

# Cryptanalysis of Lattice-Based Sequentiality Assumptions and Proofs of Sequential Work

Chris Peikert\*      Yi Tang†

February 13, 2024

## Abstract

This work *completely breaks* the sequentiality assumption (and broad generalizations thereof) underlying the candidate lattice-based proof of sequential work (PoSW) recently proposed by Lai and Malavolta at CRYPTO 2023. In addition, it breaks an essentially identical variant of the PoSW, which differs from the original in only an arbitrary choice that is immaterial to the design and security proof (under the falsified assumption). This suggests that whatever security the original PoSW may have is fragile, and further motivates the search for a construction based on a sound lattice-based assumption.

Specifically, for sequentiality parameter  $T$  and SIS parameters  $n, q, m = n \log q$ , the attack on the sequentiality assumption finds a solution of quasipolynomial norm  $m^{\lceil \log T \rceil}$  (or norm  $O(\sqrt{m})^{\lceil \log T \rceil}$  with high probability) in only *logarithmic*  $\tilde{O}_{n,q}(\log T)$  depth; this strongly falsifies the assumption that finding such a solution requires depth *linear* in  $T$ . (The  $\tilde{O}$  notation hides polylogarithmic factors in the variables appearing in its subscript.) Alternatively, the attack finds a solution of polynomial norm  $m^{1/\varepsilon}$  in depth  $\tilde{O}_{n,q}(T^\varepsilon)$ , for any constant  $\varepsilon > 0$ . Similarly, the attack on the (slightly modified) PoSW constructs a valid proof in *polylogarithmic*  $\tilde{O}_{n,q}(\log^2 T)$  depth, thus strongly falsifying the expectation that doing so requires linear sequential work.

## 1 Introduction

The notion of *timed* (or *timed-release*) cryptography was formally introduced and realized in 1996 by Rivest, Shamir, and Wagner [RSW96], following initial concepts due to May [May93] and related ideas of Cai, Lipton, Sedgewick, and Yao [CLSY93]. The general thrust of this area is to devise “puzzles” that require (roughly) a prespecified amount of time to solve—even for solvers that have a large amount of computing power. More precisely, solving the puzzle should be *inherently sequential* (i.e., high computation depth) in nature, so that using *many processors in parallel* does not lead to any major speedup in finding a solution, versus using just one processor. (Of course, using a *faster* sequential processor will unavoidably result in a speedup, but the range of available processor speeds is substantially narrower than the ability to purchase huge numbers of parallel processors.)

Several variations on this theme have emerged in the literature. The focus of this work is on one of the most basic timed primitives, called a *proof of sequential work* (PoSW) [MMV13, CP18, AKK<sup>+</sup>19]. Here the goal is simply to quickly convince a skeptical verifier that a sequential computation of some significant desired length has been performed. In other words, the computation should be inherently sequential and

---

\*University of Michigan, cpeikert@umich.edu.

†University of Michigan, yit@umich.edu.

tunable, but the result of such a computation should be publicly and very quickly verifiable. (No secret message is encrypted or decrypted, however.) Applications of PoSW include anti-spam and denial-of-service measures [DN92], and reducing the wasteful energy consumption of proof-of-work blockchains like Bitcoin (because using large-scale parallel computation is of little marginal benefit).

## 1.1 (In)Security in a Quantum World

Unfortunately, most prior timed-cryptography constructions will become *completely insecure* in the presence of general-purpose *quantum computers*. This is because the security of these constructions relies on the conjectured hardness of problems that are in fact *easy* for quantum computers. As some of the most notable examples, the constructions of [RSW96, Pie19, Wes19] rely on the presumed hardness of factoring integers, but the breakthrough work of Shor [Sho94] gave an efficient quantum algorithm for this problem. Therefore, quantum computers would be able to completely circumvent the (conjectured) classical sequentiality of these puzzles, and solve them in relatively low quantum depth.

In this light, it is important to find constructions of timed cryptography that are quantum-secure, or “post quantum”—i.e., that can be run on today’s computers, but are believed to remain secure against attacks by future quantum computers. In other parts of cryptography, the most promising post-quantum systems are based on *lattice* problems, particularly *short integer solution* (SIS) [Ajt96] and *learning with errors* (LWE) [Reg05], and their variants. Over the past two decades, countless efficient and powerful cryptographic concepts have been realized from these lattice foundations.

Yet despite so much progress in general, post-quantum *timed* cryptography is still in its infancy, with very few and limited constructions. In particular, for proofs of sequential work we know of only two types of post-quantum constructions: ones in the idealized random-oracle model (or under a closely related sequential-hashing assumption) [MMV13, CP18], and a very interesting “algebraic” proposal by Lai and Malavolta [LM23] from CRYPTO 2023 that is related to lattices, and does not require random oracles.

As a foundation for their PoSW candidate, Lai and Malavolta introduced a new SIS-related problem and sequentiality assumption, for which they gave some credible evidence. Essentially, the problem is to evaluate a long chain of iterated SIS hash functions, and the assumption is that doing so requires computation depth roughly proportional to the length of the chain. Their elegant PoSW protocol works analogously to the factoring-based construction of [Pie19] by exploiting the homomorphic properties of the SIS hash function, and they proved its security under the new sequentiality assumption.

## 1.2 Contributions

Our first main contribution is to *strongly falsify* the lattice-based sequentiality assumption proposed in [LM23], and broad generalizations thereof. In a bit more detail, the conjecture is that finding a “somewhat short” solution to a certain regular linear system (corresponding to iterated hash evaluations) requires *nearly linear depth* in the sequentiality parameter  $T$ ; see Section 2.2 for details. For typical parameters, we instead solve this problem in depth only *polylogarithmic* in  $T$ . Also, other parameterizations of our attack find asymptotically *much shorter* solutions in *small polynomial depth*  $T^\epsilon$ , for any constant  $\epsilon > 0$ . So, tightening the quantitative definition of “short” offers limited hope for salvaging the assumption, unless the norm bound is made quite small; see Section 1.3 below for a discussion.

Interestingly, while our attack breaks the *assumption* underlying the security proof for the PoSW protocol of [LM23], it does not break the *PoSW itself* as originally defined—and so far we have not found an attack that does so. However, we do manage to break *two slight variants* of the PoSW from [LM23], by employing our core techniques in more sophisticated ways (see Figure 4 for an illustration). These variants are supported

by essentially identical security proofs (under the same kind of falsified assumption) as the original PoSW. Indeed, one of the variants differs from the original in *only an arbitrary choice* in the core “folding” operation, so we consider it to be effectively identical to the original.

Our specific contributions are organized as follows.

- In [Section 3](#) we give a suite of very general and modular tools for efficiently computing, combining, and using lattice “trapdoors” in low computation depth. As we showcase in the rest of the paper, these tools can be combined in various ways to yield attacks on lattice-based timed cryptography proposals, and may also be of independent interest for other applications.
- In [Section 4](#) we use our tools to give a low-depth recursive attack that strongly falsifies the sequentiality assumption from [\[LM23\]](#).
- Finally, in [Section 5](#) we extend the attack to break two slight variants of the PoSW protocol from [\[LM23\]](#).

We refer to each individual section for further background and context for the results and techniques given therein.

### 1.3 Discussion and Future Work

As noted above, while we have not managed to break the PoSW from [\[LM23\]](#) exactly as it is written, we did break a variant that differs only in one minor and arbitrary choice, which has no effect on its underlying assumption or security proof. This state of affairs suggests that whatever security the original PoSW may have is quite fragile, and not due to any intentional design choice or technique in the security proof. Additional ideas might lead to a successful attack against the original PoSW; alternatively, it might actually be secure, perhaps with a different security proof under some other plausible assumption. We leave these topics for future work.

The effectiveness of our attacks hinges on the following key feature of the assumption and PoSW protocols following [\[LM23\]](#): solutions are allowed to be “somewhat short,” even though the “honest” computation generates a “very short” solution. More specifically, the norm (in  $\ell_\infty$ , say) of a solution is allowed to be *quasipolynomial* in the sequentiality parameter  $T$ , whereas the honestly computed solution has *constant*  $\ell_\infty$  norm. The reason for this gap is that in the protocol, the honest solution is repeatedly “folded” into one having smaller dimension but larger norm, so the verifier needs to use more permissive norm checks. (Additionally, the knowledge extractor in the security proof incurs another quasipolynomial blowup in the norm of its extracted solution, relative to the norm bounds used by the verifier.)

Our attacks crucially exploit this gap by computing a “somewhat short” solution of quasipolynomial norm in only polylogarithmic depth  $\text{polylog}(T)$ , or even a “short” one of some polynomial norm in small polynomial depth  $T^\varepsilon$ , for any constant  $\varepsilon > 0$ . However, we do not see how to compute a “very short” solution having constant  $\ell_\infty$  norm (like the honestly computed one) in depth sublinear in  $T$ . So, a much weaker version of the sequentiality assumption from [\[LM23\]](#) corresponding to these parameters, still seems plausible. Constructing a proof system that quickly proves knowledge of such a short solution is an interesting and worthwhile open problem.

## 2 Preliminaries

### 2.1 Vector and Matrix Norms

For a real vector  $\mathbf{x}$  and  $p \geq 1$ , define its  $\ell_p$  norm as  $\|\mathbf{x}\|_p := (\sum_i |x_i|^p)^{1/p}$ , and its  $\ell_\infty$  norm as  $\|\mathbf{x}\|_\infty := \max_i |x_i|$ . Observe that the Euclidean norm is simply the  $\ell_2$  norm. For any  $n$ -dimensional vector  $\mathbf{x}$  and any

$1 \leq p \leq r \leq \infty$ , a standard bound is

$$\|\mathbf{x}\|_r \leq \|\mathbf{x}\|_p \leq \|\mathbf{x}\|_r \cdot n^{1/p-1/r},$$

where we adopt the convention that  $a^{1/\infty} = 1$  for any  $a > 0$ . We extend any  $\ell_p$  norm  $\|\cdot\|_p$  on vectors to matrices  $\mathbf{X}$  by taking the maximum over its columns  $\mathbf{x}_j$ , i.e.,  $\|\mathbf{X}\|_p := \max_j \|\mathbf{x}_j\|_p$ .

For matrices it will be convenient to use the *operator* norm  $\|\mathbf{X}\|_{r \leftarrow p} := \max_{\mathbf{y} \neq \mathbf{0}} \|\mathbf{X}\mathbf{y}\|_r / \|\mathbf{y}\|_p$  for  $p, r \in [1, \infty]$ . In words, the operator norm bounds the factor by which left-multiplication by  $\mathbf{X}$  can expand norms, going from  $\ell_p$  to  $\ell_r$ :

$$\|\mathbf{X} \cdot \mathbf{Y}\|_r \leq \|\mathbf{X}\|_{r \leftarrow p} \cdot \|\mathbf{Y}\|_p. \quad (2.1)$$

The operator norm is sub-additive and sub-multiplicative, i.e.,  $\|\mathbf{X} + \mathbf{Y}\|_{r \leftarrow p} \leq \|\mathbf{X}\|_{r \leftarrow p} + \|\mathbf{Y}\|_{r \leftarrow p}$  and  $\|\mathbf{X} \cdot \mathbf{Y}\|_{r \leftarrow p} \leq \|\mathbf{X}\|_{r \leftarrow p} \cdot \|\mathbf{Y}\|_{r \leftarrow p}$ , and it is immediate that  $\|\mathbf{X}\mathbf{Y}\|_{r \leftarrow p} \leq \|\mathbf{X}\|_{r \leftarrow q} \cdot \|\mathbf{Y}\|_{q \leftarrow p}$ . For any  $p, r \in [1, \infty]$ , observe that the identity matrix  $\mathbf{I}_n$  satisfies  $\|\mathbf{I}_n\|_{r \leftarrow p} = \max(n^{1/r-1/p}, 1)$ , and in particular  $\|\mathbf{I}_n\|_{p \leftarrow p} = 1$ ; and for any  $\mathbf{X} \in \mathbb{R}^{n \times m}$ ,

$$\|\mathbf{X}\|_{r \leftarrow p} \leq n^{1/r} \cdot \|\mathbf{X}\|_{\infty \leftarrow p} \leq n^{1/r} \cdot \|\mathbf{X}\|_{\infty} \cdot m^{1-1/p}. \quad (2.2)$$

**Random matrices.** For tighter bounds, it is convenient in some cases to rely on standard results from random matrix theory; see, e.g., [Ver12]. For example, if  $\mathbf{X} \in \mathbb{R}^{m \times m}$  is a random matrix with independent columns drawn from *subgaussian* distributions (which may vary from column to column), then it will satisfy  $\|\mathbf{X}\|_{2 \leftarrow 2} = O(\sqrt{m})$  except with probability  $2^{-\Omega(m)}$ . (We refer to [Ver12] for the definition of subgaussian, which we do not need in this work.) For simplicity, we simply say “high probability” to represent  $1 - 2^{-\Omega(m)}$ . Observe that by the union bound, any  $T = 2^{o(m)}$  events that *individually* occur with high probability will also *all* occur with high probability; we often implicitly use this fact in our high-probability statements.

