 Incremental Completeness

The proof for the next proposition is in the full version of this paper.

► **Proposition 35 (Incremental Completeness of Modified DLM).** *For every security parameter $\lambda \in \mathbb{N}$, time $t \in [T - 1]$, candidate proof $\pi \in \{0, 1\}^*$ and statement $\chi \in \{0, 1\}^\lambda$:*

$$\text{If } \text{DLM.V}^{\text{H}, \text{H}'}(\chi, t, \pi) = \text{ACC} \text{ then } \text{DLM.V}^{\text{H}, \text{H}'}(\chi, t + 1, \pi') = \text{ACC} ,$$

where $\pi' = \text{DLM.P}^{\text{H}, \text{H}'}(\chi, t, \pi)$.

6.2.3 Soundness

In this section, we state the soundness property of the modified DLM system. While we altered the DLM construction, we did so by enforcing stricter verification conditions. Accordingly, we inherit the soundness of the original scheme – every valid proof in the modified scheme corresponds to a valid proof in the original scheme without the use of the random oracle in the correspondence.

► **Theorem 36 (DLM Soundness, Follows From [20]).** *Let $\mathcal{A}^{\text{H}, \text{H}'} = \{\mathcal{A}_\lambda^{\text{H}, \text{H}'}\}$ be an oracle aided adversary that makes at most $q = q(\lambda)$ total queries to both H and H' . For every $\lambda, d \in \mathbb{N}$ with $d \leq \lambda$, such that \mathcal{A} makes less than 2^d sequential queries to H we have:*

$$\Pr \left[\text{DLM.V}^{\text{H}, \text{H}'}(\chi, T_d, \pi) = \text{ACC} \mid \begin{array}{l} \chi \leftarrow \{0, 1\}^\lambda \\ \pi \leftarrow \mathcal{A}^{\text{H}, \text{H}'}(\chi) \end{array} \right] \leq O(q^2)2^{-\Omega(\lambda)} ,$$

where $T_d = 2^{d+1} - 1$.

► **Corollary 37 (Sequential Hardness).** *Let $d = d(\lambda) \leq \lambda$ be a depth parameter. For every oracle aided adversary $\mathcal{A}^{\text{H}, \text{H}'} = \{\mathcal{A}_\lambda^{\text{H}, \text{H}'}\}_{\lambda \in \mathbb{N}}$ of depth 2^d and size $2^{o(\lambda)}$, we have for every $\lambda \in \mathbb{N}$:*

$$\Pr \left[\text{DLM.V}^{\text{H}, \text{H}'}(\chi, T_d, \pi) = \text{ACC} \mid \begin{array}{l} \chi \leftarrow \{0, 1\}^\lambda \\ \pi \leftarrow \mathcal{A}^{\text{H}, \text{H}'}(\chi) \end{array} \right] \leq 2^{-\Omega(\lambda)} ,$$

where $T_d = 2^{d+1} - 1$.

6.2.3.1 Single Oracle

The DLM construction is described in terms of two random oracles H, H' (which is done mostly to simplify the analysis). From hereon, it will be simpler to think of a single random oracle $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}$. The algorithms $\text{DLM.P}^{\mathcal{O}}, \text{DLM.V}^{\mathcal{O}}$ simulate oracle calls to H, H' in the above construction by setting an appropriate prefix for every output bit in H and H' . The security reduction incurs a multiplicative polynomial loss in efficiency and a multiplicative logarithmic loss in depth (both in terms of λ).

Indeed, consider an adversary \mathcal{A} attempting to break the scheme with \mathcal{O} . In order reduce this adversary to \mathcal{A}' attacking the scheme using H, H' we modify the oracle gates to \mathcal{O} in \mathcal{A} by examining the prefix and simulating the output of the associate output bit in H or H' . In this reduction, comparing the prefix, incurs a logarithmic loss in depth and polynomial loss in size of the new adversary \mathcal{A}' . We conclude this in the following corollary.

