

# Attacking trapdoors from matrix products

Thomas Decru<sup>1</sup>, Tako Boris Fouotsa<sup>2</sup>, Paul Frixons<sup>1</sup>, Valerie Gilchrist<sup>1</sup>,  
Christophe Petit<sup>1,3</sup>

<sup>1</sup> Université Libre de Bruxelles, Brussels, Belgium

<sup>2</sup> EPFL, Lausanne, Switzerland

<sup>3</sup> University of Birmingham, Birmingham, United Kingdom

**Abstract.** Recently, Geraud-Stewart and Naccache proposed two trapdoors based on matrix products. In this paper, we answer the call for cryptanalysis. We explore how using the trace and determinant of a matrix can be used to attack their constructions. We fully break their first construction in a polynomial-time attack. We show an information leak in the second construction using characteristic polynomials, and provide an attack using traces that decreases the bit security by about half.

**Keywords:** matrix, cryptanalysis, post-quantum

## 1 Introduction

In 2023, Geraud-Stewart and Naccache proposed a new trapdoor based on matrix products [3]. From a set of public invertible matrices  $A_1, \dots, A_k$ , it sends the permutation  $\sigma$  to the product  $\prod_{i=1}^k A_{\sigma(i)}$ . Even if the practical encryption does not become competitive to already known procedures, the prospect of a new family of trapdoors is enriching for cryptography in general as it could lead to new applications. The simplicity of the scheme and the concepts are also attractive in itself when thinking about widespread implementations.

That being said, trust can only be given to a scheme once it has received careful inspection and resisted the various attempts of breaking from the cryptographic community. In the last year, no work assessing the security of these schemes has been released. This paper is inscribed in this procedure, and in fact shows that the constructions, as they are, do not meet their claimed security levels.

*Related work.* These trapdoors can be seen as an instance of Dehn’s “word search” problem. That is, from a group  $G$ , a set of generators  $g_1, \dots, g_k$  and a target  $t$ , determine the sequence  $\{m_i\}$  given a writing  $t = \prod g_{m_i}$ . As such it

---

\* Authors listed in alphabetical order: see <https://www.ams.org/profession/leaders/CultureStatement04.pdf>. Valerie Gilchrist and Thomas Decru are supported by grants from the National Fund for Scientific Research (F.N.R.S.) of Belgium; Christophe Petit is partly supported by EPSRC through grant number EP/V011324/1.

Date of this document: 2024-08-26.

is linked to other cryptosystems based on similar problems. The first one is a trapdoor based directly on Dehn’s problem from Wagner and Magyarik [7], but broken later by Levy-dit-Vehel and Perret [6]. Another one is the family of Cayley hashes initiated by Tillich and Zémor [5] and Charles, Goren, and Lauter [2].

*Organization of the paper.* In Section 2, we recall the relevant concepts and the trapdoor constructions. In Sections 3 and 4, we investigate some properties of dwarf matrices from the trapdoor constructions. Concrete attacks are proposed in Sections 5 and 6, and we conclude with some future work in Section 7.

## 2 Background

We review some of the key concepts put forward in [3] for using matrix products as trapdoors. For the rest of the paper we let  $S_k$  denote the set of permutations on the integers  $\{1, \dots, k\}$ .

Consider some  $\sigma \in S_k$ . Given some set of  $k$  matrices,  $\mathbf{A} = \{A_1, \dots, A_k\} \subset GL_n(\mathbb{F}_p)$ , we can map  $\sigma$  to a matrix product as

$$\sigma \mathbf{A} := \prod_{i=1}^k A_{\sigma(i)}.$$

By choosing the parameters  $k, n, p$  carefully, and sampling  $\mathbf{A}$  uniformly at random, we can ensure that with high probability this map is injective. That is, that the products we are mapping into are distinct for each distinct  $\sigma$ .

In order to build cryptography on top of this mapping, we would hope for the inverse computation to be cryptographically hard. To this end, the authors of [3] first provide a very specific instance where the inversion of this map is easy, and then take advantage of these easy instances to build a trapdoor. To explore their construction, we first establish a way of determining the “size” of a matrix. To do so we define a partial order relation on the set of matrices.

**Definition 1.** Let  $M := \{m_{i,j}\}_{i,j=1}^n, M' := \{m'_{i,j}\}_{i,j=1}^n$  be two matrices in  $GL_n(\mathbb{F}_p)$ . We say  $M < M'$  if for all  $i, j \in [1, \dots, n]$ , we have that  $m_{i,j} < m'_{i,j}$  when expressed as integers in  $[0, p-1]$ .

We will use the notation  $M < \alpha$ , to mean that every entry in  $M$ , when expressed as an integer in  $[0, p-1]$ , is less than  $\alpha$ .

Now, if the entries of the matrices of  $\mathbf{A}$  are very small, then we can expect the entries of  $\sigma \mathbf{A}$  to also be relatively small. In particular, we could choose our prime  $p$  so that when considering the product over the integers, it holds that  $\sigma \mathbf{A} < p$ . Thus we can work with these products as if they were over the integers.

In particular, barring few exceptions such as the identity matrix, we get that as we multiply more matrices together, the entries are strictly increasing. This means we expect that

$$A_{\sigma(1)} \cdots A_{\sigma(k)} > A_{\sigma(1)} \cdots A_{\sigma(k)} \cdot A_{\sigma(k)}^{-1} = A_{\sigma(1)} \cdots A_{\sigma(k-1)}.$$

This provides us with a method of recovering the permutation from the product. We give the details of this in Algorithm 1, which we call `Decompose`, where we assume each matrix  $A_i \in \mathbf{A}$  is such that  $A_i \leq \alpha$ , and  $p$  is chosen such that  $n^{k-1}\alpha^k < p$ . This will ensure that we are working over the integers.

---

**Algorithm 1: Decompose**


---

```

Input :  $\mathbf{A} \subset \mathcal{D}$ ,  $C = \sigma\mathbf{A}$ 
Output:  $\sigma$ 

1 if  $|\mathbf{A}| = 0$  /* if  $\mathbf{A}$  is an empty list */
2 then
3      $\perp$  return  $\perp$ 
4 if  $|\mathbf{A}| = 1$  and  $C \notin \mathbf{A}$  then
5      $\perp$  return  $\perp$ , else return 1
6 for  $A \in \mathbf{A}$  do
7      $M \leftarrow C \cdot A^{-1}$ 
8     if  $M < C$  then
9          $\mathbf{L} \leftarrow \text{Decompose}(M, \mathbf{A} \setminus \{A\})$ 
10        if  $\mathbf{L} = \perp$  then
11             $\perp$  return  $\perp$ 
12        else
13             $\mathbf{L} = \mathbf{L} \parallel [A]$ 
14            return  $\sigma$  that maps  $\mathbf{A}$  to  $\mathbf{L}$  (as ordered lists)
15 return  $\perp$ 

```

---

Notice that for `Decompose` (Algorithm 1) to work, *all* of the matrices in  $\mathbf{A}$  had to be bounded in size. Thus, we will make a distinction between these special matrices and the rest of  $GL_n(\mathbb{F}_p)$ . Note, we make a point to exclude the center of  $GL_n(\mathbb{F}_p)$ , denoted  $Z(GL_n(\mathbb{F}_p))$ , as they will not be usable in the trapdoors later on. In particular,  $Z(GL_n(\mathbb{F}_p)) = \{sI : s \in \mathbb{F}_p\}$ .

$$\mathfrak{E} := GL_n(\mathbb{F}_p) \setminus Z(GL_n(\mathbb{F}_p)),$$

$$\mathfrak{D} := \{M \in \mathfrak{E} : M \leq \alpha\}.$$

Following the terminology introduced in [3], we refer to these sets as *elves* and *dwarves* respectively. In summary, while `Decompose` can be run on a set of dwarf matrices, the work of [3] conjectures that for random elf matrices it is not easy to recover the same information. Specifically, they claim that the following problem is hard for sufficiently large  $n, k, p$ .

