A Practical Attack on the Root Problem in Braid Groups

Anja Groch, Dennis Hofheinz, and Rainer Steinwandt[*]

ABSTRACT. Using a simple heuristic approach to the *root problem* in braid groups, we show that cryptographic parameters proposed in this context must be considered as insecure. In our experiments we can, often within seconds, extract the secret key of an authentication system of Sibert et al. that is based on the root problem in braid groups.

1. Introduction

In recent years several proposals for asymmetric cryptographic schemes building on braid groups have been made. For an excellent introduction to the subject we refer to Dehornoy's survey article [**Deh04**]. Next to proposals based on (variants of) the *conjugacy problem*, more recently constructions for making use of the *root problem* in braid groups gained attention, and in the sequel we show that parameters proposed in this context succumb to a simple heuristic attack. Our approach reduces a root problem in the braid group B_n to a root problem in the symmetric group S_n, and despite its simplicity turns out as rather effective for proposed instances. To describe the details, we start by recalling some basic terminology on braid groups.

1.1. Braid Groups. For our purposes it suffices to interpret the *braid group* B_n ($n \in \mathbb{N}$) as a finitely presented group defined through the presentation (cf. [**Art25**])

$$\left\langle \ \sigma_1, \ldots, \sigma_{n-1} \ \left| \ \begin{array}{rcll} \sigma_i\sigma_j\sigma_i & = & \sigma_j\sigma_i\sigma_j & \text{if } |i-j| = 1 \\ \sigma_i\sigma_j & = & \sigma_j\sigma_i & \text{if } |i-j| > 1 \end{array} \ \right\rangle \right. .$$

As usual, we refer to $\sigma_1, \ldots, \sigma_{n-1}$ as *(Artin) generators* and to arbitrary elements of B_n as *braids*. To refer to a specific representation of a braid in terms of Artin generators, we use the term *braid word*. A braid is said to be *positive* if and only if it can be written as a product of generators σ_i, i.e., without involving negative powers of the σ_i. Here the identity $\varepsilon \in B_n$ is regarded as positive, too, and it can be shown that the positive braids in B_n form a *monoid* B_n^+ which embeds into B_n

(cf. [**Gar69**]). Moreover, we call a braid $u \in B_n^+$ a *tail* of $w \in B_n^+$ if w can be written as $w = vu$ for some $v \in B_n^+$. Analogously, $u \in B_n^+$ is referred to as a *head* of $w \in B_n^+$ if $w = uv$ for some $v \in B_n^+$.

1.2. Δ-Normal Form. Setting inductively $\Delta_1 := \sigma_1$ and $\Delta_i = \sigma_1 \cdots \sigma_i \cdot \Delta_{i-1}$ for $1 < i < n$, we define the *fundamental braid* $\Delta \in B_n$ as $\Delta := \Delta_{n-1}$. Next, we establish a partial ordering $\leq$ on the elements of B_n: For $u, v \in B_n$ we set $v \leq w$ if and only if there exist positive braids $\alpha, \beta \in B_n^+$ such that $w = \alpha v \beta$. Now any braid $\alpha \in B_n$ satisfying $\varepsilon \leq \alpha \leq \Delta$ is referred to as *canonical factor*. Further on, we have a *canonical homomorphism* $\pi : B_n \longrightarrow S_n$ from the braid group B_n into the symmetric group S_n such that

- $\pi(\sigma_i) = (i, i+1)$ and
- restricting π to the set of canonical factors in B_n induces a bijection

(see [**EM94**]). We denote by π^{-1} the map that maps permutations to canonical factors such that $\pi \circ \pi^{-1} = \mathrm{id}$.

A factorization $\gamma = \alpha\beta$ of a positive braid γ into a canonical factor α and a positive braid β is said to be *left-weighted* if and only if α has the maximal length among all such decompositions. A *right-weighted* factorization is defined analogously. Moreover, for any braid $w \in B_n$ we denote the greatest $i \in \mathbb{Z}$ with $\Delta^i \leq w$ by $\inf w$; analogously, $\sup w$ stands for the smallest $i \in \mathbb{Z}$ with $w \leq \Delta^i$. With this notation every braid $w \in B_n$ can be written *uniquely* as

$$(1) \qquad\qquad w = \Delta^r W_1 \cdots W_s$$

with $r = \inf w$, $s = \sup w - \inf w$ and canonical factors $\varepsilon < W_i < \Delta$ such that $W_i W_{i+1}$ is left-weighted for $1 \leq i < s$ (cf. [**Gar69, EM94**]). In this context, we say that s is the *canonical length* of w, and the explicit decomposition (1) is called the Δ-*normal form* of w.

Note that the W_i are canonical factors, so they can be represented uniquely by the corresponding permutations $\pi(W_i) \in S_n$. For a given braid word $w \in B_n$, its Δ-normal form can be computed in time $\mathcal{O}(|w|^2 n \log n)$ with $|\cdot|$ denoting the word length (see [**CKL$^+$01**] for details). Sometimes the Δ-normal from is called *left-normal form* because of the power of Δ being written on the left and the decompositions of canonical factors being left-weighted. Writing the power of Δ on the right and demanding right-weightedness of the decompositions of canonical factors, one obtains the *right-normal form*, which can also be computed efficiently.