► **Corollary 38** (Single Oracle Sequential Hardness). *Let $d = d(\lambda) \leq \lambda$ be a depth parameter. For every oracle aided adversary $\mathcal{A}^{\mathcal{O}} = \{\mathcal{A}_{\lambda}^{\mathcal{O}}\}_{\lambda \in \mathbb{N}}$ of depth $o(2^d / \log \lambda)$ and size $2^{o(\lambda)}$, we have for every $\lambda \in \mathbb{N}$:*

$$\Pr \left[\text{DLM.V}^{\mathcal{O}}(\chi, T_d, \pi) = \text{ACC} \mid \begin{array}{l} \chi \leftarrow \{0, 1\}^{\lambda} \\ \pi \leftarrow \mathcal{A}^{\mathcal{O}}(\chi) \end{array} \right] \leq 2^{-\Omega(\lambda)},$$

where $T_d = 2^{d+1} - 1$.

By considering $d = \lambda$, we derive the following corollary:

► **Corollary 39** (Exponential Hardness). *For every oracle aided adversary $\mathcal{A}^{\mathcal{O}} = \{\mathcal{A}_{\lambda}^{\mathcal{O}}\}_{\lambda \in \mathbb{N}}$ of size $2^{o(\lambda)}$ and $\lambda \in \mathbb{N}$:*

$$\Pr \left[\text{DLM.V}^{\mathcal{O}}(\chi, T_{\lambda}, \pi) = \text{ACC} \mid \begin{array}{l} \chi \leftarrow \{0, 1\}^{\lambda} \\ \pi \leftarrow \mathcal{A}^{\mathcal{O}}(\chi) \end{array} \right] \leq 2^{-\Omega(\lambda)},$$

where $T_{\lambda} = 2^{\lambda+1} - 1$.

6.3 Hard instances

In this section, we consider the problem $R_d^{\mathcal{O}}$, where instances are strings χ and a witness for χ is a corresponding proof of sequential work in the modified DLM scheme for time $T_d = 2^{d+1} - 1$. We show that $R_d^{\mathcal{O}}$ lies in $\text{PLS}^{\mathcal{O}}$.

► **Remark 40** (Efficiently Computable Depth Parameter). In the following, a depth parameter $d(\lambda)$ is a polynomial-time computable function. We require this for the efficiency of the reduction.

► **Definition 41.** *Let $d = d(\lambda) \leq \lambda$ be a depth parameter. The search problem $R_d^{\mathcal{O}}$ is defined as follows:*

$$R_d^{\mathcal{O}} = \left\{ (\chi, \pi) \mid \text{DLM.V}^{\mathcal{O}}(\chi, T_d(|\chi|), \pi) = \text{ACC} \right\},$$

where $T_d(\lambda) = 2^{d(\lambda)+1} - 1$.

Note that $R_d^{\mathcal{O}}$ is in $\text{TFNP}^{\mathcal{O}}$. It is in $\text{FNP}^{\mathcal{O}}$ by the efficiency of DLM. It is total by the completeness of DLM, since for every χ , there exists a proof π that passes verification – the honestly generated proof.

6.3.1 The Reduction to Local Search

Fix a depth parameter $d = d(\lambda) \leq \lambda$. We construct a polynomial-time search reduction (f_d, g_d) from $R_d^{\mathcal{O}}$ to $\text{LS}^{\mathcal{O}}$.

Instance Translation. Let $\chi \in \{0, 1\}^\lambda$ be an instance for the problem $R_d^{\mathcal{O}}$. We construct an $\text{LS}^{\mathcal{O}}$ instance $f_d(\chi) = (\mathcal{S}^{\mathcal{O}}, \mathcal{F}^{\mathcal{O}})$ in the following. Intuitively, we embed the computation of the modified DLM PoSW in an $\text{LS}^{\mathcal{O}}$ instance such that local maximums correspond to valid proofs for time $T_d = 2^{d+1} - 1$. Concretely, every node will be of a fixed polynomial length $\ell = \ell(\lambda) \geq \lambda$. Every vertex is of the form $(t, \pi) \in \{0, 1\}^\ell$ where t is parsed as an integer $t \in [T_d]$ and π as a proof of the modified DLM scheme with padding to length ℓ if needed.

Let π_1 be a valid proof for time $t = 1$ in the modified DLM scheme. For the node (t, π) , if the proof π is verified with respect to time t , the node will be given value t and be connected to $(t + 1, \pi')$ where π' is the proof the prover generates given π . If π is rejected, it is given value -1 and is connected to $(1, \pi_1)$. A formal definition of the circuits is given below.