*Problem 1.* Let  $\mathbf{A} \subset \mathfrak{E}$  be a uniformly sampled set of elf matrices. Given  $\sigma\mathbf{A}$  over  $\mathbb{F}_p$ , compute  $\sigma$ .

This leads [3] to two trapdoor constructions, a *direct construction* and an *alternating construction*. The core idea to both constructions being that we will define secret maps sending a set of dwarf matrices to a set of elf matrices. By making the elf matrices public, anyone can create a product and publish it, where adversaries will not be able to recover the permutation. With knowledge of the secret mapping, however, we will be able to send the product of elves to a product of dwarves, where we can apply Decompose to recover the permutation.

## 2.1 Direct construction

Let  $\mathbf{A} := \{A_1, \dots, A_k\} \subset \mathfrak{D}$  be a set of (secret) dwarf matrices. Choose a (secret) elf matrix,  $E \in \mathfrak{E}$ . We will use  $E$  to mask  $\mathbf{A}$  as

$$\bar{\mathbf{A}} := \{\bar{A}_i := EA_iE^{-1}\}_{i=1}^k.$$

The set  $\bar{\mathbf{A}}$  will serve as the public key. Now an external party can choose some permutation  $\sigma \in S_k$ , and compute the ciphertext

$$C := \sigma \bar{\mathbf{A}} = E(\sigma \mathbf{A})E^{-1}.$$

We can map this ciphertext back to a product of dwarf matrices since

$$E^{-1}CE = \sigma \mathbf{A}.$$

From which we can apply the Decompose algorithm to recover  $\sigma$ .

*Parameters.* The authors from [3] give parameters for varying security levels, which they refer to as “toy”, “challenge”, “recommended”, and “large”. We list them in the table below, where  $\lambda$  is their estimated security level in bits.

	$\lambda$	$k$	$n$	$\alpha$	$p$
toy	16	9	4	2	$2^{53} + 5$
challenge	64	21	8	2	$2^{167} + 83$
recommended	128	35	10	2	$2^{302} + 307$
large	512	99	24	2	$2^{1105} - 1335$

Table 1: Parameters from [3] for the direct construction.

*Extra masking.* In the original work of [3] the authors suggest that including an extra dwarf matrix in the masking could improve security, though they do not provide concrete reasons for using it. In this variant they define their public key matrices as  $\bar{A}_i := EA_iDE^{-1}$ , where  $D \in \mathfrak{D}$ . For simplicity of exposition, we exclude this masking matrix in the proceeding sections, but will justify why including a non-trivial  $D$  does not avoid the attack later on.

## 2.2 Alternating construction

Let  $\mathbf{A}^b := \{A_1^b, \dots, A_k^b\} \subset \mathfrak{D}$ , be two sets of (secret) dwarf matrices, one for each bit  $b \in \{0, 1\}$ . Choose a set of (secret) elf matrices,  $\{E_i\}_{i=0}^k \in \mathfrak{E}$ . Define

$$\bar{\mathbf{A}}^b := \{\bar{A}_i^b := E_{i-1} A_i^b E_i^{-1}\}_{i=1}^k.$$

The sets  $\bar{\mathbf{A}}^b$  will serve as the public keys. Now for a binary string  $m \in \{0, 1\}^k$ , we can compute the ciphertext as

$$C = \prod_{i=1}^k \bar{A}_i^{m_i} = E_0 \left( \prod_{i=1}^k A_i^{m_i} \right) E_k^{-1}.$$

We will refer to this computation with the notation  $m\bar{\mathbf{A}}$  for ease of notation. Thus to map  $C$  to a product of dwarf matrices we compute  $E_0^{-1} C E_k$ , and then apply the Decompose algorithm to recover  $m$ .

*Parameters.* Similarly to the direct construction, the authors from [3] give toy, challenge, recommended, and large parameters for the alternating construction, where  $\lambda$  is the estimated security level in bits. We list them in the table below.

	$\lambda$	$k$	$n$	$\alpha$	$p$
toy	16	16	4	2	$2^{47} + 5$
challenge	64	64	8	2	$2^{255} - 19$
recommended	128	128	10	2	$2^{553} + 549$
large	512	512	24	2	$2^{2859} + 641$

Table 2: Parameters from [3] for the alternating construction.

## 3 Properties of dwarf products

We will explore some properties of dwarf matrices and dwarf matrix products. We continue to use the notation outlined in Section 2.

### 3.1 Dwarf determinants

We begin by considering some properties of the determinants of dwarf matrices.

**Lemma 1.** *For a dwarf matrix  $A$  it holds that  $|\det A| \leq (\alpha\sqrt{n})^n$ , when  $\det A$  is seen as an integer in  $(-p/2, p/2)$ .*

*Proof.* This follows from Hadamard’s inequality.

In Lemma 2, we will get concrete upper bounds on the determinants of dwarf matrices, but only for the smallest values of  $\alpha$ , namely,  $\alpha = 1$  or  $2$ . This is convenient since these choices of  $\alpha$  coincide with the suggested parameters.

**Lemma 2.** *For an  $n \times n$  dwarf matrix  $A$  with  $\alpha \in \{1, 2\}$  and  $n$  even, the maximal value of  $\det A$  is the maximal value of the determinant of the  $n \times n$  Hadamard matrices with entries  $\{0, 1\}$ , respectively  $\{-1, 1\}$ .*

*Proof.* For  $\alpha = 1$  this is immediate. For  $\alpha = 2$ , it suffices to see that in both cases we are computing the volume of the maximal  $n$ -dimensional hypercube with lattice points at most distance two away.

Since not all dwarves are Hadamard matrices, the determinants of the dwarves will be much smaller than these maximal values, as can be seen from the experiments in [4]. Studying them, however, gives us an effective upper bound on dwarf determinants. The work of Tao and Vu [4, Theorem 1.1] shows that with probability tending to 1 (as  $n$  tends to infinity), the absolute value of the determinant of a random  $\{1, -1\}$  Hadamard matrix is close to  $\sqrt{n!}$ . For  $n = 8$ , this means the determinant is expected to be close to  $2^8$ .