**Block-wise norms.** For some of our purposes it will be important to have finer-grained bounds on the (operator) norms of the *row blocks* of a vector or matrix. Let  $\|\cdot\|$  be a norm on vectors or matrices, such as the ones defined above. When  $\mathbf{X}$  is seen as being made up of row blocks  $\mathbf{X}_i$  (whose definition will always be clear from context), we take  $\|\mathbf{X}\|$  to be the column *vector* of norms  $\|\mathbf{X}_i\|$ , i.e.,  $\|\mathbf{X}\| := (\|\mathbf{X}_i\|)_i$ . For vectors of norms having the same dimension, we use the partial ordering given by  $\mathbf{a} \leq \mathbf{b}$  if  $a_i \leq b_i$  for all  $i$ . So, for any matrices  $\mathbf{X}, \mathbf{Y}$  where now  $\mathbf{X}$  is seen as being made up of row blocks (but  $\mathbf{Y}$  is not), Equation (2.1) still holds, but now both sides are vectors (and  $\|\mathbf{Y}\|_p$  is a scalar).

## 2.2 Sequentiality Assumption

Here we recall the lattice-based candidate sequentiality assumption recently proposed by Lai and Malavolta [LM23] (with some slight differences in the notation). This assumption was used as the foundation for a candidate *proof of sequential work* (PoSW); see Section 5 for further details.

Let  $q$  be a positive integer modulus and  $\mathbf{g} \in \mathbb{Z}_q^\ell$  be a suitable “gadget” vector; for concreteness, we use the standard powers-of-two gadget  $\mathbf{g} = (1, 2, 4, \dots, 2^{\ell-1})^t$  for  $\ell = \lceil \log_2 q \rceil$ , but all of our results easily adapt to other choices of gadgets (see [MP12] for further details). Let  $\mathbf{g}^{-1}: \mathbb{Z}_q \rightarrow \mathbb{Z}^\ell$  be the corresponding “(bit) decomposition” function:  $\mathbf{g}^{-1}(u)$  is binary and hence “short,” and  $\langle \mathbf{g}, \mathbf{g}^{-1}(u) \rangle = \mathbf{g}^t \cdot \mathbf{g}^{-1}(u) = u$  for any  $u \in \mathbb{Z}_q$ . Finally, extend the gadget and its decomposition operation to work on matrices (including vectors) as follows: for any positive integer  $n$ , let  $\mathbf{G}_n := \mathbf{I}_n \otimes \mathbf{g}^t \in \mathbb{Z}_q^{n \times n\ell}$  denote the block-wise application of  $\mathbf{g}^t$  to each  $\ell$ -dimensional column block, and let  $\mathbf{G}_n^{-1}(\cdot)$  denote the entry-wise application of  $\mathbf{g}^{-1}$  on any matrix  $\mathbf{U}$  having  $n$  rows, so that  $\mathbf{G}_n \cdot \mathbf{G}_n^{-1}(\mathbf{U}) = \mathbf{U}$ .

For dimensions  $n$  and  $m = n\ell$ , matrix  $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times m}$ , vector  $\mathbf{u}_0 \in \mathbb{Z}_q^n$ , and sequentiality parameter  $T$ , Lai and Malavolta [LM23] consider the following linear system, where  $\mathbf{G} = \mathbf{G}_n \in \mathbb{Z}_q^{n \times m}$ :<sup>1</sup>

$$\underbrace{\begin{pmatrix} \mathbf{G} & & & & & \\ \bar{\mathbf{A}} & \mathbf{G} & & & & \\ & \bar{\mathbf{A}} & \ddots & & & \\ & & \ddots & \ddots & & \\ & & & \mathbf{G} & & \\ & & & \bar{\mathbf{A}} & \mathbf{G} & \end{pmatrix}}_{\mathbf{A}_T \in \mathbb{Z}_q^{Tn \times Tm}} \cdot \underbrace{\begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_T \end{pmatrix}}_{\mathbf{x} \in \mathbb{Z}^{Tm}} = \begin{pmatrix} \mathbf{u}_0 \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix} \in \mathbb{Z}_q^{Tn}. \quad (2.3)$$

They conjectured that for sufficiently large  $q$ , given *uniformly random*  $(\bar{\mathbf{A}}, \mathbf{u}_0)$ , computing a “somewhat short” solution  $\mathbf{x} \in \mathbb{Z}^{Tm}$  to Equation (2.3)—specifically, one having  $\ell_\infty$  norm  $\|\mathbf{x}\|_\infty \leq O(n)^{2 \log T}$ —requires  $\Omega(T)$  sequential time. (Naturally, one can consider other norms as well, like the Euclidean  $\ell_2$  norm.) Notice that a “very short” *binary* solution, which has  $\ell_\infty$  norm  $\|\mathbf{x}\|_\infty \leq 1$ , can be computed in depth proportional to  $T$ , as  $\mathbf{x}_1 = \mathbf{G}_n^{-1}(\mathbf{u}_0) \in \mathbb{Z}^m$  and then  $\mathbf{x}_i = -\mathbf{G}_n^{-1}(\bar{\mathbf{A}}\mathbf{x}_{i-1}) \in \mathbb{Z}^m$  for  $i = 2, \dots, T$ . This works because

$$\bar{\mathbf{A}}\mathbf{x}_{i-1} + \mathbf{G}\mathbf{x}_i = \bar{\mathbf{A}}\mathbf{x}_{i-1} - \bar{\mathbf{A}}\mathbf{x}_{i-1} = \mathbf{0}.$$

The reason for the gap between the “very short” bound obtained by the above computation, versus the “somewhat short” bound in the assumption, is the  $O(n)^{2 \log T}$  “slack factor” in the proof of sequential work from [LM23]. More specifically, the honest prover can use a “very short” solution to convince the verifier, but the knowledge extractor can extract only a “somewhat short” solution from any (possibly malicious) prover that manages to convince the verifier. Looking ahead, the attacks we give in Section 4 crucially exploit this gap, using a low-depth computation to find a solution that is significantly longer than the “very short” one, but still below the “somewhat short” threshold.

### 3 Attack Framework

Here we develop a suite of general tools that can be combined in various ways to yield attacks on sequentiality assumptions and proofs of sequential work.

#### 3.1 Gadget Trapdoors

We first recall from [MP12] the notion of a (gadget) *trapdoor* for a matrix  $\mathbf{A} \in \mathbb{Z}_q^{N \times W}$ , for any dimensions  $N, W$ . This is any “short” matrix  $\mathbf{R} \in \mathbb{Z}^{W \times M}$ , where  $M = N\ell$ , for which

$$\mathbf{A}\mathbf{R} = \mathbf{G}_N = \mathbf{I}_N \otimes \mathbf{g}^t \in \mathbb{Z}_q^{N \times M}. \quad (3.1)$$

More precisely,  $\mathbf{R}$  should have suitably bounded operator norm  $\|\mathbf{R}\|_{r \leftarrow p}$  for whatever  $p, r$  are most appropriate for the application. For example, when bounds on the Euclidean  $\ell_2$  norm are desired, the spectral norm

<sup>1</sup>For convenience, we have made a slight but immaterial tweak to the system appearing in [LM23], by dropping the  $\bar{\mathbf{A}}$  matrix in the bottom-right block (that appears below our bottom-right  $\mathbf{G}$  matrix) and dropping the bottom-most block  $\mathbf{u}_T$  on the right-hand side (below our zero blocks). It is easy to see that the two systems are equivalent. In addition, [LM23] considers a more compact and “algebraically structured” version of the system over a certain polynomial ring, where each  $n$ -by- $n$  block of  $\bar{\mathbf{A}}$  is the (structured) multiplication matrix of a random ring element. Our attack works for arbitrary  $\bar{\mathbf{A}}$ , so for simplicity and generality we adopt the above presentation.

$\|\mathbf{R}\|_{2\leftarrow 2}$  is usually most useful. In this work we will also frequently use the  $\ell_\infty$  norm, and also finer-grained *block-wise* norm bounds, as defined in [Section 2](#). Observe that  $\mathbf{G}_N$  itself has  $\mathbf{I}_M$  as a trapdoor.

Using such a trapdoor, it is easy to compute, in low depth, a comparably short solution  $\mathbf{x} \in \mathbb{Z}^W$  to  $\mathbf{A}\mathbf{x} = \mathbf{u}$ , for any syndrome  $\mathbf{u} \in \mathbb{Z}_q^N$ : simply let  $\mathbf{x} = \mathbf{R} \cdot \mathbf{G}_N^{-1}(\mathbf{u})$ . This works because both  $\mathbf{R}$  and  $\mathbf{G}_N^{-1}(\mathbf{u})$  are short, and

$$\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{R} \cdot \mathbf{G}_N^{-1}(\mathbf{u}) = \mathbf{G}_N \cdot \mathbf{G}_N^{-1}(\mathbf{u}) = \mathbf{u} .$$

Recall that  $\mathbf{G}_N^{-1}$  applies  $\mathbf{g}^{-1}$  to each entry of  $\mathbf{u}$  independently, so  $\mathbf{x}$  can be computed in depth  $\tilde{O}_{N,q}(1)$ , i.e., polylogarithmic in  $N$  and  $q$ . And because  $\mathbf{G}_N^{-1}(\mathbf{u})$  is binary, for any  $r \in [1, \infty]$  we have that  $\mathbf{x}$  satisfies the (potentially block-wise) norm bound

$$\|\mathbf{x}\|_r \leq \|\mathbf{R}\|_{r\leftarrow \infty} . \quad (3.2)$$

For example, by [Equation \(2.2\)](#) this is at most  $W^{1/r} \cdot \|\mathbf{R}\|_\infty \cdot M$ , but tighter bounds may be available in specific circumstances.

**Subspace trapdoors.** More generally, for some relations (including the specific one from [\[LM23\]](#)) it suffices, and will yield better bounds, to have a limited trapdoor with respect to just the top  $s$  rows for some  $s \ll N$ , which we call an *s-subspace* trapdoor.<sup>2</sup> For convenience of usage later on, we allow  $s$  to be arbitrary, including  $s > N$ . An *s-subspace* trapdoor of  $\mathbf{A}$  is a short matrix  $\mathbf{R} \in \mathbb{Z}^{W \times s\ell}$  for which

$$\mathbf{A}\mathbf{R} = \mathbf{G}_{N,s\ell} := \mathbf{I}_{N,s} \otimes \mathbf{g}^t \in \mathbb{Z}_q^{N \times s\ell} , \quad (3.3)$$

where  $\mathbf{I}_{a,b}$  is the *a-by-b* matrix obtained by padding  $\mathbf{I}_{\min(a,b)}$  with zeros on the right or the bottom, as appropriate. We have that  $\|\mathbf{I}_{a,b}\|_{r\leftarrow p} = \|\mathbf{I}_{\min(a,b)}\|_{r\leftarrow p}$ . Notice that the gadget matrix  $\mathbf{G}_N$  itself has the matrix  $\mathbf{I}_{M,s\ell}$  as an *s-subspace* trapdoor.

**Definition 3.1 (*s-admissible vector*).** For a positive integer  $s$ , a vector  $\mathbf{u} \in \mathbb{Z}_q^N$  is *s-admissible* if all its entries below the first  $\min(s, N)$  are zero.