Successor Circuit $\mathcal{S}^{\mathcal{O}}(t, \pi)$

Hardwired: The instance χ .

Algorithm:

1. If $\text{DLM.V}^{\mathcal{O}}(\chi, t, \pi) = \text{REJ}$, output $(1, \pi_1)$.
2. If $t = T_d$, output (t, π) .
3. Compute $\pi' = \text{DLM.P}^{\mathcal{O}}(\chi, t, \pi)$ and output $(t + 1, \pi')$.

Value Circuit $\mathcal{F}^{\mathcal{O}}(t, \pi)$

Hardwired: The instance χ .

Algorithm:

1. If $\text{DLM.V}^{\mathcal{O}}(\chi, t, \pi) = \text{REJ}$, output -1 .
2. Else, output t .

Witness Translation. The function $g_d(t, \pi)$ outputs π .

Efficiency. By the efficiency of the modified DLM scheme and the efficiency of the depth parameter $d(\lambda)$, the instance translation f_d is done in $\text{poly}(\lambda)$ time. The witness translation g_d is also done in $\text{poly}(\lambda)$ time since the length of nodes ℓ is of size $\text{poly}(\lambda)$. This further implies that $\mathcal{S}^{\mathcal{O}}, \mathcal{F}^{\mathcal{O}}$ are of size $\text{poly}(\lambda)$. Since the input length is a polynomial such that $\ell(\lambda) \geq \lambda$, this implies that $\mathcal{S}^{\mathcal{O}}, \mathcal{F}^{\mathcal{O}}$ are polynomial-size circuits.

6.3.2 Security Analysis

In the following claim we rely on the incremental completeness of the modified DLM Section 6.2.2.

▷ **Claim 42.** Let $d = d(\lambda) \leq \lambda$ be a depth parameter. Fix $\lambda \in \mathbb{N}$ and $\chi \in \{0, 1\}^\lambda$. Let $w = (t, \pi)$ be a local maximum of $(\mathcal{S}^{\mathcal{O}}, \mathcal{F}^{\mathcal{O}}) = f_d(\chi)$. We have $\text{DLM.V}^{\mathcal{O}}(\chi, t, \pi) = \text{ACC}$ and $t = T_d$.

Proof. We have $\text{DLM.V}^{\mathcal{O}}(\chi, t, \pi) = \text{ACC}$ as otherwise w is given value -1 and his successor is $(1, \pi_1)$ of value 1. If $t < T_d$ then w is given value t and is connected to $(t + 1, \pi')$ for $\pi' = \text{DLM.P}^{\mathcal{O}}(\chi, t, \pi)$. By incremental completeness Proposition 35, we have that $\text{DLM.P}(\chi, t + 1, \pi') = \text{ACC}$. Therefore, the value of $(t + 1, \pi')$ is $t + 1$, and w is not a local maximum. We are left with $t = T_d$. This concludes the claim. ◁

► **Theorem 43** (Hard Problems in $\text{PLS}^\mathcal{O}$). *For every depth parameter $d = d(\lambda) \leq \lambda$, the search problem $R_d^\mathcal{O}$ lies in $\text{PLS}^\mathcal{O}$.*

Proof. The class $\text{PLS}^\mathcal{O}$ consists of all $\text{TFNP}^\mathcal{O}$ search problems that are polynomial-time reducible to $\text{LS}^\mathcal{O}$. The search problem $R_d^\mathcal{O} \in \text{TFNP}^\mathcal{O}$ and by Claim 42, the tuple (f_d, g_d) is a valid polynomial-time reduction to $\text{PLS}^\mathcal{O}$. Indeed, for every string χ , we have that all witnesses w' for the instance $(\mathcal{S}^\mathcal{O}, \mathcal{F}^\mathcal{O}) = f_d(\chi)$ are mapped under g_d to $w = \pi$ such that $\text{DLM.V}^\mathcal{O}(\chi, T_d, \pi) = \text{ACC}$ – a valid witness to $R_d^\mathcal{O}$. Hence $R_d^\mathcal{O} \in \text{PLS}^\mathcal{O}$. A proof is given in the full version. ◀