For  $\alpha = 2$  and  $n$  ranging from 1 to 10 we then get that  $|\det A|$  is maximally equal to

$$2, 4, 16, 48, 160, 576, 4096, 14336, 73728, 327680.$$

There is no known closed-form expression for these values, but the first 22 can be found as A003433 in The On-Line Encyclopedia of Integer Sequences [1]. Some of the larger values of determinants cannot be attained; e.g. for  $n = 3$  we cannot construct a matrix with determinant  $\pm 13, \pm 14, \pm 15$ . The maximum values are typically achieved by incorporating a lot of structure. The following is an example for  $n = 8$  with maximal determinant:

$$\begin{pmatrix} 2 & 0 & 2 & 0 & 0 & 2 & 2 & 0 \\ 2 & 2 & 0 & 2 & 0 & 0 & 2 & 2 \\ 2 & 2 & 2 & 0 & 2 & 0 & 0 & 2 \\ 0 & 2 & 2 & 2 & 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 2 & 2 & 0 & 2 & 0 \\ 2 & 0 & 0 & 2 & 2 & 2 & 0 & 2 \\ 0 & 2 & 0 & 0 & 2 & 2 & 2 & 0 \\ 0 & 0 & 2 & 0 & 0 & 2 & 2 & 2 \end{pmatrix}.$$

### 3.2 Dwarf traces

The trace of a matrix is invariant under conjugation, but is not multiplicative. Thus  $\text{Tr}(E\sigma\mathbf{A}E^{-1}) = \text{Tr}(\sigma\mathbf{A})$ , and will change depending on the choice of  $\sigma$ . For dwarf products in general we expect that as more dwarves are multiplied to it, the trace should strictly increase. This will provide us with a test to use in our attack later on. We formalize this idea in Heuristic 1 and give both heuristic arguments and thorough experimental evidence to support it.

**Heuristic 1.** Choose parameters  $(n, k, p, \alpha)$  as described in Section 2.1.

Let  $\{A_1, \dots, A_k\} \subset \mathfrak{D}$ ,  $\sigma \in S_k$ , and  $m \leq k$ .

Then we have that

$$\text{tr} \left( \prod_{i=1}^{m-1} A_{\sigma(i)} \right) < \text{tr} \left( \prod_{i=1}^m A_{\sigma(i)} \right)$$

with overwhelming probability.

Recall that every entry in a dwarf matrix is at most  $\alpha$ . By considering a set of matrices, where each entry is exactly the maximum value  $\alpha$ , we can upper bound the trace of a product of  $m$  dwarf matrices by  $n^m \alpha^m$ . Note that the prime is chosen to be larger than  $n^{k-1} \alpha^k$ . However, this upper bound is significantly larger than what occurs in practice since the bound is computed from considering a matrix with only entries equal to  $\alpha$ , which is not itself an invertible matrix. So in practice, we expect the trace of the product to always be less than  $p$ , meaning we are working strictly over the integers.

Now, let  $D_\alpha$  be the distribution of a variable uniformly randomly sampled in  $\{0, \dots, \alpha\}$ . Though dwarf matrices are required to be invertible and so there may be some bias introduced, we estimate that this bias does not have a significant impact on the results of this analysis, and we will support this claim with experimental evidence later on. Thus, we assume the elements of a dwarf matrix are sampled according to  $D_\alpha$ . Their expected value is

$$\mu_\alpha := \frac{1}{\alpha + 1} \sum_{i=0}^{\alpha} i = \alpha/2.$$

Their variance is

$$\sigma_\alpha^2 = \frac{1}{\alpha + 1} \left( \sum_{i=0}^{\alpha} i^2 \right) - \mu_\alpha^2 = \frac{\alpha(2\alpha + 1)}{6} - (\alpha/2)^2 = \frac{\alpha(\alpha + 2)}{12}.$$

We will use the following fact. Let  $X$  and  $Y$  be two random variables, and let  $Z$  be their product. Assume  $X$  and  $Y$  are independent. Then we have

$$\mu_Z = \mu_X \mu_Y \quad \text{and} \quad \sigma_Z^2 = (\sigma_X^2 + \mu_X^2)(\sigma_Y^2 + \mu_Y^2) - \mu_X^2 \mu_Y^2.$$

Entries in the product of two dwarf matrices can now be estimated as follows. Each entry is the sum of  $n$  terms, where each term is the product of two independent variables, each one distributed as  $D_\alpha$ . Their average value is then

$$\mu_\alpha^{(2)} = n \mu_\alpha^2 = n \frac{\alpha^2}{4},$$

and their variance is

$$(\sigma_\alpha^{(2)})^2 = n [(\sigma_\alpha^2 + \mu_\alpha^2)(\sigma_\alpha^2 + \mu_\alpha^2) - \mu_\alpha^2 \mu_\alpha^2] = n \alpha^2 \frac{7\alpha^2 + 16\alpha + 4}{144}.$$

Entries in the product of three or more dwarf matrices are harder to estimate rigorously. We can try to iterate the previous argument, by heuristically ignoring correlations between elements of partial products. We then obtain formulae for the average as

$$\mu_\alpha^{(k)} = n\mu_\alpha^{(k-1)}\mu_\alpha = n^{k-1}\mu_\alpha^k \quad (1)$$

and for the variances as

$$\begin{aligned} (\sigma_\alpha^{(k)})^2 &= n \left[ ((\sigma_\alpha^{(k-1)})^2 + (\mu_\alpha^{(k-1)})^2)(\sigma_\alpha^2 + \mu_\alpha^2) - (\mu_\alpha^{(k-1)})^2 \mu_\alpha^2 \right] \\ &= n \left[ (\sigma_\alpha^{(k-1)})^2 (\sigma_\alpha^2 + \mu_\alpha^2) + (\mu_\alpha^{(k-1)})^2 \sigma_\alpha^2 \right]. \end{aligned} \quad (2)$$

Consider the probability that the trace of a product of  $k+1$  dwarves is smaller than the trace of the product of the first  $k$  factors. In other words, letting  $P$  be the product of  $k$  dwarf matrices, and  $A$  a dwarf matrix, we are interested in when  $\text{Tr}(PA) - \text{Tr}(P) = \text{Tr}((P(A - I)))$  is negative. Ignoring correlations, the expected value of the trace difference can be approximated as follows

$$\mu := \text{Tr}(PA) - \text{Tr}(P) = n(\mu_\alpha^{(k+1)} - \mu_\alpha^{(k)}) = n\mu_\alpha^{(k)}(n\alpha/2 - 1).$$

This average is of course positive, but we also need the variance to argue about the probability to obtain a negative value. We estimate this as follows

$$\sigma^2 = n((\sigma_\alpha^{(k+1)})^2 + (\sigma_\alpha^{(k)})^2).$$

We can then bound the probability that the trace difference is negative using Tchebychev inequality

$$\Pr[\text{Tr}(PA) - \text{Tr}(P) \leq 0] \leq \left( \frac{\sigma}{\mu} \right)^2. \quad (3)$$

This implies that  $\epsilon(\alpha, n, k) := 1 - (\sigma/\mu)^2$  is a lower bound on the probability that  $\text{Tr}(PA) > \text{Tr}(P)$  holds. For each set  $(n, k, \alpha)$  of the parameters from [3] for the direct construction, we computed  $\epsilon(\alpha, n, k')$  where  $1 \leq k' \leq k$ . This data is given in Figure 1. When verifying these figures experimentally, we checked Heuristic 1 for  $k = 2, 3$  using  $10^6$  samples for each of the toy, challenge, and recommended parameters, and found that it was true 100% of the time. Thus, these results support the claim from Heuristic 1.

