REMARKS ON A TROPICAL KEY EXCHANGE SYSTEM

The Texas Tech community has made this publication openly available. Please share how this access benefits you. Your story matters to us.

Research Article

Dylan Rudy and Chris Monico*

Remarks on a Tropical Key Exchange System

Abstract: We consider a key-exchange protocol based on matrices over a tropical semiring which was recently proposed in [2]. We show that a particular private parameter of that protocol can be recovered with a simple binary search, rendering it insecure.

1 Introduction

Let S be any nonempty subset of $\mathbb{R}$ which is closed under addition. Define two operations $\oplus$ and $\otimes$ on S by

$$a \oplus b = \min\{a, b\},$$
$$a \otimes b = a + b.$$

Both operations are associative and commutative and $\otimes$ distributes over $\oplus$, and hence S is a commutative semiring, called a *tropical semiring*. The set $\mathcal{M} = \mathrm{Mat}_{k \times k}(S)$ of $k \times k$ matrices over S is therefore a semiring with the induced operations

$$(a_{ij}) \oplus (b_{ij}) = (a_{ij} \oplus b_{ij}),$$
$$(a_{ij}) \otimes (b_{ij}) = (c_{ij}), \quad \text{where } c_{ij} = (a_{i1} \otimes b_{1j}) \oplus (a_{i2} \otimes b_{2j}) \oplus \cdots \oplus (a_{ik} \otimes b_{kj}).$$

In [1], the authors proposed two key exchange protocols based on the structure $\mathcal{M}$. Shortly after, an effective attack was given on one of those protocols in [3]. Subsequently, a new key exchange protocol was proposed in [2] (in fact, two new protocols, but they are very closely related to each other). It is this protocol that we consider in this paper.

In [2], the authors give two semigroup operations on $\mathcal{M} \times \mathcal{M}$ each arising as a semidirect product induced by a specified action of these matrices on themselves. The two semigroup operations are given by

$$(M, G) \circ (S, H) = \Big(M \oplus S \oplus H \oplus (M \otimes H), \quad G \oplus H \oplus (G \otimes H) \Big), \tag{1}$$

$$(M, G) \star (S, H) = \Big((H \otimes M^T) \oplus (M^T \otimes H) \oplus S, \quad G \otimes H \Big). \tag{2}$$

Note that for each of these operations, the first component of the product does not depend on G. This fact plays a key role in the two key exchange protocols they then propose (one corresponding to each operation):

1. Alice and Bob agree on public matrices $M, H \in \mathcal{M}$ whose entries are integers in the range $[-N, N]$, and they agree on a positive integer K. Alice selects a private positive integer $m < 2^K$ and Bob selects a private positive integer $n < 2^K$.

*** Corresponding Author: Chris Monico:** Department of Mathematics and Statistics, Texas Tech University, United States of America; Email: c.monico@ttu.edu
Dylan Rudy: Department of Mathematics and Statistics, Texas Tech University, United States of America

2. Alice computes $(M, H)^m = (A, P_A)$ and sends A to Bob.
3. Bob computes $(M, H)^n = (B, P_B)$ and sends B to Alice.
4. Alice determines the first component of $(M, H)^{m+n} = (M, H)^n(M, H)^m = (B, P_B)(A, P_A)$ from her knowledge of A, P_A, and B (knowledge of P_B is not necessary for either of the operations (1) or (2).
5. Bob similarly determines the first component of $(M, H)^{m+n} = (M, H)^m(M, H)^n = (A, P_A)(B, P_B)$ from his knowledge of B, P_B, and A.

In the next section, we show that an eavesdropper can find a positive integer m' for which the first component of $(M, H)^{m'}$ is A; she can then use this m' to compute the shared secret key in essentially the same way as Alice. Furthermore, such an m' can be found using $\mathcal{O}(K^2)$ operations (1) or (2).

2 The attack

Since addition of matrices in $\mathcal{M}$ is idempotent, i.e., $G \oplus G = G$, we have a partial order on $\mathcal{M}$ defined by

$$X \leq Y \quad \text{if } X \oplus Y = X.$$

Clearly we have that $X \leq Y$ iff $x_{ij} \leq y_{ij}$ for all $i, j \in \{1, 2, \ldots, k\}$. Furthermore, this partial order respects both operations on $\mathcal{M}$; if $X \leq Y$ and $Z \in \mathcal{M}$, then $X \oplus Z \leq Y \oplus Z$ and $X \otimes Z \leq Y \otimes Z$.