Using an *s-subspace* trapdoor  $\mathbf{R}$  for  $\mathbf{A}$ , for any *s-admissible* syndrome  $\mathbf{u} \in \mathbb{Z}_q^N$ , it is easy to compute, in depth  $\tilde{O}_{s,q}(1)$ , a solution to  $\mathbf{A}\mathbf{x} = \mathbf{u}$  that satisfies the norm bound in [Equation \(3.2\)](#). Let  $\mathbf{u}_s \in \mathbb{Z}_q^s$  be the truncation or padding by zeros (as appropriate) of  $\mathbf{u}$ ; observe that  $\mathbf{I}_{N,s}\mathbf{u}_s = \mathbf{u}$  because  $\mathbf{u}$  is *s-admissible*. Output  $\mathbf{x} = \mathbf{R} \cdot \mathbf{G}_s^{-1}(\mathbf{u}_s) \in \mathbb{Z}^W$ . Then we have that, as needed,

$$\mathbf{A}\mathbf{x} = \mathbf{G}_{N,s\ell} \cdot \mathbf{G}_s^{-1}(\mathbf{u}_s) = \mathbf{I}_{N,s} \cdot \mathbf{u}_s = \mathbf{u} .$$

## 3.2 Trapdoor Combiners

The core idea underlying our attacks is to *recursively* compute, in fairly low depth, a trapdoor for the block-triangular matrix of the linear system in question, using trapdoors for sub-matrices of the system. This trapdoor can then be used to find a short solution for any desired syndrome, as described above. In this section we give a variety of low-depth, non-recursive *combiner algorithms* for constructing trapdoors and finding short solutions using sub-trapdoors. [Sections 4](#) and [5](#) then give *parallel* recursive “driver” algorithms that use these combiners for specific attacks.

The base case is where the system’s matrix is simply the gadget matrix  $\mathbf{G}_n$  for some (typically small)  $n$ —which trivially has the identity matrix  $\mathbf{I}_{n\ell}$  as a trapdoor—or some other matrix having a known trapdoor. For

<sup>2</sup>This definition and the associated techniques naturally generalize to any particular set of rows, or even to any subspace of the column space.



the recursive case, suppose that the system's matrix  $\mathbf{A}$  has  $N = N_0 + N_1$  rows and block lower-triangular form

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_0 & \\ \mathbf{W} & \mathbf{A}_1 \end{pmatrix}, \quad (3.4)$$

where  $\mathbf{A}_0$  and  $\mathbf{A}_1$  respectively have  $N_0$  and  $N_1$  rows; typically, one would have  $N_0 \approx N_1$ . Note that the matrix  $\mathbf{A}_T$  from the system in Equation (2.3) has this form, where  $\mathbf{A}_0 = \mathbf{A}_1 = \mathbf{A}_{T/2}$  for even  $T$  (and similarly for odd  $T$ ) and  $\mathbf{W}$  is all zeros except in its upper-rightmost  $n$ -by- $m$  block.

The following lemma is the main result of this section.

**Lemma 3.2.** *Let  $\mathbf{A}$  have the form given in Equation (3.4). There is a polynomial-time, depth- $\tilde{O}_{N,q}(1)$  algorithm CombTrap that, given  $\mathbf{A}$  and a trapdoor  $\mathbf{R}_i$  (having  $M_i = N_i \ell$  columns) for each  $\mathbf{A}_i$ , outputs a trapdoor  $\mathbf{R}$  (having  $M = N \ell$  columns) for  $\mathbf{A}$  satisfying for all  $p, r \in [1, \infty]$  the block-wise bound*

$$\|\mathbf{R}\|_{r \leftarrow p} \leq \left( \begin{array}{l} \|\mathbf{R}_0\|_{r \leftarrow p} \\ \|\mathbf{R}_1\|_{r \leftarrow p} \cdot (M_0^{1-1/p} \cdot M_1^{1/p} + 1) \end{array} \right) \quad (3.5)$$

Moreover, for  $p = 2$ , the bound holds with  $O(\sqrt{M})$  in place of  $(M_0^{1-1/p} \cdot M_1^{1/p} + 1)$ , with high probability.

*Proof.* Observe that

$$\mathbf{A} \cdot \begin{pmatrix} \mathbf{R}_0 & \\ & \mathbf{R}_1 \end{pmatrix} = \begin{pmatrix} \mathbf{G}_{N_0} & \\ \mathbf{W}\mathbf{R}_0 & \mathbf{G}_{N_1} \end{pmatrix}. \quad (3.6)$$

Therefore, a trapdoor  $\mathbf{R}$  for  $\mathbf{A}$  is

$$\mathbf{R} = \begin{pmatrix} \mathbf{R}_0 & \\ & \mathbf{R}_1 \end{pmatrix} \begin{pmatrix} \mathbf{I}_{M_0} & \\ \mathbf{R}' & \mathbf{I}_{M_1} \end{pmatrix} = \begin{pmatrix} \mathbf{R}_0 & \\ \mathbf{R}_1 \mathbf{R}' & \mathbf{R}_1 \end{pmatrix} \text{ where } \mathbf{R}' = \mathbf{G}_{N_1}^{-1}(-\mathbf{W}\mathbf{R}_0) \in \mathbb{Z}^{M_1 \times M_0}. \quad (3.7)$$

This is because

$$\mathbf{A}\mathbf{R} = \begin{pmatrix} \mathbf{G}_{N_0} & \\ \mathbf{W}\mathbf{R}_0 & \mathbf{G}_{N_1} \end{pmatrix} \begin{pmatrix} \mathbf{I}_{M_0} & \\ \mathbf{R}' & \mathbf{I}_{M_1} \end{pmatrix} = \begin{pmatrix} \mathbf{G}_{N_0} & \\ \mathbf{W}\mathbf{R}_0 - \mathbf{W}\mathbf{R}_0 & \mathbf{G}_{N_1} \end{pmatrix} = \mathbf{G}_N.$$

Note that  $\mathbf{R}$  can be computed (given  $\mathbf{R}_0, \mathbf{R}_1, \mathbf{W}$ ) in depth  $\tilde{O}_{N,q}(1)$ , by first computing  $\mathbf{W}\mathbf{R}_0$ , then  $\mathbf{R}'$ , then  $\mathbf{R}_1 \mathbf{R}'$ .

For the norm bound, first observe that  $\mathbf{R}' \in \mathbb{Z}^{M_1 \times M_0}$  has binary entries, so  $\|\mathbf{R}'\|_\infty \leq 1$ . From Equation (3.7), and sub-additivity and transitivity of the operator norm, we get the block-wise bound

$$\|\mathbf{R}\|_{r \leftarrow p} \leq \left( \begin{array}{l} \|\mathbf{R}_0\|_{r \leftarrow p} \\ \|\mathbf{R}_1\|_{r \leftarrow p} \cdot (\|\mathbf{R}'\|_{p \leftarrow p} + 1) \end{array} \right).$$

The bound  $\|\mathbf{R}'\|_{p \leftarrow p} \leq M_0^{1-1/p} \cdot M_1^{1/p}$  follows from Equation (2.2). Moreover, if we use a randomized, subgaussian variant of  $\mathbf{g}^{-1}$ , we get that  $\|\mathbf{R}'\|_{2 \leftarrow 2} = O(\sqrt{M})$  with high probability.  $\square$

### 3.2.1 Subspace Trapdoors

To attack the specific relation from [LM23] (see Equation (2.3)), because the syndrome is  $n$ -admissible it suffices to compute an  $n$ -subspace trapdoor, regardless of how large  $N = Tn$  is. This can be done by a simple optimization of the above approach, which yields much better matrix-norm bounds, and somewhat better computation-depth bounds.

**Definition 3.3 ( $s$ -admissible matrix).** Let  $\mathbf{A}$  have the form given in Equation (3.4). For a positive integer  $s$ , we say that  $\mathbf{A}$  is  $s$ -admissible if all the columns of  $\mathbf{W}$  are  $s$ -admissible vectors (Definition 3.1), i.e., all their entries below the first  $\min(s, N_1)$  are zero.

For example, observe that the matrix  $\mathbf{A}_T$  from [LM23] (see Equation (2.3)) is  $n$ -admissible.

For  $s$ -admissible matrices, the CombTrap algorithm from Lemma 3.2 adapts to an algorithm CombTrap $_s$  having the following differences:<sup>3</sup>

- It takes as input, and produces as output,  $s$ -subspace trapdoors of their respective matrices; note that these trapdoors have just  $s\ell$  columns (not  $M = N\ell$  or  $M_i = N_i\ell$ ).
- The output can be computed in depth  $\tilde{O}_{s,q}(1)$ ; note that in the subscript,  $s$  appears in place of  $N$ .
- The operator-norm bound from Equation (3.5) becomes

$$\|\mathbf{R}\|_{r\leftarrow p} \leq \begin{pmatrix} \|\mathbf{R}_0\|_{r\leftarrow p} \\ \|\mathbf{R}_1\|_{r\leftarrow p} \cdot s\ell \end{pmatrix}, \quad (3.8)$$

and for  $p = 2$ , the bound holds with  $O(\sqrt{s\ell})$  in place of  $s\ell$ , with high probability.

Specifically, CombTrap $_s$  works as follows:

1. Write  $\mathbf{G}_{N,s\ell} = \begin{pmatrix} \mathbf{G}_{N_0,s\ell} \\ \mathbf{H}_{N_1,s\ell} \end{pmatrix}$ . Observe that the columns of  $\mathbf{G}_{N,s\ell}$  are  $s$ -admissible vectors, hence the same goes for  $\mathbf{H}_{N_1,s\ell}$ .
2. Let  $\mathbf{W}_s$  and  $\mathbf{H}_s$  respectively be the truncations or zero-paddings (as appropriate) of  $\mathbf{W}$  and  $\mathbf{H}_{N_1,s\ell}$  to  $s$  rows. Because the columns of  $\mathbf{W}$  and  $\mathbf{H}_{N_1,s\ell}$  are  $s$ -admissible vectors,  $\mathbf{I}_{N_1,s} \cdot \mathbf{W}_s = \mathbf{W}$  and  $\mathbf{I}_{N_1,s} \cdot \mathbf{H}_s = \mathbf{H}_{N_1,s\ell}$ .
3. Output

$$\mathbf{R} = \begin{pmatrix} \mathbf{R}_0 & \\ & \mathbf{R}_1 \end{pmatrix} \begin{pmatrix} \mathbf{I}_{s\ell} \\ \mathbf{R}' \end{pmatrix} = \begin{pmatrix} \mathbf{R}_0 \\ \mathbf{R}_1 \mathbf{R}' \end{pmatrix} \text{ where } \mathbf{R}' = \mathbf{G}_s^{-1}(\mathbf{H}_s - \mathbf{W}_s \mathbf{R}_0) \in \mathbb{Z}^{s\ell \times s\ell}. \quad (3.9)$$

The output is an  $s$ -subspace trapdoor of  $\mathbf{A}$  due to the subspace-trapdoor relation (Equation (3.3)), the above parsing of  $\mathbf{G}_{N,s\ell}$ , and a calculation similar to the one in the proof of Lemma 3.2, using the fact that

$$\mathbf{A}_1 \mathbf{R}_1 \cdot \mathbf{R}' = \mathbf{G}_{N_1,s\ell} \cdot \mathbf{G}_s^{-1}(\mathbf{H}_s - \mathbf{W}_s \mathbf{R}_0) = \mathbf{I}_{N_1,s} \cdot (\mathbf{H}_s - \mathbf{W}_s \mathbf{R}_0) = \mathbf{H}_{N_1,s\ell} - \mathbf{W} \mathbf{R}_0.$$

The depth bound of  $\tilde{O}_{s,q}(1)$  holds because  $\mathbf{R}_1$  has only  $s\ell$  columns. And bounds on  $\|\mathbf{R}'\|_{p\leftarrow p}$  analogous to the ones in the proof of Lemma 3.2 hold, but with  $s\ell$  in place of  $M_0$  and  $M_1$ .

### 3.2.2 Combining Solver

Here we describe a slightly optimized solution finder, which is important for our attack on the PoSW protocol with its original norm bounds (see Section 5). Instead of building a full trapdoor from sub-trapdoors and then using it to find a comparably short solution for a desired syndrome, we can instead *directly* use the sub-trapdoors to get a somewhat shorter solution. Essentially, this optimization corresponds to solving the system “slightly honestly,” where we sequentially compute a solution one block at a time, using the sub-trapdoors to compute each block of the solution in low depth. Comparing the combination of Equations (3.2) and (3.5) above with Equation (3.10) below, this method yields better block-wise norm bounds for (say) the  $\ell_\infty$  norm by at least an  $M_0$  factor for full trapdoors, and an  $s\ell$  factor for an  $s$ -subspace trapdoor.

<sup>3</sup>As with subspace trapdoors themselves, this modification naturally generalizes to any linear subspace of the column space.