*Experiments.* In Figure 2, we computed the average trace of a product of  $i$  dwarf matrices, for  $i \in [1, k]$ . The averages were taken from 10,000 samples for several parameter sets. As can be seen, the traces follow a logarithmic line, further supporting the claim in Heuristic 1. The coloured bars indicate the standard deviation.

The experimental results shown in Figures 3 and 4 show how varying the values of  $n$  and  $\alpha$  will affect the trace. In general, increasing these values will increase the average trace of the dwarf products. In Figure 3 the slopes of the lines change by a constant value close to 1, with values close to 1,2, and 3. This



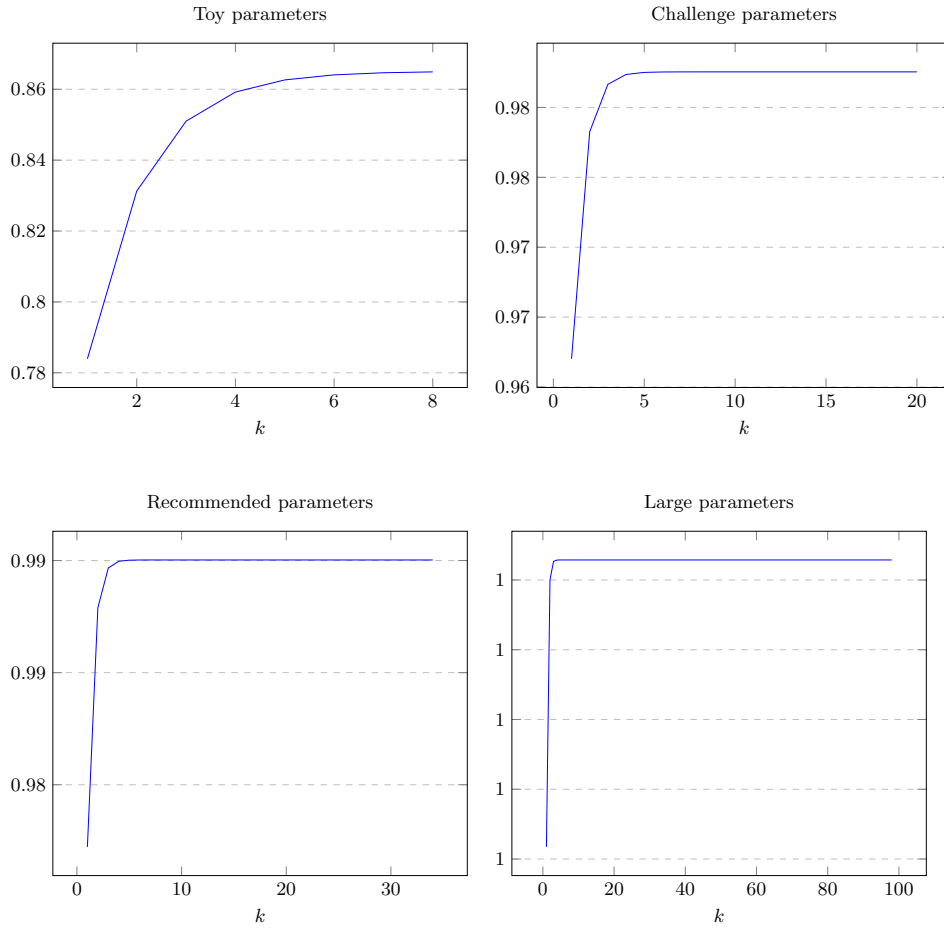


Fig. 1: We plot the probabilities that the trace differences are non-negative according to Equation 3 for each set of parameters from Section 2.1.

means that increasing  $n$  will linearly increase the average trace, as would be expected. In Figure 4 we see that the lines seem to be converging to each other as  $\alpha$  increases, with the most dramatic difference between  $\alpha = 1$  and  $\alpha = 2$ . This signifies that we get the most dramatic differences in trace for small  $\alpha$  values. We also computed these values from 10,000 samples, which were generated using the challenge parameters.

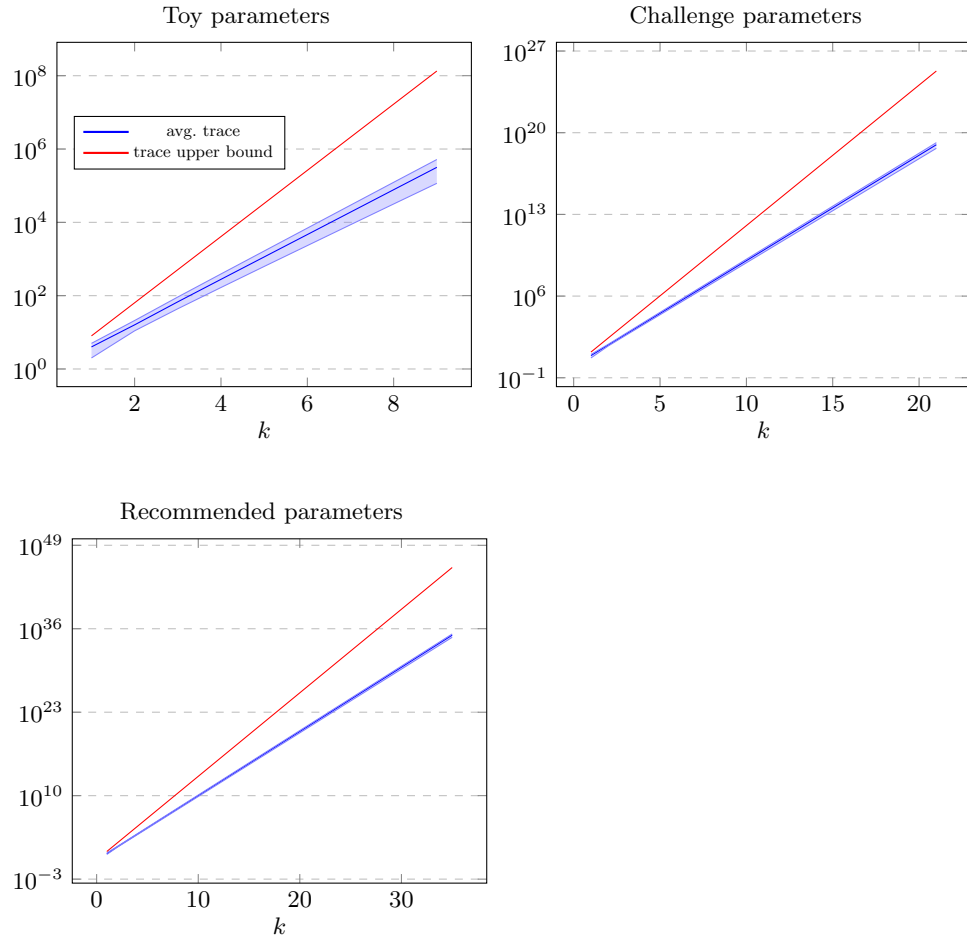


Fig. 2: We plot the average trace values for products of dwarves of various parameter sets together with the maximum possible trace value.