► **Corollary 44** (Exponential Hardness in $\text{PLS}^\mathcal{O}$). *There exists search problems in $\text{PLS}^\mathcal{O}$ that are average-case exponentially hard.*

Proof. By Theorem 43, for every depth parameter $d = d(\lambda) \leq \lambda$, the search problem $R_d^\mathcal{O}$ is in $\text{PLS}^\mathcal{O}$. By Corollary 39, for $d = \lambda$ the search problem $R_\lambda^\mathcal{O}$ is exponentially average-case hard. ◀

This implies that $\text{LS}^\mathcal{O}$ is sub-exponentially hard due to the polynomial blowup the reduction incurs. We formulate it in the following corollary. In the following, n stands for the size of the $\text{LS}^\mathcal{O}$ instance $(\mathcal{S}^\mathcal{O}, \mathcal{F}^\mathcal{O})$ and λ stands for the security parameter that is given as input to the sampler. A proof is given in the full version of this paper.

► **Corollary 45.** *There exists an efficient sampler HARD of $\text{LS}^\mathcal{O}$ instances and a constant $\delta > 0$ such that for every oracle aided adversary $\mathcal{A}^\mathcal{O} = \{\mathcal{A}_n^\mathcal{O}\}_{n \in \mathbb{N}}$ of size at most 2^{n^δ} and for every $\lambda \in \mathbb{N}$:*

$$\Pr \left[\mathcal{F}^\mathcal{O}(\mathcal{S}^\mathcal{O}(w)) \leq \mathcal{F}^\mathcal{O}(w) \mid \begin{array}{l} (\mathcal{S}^\mathcal{O}, \mathcal{F}^\mathcal{O}) \leftarrow \text{HARD}(1^\lambda) \\ w \leftarrow \mathcal{A}(\mathcal{S}^\mathcal{O}, \mathcal{F}^\mathcal{O}) \end{array} \right] \leq 2^{-\Omega(\lambda)} ,$$

where $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}$ is a random oracle.

6.4 Depth Robust Instances

In this section, we deduce the existence of depth-robust, moderately hard, search problems in $\text{PLS}^\mathcal{O}$. That is, search problems with instances that can be solved in polynomial time yet also require high sequential time.

► **Proposition 46** (Depth Robust Problems in $\text{PLS}^\mathcal{O}$). *Fix any constant $\varepsilon > 0$. For any large enough constant α , consider $R_d^\mathcal{O}$ for $d(\lambda) = \alpha \log_2 \lambda$ a depth parameter. The following properties hold for $\lambda > 1$:*

- **Moderately Hard:** *The search problem $R_d^\mathcal{O}$ can be solved in time $\lambda^{(1+\varepsilon)\alpha}$.*
- **Depth Robust:** *If \mathcal{A} is an adversary of depth $\lambda^{(1-\varepsilon)\alpha}$ and size $2^{o(\lambda)}$ then:*

$$\Pr \left[(\chi, \pi) \in R_d^\mathcal{O} \mid \begin{array}{l} \chi \leftarrow \{0, 1\}^\lambda \\ \pi \leftarrow \mathcal{A}^\mathcal{O}(\chi) \end{array} \right] \leq 2^{-\Omega(\lambda)} .$$

Proof. Fix $\varepsilon > 0$ throughout the proof. We start by proving that $R_d^\mathcal{O}$ is moderately hard.

▷ **Claim 47 (Moderately Hard).** *The search problem $R_d^\mathcal{O}$, where $d(\lambda) = \alpha \log_2 \lambda$, can be solved in time $\lambda^{(1+\varepsilon)\alpha}$ when α is a large enough constant.*

6:26 On the Cryptographic Hardness of Local Search

Proof. Fix $\lambda > 1$ and let $\chi \in \{0, 1\}^\lambda$ be an instance of $R_d^{\mathcal{O}}$. We have $T_d = 2^{\alpha \log_2 \lambda + 1} - 1 \leq 2\lambda^\alpha \leq \lambda^{\alpha+1}$. By efficiency of DLM, the run-time of $\text{DLM.P}^{\mathcal{O}}$ is at most λ^c for some fixed constant c , independent of ε and α . The proof π_1 for time $t = 1$ can also be generated in $\text{poly}(\lambda)$ time. We assume w.l.o.g that it can be generated in time λ^c . By completeness of DLM, generating π_1 and applying the prover $T_d - 1$ times on it, results in a valid proof for time $t = T_d$. This is a valid witness for the instance χ in $R_d^{\mathcal{O}}$. The above algorithm takes at most

$$T_d \cdot \lambda^c \leq \lambda^{\alpha+1} \cdot \lambda^c \leq \lambda^{(1+\varepsilon)\alpha} ,$$

where the last inequality holds for large enough α , concretely $\alpha > (c + 1)/\varepsilon$. \triangleleft

We proceed to show that $R_d^{\mathcal{O}}$ is depth robust.