**Lemma 3.4.** *Let  $\mathbf{A}$  have the form given in Equation (3.4). There is a polynomial-time, depth- $\tilde{O}_{N,q}(1)$  algorithm CombSolve that, given  $\mathbf{A}$ , any syndrome  $\mathbf{u} \in \mathbb{Z}_q^N$ , and a trapdoor  $\mathbf{R}_i$  (having  $M_i = N_i \ell$  columns) for each  $\mathbf{A}_i$ , outputs a solution to  $\mathbf{A}\mathbf{x} = \mathbf{u}$  satisfying for any  $p, r \in [1, \infty]$  the block-wise norm bound*

$$\|\mathbf{x}\|_r \leq \left( \begin{array}{l} \|\mathbf{R}_0\|_{r \leftarrow p} \cdot M_0^{1/p} \\ \|\mathbf{R}_1\|_{r \leftarrow p} \cdot M_1^{1/p} \end{array} \right). \quad (3.10)$$

Similarly, for any positive integer  $s$ , there is a polynomial-time, depth- $\tilde{O}_{s,q}(1)$  algorithm CombSolve $_s$  that, given any  $s$ -admissible  $\mathbf{A}$  and  $\mathbf{u} \in \mathbb{Z}_q^N$ , and an  $s$ -subspace trapdoor  $\mathbf{R}_i$  for each  $\mathbf{A}_i$ , outputs a solution to  $\mathbf{A}\mathbf{x} = \mathbf{u}$  satisfying the same block-wise norm bound as in Equation (3.10), but with  $s\ell$  in place of each  $M_i$ .

*Proof.* For our purposes in Section 5.3 it is convenient to define CombSolve in terms of a ‘‘helper’’ function CombSolveHelper that has a slightly different interface. Given:

- $\mathbf{A}$  and an initial solution block  $\mathbf{x}_0$  (that is meant to satisfy  $\mathbf{A}_0\mathbf{x}_0 = \mathbf{u}_0$  for some initial syndrome block  $\mathbf{u}_0 \in \mathbb{Z}_q^{N_0}$ ),
- a remaining syndrome  $\mathbf{u}_1 \in \mathbb{Z}_q^{N_1}$ , and
- a trapdoor  $\mathbf{R}_i$  for  $\mathbf{A}_i$  for each  $i > 0$  (so  $\mathbf{R}_0$  is not needed),

it outputs a *full* solution  $\mathbf{x} = \begin{pmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \end{pmatrix}$  to  $\mathbf{A}\mathbf{x} = \mathbf{u}$ , by computing a solution to  $\mathbf{A}_1\mathbf{x}_1 = \mathbf{u}'_1 := \mathbf{u}_1 - \mathbf{W}\mathbf{x}_0$ , namely,  $\mathbf{x}_1 = \mathbf{R}_1 \cdot \mathbf{G}_{N_1}^{-1}(\mathbf{u}'_1)$ .

We now define CombSolve. It parses  $\mathbf{u} = \begin{pmatrix} \mathbf{u}_0 \\ \mathbf{u}_1 \end{pmatrix} \in \mathbb{Z}_q^N$  where  $\mathbf{u}_i \in \mathbb{Z}_q^{N_i}$ , computes the initial solution block as  $\mathbf{x}_0 = \mathbf{R}_0 \cdot \mathbf{G}_{N_0}^{-1}(\mathbf{u}_0)$ , and finally outputs the full solution CombSolveHelper( $\mathbf{A}, \mathbf{x}_0, \mathbf{u}_1, \mathbf{R}_1$ ). Observe that the output is computed in depth  $\tilde{O}_{N,q}(1)$ , and satisfies  $\|\mathbf{x}_i\|_r \leq \|\mathbf{R}_i\|_{r \leftarrow p} \cdot M_i^{1/p}$  because the output of  $\mathbf{G}_{N_i}^{-1}$  is binary.

The claim about CombSolve $_s$  follows similarly, via a corresponding helper CombSolveHelper $_s$ , by adapting the approach for combining subspace trapdoors from Section 3.2.1 above.  $\square$

### 3.2.3 Larger Arity

Finally, all of the above easily generalizes to the case where  $\mathbf{A}$  can be split into larger numbers of blocks. This will be needed for our attacks on PoSW variants in Section 5.

Specifically, suppose that for some  $k \geq 2$ , the matrix  $\mathbf{A}$  has the form

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_0 & & & \\ \mathbf{W}_{1,0} & \mathbf{A}_1 & & \\ \vdots & \vdots & \ddots & \\ \mathbf{W}_{k-1,0} & \mathbf{W}_{k-1,1} & \cdots & \mathbf{A}_{k-1} \end{pmatrix}, \quad (3.11)$$

where  $\mathbf{A}_i$  has  $N_i$  rows. Then everything in Lemmas 3.2 and 3.4 and Section 3.2.1 generalizes in the natural way, with algorithms having depth  $\tilde{O}_{N,q}(k)$ , using techniques analogous to Gaussian elimination (run in parallel over columns).

In a bit more detail, CombTrap from Lemma 3.2 generalizes as follows. It can be verified that if  $\mathbf{R}_i$  is a trapdoor (having  $M_i = N_i \ell$  columns) for  $\mathbf{A}_i$ , then a trapdoor for  $\mathbf{A}$  is

$$\mathbf{R} = \begin{pmatrix} \mathbf{R}_0 & & & \\ & \mathbf{R}_1 & & \\ & & \ddots & \\ & & & \mathbf{R}_{k-1} \end{pmatrix} \begin{pmatrix} \mathbf{I}_{M_0} & & & \\ \mathbf{R}'_{1,0} & \mathbf{I}_{M_1} & & \\ \vdots & \vdots & \ddots & \\ \mathbf{R}'_{k-1,0} & \mathbf{R}'_{k-1,1} & \cdots & \mathbf{I}_{M_{k-1}} \end{pmatrix}$$

where

$$\mathbf{R}'_{i,j} = \mathbf{G}_{N_i}^{-1} \left( -\mathbf{W}_{i,j} \mathbf{R}_j - \sum_{j < j' < i} \mathbf{W}_{i,j'} \mathbf{R}_{j'} \mathbf{R}'_{j',j} \right).$$

These blocks can be computed in  $k - 1$  sequential stages, in parallel over the columns  $j = 0, \dots, k - 1$ , by computing  $\mathbf{R}'_{i,j}$  for  $i = j + 1, \dots, k - 1$  in sequence. So,  $\mathbf{R}$  can be computed in depth  $\tilde{O}_{N,q}(k)$ . A block-wise operator-norm bound on  $\mathbf{R}$  can easily be obtained from operator-norm bounds on the  $\mathbf{R}'_{i,j}$  and the triangle inequality, as in the proof of [Lemma 3.2](#). (We will not need such a bound in this work, so we omit its rather large expression.)

Similarly, the optimized solver CombSolve from [Lemma 3.4](#) generalizes to this setting as well, in an analogous way.

**Subspace trapdoors.** All the optimizations relating to *subspace* trapdoors also generalize similarly. First, we generalize the notion of  $s$ -admissible matrix ([Definition 3.3](#)) to apply to matrices having the form given in [Equation \(3.11\)](#), where we require that the columns of *all* the  $\mathbf{W}_{i,j}$  are  $s$ -admissible vectors. The combiners and solvers for  $s$ -admissible matrices and subspace trapdoors given in [Sections 3.2.1](#) and [3.2.2](#) then generalize straightforwardly, and have depth  $\tilde{O}_{s,q}(k)$ . More specifically, the output of  $\text{CombTrap}_s$  satisfies an operator-norm bound analogous to [Equation \(3.8\)](#), where *all blocks but the first* incur an expansion, by the same factor. Finally, the optimized solver  $\text{CombSolve}_s$  from [Lemma 3.4](#) and its norm bound also generalize in the natural way.

## 4 Attack on the Assumption

Here we use the tools from [Section 3](#) to give a recursive attack on the sequentiality assumption from [[LM23](#)] (see [Section 2.2](#)). Our ultimate result in this respect is as follows.

**Theorem 4.1 (Attack on the sequentiality assumption).** *Let  $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times m}$ ,  $T$  and  $k \in [2, T]$  be positive integers, and  $\mathbf{A}_T$  be the matrix given in [Equation \(2.3\)](#). Given  $\bar{\mathbf{A}}, T, k$ , an  $n$ -subspace trapdoor  $\mathbf{R}_T$  for  $\mathbf{A}_T$  satisfying the following (scalar) norm bound for any  $p, r \in [1, \infty]$  can be computed in depth  $\tilde{O}_{n,q}(k \log_k T)$ :*

$$\|\mathbf{R}_T\|_{r \leftarrow p} \leq (k^{1/r} \cdot m)^{\lceil \log_k T \rceil} \cdot \|\mathbf{I}_m\|_{r \leftarrow p}.$$

*Moreover, for  $p = 2$  and  $T = 2^{o(m)}$ , the bound holds with  $\sqrt{m}$  in place of  $m$  in the factor on the left, with high probability.*

*As a result, a solution  $\mathbf{x}$  to [Equation \(2.3\)](#) satisfying the following (scalar) norm bound for any  $p \in [1, \infty]$  can be computed in depth  $\tilde{O}_{n,q}(k \log_k T)$ :*

$$\|\mathbf{x}\|_p \leq (k^{1/p} \cdot m)^{\lceil \log_k T \rceil} \cdot m^{1/p},$$

*and for  $p = 2$  and  $T = 2^{o(m)}$  the bound holds with  $\sqrt{m}$  in place of  $m$  in the factor on the left, with high probability.*

We prove this theorem below as a corollary of a much more general setup and attack. As a couple of examples, we can get the following parameterizations of [Theorem 4.1](#):

- For  $k = 2$ , in depth just  $\tilde{O}_{n,q}(\log T)$  we get a solution having  $\ell_\infty$  norm bounded by  $m^{\lceil \log_2 T \rceil}$ , or  $\ell_2$  norm (and hence  $\ell_\infty$  norm as well) bounded by  $(2m)^{(\lceil \log_2 T \rceil + 1)/2}$  with high probability.

```

RecTraps(A)


---


if A = Gn : return In,sℓ
foreach i in parallel : Ri = RecTraps(Ai)
return CombTraps(A, [Ri]i)

```

Figure 1: Algorithm that computes a subspace trapdoor for any recursively  $s$ -admissible matrix  $\mathbf{A}$ .

- By setting  $k \approx T^\varepsilon$  for an arbitrarily small constant  $\varepsilon > 0$ , in depth  $\tilde{O}_{n,q}(T^\varepsilon)$  we get a solution of just polynomially bounded  $\ell_\infty$  norm  $m^{1/\varepsilon}$ .

In a typical instantiation where  $\log q = \tilde{O}_{n,T}(1)$ , the  $q$  subscripts in the above  $\tilde{O}$  expressions can be replaced by  $T$ . In general, for  $\log q = o(n)$  and hence  $m = o(n^2)$ , the above falsifies (even a major weakening of) the assumption made in [LM23], which posits that computing a solution of  $\ell_\infty$  norm  $O(n)^{2 \log_2 T}$  requires depth  $\tilde{\Omega}_{n,q}(T)$  (or in weaker form, depth  $\tilde{\Omega}_{n,q}(T^\varepsilon)$  for some constant  $\varepsilon > 0$ ).

**General attack.** In fact, our attack applies to a *broad generalization* of the system in Equation (2.3), namely, any one in which the matrix is block lower-triangular and has “gadget” matrices  $\mathbf{G}$  (or any other matrices having known trapdoors) as the diagonal blocks. For simplicity of presentation, throughout this section and Section 5, our general attack algorithms assume that the system’s matrix  $\mathbf{A}$  is *implicitly* given in block form following Equation (3.4) (or more generally, Equation (3.11)), and similarly for its component  $\mathbf{A}_i$  matrices, *recursively*. Then, an attack against a specific assumption or protocol is obtained as an instantiation of the general attack, by specifying the recursive block structure for the proposal in question.

**Definition 4.2 (Recursively  $s$ -admissible).** For a positive integer  $s$ , a matrix  $\mathbf{A}$  is *recursively  $s$ -admissible* if  $\mathbf{A} = \mathbf{G}_n$  for some  $n$  (the base case), or if it has the form given in Equation (3.4) (or more generally, Equation (3.11)), it is  $s$ -admissible, and all its  $\mathbf{A}_i$  submatrices are recursively  $s$ -admissible.

For example, notice that the matrix  $\mathbf{A}_T$  from Equation (2.3) is recursively  $n$ -admissible for any choice of arity  $k \in [2, T]$ , with a recursive block structure of depth  $\lceil \log_k T \rceil$ .