## 4 Trace of inverse dwarf products

### 4.1 Dwarf products and inverse dwarves

Now instead of looking at how the trace changes in dwarf products, we consider how it might change when multiplying by a non-dwarf matrix in the following heuristic. Note that the inverse of a dwarf is not expected to be a dwarf.

**Heuristic 2.** Let  $C = \prod_{i=1}^m A_{\sigma(i)}$  be a product of dwarf matrices. The expected value of  $\text{tr}(C)$  is less than the expected value of  $\text{tr}(A_{\sigma(j)}^{-1}C)$  where  $j \neq 1, m$ .




Fig. 3: Average trace value for dwarf products of varying  $k$  values for the challenge parameters. The different lines indicate different  $n$  values.

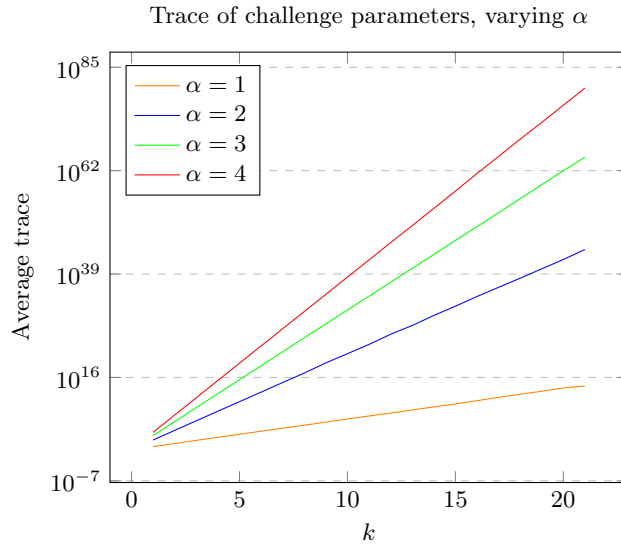


Fig. 4: Average trace value for dwarf products of varying  $k$  values for the challenge parameters. The different lines indicate different  $\alpha$  values.

In our experiments, after running 10,000 samples, we found that for the recommended parameters from the direct construction, Heuristic 2 was true 71% of the time. For the large parameter set this probability increased to 99%.

## 4.2 Estimating entries of inverse dwarves

Let  $D$  be an  $n \times n$  dwarf matrix, then  $D^{-1} = 1/\det(D) \cdot \text{Adj}(D)$ , where  $\text{Adj}(D)$  denotes the adjugate of  $D$ . Recall that the entries of  $\text{Adj}(D)$  are minors of  $D$  i.e. determinants of the  $(n-1) \times (n-1)$  submatrices of  $D$ . Suppose the coefficients of  $D$  are distributed with mean  $\mu$  and variance  $\sigma^2$ , we now study the expected distribution of the minors of  $D$ . A minor of size  $r$  is a sum (with signs) of  $r!$  terms, with each term a product of  $r$  coefficients in  $D$ . If these behave as independent, we then expect

$$\mu_r = \begin{cases} 0 & \text{if } r \text{ is even} \\ \mu^r & \text{if } r \text{ is odd} \end{cases}, \quad \sigma_r^2 = (r!)((\sigma^2 + \mu^2)^r - \mu^{2r}).$$

In particular, we expect the determinant to be a random variable with average and variance  $\mu_n$  and  $\sigma_n^2$ . Coefficients of an inverse dwarf matrix times its determinant should have average  $\mu_{n-1}$  and variance  $\sigma_{n-1}^2$ .

This analysis indicates that though the inverse of a dwarf matrix over  $\mathbb{F}_p$  is not itself a dwarf, they still behave differently from a randomly sampled elf matrix. For example, we know that when the entries of  $\text{Adj}(D)$  are considered as integers in  $[0, p-1]$ , they will not be small. When they are considered as integers in  $[-(p-1)/2, (p-1)/2]$ , however, they are in fact close to zero. In what follows, we will be looking at the *absolute value* of these matrix entries, where we consider the entry in  $[-(p-1)/2, (p-1)/2]$ , and multiply it by  $-1$  if it is negative to ensure a positive value. Thus, we claim that the absolute value of the trace of  $\text{Adj}(D)$  will also be small in general. We compute an upper bound for it in Lemma 3.

**Lemma 3.** *Let  $D \in \mathfrak{D}$ , so  $D < \alpha$  and has dimensions  $n \times n$ . Then  $|\text{Tr}(\text{Adj}(D))| \leq n\alpha^{n-1}(n-1)^{(n-1)/2}$ .*

*Proof.* The elements of  $\text{Adj}(D)$  are determinants of the  $(n-1) \times (n-1)$  submatrices of  $D$ . Thus,  $|\text{Tr}(\text{Adj}(D))| \leq |n \cdot d_{n-1}|$ , where  $d_{n-1}$  bounds the determinant of dwarves of dimensions  $n-1$ . From Hadamard's inequality we get that

$$d_{n-1} \leq \alpha^{n-1}(n-1)^{(n-1)/2}.$$

We now give experimental evidence to support these claims.

*Experiments.* We considered  $10^6$  dwarf matrices,  $D$ , for each set of parameters proposed for the alternating construction. We then computed the absolute value of  $\text{Tr}(\text{Adj}(D)) = \text{Tr}(\det(D) \cdot D^{-1})$  and list the average of the results in Table 3. Remark that the trace for each  $D$  is roughly  $n$  times the average coefficient of  $\text{Adj}(D)$ .

	$\lambda$	$n$	$\alpha$	$p$	$\log(\text{max bound})$	$\log( \text{trace} )$	$\log(\text{stand. dev.})$
toy	16	4	2	$2^{47} + 5$	7.3	2.6	2.3
challenge	64	8	2	$2^{255} - 19$	19.8	6.9	7.0
recommended	128	10	2	$2^{553} + 549$	26.6	9.6	9.9
large	512	24	2	$2^{2859} + 641$	79.6	34.5	35.1

Table 3: Toy, challenge, recommended, and large parameters from [3] for the alternating construction, together with the average and standard deviation values for  $\log_2(|\text{tr}(\text{Adj}(D))|)$  taken from  $10^6$  random samples. We also include the maximum bound from Lemma 3.

## 5 Trace attack on the direct construction

In what follows, we outline an attack on the direct construction from [3], and so assume the notation and structure summarized in Section 2.1.