\triangleright **Claim 48 (Depth Robust).** Let α be a large enough constant and let $d = \alpha \log_2 \lambda$. For every adversary \mathcal{A} of depth $\lambda^{(1-\varepsilon)\alpha}$ and size $2^{o(\lambda)}$ we have for every $\lambda \in \mathbb{N}$:

$$\Pr \left[(\chi, \pi) \in R_d^{\mathcal{O}} \mid \begin{array}{l} \chi \leftarrow \{0, 1\}^\lambda \\ \pi \leftarrow \mathcal{A}^{\mathcal{O}}(\chi) \end{array} \right] \leq 2^{-\Omega(\lambda)} .$$

Proof. This is a direct application of single oracle sequential hardness Corollary 38. We have $T_d = 2^{\alpha \log_2 \lambda + 1} - 1 \geq \lambda^\alpha$. Hence $\lambda^{(1-\varepsilon)\alpha} = o(T_d / \log \alpha)$ and the claim follows by Corollary 38. \triangleleft

The above two claims conclude the proof of the proposition. \blacktriangleleft

References

- 1 Tim Abbot, Daniel Kane, and Paul Valiant. On algorithms for nash equilibria. *Unpublished*, 2004.
- 2 Omer Angel, Sébastien Bubeck, Yuval Peres, and Fan Wei. Local max-cut in smoothed polynomial time. *CoRR*, abs/1610.04807, 2016. [arXiv:1610.04807](https://arxiv.org/abs/1610.04807).
- 3 Saikrishna Badrinarayanan, Yael Tauman Kalai, Dakshita Khurana, Amit Sahai, and Daniel Wichs. Succinct delegation for low-space non-deterministic computation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 709–721, 2018. [doi:10.1145/3188745.3188924](https://doi.org/10.1145/3188745.3188924).
- 4 Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (Im)Possibility of Obfuscating Programs. *J. ACM*, 59(2):6:1–6:48, May 2012. [doi:10.1145/2160158.2160159](https://doi.org/10.1145/2160158.2160159).
- 5 Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive Composition and Bootstrapping for SNARKS and Proof-carrying Data. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, pages 111–120, New York, NY, USA, 2013. ACM. [doi:10.1145/2488608.2488623](https://doi.org/10.1145/2488608.2488623).
- 6 Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the Existence of Extractable One-Way Functions. *SIAM J. Comput.*, 45(5):1910–1952, 2016. [doi:10.1137/140975048](https://doi.org/10.1137/140975048).
- 7 Nir Bitansky, Omer Paneth, and Alon Rosen. On the Cryptographic Hardness of Finding a Nash Equilibrium. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1480–1498, October 2015. [doi:10.1109/FOCS.2015.94](https://doi.org/10.1109/FOCS.2015.94).
- 8 Dan Boneh, Xavier Boyen, and Eu-Jin Goh. Hierarchical Identity Based Encryption with Constant Size Ciphertext. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, pages 440–456, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- 9 Dan Boneh, Craig Gentry, and Brent Waters. Collusion Resistant Broadcast Encryption with Short Ciphertexts and Private Keys. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, pages 258–275, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