In the rest of this section, we define and analyze a simple recursive “driver” of the (general arity, subspace) trapdoor combiner CombTrap from Section 3.2; see Figure 1 for its formal definition. Then we show that a straightforward instantiation of this driver yields Theorem 4.1 as corollary.

**Lemma 4.3.** *Let  $\mathbf{A}$  be a recursively  $s$ -admissible matrix that recursively has the form given in Equation (3.11) for some arity  $k$  and recursion depth at most  $d$  (where  $d = 0$  is the base case). Then  $\text{RecTrap}_s(\mathbf{A})$  computes, in depth  $\tilde{O}_{s,q}(kd)$ , an  $s$ -subspace trapdoor  $\mathbf{R}$  for  $\mathbf{A}$  satisfying for any  $p, r \in [1, \infty]$  the (scalar) norm bound*

$$\|\mathbf{R}\|_{r \leftarrow p} \leq (k^{1/r} \cdot s\ell)^d \cdot \|\mathbf{I}_{s\ell}\|_{r \leftarrow p}.$$

*For  $p = 2$  and  $k^d = 2^{o(s\ell)}$ , the bound holds with  $\sqrt{s\ell}$  in place of  $s\ell$  in the factor on the left, with high probability. In particular, for  $r = \infty$  we have that  $\|\mathbf{R}\|_{\infty \leftarrow p} \leq (s\ell)^d$ , and  $\|\mathbf{R}\|_{\infty \leftarrow 2} \leq (s\ell)^{d/2}$  with high probability.*

*Proof.* It is easy to see that at each level of the recursion, by (the general arity, subspace version of) Lemma 3.2, CombTrap<sub>s</sub> has depth  $\tilde{O}_{s,q}(k)$ . Hence the overall computation depth of RecTrap<sub>s</sub>( $\mathbf{A}$ ) is  $\tilde{O}_{s,q}(kd)$ .

For the operator norm of  $\mathbf{R}$ , again by (the general arity, subspace version of) [Lemma 3.2](#), we have the recurrence

$$\|\mathbf{R}\|_{r \leftarrow p} \leq \left\| \begin{pmatrix} \|\mathbf{R}_0\|_{r \leftarrow p} \\ \|\mathbf{R}_1\|_{r \leftarrow p} \cdot s\ell \\ \vdots \\ \|\mathbf{R}_{k-1}\|_{r \leftarrow p} \cdot s\ell \end{pmatrix} \right\|_r \leq k^{1/r} \cdot \max_i \|\mathbf{R}_i\|_{r \leftarrow p} \cdot s\ell.$$

For  $p = 2$ , the recurrence holds with  $\sqrt{s\ell}$  in place of  $s\ell$ , with high probability. The base case of this recurrence is  $\|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p} \leq \|\mathbf{I}_{s\ell}\|_{r \leftarrow p}$ , and the recursion has depth at most  $d$ . This yields the claimed bound on  $\|\mathbf{R}\|_{r \leftarrow p}$ .  $\square$

*Proof of [Theorem 4.1](#).* Recall that the matrix  $\mathbf{A}_T$  from [Equation \(2.3\)](#) is recursively  $n$ -admissible for any choice of arity  $k \in [2, T]$ , with a corresponding depth of  $d = \lceil \log_k T \rceil$ . The computation of  $\mathbf{R}_T$  and the bound on  $\|\mathbf{R}_T\|_{r \leftarrow p}$  then follows immediately from [Lemma 4.3](#) and the fact that  $m = n\ell$ . For the particular solution  $\mathbf{x}$ , we use the subspace-trapdoor solver (see [Section 3.1](#)), letting  $\mathbf{x} = \mathbf{R}_T \cdot \mathbf{G}_n^{-1}(\mathbf{u}_0)$ . Then

$$\|\mathbf{x}\|_p \leq \|\mathbf{R}_T\|_{p \leftarrow p} \cdot \|\mathbf{G}_n^{-1}(\mathbf{u}_0)\|_p \leq (k^{1/p} \cdot m)^{\lceil \log_k T \rceil} \cdot m^{1/p}. \quad \square$$

## 5 Attacks on Proofs of Sequential Work

Lai and Malavolta [[LM23](#)] also gave a candidate *proof of sequential work* (PoSW) protocol and proved its security based on the sequentiality assumption stated in [Section 2.2](#). While the attack from [Section 4](#) strongly *falsifies the assumption*, and thus renders the security proof vacuous, it does not immediately follow that the *PoSW itself* is broken. Indeed, the attack as stated does not break the protocol, because it produces a solution vector whose components have much larger norms than the “honestly computed” ones, so running the (honest) prover with this vector would not yield short enough component vectors to convince the verifier.

In this section we give attacks that use the tools from [Section 3](#) in a more sophisticated way, to break two PoSWs that are very similar (but not quite identical) to the one given in [[LM23](#)]; see [Figure 2](#) for the formal definition. Both variants tweak the core “folding” operation to multiply the verifier’s random challenge by the *first*, rather than the *second*, half of the prover’s solution vector. In the design of the protocol the choice is arbitrary, since the analysis and security proof apply equally well to either version, but the choice seems to have significant implications for attacks, as we explain below. In addition, in our first variant protocol, the verifier checks the prover’s responses using *somewhat relaxed norm bounds* (which are polynomially related to the original ones). Our second variant uses the original norm bounds from [[LM23](#)]; the only change is the tweaked folding operation.

The rest of this section is organized as follows. In [Section 5.1](#) we describe some of the challenges that arise in attacking the PoSW protocol of [[LM23](#)]. In [Section 5.2](#) we break the first variant protocol using the simple RecTrap algorithm from [Section 4](#), but with a slightly different recursive block decomposition of the system’s matrix. Then in [Section 5.3](#) we break the second variant using a much more sophisticated recursive strategy.

### 5.1 Challenges in Attacking PoSWs

We start by describing some of the additional challenges that arise in attacking the PoSW protocol of [[LM23](#)], which motivate our protocol tweaks and enhanced attacks. See [Figure 2](#) for a formal definition of the protocol, with our tweaks highlighted.

$\text{Prover}(\bar{\mathbf{A}}, T, \mathbf{u}_0, \{\mathbf{u}_i\}_{i \in [T]}, \{\mathbf{x}_i\}_{i \in [T]})$	$\text{Verifier}(\bar{\mathbf{A}}, T, \mathbf{u}_0, \mathbf{u}_T, \beta)$
$\text{if } T = 1 :$ send $\mathbf{x}_1$ <b>return</b>	$\text{if } T = 1 :$ receive $\mathbf{x}_1$ <b>return</b> $[\ \mathbf{x}_1\  \leq \beta$ $\wedge \mathbf{G}\mathbf{x}_1 = \mathbf{u}_0 \wedge -\bar{\mathbf{A}}\mathbf{x}_1 = \mathbf{u}_1]$
$T' = (T - 1)/2$ send $\mathbf{x}_{T'+1}$	$T' = (T - 1)/2$ receive $\mathbf{x}_{T'+1}$
receive $c$ $\mathbf{u}'_i = [c \cdot] \mathbf{u}_i + \mathbf{u}_{i+T'+1}$ $\mathbf{x}'_i = [c \cdot] \mathbf{x}_i + \mathbf{x}_{i+T'+1}$	<b>if</b> $\ \mathbf{x}_{T'+1}\  > \beta$ : <b>return</b> 0 send random $c$ with $ c  \leq \gamma$ $\mathbf{u}'_0 = [c \cdot] \mathbf{u}_0 + (-\bar{\mathbf{A}}\mathbf{x}_{T'+1})$ $\mathbf{u}'_{T'} = [c \cdot] \mathbf{G}\mathbf{x}_{T'+1} + \mathbf{u}_T$
$\text{Prover}(\bar{\mathbf{A}}, T', \mathbf{u}'_0, \{\mathbf{u}'_i\}_{i \in [T']}, \{\mathbf{x}'_i\}_{i \in [T']})$	<b>return</b> $\text{Verifier}(\bar{\mathbf{A}}, T', \mathbf{u}'_0, \mathbf{u}'_{T'}, 2[\gamma'] \cdot \beta)$

Figure 2: The core algorithms of the PoSW from [LM23], with boxed modifications. For simplicity, we give the code only for  $T$  of the form  $T = 2^t - 1$ . The original PoSW applies the challenge factor  $c$  to the *second* term in each sum (instead of the first), and it uses  $\gamma' = \gamma$ .

**Structure of the PoSW.** In the PoSW, the honest prover computes a short solution  $\mathbf{x}$  to Equation (2.3), then engages in an interactive public-coin protocol to convince the verifier that it knows such a solution. (The protocol can be made non-interactive in the usual way via the Fiat–Shamir transform.) To do this, it uses the verifier’s small random challenge to linearly “fold” the first and second halves of  $\mathbf{x}$  together, which yields a somewhat longer solution of half the dimension, then recursively proves knowledge of this folded solution.

More precisely, the prover first announces  $\mathbf{u}_T = -\bar{\mathbf{A}}\mathbf{x}_T$  (or equivalently,  $\mathbf{x}_T$ ) as its claimed result of the sequential computation. Then it gives a proof of knowledge of its solution  $\mathbf{x}$  to Equation (2.3): it announces  $\mathbf{x}_{(T+1)/2}$  (assume that  $T$  is odd for simplicity), which the verifier checks is short enough, and the verifier announces a small random challenge  $c$ .<sup>4</sup> Observe that the remaining halves of  $\mathbf{x}$  form two solutions to reduced-dimension instances of Equation (2.3), with known right-hand sides of the appropriate form. So, the prover linearly combines these solutions using  $c$ , and recursively proves knowledge of the resulting (somewhat longer) solution in the same manner, until the dimension is small enough to simply reveal and check the solution. Note that with each successive stage of the recursion, the verifier must apply a more relaxed norm check on the prover’s announced value, because folding increases the norm of the solution by some fixed factor (independent of  $T$ ).

**The difficulty.** The key challenge in attacking the PoSW seems to be as follows: (1) the prover must first announce some  $\mathbf{u}_T$  to the verifier, so (2) the prover can know *at most one* short solution for whatever value it announces, and (3) the only way we see to convince the verifier is by knowing a solution whose middle component  $\mathbf{x}_{(T+1)/2}$ , along with all subsequent middle components under the recursive folding, are *nearly as short* as what the honest prover would compute. We elaborate on each of these points next.

The announced value of  $\mathbf{u}_T$  represents the claimed result of the sequential computation that the prover

---

<sup>4</sup>As mentioned before, [LM23] considers an “algebraically structured” version of the system Equation (2.3) over a certain polynomial ring, and the challenge  $c$  there is also a ring element. Hence in our language, we actually need to model the challenge  $c$  as a matrix, and measure its shortness by some suitable operator norm. For simplicity we still write  $c$  as a scalar, and we give more details about the norm of  $c$  as a matrix when it comes to the actual calculations later on.

supposedly performed to get a solution to Equation (2.3). It is straightforward to show that computing *distinct* short solutions, for *any* fixed and possibly adversarially chosen right-hand side, is at least as hard as solving the corresponding SIS problem for the random matrix  $\bar{\mathbf{A}}$ . So, under the standard assumption that SIS is intractable (see, e.g., [Ajt96, MR04]), once an efficient prover (of any depth) reveals some  $\mathbf{u}_T$ , it can know *at most one* short solution for it. The same goes for the later stages of the protocol with lower-dimensional instances of Equation (2.3), where the first and last components of the right-hand side are determined by the previous stages.<sup>5</sup>

With these constraints, the only way we see to convince the verifier is by proceeding exactly as the honest (specified) prover would, using a single known solution. This means that the middle component of the solution vector, and all subsequent middle components under recursive folding, need to satisfy the verifier’s norm checks. In particular, it seems that *whichever half of the solution is multiplied by the verifier’s challenge needs to be nearly as short as what the honest prover would compute*. The PoSW from [LM23] multiplies by the *second* half, and we do not see a way to generate such a short second half in depth significantly less than  $T$ . This is because our techniques yield trapdoors and solutions whose second halves are larger than their first halves by some polynomial factor. Again we stress that the attack from Section 4 achieves low depth by exploiting the moderately large slack factor: the solution is merely “somewhat short” in its second half (and fourth quarter, etc.), but not nearly as short as what the honest computation produces.