Given some ciphertext  $C = E(\sigma \mathbf{A})E^{-1}$ , and a public key  $\bar{\mathbf{A}} := \{\bar{A}_1, \dots, \bar{A}_k\}$ , we can iteratively multiply by  $(\bar{A}_i)^{-1}$ , while checking the trace. If the trace is less than  $\text{Tr}(C)$ , then we will assume that we have correctly guessed the first matrix in the product. Heuristic 2 asserts that this will likely be the case, we estimate that the probability of this being a correct guess is at least 0.99, as seen in Section 3. The issue here is that trace is invariant under conjugation, thus when we left-multiply by the inverse of the last matrix we get that

$$\text{Tr}\left((\bar{A}_{\sigma(k)})^{-1} \cdot \bar{A}_{\sigma(1)} \cdots \bar{A}_{\sigma(k)}\right) = \text{Tr}\left(\bar{A}_{\sigma(1)} \cdots \bar{A}_{\sigma(k-1)}\right).$$

This trace will also be smaller than  $\text{Tr}(C)$ .

As it turns out, if we make an incorrect guess of the first matrix using this trace check, we can quickly flag it as incorrect by continuing to left-multiply by inverses. If the guess had been correct, there will always exist at least one matrix from the public key such that left-multiplying by its inverse gives a smaller trace. In the case of an incorrect guess, this quickly stops being the case, since the odds of being in one of these special cases, described above, becomes less and less likely as the product grows. Hence we can discard it. This approach is summarized in Algorithm 2.

Note that the set  $T_k$  from Algorithm 2 will contain two permutations that are equal except for the final two indices, which are permuted. This is due to the fact that once we get down to a product of two matrices, we get that  $\text{tr}(\bar{A}_{k-1}\bar{A}_k) = \text{tr}(\bar{A}_k\bar{A}_{k-1})$ . Determining which permutation was the correct one can be done easily by computing the corresponding products, and comparing them to the ciphertext.

We coded Algorithm 2 in Magma, which can be found at the following link:

<https://github.com/vgilchri/matrix-product-attack>.

**Algorithm 2:** DirectTraceAttack

---

**Input** :  $C = \sigma \bar{\mathbf{A}}, \bar{\mathbf{A}}$   
**Output:**  $\sigma$

```

1  $T_0 \leftarrow \{I\}$ 
2 for  $i = 1$  to  $k$  do
3    $T_i \leftarrow \{\}$  /* trace-decreasing products of length  $i$  */
4   for  $M \in T_{i-1}$  do
5     for  $\bar{A}_j \in \bar{\mathbf{A}}$  do
6       if  $\text{Tr}(\bar{A}_j^{-1} \cdot M^{-1} \cdot C) < \text{Tr}(M^{-1} \cdot C)$  then
7          $T_i \leftarrow T_i \cup \{M \cdot \bar{A}_j\}$ 
8 return  $\sigma \in T_k$ 

```

---

We give timings for this attack on the “challenge”, “recommended”, and “large” parameters from [3] in Table 4. The experiments were run using Magma V2.27-7 on a laptop with an Intel Dual-Core i3 at 1.1 GHz.

	$\lambda$	$k$	$n$	$\alpha$	$p$	time (s)
challenge	64	21	8	2	$2^{167} + 83$	0.2
recommended	128	35	10	2	$2^{302} + 307$	2.8
large	512	99	24	2	$2^{1105} - 1335$	915

Table 4: Timings in seconds for attacking the direct construction, where  $\lambda$  was the previously estimated security in bits.

**Theorem 1.** *Given a ciphertext  $C = \sigma \bar{\mathbf{A}}$ , and a public key,  $\bar{\mathbf{A}}$ , as described in Section 2.1, Algorithm 2 can recover  $\sigma$  in complexity  $O(k^2 n^\omega)$  under Heuristic 2, where  $O(n^\omega)$  is the cost of inverting an  $n \times n$ -matrix.*

*Proof.* Looking at Algorithm 2, we see that the first loop has length  $k$ . From Heuristic 2 we assume the second loop will require  $O(1)$  iterations. The last loop starts at length  $k$ , but decreases by one each time it is called (since we do not need to check matrices already in the product  $M$ ). The dominating cost of line 6 is from the matrix multiplications and inversions, thus we get a total complexity of  $O(k^2 n^\omega)$ .  $\square$

Note, our experiments showed that in practice  $|T_i|$  never exceeded 3 for any of the parameter sets, thereby supporting our use of Heuristic 2.

*Extra masking.* Recall from Section 2.1 that the authors of [3] suggested that including an extra masking matrix could improve the security of the scheme. They defined their public key matrices as  $\bar{A}_i := EA_iDE^{-1}$ , where  $D \in \mathfrak{D}$ . Note, this would require a larger prime to ensure `Decompose` still runs, since the length of the dwarf product essentially doubles. This alternate construction, with a non-trivial choice of  $D$ , does not avoid the attack from this section. Though  $A_iD$  is not itself a dwarf matrix, the properties of the trace function being an increasing function in our context remains true. Though  $D$  can be chosen to have as large a trace as possible in the hopes of making the trace of the products larger than  $p$ , we argue this would not be an effective countermeasure since you also risk the `Decompose` algorithm, that is central to the trapdoor function, failing.

## 6 Attacking the alternating construction

Recall that in the alternating construction we have two sets of dwarf matrices, say the  $\{A_i^0\}_{i=1}^k$  and the  $\{A_i^1\}_{i=1}^k$ . We also have elves  $E_0, \dots, E_k$ . For each  $i$ , the matrix  $\bar{A}^b := E_{i-1}A_i^bE_i^{-1}$  is computed. When the sender wants to encrypt a bitstring  $m = m_1 \dots m_k$ , they compute the following product:

$$C := \prod E_{i-1}A_i^{m_i}E_i^{-1} = E_0 \left( \prod A_i^{m_i} \right) E_k^{-1}.$$

In Section 6.1 we outline an information leak that impacts the entropy of the scheme, and in Section 6.2 we give an attack that significantly decreases the bit security.

### 6.1 Recovering the determinant of dwarf matrices

In what follows we will assume that  $n$  is much smaller than  $k$ , which is true for every suggested parameter set of the protocol. Recall, we have matrices  $\bar{A}_i^b$  which – when combined – we can turn into conjugate matrices as follows:

$$\bar{A}_i^0(\bar{A}_i^1)^{-1} = E_{i-1}A_i^0(A_i^1)^{-1}E_{i-1}^{-1}.$$

Following [3, Remark 7],  $A_i^0$  and  $A_i^1$  are chosen to have the same determinant as to not leak information. We know that  $A_i^0$  is a dwarf such that it has small entries. This is unfortunately not true for  $(A_i^1)^{-1}$ , since it is the inverse of a matrix with small entries, so when seen as a matrix over  $\mathbb{Z}$ , its entries can be huge. However, as described in Section 4, we can construct the inverse of a matrix through its determinant and its adjugate; i.e.  $(A_i^1)^{-1} = \frac{1}{\det A_i^1} \cdot \text{adj} A_i^1$ . We also know that  $\det A_i^1$  is small since we can upperbound it in terms of  $n$  and  $\alpha$  like in Lemma 1.

Now consider

$$\det(A_i^1) \cdot \bar{A}_i^0(\bar{A}_i^1)^{-1} = E_{i-1}A_i^0 \cdot \text{adj}(A_i^1)E_{i-1}^{-1}.$$

Even though the adjugate will not have coefficients bounded by  $\alpha$ , it *will* have coefficients much smaller than  $p$ ; i.e. the entries are bounded from above by entry

$n-1$  of the A003433 sequence, as seen in Section 3.1. Recall, we also assume that  $n$  is much smaller than  $k$ , which is true for every suggested parameter set of the protocol. Furthermore, conjugate matrices keep their characteristic polynomial (which includes the trace as well as the determinant).