- 10 Shant Boodaghians, Rucha Kulkarni, and Ruta Mehta. Nash Equilibrium in Smoothed Polynomial Time for Network Coordination Games, September 2018. [arXiv:1809.02280](https://arxiv.org/abs/1809.02280).
- 11 Zvika Brakerski, Justin Holmgren, and Yael Tauman Kalai. Non-interactive delegation and batch NP verification from standard computational assumptions. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 474–482, 2017. doi:10.1145/3055399.3055497.
- 12 Buresh-Oppenheim. On the TFNP complexity of factoring. *Unpublished*, 2006.
- 13 Ran Canetti, Yilei Chen, and Leonid Reyzin. On the Correlation Intractability of Obfuscated Pseudorandom Functions. In *Proceedings, Part I, of the 13th International Conference on Theory of Cryptography - Volume 9562, TCC 2016-A*, pages 389–415, Berlin, Heidelberg, 2016. Springer-Verlag. doi:10.1007/978-3-662-49096-9_17.
- 14 Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the Complexity of Computing Two-player Nash Equilibria. *J. ACM*, 56(3):14:1–14:57, May 2009. doi:10.1145/1516512.1516516.
- 15 Arka Rai Choudhuri, Pavel Hubacek, Chethan Kamath, Krzysztof Pietrzak, Alon Rosen, and Guy N. Rothblum. PPAD-Hardness via Iterated Squaring Modulo a Composite. *Cryptology ePrint Archive*, Report 2019/667, 2019. URL: <https://eprint.iacr.org/2019/667>.
- 16 Arka Rai Choudhuri, Pavel Hubáček, Chethan Kamath, Krzysztof Pietrzak, Alon Rosen, and Guy N. Rothblum. Finding a Nash Equilibrium is No Easier Than Breaking Fiat-Shamir. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, pages 1103–1114, New York, NY, USA, 2019. ACM. doi:10.1145/3313276.3316400.
- 17 Bram Cohen and Krzysztof Pietrzak. Simple Proofs of Sequential Work. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018*, pages 451–467, Cham, 2018. Springer International Publishing.
- 18 Constantinos Daskalakis, Paul Goldberg, and Christos H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *SIAM J. Comput.*, 39:195–259, February 2009. doi:10.1137/070699652.
- 19 Constantinos Daskalakis and Christos Papadimitriou. Continuous Local Search. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '11*, pages 790–804, Philadelphia, PA, USA, 2011. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=2133036.2133098>.
- 20 Nico Döttling, Russell W. F. Lai, and Giulio Malavolta. Incremental Proofs of Sequential Work. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019*, pages 292–323, Cham, 2019. Springer International Publishing.
- 21 Naomi Ephraim, Cody Freitag, Ilan Komargodski, and Rafael Pass. Continuous Verifiable Delay Functions. *Cryptology ePrint Archive*, Report 2019/619, 2019. URL: <https://eprint.iacr.org/2019/619>.
- 22 Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The Complexity of Pure Nash Equilibria. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing, STOC '04*, pages 604–612, New York, NY, USA, 2004. ACM. doi:10.1145/1007352.1007445.
- 23 Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. Revisiting the Cryptographic Hardness of Finding a Nash Equilibrium. In *Advances in Cryptology – CRYPTO 2016*, pages 579–604, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- 24 Craig Gentry and Daniel Wichs. Separating Succinct Non-interactive Arguments from All Falsifiable Assumptions. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing, STOC '11*, pages 99–108, New York, NY, USA, 2011. ACM. doi:10.1145/1993636.1993651.
- 25 Paul W. Goldberg and Christos H. Papadimitriou. Towards a unified complexity theory of total functions. *J. Comput. Syst. Sci.*, 94:167–192, 2018. doi:10.1016/j.jcss.2017.12.003.
- 26 Oded Goldreich and Johan Håstad. On the Complexity of Interactive Proofs with Bounded Communication. *Inf. Process. Lett.*, 67(4):205–214, 1998. doi:10.1016/S0020-0190(98)00116-1.