Proposition 2.1. *Consider the semigroup $\mathcal{M} \times \mathcal{M}$ equipped with either of the two operations defined by* (1) *and* (2). *Let $(M, H) \in \mathcal{M} \times \mathcal{M}$, and for each positive integer ℓ let $(M_\ell, H_\ell) = (M, H)^\ell$. Then the sequence $\{M_\ell\}$ is monotonically decreasing: $M_1 \geq M_2 \geq M_3 \geq \ldots$.*

Proof. Let $\ell \geq 2$. For the operation $\circ$ we have

$$(M_\ell, H_\ell) = (M_{\ell-1}, H_{\ell-1}) \circ (M, H)$$
$$= \Big(M_{\ell-1} \oplus M \oplus H \oplus (M_{\ell-1} \otimes H), H_{\ell-1} \oplus H \oplus (H_{\ell-1} \otimes H)\Big),$$

so that $M_\ell = M_{\ell-1} \oplus M \oplus H \oplus (M_{\ell-1} \otimes H)$. In particular, $M_\ell \oplus M_{\ell-1} = M_\ell$, and hence $M_\ell \leq M_{\ell-1}$.

Similarly, for the operation $\star$ we have that

$$(M_\ell, H_\ell) = (M, H) \star (M_{\ell-1}, H_{\ell-1})$$
$$= \Big((H_{\ell-1} \otimes M^T) \oplus (M^T \otimes H_{\ell-1}) \oplus M_{\ell-1}, H \otimes H_{\ell-1}\Big),$$

and hence $M_\ell = (H_{\ell-1} \otimes M^T) \oplus (M^T \otimes H_{\ell-1}) \oplus M_{\ell-1}$. Again, $M_\ell \oplus M_{\ell-1} = M_\ell$, so that $M_\ell \leq M_{\ell-1}$. □

The problem alluded to at the end of the introduction is now easily solved with a binary search. Let $M, H \in \mathcal{M}$ and $(M, H)^\ell = (M_\ell, H_\ell)$. Suppose $A \in \mathcal{M}$ satisfies $A = M_m$ for some positive integer $m < 2^K$. First, obtain an upper bound on m by computing successive squares

$$M_1, M_2, M_4, M_8, \ldots$$

until finding a positive integer t for which $A \leq M_{2^t}$. Since it is then known that $1 \leq m \leq 2^t$, a simple binary search will find an integer m' for which $M_{m'} = A$. The sequence $M_1, M_2, \ldots$ is generally strictly decreasing, in which case $m' = m$. However, even if $m' \neq m$, finding such an integer m' is enough for the eavesdropper to recover the shared secret key. Let $\pi_1 : \mathcal{M} \times \mathcal{M} \longrightarrow \mathcal{M}$ be the map $\pi_1(C, D) = C$. Suppose $(M, H)^n = (B, P_B)$, $(M, H)^m = (A, P_A)$ and $(M, H)^{m'} = (A, P_E)$. Then for each of the operations (1) and (2), the shared secret key satisfies

$$\pi_1((M, H)^{m+n}) = \pi_1((M, H)^{m'+n}).$$

This is clear, since this shared secret key can be expressed in terms of A, B, and P_B only, but it may also be explicitly verified. For example, with the operation (1),

$$\pi_1((M, H)^{m+n}) = \pi_1((A, P_A) \circ (B, P_B))$$

$$= A \oplus B \oplus P_B \oplus (A \otimes P_B)$$
$$= \pi_1((A, P_E) \circ (B, P_B))$$
$$= \pi_1((M, H)^{m'+n}).$$

In particular, the eavesdropper may recover the shared secret key via

$$\pi_1((M, H)^{m+n}) = \pi_1((M, H)^n \circ (M, H)^{m'})$$
$$= \pi_1((B, P_B) \circ (A, P_E))$$
$$= B \oplus A \oplus P_E \oplus (B \otimes P_E).$$

Finding t as described above requires at most K semigroup operations in $\mathcal{M} \times \mathcal{M}$. The binary search, done in the most obvious way, would compute K powers of (M, H), each of which requires no more than $2K$ semigroup operations in $\mathcal{M} \times \mathcal{M}$, for a total complexity of at most $2K^2 + K$ operations in $\mathcal{M} \times \mathcal{M}$. This can be reduced to $K^2 + K$ by storing the successive squares $(M_1, H_1), (M_2, H_2), (M_4, H_4), \ldots$ and using them to compute each power of (M, H) during the binary search phase.