With the public information, we can compute the characteristic polynomial of the matrix  $\bar{A}_i^0(\bar{A}_i^1)^{-1}$  (as an element of  $\mathbb{F}_p[x]$ ) and use lattice techniques to find the coefficients of this polynomial (as elements of  $\mathbb{Z}$ ). Indeed, write  $\sum_{i=0}^n c_i x^i$  for the characteristic polynomial (seen with coefficients in  $[0, \dots, p-1]$ ), and consider the lattice generated by the following basis:

$$A = \begin{pmatrix} c_0 & c_1 & c_2 & \dots & c_{n-1} \\ p & 0 & 0 & \dots & 0 \\ 0 & p & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & p \end{pmatrix}.$$

Remark that  $c_0 = 1$  since  $\bar{A}_i^0$  and  $\bar{A}_i^1$  have the same determinant. When multiplying  $\bar{A}_i^0(\bar{A}_i^1)^{-1}$  with  $\det(\bar{A}_i^1)$ , we get the matrix  $\bar{A}_i^0 \text{Adj}(\bar{A}_i^1)$  with characteristic polynomial  $\det(\bar{A}_i^1) \sum_{i=0}^n c_i x^i$ , whose entries are expected to be small relative to  $p$ , given that the coefficients of both  $\bar{A}_i^0$  and  $\text{Adj}(\bar{A}_i^1)$  are small relative to  $p$ . Now, the vector  $\mathbf{v} = (\det(\bar{A}_i^1)c_0, \dots, \det(\bar{A}_i^1)c_{n-1})$  is in  $A$ , so we can expect it to be a relatively short vector. The first entry of this short vector will be the determinant of  $A_i^1$  up to sign, since  $c_0 = 1$ . From this, one can also easily deduce the coefficients of the characteristic polynomials of  $A_i^0 \cdot (A_i^1)^{-1}$  and  $A_i^1 \cdot (A_i^0)^{-1}$ , seen as elements in  $\mathbb{Q}$ .

How likely is the vector  $\mathbf{v}$  to be the shortest vector in  $A$ ? Due to the construction of the characteristic polynomial, its first entry can achieve the largest possible value so if we bound this from above by the aforementioned maximal value we get

$$\begin{aligned} \|\mathbf{v}\| &= \|(\det(\bar{A}_i^1)c_0, \dots, \det(\bar{A}_i^1)c_{n-1})\| \\ &\leq \|(\det(\bar{A}_i^1)c_0, \dots, \det(\bar{A}_i^1)c_0)\| \\ &= \|(\det(\bar{A}_i^1), \dots, \det(\bar{A}_i^1))\| \\ &= \sqrt{n} \cdot \det(\bar{A}_i^1) \\ &\leq \sqrt{n} \cdot (\alpha\sqrt{n})^n \\ &= \alpha^n \sqrt{n}^{n+1}. \end{aligned}$$



On the other hand, we can consider the Minkowski upper bound for the shortest vector  $\lambda_1$  in  $\Lambda$ :

$$\begin{aligned} \|\lambda_1\| &\leq \sqrt{n} \cdot \left( \sqrt{\det(A^T A)} \right)^{1/n} \\ &= \sqrt{n} \cdot \left( p^{n-1} \sqrt{p^2 + \left( \sum_{i=0}^{n-1} c_i^2 \right)} \right)^{1/n} \\ &\leq \sqrt{n} \cdot \left( p^{n-1} \sqrt{p^2 + np^2} \right)^{1/n} \\ &= \sqrt{n} \cdot p \cdot (\sqrt{n+1})^{1/n}. \end{aligned}$$

Now if we assume  $p \approx \alpha^k n^{k-1}$  as in the protocol, as well as  $n$  being much smaller than  $k$ , we can simplify this to

$$\|\lambda_1\| \leq \alpha^k n^{k-1/2}.$$

Even though it is clear that  $\mathbf{v}$  is short compared to this bound, the lattice  $\Lambda$  is not constructed at random so we cannot conclude that  $\mathbf{v}$  is likely to be the shortest vector. In fact, for small values of  $p$  and  $n$  one can easily construct counterexamples to this statement. For realistic parameters however, only one exceptional case appears heuristically. This is the case where all entries of  $\mathbf{v}$  share a common factor over  $\mathbb{Z}$ , in which case a shortest vector algorithm will only get  $\mathbf{v}$  divided by this common factor. It is hard to determine what the odds of this happening are exactly, given that the  $c_i$  are not drawn uniformly at random and will depend on the (small) parameter  $\alpha$ . Generically, we expect every  $c_i$  to be divisible by a prime  $\ell_j$  with probability  $\ell_j^{-1}$ , so all of them will be divisible by  $\ell_j$  simultaneously with a probability of  $\ell_j^{-n}$ . This probability is heavily dominated by  $\ell_j = 2$  of course, and in practice we see that this is the case as well for all realistic parameters. Even for the challenge parameters, it means we can determine the determinant of the used dwarf matrices in the public key with over 99% accuracy. When swapping to the recommended parameters, this turns into 99.9% already. This can be seen from the code provided at

<https://github.com/vgilchri/matrix-product-attack>.

Since the determinant of the dwarf matrices leaks, this means the key generation can be unlucky and have a pair of dwarf matrices  $A_i^0, A_i^1$  with very large determinant, although the chances of this happening are rather slim.<sup>4</sup> Given how there are significantly less matrices with these large determinants, their entropy can easily be too small which may allow brute force guessing.

<sup>4</sup> Remark that this is also noticeable in the timing of the key generation: sampling random elements from  $\mathfrak{D}$  until a second one with the same huge determinant is found can take an egregious amount of time.

*Countermeasures.* One possible countermeasure would be to only sample dwarves from a set  $\mathfrak{D} \cap \mathfrak{F}$ , where  $\mathfrak{F}$  is a set of matrices with fixed determinant (up to sign). The choice of  $\mathfrak{F}$  would of course require enough entropy for  $\mathfrak{D} \cap \mathfrak{F}$ . One suggestion would be  $\mathfrak{F} = SL_n(\mathbb{F}_p)$ , although depending on  $\alpha$  and  $n$ , better options may be available. Apart from this, any type of security reduction would still need to take into account information leaking from the characteristic polynomials of  $A_i^0 \cdot (A_i^1)^{-1}$  and  $A_i^1 \cdot (A_i^0)^{-1}$ , making it a lot harder to argue why they would be indistinguishable from random for example.