## 5.2 Breaking a Relaxed PoSW

The above discussion motivates a natural alternative folding operation: multiply the verifier’s random challenge  $c$  by the *first*, rather than the *second*, half of the solution. That is, instead of folding a solution  $\mathbf{x}$  into the lower-dimensional one  $\mathbf{x}_{\text{first}} + c \cdot \mathbf{x}_{\text{last}}$ , use  $c \cdot \mathbf{x}_{\text{first}} + \mathbf{x}_{\text{last}}$ . The security proof from [LM23] (modified in the obvious way) holds equally well for this option, because the two halves are treated symmetrically. (As far as we can tell, in [LM23] the choice of folding operation between these two options was arbitrary.)

Interestingly, this trivial modification makes the PoSW breakable using our attack framework. In this subsection, as a warmup we break this tweaked protocol with a verifier that also uses somewhat relaxed norm checks. The precise statement is as follows; we prove it below after introducing the definitions and tools needed to do so.

**Theorem 5.1 (Attack on the PoSW with relaxed parameters).** *There is a depth- $\tilde{O}_{n,q}(\log T)$  malicious prover that, given any  $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times m}$ ,  $\mathbf{u}_0 \in \mathbb{Z}_q^n$ , and positive integer  $T$ , convinces the verifier in the modified PoSW, where in the folding operation the challenge is multiplied by the first half of the solution, and parameters  $\gamma' \geq (\gamma + m)/2$  and  $\beta \geq m$  are used.*

The attack simply uses the RecTrap algorithm from Section 4, but with a slightly different recursive block decomposition of the matrix  $\mathbf{A}_T$ . Essentially, this works because the  $\mathbf{x}_{\text{last}}$  constructed by the attack is a small factor longer than  $\mathbf{x}_{\text{first}}$ , so the two summands  $c \cdot \mathbf{x}_{\text{first}}$ ,  $\mathbf{x}_{\text{last}}$  in the tweaked folding operation are more “balanced,” hence their sum passes the appropriately relaxed norm checks. However, we also need to ensure that the revealed middle component is sufficiently short, at every stage of the recursion. For this we impose a suitable recursive block structure on  $\mathbf{A}_T$ , which treats the middle components specially, as base cases.

**Definition 5.2 (PoSW topology).** For a positive integer  $r$ , a block vector or matrix has *PoSW topology* of depth  $d$  and base rows  $r$  if for the base case  $d = 0$  it has  $r$  rows, and if for  $d > 0$  it has three row blocks,

---

<sup>5</sup>This state of affairs is quite different from the context of the attack from Section 4, where the adversary is not bound to any particular right-hand side of Equation (2.3), and knows short solutions to many different right-hand sides via its trapdoor.



respectively having PoSW topologies of depths  $d - 1, 0, d - 1$  (all with base rows  $r$ ). In addition, for a recursively  $s$ -admissible matrix  $\mathbf{A}$  (Definition 4.2), we also require that  $\mathbf{A} = \mathbf{G}_r$  in the base case  $d = 0$ .

Observe that by induction, a vector or matrix having PoSW topology of depth  $d$  has  $T = 2^{d+1} - 1$  rows blocks, of  $r$  rows each. Also observe that for such  $T$  and for  $r = n$ , the recursively  $n$ -admissible matrix  $\mathbf{A}_T$  from the system in Equation (2.3) can be given this PoSW topology, and the base case  $\mathbf{G}_n$  has an ( $n$ -subspace) trapdoor  $\mathbf{I}_{n\ell}$ .

*Remark 5.3.* In Definition 5.2, for simplicity of presentation, we define the PoSW topology only for matrices having exactly  $T = 2^{d+1} - 1$  row blocks. Our attacks, analyses, and results in Sections 5.2 and 5.3 rely on the PoSW topology, and consequently work for  $T$  of this form. However, it is easy to generalize the recursive definition to work for an even number of blocks, following what is done in the original definition of the PoSW of [LM23]. This defines a PoSW topology for any positive integer  $T$ , and our attacks, analyses, and results immediately generalize analogously.

To aid in the analysis, we define the following operations on a vector or matrix  $\mathbf{M}$  having PoSW topology of depth  $d > 0$  and base rows  $r$ , whose three row blocks are denoted  $\mathbf{M}_0, \mathbf{M}_1, \mathbf{M}_2$ . The middle block, and folding operation with multiplier  $c$ , are respectively defined as

$$\text{mid}(\mathbf{M}) := \mathbf{M}_1 = \begin{pmatrix} \mathbf{0}_{r,N'} & \mathbf{I}_r & \mathbf{0}_{r,N'} \end{pmatrix} \cdot \mathbf{M} \quad (5.1)$$

$$\text{fold}_c(\mathbf{M}) := c\mathbf{M}_0 + \mathbf{M}_2 = \begin{pmatrix} c\mathbf{I}_{N'} & \mathbf{0}_{N',r} & \mathbf{I}_{N'} \end{pmatrix} \cdot \mathbf{M}, \quad (5.2)$$

where  $N' = T'r$  for  $T' = (T - 1)/2 = 2^d - 1$  is the number of rows in each of  $\mathbf{M}_0, \mathbf{M}_2$ . For convenience, for  $\mathbf{M}$  having depth  $d = 0$  also define  $\text{mid}(\mathbf{M}) := \mathbf{M} = \mathbf{I}_r \cdot \mathbf{M}$ .

Observe that for a vector or matrix having a PoSW topology, its corresponding recursive block-wise vector of (vector or matrix) norms also has a PoSW topology of the same depth, with one row in the base case. In particular, if  $|c| \leq \gamma$ , then by the triangle inequality,

$$\|\text{fold}_c(\mathbf{M})\|_{r \leftarrow p} \leq \text{fold}_\gamma(\|\mathbf{M}\|_{r \leftarrow p}). \quad (5.3)$$

This can be applied iteratively: for any  $c_i$  with  $|c_i| \leq \gamma$  for all  $1 \leq i \leq \tau$  where  $\tau \leq d$ , we have that

$$\|\text{fold}_{c_\tau}(\cdots(\text{fold}_{c_1}(\mathbf{M}))\|_{r \leftarrow p} \leq \text{fold}_\gamma^{(\tau)}(\|\mathbf{M}\|_{r \leftarrow p}). \quad (5.4)$$

More generally, our treatment (which expands polynomial-ring elements as vectors and matrices) actually models the multiplier  $c$  as a square matrix having the same dimension as the base-case number of rows. Then the folding is defined as

$$\text{fold}_c(\mathbf{M}) := \begin{pmatrix} \mathbf{I}_{T'} \otimes c & \mathbf{0}_{N',r} & \mathbf{I}_{N'} \end{pmatrix} \cdot \mathbf{M},$$

and for any  $c$  with  $\|c\|_{r \leftarrow r} \leq \gamma$ , Equation (5.3) still holds (and similarly for its iterated version Equation (5.4)).

**Lemma 5.4.** *Let  $\mathbf{A}$  be a recursively  $s$ -admissible matrix with PoSW topology of depth  $d$  and base rows  $n$ . Then  $\text{RecTrap}_s(\mathbf{A})$  computes, in depth  $\tilde{O}_{s,q}(d)$ , an  $s$ -subspace trapdoor  $\mathbf{R}$  for  $\mathbf{A}$  satisfying the following bound for any  $p, r \in [1, \infty]$  and any integer  $\tau \in [0, d]$ :*

$$\text{mid}(\text{fold}_\gamma^{(\tau)}(\|\mathbf{R}\|_{r \leftarrow p})) \leq (\gamma + s\ell)^\tau \cdot \|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p} \cdot s\ell.$$

*Proof.* At each level of the recursion, by (the general arity, subspace version of) [Lemma 3.2](#),  $\text{CombTrap}_s$  has computation depth  $\tilde{O}_{s,q}(1)$ . Hence the overall computation depth for  $\text{RecTrap}_s(\mathbf{A})$  is  $\tilde{O}_{s,q}(d)$ .

For the operator norm of  $\mathbf{R}$ , again by (the general arity, subspace version of) [Lemma 3.2](#), we have the recurrence

$$\|\mathbf{R}\|_{r \leftarrow p} \leq \begin{pmatrix} \|\mathbf{R}_0\|_{r \leftarrow p} \\ \|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p} \cdot s\ell \\ \|\mathbf{R}_2\|_{r \leftarrow p} \cdot s\ell \end{pmatrix},$$

and the base case  $d = 0$  has  $\|\mathbf{R}\|_{r \leftarrow p} = \|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p}$ . Hence  $\|\mathbf{R}\|_{r \leftarrow p} \leq \text{pf}(d)$ , a block vector recursively defined as

$$\text{pf}(d) := \begin{pmatrix} \text{pf}(d-1) \\ \|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p} \cdot s\ell \\ \text{pf}(d-1) \cdot s\ell \end{pmatrix}, \quad (5.5)$$

with base case  $\text{pf}(0) = \|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p}$ . Therefore,

$$\text{fold}_\gamma(\text{pf}(d)) = \begin{pmatrix} \gamma \mathbf{I} & \mathbf{0} & \mathbf{I} \end{pmatrix} \cdot \begin{pmatrix} \text{pf}(d-1) \\ \|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p} \cdot s\ell \\ \text{pf}(d-1) \cdot s\ell \end{pmatrix} = (\gamma + s\ell) \cdot \text{pf}(d-1).$$

By iterating, for any integer  $\tau \in [0, d]$ , we get that  $\text{fold}_\gamma^{(\tau)}(\text{pf}(d)) = (\gamma + s\ell)^\tau \cdot \text{pf}(d - \tau)$ . Therefore, as desired,

$$\begin{aligned} \text{mid}(\text{fold}_\gamma^{(\tau)}(\|\mathbf{R}\|_{r \leftarrow p})) &\leq \text{mid}(\text{fold}_\gamma^{(\tau)}(\text{pf}(d))) \\ &= (\gamma + s\ell)^\tau \cdot \text{mid}(\text{pf}(d - \tau)) \\ &\leq (\gamma + s\ell)^\tau \cdot \|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p} \cdot s\ell. \end{aligned}$$

(Note that  $\text{mid}(\text{pf}(d - \tau)) = \|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p}$  in the base case  $\tau = d$ , but  $\|\mathbf{I}_{n\ell, s\ell}\|_{r \leftarrow p} \cdot s\ell$  is still a valid upper bound.)  $\square$

*Proof of [Theorem 5.1](#).* For simplicity of presentation we assume that  $T$  has the form  $T = 2^{d+1} - 1$ , as required by [Definition 5.2](#); for other  $T$ , we can use a generalized PoSW topology as described in [Remark 5.3](#).

The matrix  $\mathbf{A}_T$  from [Equation \(2.3\)](#) is recursively  $n$ -admissible with PoSW topology of depth  $d = \lfloor \log_2 T \rfloor$ . The malicious prover computes an  $n$ -subspace trapdoor  $\mathbf{R}$  of  $\mathbf{A}_T$  following [Lemma 5.4](#), uses the subspace-trapdoor solver (see [Section 3.1](#)) to get a solution  $\mathbf{x} = \mathbf{R}_T \cdot \mathbf{G}_n^{-1}(\mathbf{u}_0)$  to [Equation \(2.3\)](#), and then proceeds in the same way as the honest prover. By [Lemma 5.4](#), this prover has depth  $\tilde{O}_{n,q}(\log T)$ .

Note that following the honest prover using a valid solution ensures that the verifier's "linear" checks are satisfied, so it just remains to confirm that the verifier's norm checks are also satisfied. By [Lemma 5.4](#) again, for all  $\tau \in [0, d]$  the block-wise norm of the solution  $\mathbf{x}$  satisfies

$$\text{mid}(\text{fold}_\gamma^{(\tau)}(\|\mathbf{x}\|_\infty)) \leq \text{mid}(\text{fold}_\gamma^{(\tau)}(\|\mathbf{R}\|_{\infty \leftarrow \infty})) \cdot 1 \leq (\gamma + m)^\tau \cdot m \leq (2\gamma')^\tau \cdot \beta.$$

So, considering [Equation \(5.4\)](#) and the fact that the verifier's challenges  $c$  all satisfy  $|c| \leq \gamma$  (or more precisely,  $\|c\|_{\infty \leftarrow \infty} \leq \gamma$ ), the verifier's norm checks on the folded solution  $\mathbf{x}$  itself are indeed satisfied.  $\square$

### 5.3 Breaking a PoSW with the Original Norm Bounds

The prior subsection breaks the PoSW with tweaked folding operation and relaxed norm bounds. Here we break a PoSW whose only modification from [\[LM23\]](#) is the tweaked folding operation (the norm bounds

remain unchanged). This attack uses much more sophisticated recursive strategy, which we summarize below, and present formally in [Figure 3](#); see also [Figure 4](#) for a visual illustration. Our main result in this subsection is as follows; as in the previous subsection, we prove it below after introducing the needed tools.