Addition of $k \times k$ matrices can be accomplished with $\mathcal{O}(k^2)$ integer max operations, and multiplication accomplished using $\mathcal{O}(k^3)$ integer addition and max operations. Therefore this attack requires $\mathcal{O}(K^2 k^3)$ integer operations. We argue below that the typical entry of A has about K bits. In that case, each integer addition and max operation requires no more than K bit operations, for a total of $\mathcal{O}(K^3 k^3)$ bit operations. If we let α denote the number of bits required to represent A (i.e., the key size) it follows that $\alpha \approx K k^2$, and this attack requires $\mathcal{O}(\alpha^3)$ bit operations, a polynomial-time function of the input size. If K is fixed, as in our experiments, then it requires $\mathcal{O}(\alpha^{1.5})$ bit operations.

We coded this method in C, and performed some experiments on a single core of an i7 CPU at 3.10GHz. Using $\mathcal{M} = \mathrm{Mat}_{k \times k}(S)$ for various values of k, and the parameters $N = 1000$, $K = 200$ suggested in [2], we performed 40 experiments for each value of k. In each experiment, we generated random matrices M, H and chose random positive integers $m, n < 2^K$ and measured the time to recover an m' as described above. The results of these experiments are summarized in Table 1. For reference, we also report the average number of bits α in the matrix A that would be shared by Alice, and the values t/k^3 and $t/\alpha^{1.5}$ for comparison with the asymptotic runtime estimates given above.

Table 1: Average number of bits α to represent A (Alice's matrix, from Section 1), and average time t (in seconds) to recover m' for various sized ($k \times k$) matrices, with $N = 1000$ and $K = 200$.

k	α	t	t/k^3	$t/\alpha^{1.5}$
5	5222	0.12	0.00096	3.2e−7
10	20885	0.66	0.00066	2.2e−7
15	47025	2.43	0.00072	2.4e−7
20	83710	4.76	0.00060	2.0e−7
25	130594	10.53	0.00067	2.2e−7
30	188145	17.75	0.00066	2.2e−7
35	256484	24.05	0.00056	1.9e−7
40	334040	40.92	0.00064	2.1e−7
45	422111	45.80	0.00050	1.7e−7
50	523312	78.33	0.00063	2.1e−7
55	631091	98.19	0.00059	2.0e−7
60	752490	122.57	0.00057	1.9e−7

We would like to make one final remark about the key sizes in this system. With the notation as above and the operation (1), for example, we have

$$M_{\ell+1} = M_\ell \oplus M \oplus H \oplus (M_l \otimes H).$$

Since $M_2 \preceq M$ and $M_2 \preceq H$ and $M_{\ell+1} \preceq M_2$ for all $\ell \geq 2$, it follows that

$$M_{\ell+1} = M_\ell \oplus (M_\ell \otimes H), \qquad \text{for } \ell \geq 2.$$

This means that, on average, the entries of $M_{\ell+1}$ decrease from those of M_ℓ by an approximately constant amount, proportional to the size of the entries of H. With Alice's $m \approx 2^K$, this means that the entries of A are on the order of $-c \times 2^K$, or about K bits each. With the parameter sizes $K = 200$, $k = 30$, $N \approx 1000$ suggested in [2], one would have M and H consisting of about 9000 bits each and A with about $30 \times 30 \times 200 = 180,000$ bits.

3 Conclusion

The attack presented here exploits the fact that the sequence $\{(M, H)^\ell\}$ is linearly ordered. It is quite effective and practical against the protocols described in [2]. For those protocols, Alice and Bob must do approximately $\mathcal{O}(K)$ operations in the semigroup $\mathcal{M} \times \mathcal{M}$, and this attack requires about $\mathcal{O}(K^2)$ operations in that same semigroup, so an increase of parameter sizes does not help.

We thank the referees for their thoughtful reading of this manuscript and their feedback.

References

[1] Dima Grigoriev and Vladimir Shpilrain. Tropical cryptography. *Comm. Algebra*, 42(6):2624–2632, 2014.
[2] Dima Grigoriev and Vladimir Shpilrain. Tropical cryptography II: extensions by homomorphisms. *Comm. Algebra*, 47(10):4224–4229, 2019.
[3] Matvei Kotov and Alexander Ushakov. Analysis of a key exchange protocol based on tropical matrix algebra. *J. Math. Cryptol.*, 12(3):137–141, 2018.