## 6.2 Trace attack

The naïve approach to applying the trace attack from Section 5 would be to compute a new product from a bit string,  $d$ , of the form

$$D := \prod E_{i-1} A_i^{d_i} E_i^{-1} = E_0 \left( \prod A_i^{d_i} \right) E_k^{-1}.$$

Then we can compute  $CD^{-1}$  to obtain a product of the form

$$CD^{-1} = E_0 \prod A_i^{m_i} \prod (A_i^{d_i})^{-1} E_0^{-1}.$$

This is a product of dwarves and inverse dwarves, conjugated by one elf. The hope, thus, would be that since the product is conjugated by an elf, we can use the trace to gain some knowledge about the secret,  $m$ . The caveat here is that inverse dwarves are not guaranteed to have small entries, and in fact, often have seemingly very large entries as we saw in Section 4. Hence, some more care will be required for this construction.

We will be able to use the same overall idea to gain information about the secret using the trace, but will need to consider the absolute value of the trace instead. This time we have more variance in the size of the entries, unlike in Section 5 where we were guaranteed to have entries with a fixed upper bound. We will need to guess several bits at once in order to see noticeable differences in the absolute value of the trace. Another issue is that since the product of matrices now has  $2k$  matrices, it is likely that the combined trace of the product will be larger than  $p$ . So in order for the attack to run we will first need to guess some bits of  $m$ , and then we can run the attack to recover the remaining bits of  $m$ .

We will begin with some initial guess  $d$ , and take note of the absolute value of the trace  $|tr(CD^{-1})|$ . Then, we will adjust  $d$  by switching the last *step* bits of  $d$  with every possible bit string of length *step*, and compare all of the traces. We expect there to be a big difference in trace size between correct and incorrect guesses, as seen in Heuristic 2. Thus, we keep any bit strings that achieve the minimum trace value, or close to the minimum trace value (we denote this interval using  $\epsilon$ ). We iterate until we have some candidates for  $m$ .

We outline the attack in Algorithm 3, where we would like to recover  $k'$  bits of  $m$ , assuming we have already guessed the first  $k - k'$  bits. Recall that the notation  $m\bar{\mathbf{A}}$  refers to computing a product of matrices using  $\bar{\mathbf{A}}^0$  and  $\bar{\mathbf{A}}^1$  that relies on  $m$ , as outlined in Section 2.2. We use  $\epsilon$  to denote some “wobble room”, that is, we want to keep any candidates that give a trace that is *close* to the minimum, even if not exactly the minimum. We will use `AllBitStrings(length=step)` to denote a function that lists all bitstrings of length *step*. Note also we will use *prev* to track what the previous product’s trace was in order to keep only products of decreasing trace. It is initialized to be as large as possible since the first product does not have any trace to be compared to.

---

**Algorithm 3: AlternatingTraceAttack**


---

```

Input :  $\bar{\mathbf{A}}^0, \bar{\mathbf{A}}^1, step, C = m\bar{\mathbf{A}}$ 
Output:  $m$ 

1  $str \leftarrow \text{AllBitStrings}(\text{length}=step)$ 
2  $T \leftarrow \{str\}$  /* track "good" candidates */
3  $prev \leftarrow p - 1 - \epsilon$  /* track previous trace for comparison */
4 for  $i = 1$  to  $k'/step$  do
5     for  $b \in T$  do
6          $traces \leftarrow \{\}$ 
7         for  $s \in str$  do
8              $D_s \leftarrow (b||s||0\dots0)\bar{\mathbf{A}}$  /* add zeros until length  $k$  */
9              $t_s \leftarrow \text{Abs}(\text{Tr}(CD_s^{-1}))$ 
10            if  $t_s < prev$  then
11                 $traces \leftarrow traces \cup \{t_s\}$ 
12     $T \leftarrow \{s : t_s < \min(traces) + \epsilon\}$ 
13     $prev \leftarrow \min(traces) + \epsilon$ 
14 return  $m \in T$ 

```

---

We coded the attack in Magma, which can be found at the following link:

<https://github.com/vgilchri/matrix-product-attack>.

We see that the first loop has length  $k'/step$ . The second loop has length upper bounded by  $2^{step}$ , but in practice is far smaller. The third loop has length exactly  $2^{step}$ . The dominating subroutine within the loops is the matrix multiplications, which has complexity  $O(n^\omega)$ , where  $\omega$  depends on the choice of algorithm used. Based on experimental evidence, we estimate that *step* can be chosen between  $k'/8$  and  $k'/4$ , with the latter being more expensive but having a higher likelihood of success. In total this means that the complexity can range between  $O(2^{k'/4}n^\omega)$  and  $O(2^{k'/2}n^\omega)$ , which gives a time-accuracy trade-off for the attack. This reduces the bit complexity of the scheme from  $k$  to approximately  $k - 3k'/4$  bits of security.



*Countermeasures.* One possible countermeasure would be to reduce the size of the prime to as small as possible such that `Decompose` still runs. This would save back some bits of security, but does not avoid the attack entirely.

The other obvious countermeasure is to increase the size of  $k$ . As mentioned, we are estimating the new security value as  $\lambda' = k - (3/4)k'$ , so in order to reach  $\lambda$  bits of security, we would need  $k = \lambda + (3/4)k'$ . For example, this means that for 64 bits of security,  $k = 88$ .

## 7 Conclusion

We have shown that the constructions outlined in [3] are not secure by giving two full message recovery attacks, running in polynomial time, and detailing several other weaknesses. We expect these attacks cannot easily be avoided with simple countermeasures.

We believe the attack from Section 6.2 may be an interesting example where machine learning could be used to make a more adaptive attack that makes use of previous (failed) guesses, and optimizes some of the bounds used such as  $\epsilon$  and  $k'$ . It may also be combined with the ideas from Section 6.1 to decrease the size of the product and thus allow for a larger  $k'$  value. Future work could also consider how the attacks detailed in this paper may apply to other non-commutative objects, such as tensors, or other word problem based cryptosystems.

## References

1. OEIS Foundation Inc. (2024), Entry A003433 in The On-Line Encyclopedia of Integer Sequences, <https://oeis.org/A003433>
2. Charles, D.X., Lauter, K.E., Goren, E.Z.: Cryptographic hash functions from expander graphs. *Journal of CRYPTOLOGY* **22**(1), 93–113 (2009)
3. Geraud-Stewart, R., Naccache, D.: New public-key cryptosystem blueprints using matrix products in  $\mathbb{F}_p$ . *Cryptology ePrint Archive*, Paper 2023/1745 (2023), <https://eprint.iacr.org/2023/1745>
4. Tao, T., Vu, V.: On random  $\pm 1$  matrices: singularity and determinant. *Random Structures Algorithms* **28**(1), 1–23 (2006). <https://doi.org/10.1002/rsa.20109>, <https://doi.org/10.1002/rsa.20109>
5. Tillich, J.P., Zémor, G.: Hashing with sl 2. In: *Advances in Cryptology—CRYPTO'94: 14th Annual International Cryptology Conference Santa Barbara, California, USA August 21–25, 1994 Proceedings* 14. pp. 40–49. Springer (1994)
6. Levy-dit Vehel, F., Perret, L.: Security analysis of word problem-based cryptosystems. *Designs, Codes and Cryptography* **54**, 29–41 (2010)
7. Wagner, N.R., Magyarik, M.R.: A public-key cryptosystem based on the word problem. In: *Workshop on the Theory and Application of Cryptographic Techniques*. pp. 19–36. Springer (1984)