**Theorem 5.5 (Attack on the PoSW with original norm bounds).** *For  $\log q = o(n)$ , there is a depth- $\tilde{O}_{n,q}(\log^2 T)$  malicious prover that, given any  $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times m}$ ,  $\mathbf{u}_0 \in \mathbb{Z}_q^n$ , and positive integer  $T$ , convinces the verifier in the modified PoSW where in the folding operation the challenge is multiplied by the first half of the solution, and the original parameters  $\gamma' = \gamma$  and  $\beta = 1$  are used.*

Recall that our previous attack breaks the PoSW with parameters  $\gamma' = (\gamma + m)/2$  and  $\beta = m$  ([Theorem 5.1](#)). Our goal in the improved attack is to handle  $\gamma' = \gamma$  and  $\beta = 1$ . At a high level, we improve the parameter  $\beta$  by applying the “direct solution” technique from [Section 3.2.2](#), and we improve  $\gamma'$  by designing a more sophisticated block decomposition of  $\mathbf{A}_T$  (in a finer-grained version of the PoSW topology).

**Obtaining  $\beta = 1$ .** To see the idea behind our new strategy, consider the concrete example of (the core recursive part of) the PoSW for  $T = 15$ . In the first round, the prover sends the “middle” component  $\mathbf{x}_8$  to the verifier, which checks that  $\|\mathbf{x}_8\|_\infty \leq \beta$ . Our previous attack constructs the solution as  $\mathbf{x} = \mathbf{R}_T \cdot \mathbf{G}_n^{-1}(\mathbf{u}_0)$ , where  $\mathbf{R}_T$  is a  $n$ -subspace trapdoor for  $\mathbf{A}_T$  generated by RecTrap. As a result, our previous attack can ensure only that  $\|\mathbf{x}_8\|_\infty \leq \|\mathbf{R}_8\|_{\infty \leftarrow \infty}$  (where  $\mathbf{R}_8$  is the corresponding row block of  $\mathbf{R}_T$ ), and by construction, this is bounded by  $\|\mathbf{I}\|_{\infty \leftarrow \infty} \cdot m = m$ , which is why we took  $\beta = m$ .

The factor  $m$  above comes from (the general arity, subspace version of) [Lemma 3.2](#), but notice that in its “direct solution” counterpart [Lemma 3.4](#), the expansion factor is  $m^{1/\infty} = 1$  instead of  $m$ . So, if  $\mathbf{x}$  is constructed using CombSolve instead, we would have  $\|\mathbf{x}_8\|_\infty \leq 1$ , which allows us to take  $\beta = 1$ , at least for the first round. To get an attack that works for every round, we apply this idea recursively, solving for the first half of  $\mathbf{x}$  recursively, and *in parallel* computing suitable trapdoor(s) for the second-half block. We then use CombSolveHelper to construct an entire solution from these pieces. See the definition of RecSolve<sup>vec</sup> in [Figure 3](#) for the precise definition; also see [Figure 4](#) for an illustration, where the solid lines represent recursive calls to RecSolve<sup>vec</sup> (ignore the other types of lines for now).

**Obtaining  $\gamma' = \gamma$ .** Now let us resume the example. After the norm check in the first round, the prover receives a challenge  $c_1$  from the verifier. Then in the second round, the prover sends  $\mathbf{x}'_4 = c_1 \cdot \mathbf{x}_4 + \mathbf{x}_{12}$  to the verifier, which tests whether  $\|\mathbf{x}'_4\|_\infty \leq 2\gamma' \cdot \beta$ . Once the above “direct solution” idea is applied recursively, we can similarly get  $\|\mathbf{x}_4\|_\infty \leq 1$ . However, the  $\|\mathbf{x}_{12}\|_\infty$  term in the second half still picks up an  $m$  factor when the trapdoor for the second-half block is constructed using CombTrap. This results in no improvement to the parameter  $\gamma' = (\gamma + m)/2$ .

Our key new idea here is to partly “sequentialize” the solving within second half: instead of constructing a combined, longer trapdoor for the entire second-half block from trapdoors for its sub-blocks, we just directly use those trapdoors; see RecSolve<sup>list</sup> in [Figure 3](#). (Recall that CombSolveHelper works for an arbitrary block decomposition.) Specifically, the sub-blocks in the second-half block are the third-quarter block, a singleton block  $\mathbf{G}_n$ , and the last-quarter block. We then simply use this list of trapdoors in CombSolveHelper to get the final solution  $\mathbf{x}$ . This leads to  $\|\mathbf{x}_{12}\|_\infty \leq 1$ , thus allowing the use of  $\gamma' = \gamma$  in the second round.

Continuing the example, we see that the “ $(2^\tau + 1)/2^{\tau+1}$ -points” (e.g., the “3/4-point”  $\mathbf{x}_{12}$  we just considered, the “5/8-point”  $\mathbf{x}_{10}$ , etc.) are the only places where we need to apply the “sequentialization” idea in order to obtain  $\gamma' = \gamma$ . Trapdoors for other groups of sub-blocks can be safely combined into one trapdoor without violating the norm bounds; see RecSolve<sup>trap</sup> in [Figure 3](#). Importantly, therefore, the lengths of the lists of trapdoors remain linear in the recursion depth  $d$ , so the depths of the sequential solving steps

$\text{RecSolve}_s^{\text{vec}}(\mathbf{A}, \mathbf{u})$ <hr/>	
<b>if</b> $\mathbf{A} = \mathbf{G}_n$ : <b>return</b> $\text{CombSolve}_s(\mathbf{A}, \mathbf{u}, [\mathbf{I}_{m,s\ell}])$	
<i>in parallel:</i>	
$\mathbf{x}_0 = \text{RecSolve}_s^{\text{vec}}(\mathbf{A}_0, \mathbf{u}_0)$	
$L = \text{RecSolve}_s^{\text{list}}(\mathbf{A}_2)$	
<b>return</b> $\text{CombSolveHelper}_s(\mathbf{A}, \mathbf{x}_0, \mathbf{u}', [\mathbf{I}_{m,s\ell}; L])$	
$\text{RecSolve}_s^{\text{trap}}(\mathbf{A})$ <hr/>	$\text{RecSolve}_s^{\text{list}}(\mathbf{A})$ <hr/>
<b>if</b> $\mathbf{A} = \mathbf{G}_n$ : <b>return</b> $\mathbf{I}_{m,s\ell}$	<b>if</b> $\mathbf{A} = \mathbf{G}_n$ : <b>return</b> $[\mathbf{I}_{m,s\ell}]$
<i>in parallel:</i>	<i>in parallel:</i>
$\mathbf{R}_0 = \text{RecSolve}_s^{\text{trap}}(\mathbf{A}_0)$	$L = \text{RecSolve}_s^{\text{list}}(\mathbf{A}_0)$
$L = \text{RecSolve}_s^{\text{list}}(\mathbf{A}_2)$	$\mathbf{R}_2 = \text{RecSolve}_s^{\text{trap}}(\mathbf{A}_2)$
<b>return</b> $\text{CombTrap}_s(\mathbf{A}, [\mathbf{R}_0, \mathbf{I}_{m,s\ell}; L])$	<b>return</b> $[L; \mathbf{I}_{m,s\ell}, \mathbf{R}_2]$

Figure 3: Algorithm  $\text{RecSolve}_s^{\text{vec}}$  that solves  $\mathbf{Ax} = \mathbf{u}$  for any recursively  $s$ -admissible matrix  $\mathbf{A}$  with PoSW topology, and  $s$ -admissible syndrome  $\mathbf{u}$ . It parses  $\mathbf{u} = \begin{pmatrix} \mathbf{u}_0 \\ \mathbf{u}' \end{pmatrix}$  where  $\mathbf{u}_0 \in \mathbb{Z}_q^{N_0}$  for  $N_0 = (N - n)/2$ . Throughout,  $m = n\ell$ .

remain low. See Figure 4 for an illustration of the entire example, where the thick and thin dashed lines represent recursive calls to  $\text{RecSolve}_s^{\text{list}}$  and  $\text{RecSolve}_s^{\text{trap}}$ , respectively; note that these recursive calls alternate in a way that exactly traces the “ $(2^\tau + 1)/2^{\tau+1}$ -points.” Altogether, combining the “direct solution” and the “sequentialization” idea, we manage to get our new attack to work for the original parameters  $\gamma' = \gamma$  and  $\beta = 1$ .

**Lemma 5.6.** *Suppose that  $s\ell \leq \gamma^2$ . Let  $\mathbf{A}$  be a recursively  $s$ -admissible matrix with PoSW topology of depth  $d$  and base rows  $n$ , and let  $\mathbf{u}$  be an  $s$ -admissible vector. Then  $\text{RecSolve}_s^{\text{vec}}(\mathbf{A})$  computes, in depth  $\tilde{O}_{s,q}(d^2)$ , a solution to  $\mathbf{Ax} = \mathbf{u}$  satisfying the following bound for any  $p \in [1, \infty]$  and any integer  $\tau \in [0, d]$ :*

$$\text{mid}(\text{fold}_\gamma^{(\tau)}(\|\mathbf{x}\|_p)) \leq (2\gamma)^\tau \cdot (s\ell)^{1/p}.$$

*In particular, for  $p = \infty$ , we have that  $\text{mid}(\text{fold}_\gamma^{(\tau)}(\|\mathbf{x}\|_\infty)) \leq (2\gamma)^\tau$ .*

*Remark 5.7.* In Lemma 5.6 we assume that  $s\ell \leq \gamma^2$ . It is possible to relax the assumption to  $s\ell \leq \gamma^{2+a}$  for any integer  $a \geq 0$ , at the cost of increasing the depth of the computation from  $\tilde{O}_{s,q}(d^2)$  to  $\tilde{O}_{s,q}(d^{2+a})$ . Note that for any polynomial relationship between  $s\ell$  and  $\gamma$ , this  $a$  will be a constant and so the extra factor  $d^a$  in the computation depth will be polynomial in the depth of the PoSW topology.

The generalization is to generalize  $\text{RecSolve}_s^{\text{list}}$  to  $\text{RecSolve}_s^{\text{list}}[\alpha]$  (with default value  $\alpha = a$ ): it makes recursive calls to  $\text{RecSolve}_s^{\text{list}}$  and  $\text{RecSolve}_s^{\text{trap}}$  as usual for  $\alpha = 0$ , while for  $\alpha > 0$  it makes recursive calls to  $\text{RecSolve}_s^{\text{list}}[\alpha]$  and  $\text{RecSolve}_s^{\text{list}}[\alpha - 1]$  instead (and combines by simply concatenating the returned lists, still with an  $\mathbf{I}_{m,s\ell}$  in the middle).

*Proof.* Again we assume that  $T$  has the form  $T = 2^{d+1} - 1$ ; the case of general  $T$  can be handled as described in Remark 5.3.

We first bound the computation depth of  $\text{RecSolve}_s$ , in any of its three “modes” (vec, trap, list), when given a recursively  $s$ -admissible matrix having PoSW topology of depth  $d$ . It makes two *parallel* recursive calls




Figure 4: An example recursion tree for  $\text{RecSolve}^{\text{vec}}$  on a matrix with PoSW topology of depth  $d = 3$ . Recursive calls to  $\text{RecSolve}$  in the “modes”  $\text{vec}$ ,  $\text{list}$ ,  $\text{trap}$  are represented by solid lines, thick dashed lines, and thin dashed lines, respectively; singleton identity matrices are marked by thin dotted lines. The thickness of each (thick dashed) line for a  $\text{list}$ -mode call roughly depicts the number of returned trapdoors, and the width of each internal recursion node roughly depicts the amount of sequential work at that node. Each number  $i$  marks the vector/trapdoor corresponding to the block  $\mathbf{x}_i$  in the solution  $\mathbf{x}$ .

(with varying modes) on submatrices that follow the PoSW topology (with the same base case), so its total recursion depth is  $d$ . Then because  $\text{RecSolve}^{\text{trap}}$  returns a single trapdoor,  $\text{RecSolve}^{\text{list}}$  returns a list having  $O(d)$  trapdoors. So, by (the general arity, subspace versions of) [Lemmas 3.2](#) and [3.4](#), the computation depth of the non-recursive work in  $\text{RecSolve}$ —namely,  $\text{CombSolveHelper}$  or  $\text{CombTrap}$ —is  $\tilde{O}_{s,q}(d)$ , and hence the overall computation depth is  $\tilde{O}_{s,q}(d^2)$ .