- 27 Oded Goldreich, Salil P. Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1-2):1–53, 2002. doi:10.1007/s00037-002-0169-0.
- 28 Pavel Hubáček, Moni Naor, and Eylon Yogev. The Journey from NP to TFNP Hardness. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 60:1–60:21, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ITCS.2017.60.
- 29 Pavel Hubáček and Eylon Yogev. *Hardness of Continuous Local Search: Query Complexity and Cryptographic Lower Bounds*, pages 1352–1371. ACM, 2016. doi:10.1137/1.9781611974782.88.
- 30 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which Problems Have Strongly Exponential Complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.
- 31 Hisao Ishibuchi and Tadahiko Murata. A multi-objective genetic local search algorithm and its application to flowshop scheduling. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 28(3):392–403, August 1998. doi:10.1109/5326.704576.
- 32 Emil Jerábek. Integer factoring and modular square roots. *CoRR*, abs/1207.5220, 2012. arXiv:1207.5220.
- 33 David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988. doi:10.1016/0022-0000(88)90046-3.
- 34 Yael Tauman Kalai, Omer Paneth, and Lisa Yang. How to Delegate Computations Publicly. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, pages 1115–1124, New York, NY, USA, 2019. ACM. doi:10.1145/3313276.3316411.
- 35 Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. How to Delegate Computations: The Power of No-Signaling Proofs. In *In Proceedings of the 46th annual ACM symposium on Theory of computing (STOC)*, pages 485–494. ACM, January 2014. URL: <https://www.microsoft.com/en-us/research/publication/delegate-computations-power-no-signaling-proofs/>.
- 36 Ilan Komargodski, Moni Naor, and Eylon Yogev. White-Box vs. Black-Box Complexity of Search Problems: Ramsey and Graph Property Testing. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 622–632, October 2017.
- 37 Ilan Komargodski, Moni Naor, and Eylon Yogev. White-Box vs. Black-Box Complexity of Search Problems: Ramsey and Graph Property Testing. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 622–632, 2017. doi:10.1109/FOCS.2017.63.
- 38 S. Lin and Brain W. Kernighan. An Effective Heuristic Algorithm for the Traveling-Salesman Problem. *Oper. Res.*, 21(2):498–516, April 1973. doi:10.1287/opre.21.2.498.
- 39 Carsten Lund, Lance Fortnow, H Karloff, and Noam Nisan. Algebraic Methods for Interactive Proof Systems, November 1990. doi:10.1109/FSCS.1990.89518.
- 40 Mohammad Mahmoody, Tal Moran, and Salil Vadhan. Publicly Verifiable Proofs of Sequential Work. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science, ITCS '13*, pages 373–388, New York, NY, USA, 2013. ACM. doi:10.1145/2422436.2422479.
- 41 Nimrod Megiddo and Christos H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theoretical Computer Science*, 81(2):317–324, 1991. doi:10.1016/0304-3975(91)90200-L.
- 42 Ralph Merkle. *Secrecy, Authentication and Public Key Systems*. PhD thesis, Stanford University, Department of Electrical Engineering, June 1979.
- 43 John C. Nash. The (Dantzig) simplex method for linear programming. *Computing in Science Engineering*, 2(1):29–31, January 2000. doi:10.1109/5992.814654.

- 44 Omer Paneth and Guy N. Rothblum. On Zero-Testable Homomorphic Encryption and Publicly Verifiable Non-interactive Arguments. In Yael Kalai and Leonid Reyzin, editors, *Theory of Cryptography*, pages 283–315, Cham, 2017. Springer International Publishing.
- 45 Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498–532, 1994. doi: 10.1016/S0022-0000(05)80063-7.
- 46 Krzysztof Pietrzak. Simple Verifiable Delay Functions. Cryptology ePrint Archive, Report 2018/627, 2018. URL: <https://eprint.iacr.org/2018/627>.
- 47 Ronald L. Rivest, Adi Shamir, and David A. Wagner. Time-lock puzzles and timed-release crypto. Technical Report MIT/LCS/TR-684, MIT, February 2000.
- 48 Alon Rosen, Gil Segev, and Ido Shahaf. Can PPAD Hardness be Based on Standard Cryptographic Assumptions? In Yael Kalai and Leonid Reyzin, editors, *Theory of Cryptography*, pages 747–776, Cham, 2017. Springer International Publishing.
- 49 Alejandro Schaffer and Mihalis Yannakakis. Simple Local Search Problems That are Hard to Solve. *SIAM J. Comput.*, 20:56–87, February 1991. doi:10.1137/0220004.
- 50 Katerina Sotiraki, Manolis Zampetakis, and Giorgos Zirdelis. PPP-Completeness with Connections to Cryptography. *CoRR*, abs/1808.06407, 2018. arXiv:1808.06407.
- 51 Paul Valiant. Incrementally Verifiable Computation or Proofs of Knowledge Imply Time/Space Efficiency. In Ran Canetti, editor, *Theory of Cryptography*, pages 1–18, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.