For the (block-wise) norm of  $\mathbf{x}$ , we analyze more generally the norms of the outputs of  $\text{RecSolve}^{\{\text{vec}, \text{list}, \text{trap}\}}$  when given an arbitrary recursively  $s$ -admissible matrix with PoSW topology. For an (arbitrary) vector  $\mathbf{x}$  returned by  $\text{RecSolve}^{\text{vec}}$ , again by (the general arity version of) [Lemma 3.4](#), we have the recurrence

$$\|\mathbf{x}\|_p \leq \begin{pmatrix} \|\mathbf{x}_0\|_p \\ \|\mathbf{I}_{m,sl}\|_{p \leftarrow p} \cdot (sl)^{1/p} \\ \|\hat{\mathbf{R}}_2\|_{p \leftarrow p} \cdot (sl)^{1/p} \end{pmatrix} = \begin{pmatrix} \|\mathbf{x}_0\|_p & (sl)^{1/p} \\ & \|\hat{\mathbf{R}}_2\|_{p \leftarrow p} \cdot (sl)^{1/p} \end{pmatrix},$$

and the base case is  $\|\mathbf{x}\|_p = \|\mathbf{G}_s^{-1}(\star)\|_p \leq (sl)^{1/p}$ . Here  $\hat{\mathbf{R}}_2$  is the block matrix whose row blocks are the vertically stacked  $s$ -subspace trapdoors (which, to recall, all have  $sl$  columns) from the list  $L$  returned by  $\text{RecSolve}^{\text{list}}$ ; we similarly stack other outputs of  $\text{RecSolve}^{\text{list}}$  below. For a trapdoor stack  $\hat{\mathbf{R}}$  returned by  $\text{RecSolve}^{\text{list}}$ , straightforwardly,

$$\|\hat{\mathbf{R}}\|_{p \leftarrow p} = \begin{pmatrix} \|\hat{\mathbf{R}}_0\|_{p \leftarrow p} \\ \|\mathbf{I}_{m,sl}\|_{p \leftarrow p} \\ \|\mathbf{R}_2\|_{p \leftarrow p} \end{pmatrix} = \begin{pmatrix} \|\hat{\mathbf{R}}_0\|_{p \leftarrow p} \\ 1 \\ \|\mathbf{R}_2\|_{p \leftarrow p} \end{pmatrix};$$

and for a trapdoor  $\mathbf{R}$  returned by  $\text{RecSolve}^{\text{trap}}$ , again by (the general arity, subspace version of) [Lemma 3.2](#), we have that

$$\|\mathbf{R}\|_{p \leftarrow p} \leq \begin{pmatrix} \|\mathbf{R}_0\|_{p \leftarrow p} \\ \|\mathbf{I}_{m,sl}\|_{p \leftarrow p} \cdot sl \\ \|\hat{\mathbf{R}}_2\|_{p \leftarrow p} \cdot sl \end{pmatrix} = \begin{pmatrix} \|\mathbf{R}_0\|_{p \leftarrow p} & \\ & sl \end{pmatrix} \leq \begin{pmatrix} \|\mathbf{R}_0\|_{p \leftarrow p} & \\ & \gamma^2 \end{pmatrix},$$

where for the last inequality we use the hypothesis  $s\ell \leq \gamma^2$ . Both  $\|\hat{\mathbf{R}}\|_{p \leftarrow p}$  and  $\|\mathbf{R}\|_{p \leftarrow p}$  have the same base case  $\|\mathbf{I}_{m, s\ell}\|_{p \leftarrow p} = 1$ .

So, following the recurrence pattern, when the input to RecSolve has PoSW topology of depth  $d$ , we have that  $(\|\mathbf{R}\|_{p \leftarrow p} \ \|\hat{\mathbf{R}}\|_{p \leftarrow p} \ \|\mathbf{x}\|_p)$  is bounded from above by the matrix  $(\text{pf}_{\text{trap}}(d) \ \text{pf}_{\text{list}}(d) \ \text{pf}_{\text{vec}}(d))$ , which is recursively defined as follows:

$$(\text{pf}_{\text{trap}}(d) \ \text{pf}_{\text{list}}(d) \ \text{pf}_{\text{vec}}(d)) := \begin{pmatrix} \text{pf}_{\text{trap}}(d-1) & \text{pf}_{\text{list}}(d-1) & \text{pf}_{\text{vec}}(d-1) \\ & \gamma^2 & 1 \\ \text{pf}_{\text{list}}(d-1) \cdot \gamma^2 & \text{pf}_{\text{trap}}(d-1) & \text{pf}_{\text{list}}(d-1) \cdot (s\ell)^{1/p} \end{pmatrix}.$$

For convenience, we denote  $\text{profile}(d) := (\text{pf}_{\text{trap}}(d) \ \text{pf}_{\text{list}}(d) \ \text{pf}_{\text{vec}}(d))$ . The base case of the recursion is  $\text{profile}(0) = (1 \ 1 \ (s\ell)^{1/p})$ .

From these definitions it can be verified that

$$\begin{aligned} \text{fold}_\gamma(\text{profile}(d)) &= (\gamma \mathbf{I} \ \mathbf{0} \ \mathbf{I}) \cdot \begin{pmatrix} \text{pf}_{\text{trap}}(d-1) & \text{pf}_{\text{list}}(d-1) & \text{pf}_{\text{vec}}(d-1) \\ & \gamma^2 & 1 \\ \text{pf}_{\text{list}}(d-1) \cdot \gamma^2 & \text{pf}_{\text{trap}}(d-1) & \text{pf}_{\text{list}}(d-1) \cdot (s\ell)^{1/p} \end{pmatrix} \\ &= \text{profile}(d-1) \cdot \begin{pmatrix} \gamma & 1 & 0 \\ \gamma^2 & \gamma & (s\ell)^{1/p} \\ 0 & 0 & \gamma \end{pmatrix}. \end{aligned}$$

Hence by iterating, for any  $\tau$ , we get that

$$\text{fold}_\gamma^{(\tau)}(\text{profile}(d)) = \text{profile}(d-\tau) \cdot \begin{pmatrix} \gamma & 1 & 0 \\ \gamma^2 & \gamma & (s\ell)^{1/p} \\ 0 & 0 & \gamma \end{pmatrix}^\tau.$$

It can be verified by induction that this matrix power expands to

$$\begin{pmatrix} \gamma & 1 & 0 \\ \gamma^2 & \gamma & (s\ell)^{1/p} \\ 0 & 0 & \gamma \end{pmatrix}^\tau = \begin{pmatrix} (2\gamma)^{\tau-1} \cdot \gamma & (2\gamma)^{\tau-1} & (2^{\tau-1} - 1) \cdot \gamma^{\tau-2} \cdot (s\ell)^{1/p} \\ (2\gamma)^{\tau-1} \cdot \gamma^2 & (2\gamma)^{\tau-1} \cdot \gamma & (2\gamma)^{\tau-1} \cdot (s\ell)^{1/p} \\ 0 & 0 & \gamma^\tau \end{pmatrix}.$$

Then we get, as desired,

$$\begin{aligned} \text{mid}(\text{fold}_\gamma^{(\tau)}(\|\mathbf{x}\|_p)) &\leq \text{mid}(\text{fold}_\gamma^{(\tau)}(\text{pf}_{\text{vec}}(d))) \\ &= \text{mid}(\text{profile}(d-\tau)) \cdot \begin{pmatrix} (2^{\tau-1} - 1) \cdot \gamma^{\tau-2} \cdot (s\ell)^{1/p} \\ (2\gamma)^{\tau-1} \cdot (s\ell)^{1/p} \\ \gamma^\tau \end{pmatrix} \\ &\leq (\gamma^2 \ 1 \ (s\ell)^{1/p}) \cdot \begin{pmatrix} (2^{\tau-1} - 1) \cdot \gamma^{\tau-2} \cdot (s\ell)^{1/p} \\ (2\gamma)^{\tau-1} \cdot (s\ell)^{1/p} \\ \gamma^\tau \end{pmatrix} \\ &= 2^{\tau-1}(\gamma^\tau + \gamma^{\tau-1}) \cdot (s\ell)^{1/p} \\ &\leq (2\gamma)^\tau \cdot (s\ell)^{1/p}. \end{aligned}$$

Here note that  $\text{mid}(\text{pf}(d-\tau)) = (1 \ 1 \ (s\ell)^{1/p})$  in the base case  $\tau = d$ , but  $(\gamma^2 \ 1 \ (s\ell)^{1/p})$  is still a valid upper bound.  $\square$



*Proof of Theorem 5.5.* Because  $\log q = o(n)$ , we have that  $m = n \log q = o(n^2) = o(\gamma^2)$ , and thus  $m \leq \gamma^2$  holds (for all sufficiently large  $n$ ). Recall that the matrix  $\mathbf{A}_T$  from Equation (2.3) is recursively  $n$ -admissible with PoSW topology of depth  $d = \lfloor \log_2 T \rfloor$ . The malicious prover computes a solution  $\mathbf{x}$  to Equation (2.3) following Lemma 5.6, and then proceeds in the same way as the honest prover. By Lemma 5.6, this prover has depth  $\tilde{O}_{n,q}(\log^2 T)$ .

It remains to confirm that this will satisfy the verifier’s norm checks, and in particular to analyze the  $\gamma$ -folding of the block-wise norms of the solution  $\mathbf{x}$ . Similar to the proof of Theorem 5.1, by Lemma 5.6 again, for all  $\tau \in [0, d]$  the block-wise norm of the solution  $\mathbf{x}$  satisfies

$$\text{mid}(\text{fold}_\gamma^{(\tau)}(\|\mathbf{x}\|_\infty)) \leq (2\gamma)^\tau \cdot 1,$$

so the verifier’s norm checks are indeed satisfied with bound  $(2\gamma')^\tau \cdot \beta$  for parameters  $\gamma' = \gamma$  and  $\beta = 1$ .  $\square$

## References

- [Ajt96] M. Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996. Pages 2 and 14.
- [AKK<sup>+</sup>19] H. Abusalah, C. Kamath, K. Klein, K. Pietrzak, and M. Walter. Reversible proofs of sequential work. In *EUROCRYPT*, pages 277–291. 2019. Page 1.
- [CLSY93] J. Cai, R. J. Lipton, R. Sedgewick, and A. C. Yao. Towards uncheatable benchmarks. In *Structure in Complexity Theory Conference*, pages 2–11. 1993. Page 1.
- [CP18] B. Cohen and K. Pietrzak. Simple proofs of sequential work. In *EUROCRYPT*, pages 451–467. 2018. Pages 1 and 2.
- [DN92] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In *CRYPTO*, volume 740, pages 139–147. 1992. Page 2.
- [LM23] R. W. F. Lai and G. Malavolta. Lattice-based timed cryptography. In *CRYPTO*, pages 782–804. 2023. Pages 2, 3, 4, 5, 6, 7, 8, 10, 11, 12, 13, 14, 15, and 16.
- [May93] T. C. May. Timed-release crypto, February 1993. <http://cypherpunks.venona.com/date/1993/02/msg00129.html>. Page 1.
- [MMV13] M. Mahmoody, T. Moran, and S. P. Vadhan. Publicly verifiable proofs of sequential work. In *Innovations in Theoretical Computer Science (ITCS)*, pages 373–388. 2013. Pages 1 and 2.
- [MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718. 2012. Pages 4 and 5.
- [MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004. Page 14.
- [Pie19] K. Pietrzak. Simple verifiable delay functions. In *Innovations in Theoretical Computer Science Conference (ITCS)*, volume 124 of *LIPICs*, pages 60:1–60:15. 2019. Page 2.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005. Page 2.

- [RSW96] R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, Massachusetts Institute of Technology, Cambridge, MA, USA, 1996. Pages 1 and 2.
- [Sho94] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997. Preliminary version in FOCS 1994. Page 2.
- [Ver12] R. Vershynin. *Introduction to the non-asymptotic analysis of random matrices*, chapter 5, pages 210–268. Cambridge University Press, 2012. Available at <http://www-personal.umich.edu/~romanv/papers/non-asymptotic-rmt-plain.pdf>. Page 4.
- [Wes19] B. Wesolowski. Efficient verifiable delay functions. *J. Cryptol.*, 33(4):2113–2147, 2020. Preliminary version in EUROCRYPT 2019. Page 2.