1.3. Conjugacy and Root Problem. For the *conjugacy (search) problem* (short: *CSP*), we are given two braid words so that the corresponding braids $u, w \in B_n$ are related by $w = \alpha^{-1} u \alpha$ for some $\alpha \in B_n$, and the goal is to find some braid $\alpha' \in B_n$ with $w = \alpha'^{-1} u \alpha'$. Leaving aside the question of efficiency, this problem has been solved in [**Gar69**], but finding an efficient algorithm for CSP is the topic of ongoing research. In parts motivated by cryptanalytic purposes, a number of heuristic approaches to the CSP have been brought up (see, e.g., [**Hug02, HS02, GKT$^+$05**]). An experimentally very efficient Las Vegas algorithm for solving the CSP is known, too (cf. [**Geb03**]), and from the cryptographic point of view further research on how to generate hard instances of the CSP is desirable.

Most of the braid group based cryptographic schemes that have been proposed rely on the difficulty of some variation of the CSP. For instance, the public key cryptosystem of [**KLC$^+$00**] is based on a variant of the CSP, where one is given

$\alpha^{-1}u\alpha$ and $\beta^{-1}u\beta$ for α and β chosen from two commuting subgroups, and one needs to find $(\alpha\beta)^{-1}u(\alpha\beta)$ (see also [**SU04**]). Note that this particular problem can be solved in polynomial time, cf. [**CJ03**].

For the *root extraction problem* (short: *REP*), we are given a braid word representing some $w \in B_n$ along with some natural number $m \in \mathbb{N}$ such that $w = u^m$ for some $u \in B_n$. The goal is to find a braid u' with $u'^m = w$. As shown by González-Meneses [**GM03**], u' and u then in particular must be conjugate.

Using the REP for cryptography has been suggested by Ko et al. in [**KLC$^+$00**], and a concrete cryptographic system for entity authentication basing on the hardness of the REP has been put forward by Sibert et al. in [**SDG02, Sib03**]. This system is provably secure (assuming the hardness of REP *and* CSP) and is described in more detail in Section 2. Very recently, in [**LC05**], two authentication schemes, claimed to be based on the REP, were suggested, but the first of these has been attacked successfully in [**Tsa05**] without solving the REP. The second scheme is unsatisfying in the sense that is relies on the difficulty of the Diffie-Hellman decomposition problem which can be solved in polynomial time [**CJ03**]. Moreover, [**LC05**] lacks concrete parameter suggestions that could serve as testground for practical tests. Hence here we do not investigate these two proposals further.

Disregarding the question of efficiency, a result of Styshnev [**Sty79**] provides an algorithmic solution for the REP. Furthermore, Lee in [**Lee04**] shows how to reduce the REP to a conjugacy search problem in a different (Garside) group. In combination with, e.g., the techniques developed in [**Geb03**] (which apply to arbitrary Garside groups) this can be used to solve the REP. However, to our knowledge no results about the efficiency of that method are currently known. In particular, it seems unclear how efficient the algorithm of [**Geb03**] performs in the Garside groups considered in [**Lee04**]. While exploring this question certainly looks worthwhile, in the sequel we show that already a rather simple heuristic approach enables a practical attack on proposed parameters of the root problem.

2. A Proposal of Sibert, Dehornoy and Girault

To get an idea of how a cryptographic system might be founded on CSP and REP (and also to motivate the parameters used in our experiments in Section 3.6), we reproduce an authentication system described in [**SDG02, Sib03**].

2.1. System Description. The authors of [**SDG02**] refer to the system discussed here as *Scheme III*. It follows the Fiat-Shamir paradigm (cf. [**FS87**]) and involves a *prover* P who chooses a braid $s \in B_n$ along with an integer $m \in \mathbb{N}$. Then P publishes both $b := s^m$ and m, but keeps s secret. During an authentication phase, P wants to convince a *verifier* V that he knows s.

The idea is to convince the verifier V not in the trivial way (namely by sending her s), but rather in *zero-knowledge* (see [**GMR89**]), i.e., without revealing any "usable" information about s. This makes the scheme useful for authentication in the sense that P is the only one who can convince another party of knowing s: even after having convinced V, the braid s remains the prover P's individual secret. For P and V being honest, the communication in Scheme III reads as follows.

(1) P chooses a random $r \in B_n$ and sends $x = rbr^{-1}$ to V.
(2) V chooses a random $\epsilon \in \{0, 1\}$ and sends it to P.
(3)   • If $\epsilon = 0$, P sends $y = r$ to V; V then checks $x \overset{?}{=} yby^{-1}$,

- if $\epsilon = 1$, P sends $y = rsr^{-1}$ to V; V then checks $x \stackrel{?}{=} y^m$.

The verifier V accepts the proof exactly if the check in the third step succeeds. Of course, extracting an m^{th} root of b will enable anybody to authenticate as P. Also, when communicating with P, choosing $\epsilon = 1$ combined with the ability to solve the CSP (applied to $u = b$ and $w = y^m = rbr^{-1}$) enables an evil verifier to get r and therewith P's indivual secret s. On the other hand, under the assumption that REP and CSP are hard, and additionally some assumptions concerning the choice of instances, [**SDG02**] shows that the scheme indeed is zero-knowledge. That means that no computationally usable information about s is leaked.

2.2. Suggested Parameters. In [**SDG02**], the following parameters are suggested: $n = 30$, and random braids should be chosen either as products of 15 random permutations, or as products of at least 1000 random (Artin) generators. Additionally, to protect against a heuristic approach to the CSP (cf. [**HS02**]), [**Sib03**] suggests to employ cycling and decycling steps in the generation of instances (see [**Sib03**, Algorithme 3.2.2] for details).

3. A Heuristic Apprach to the REP

As outlined in the explanation of the REP, subsequently we assume that $w = u^m \in B_n$ and $m \in \mathbb{N}$ are given, and our goal is to find a $u' \in B_n$ such that $w = u'^m$. Moreover, for most cryptographic applications, infimum and canonical length of u are also either public or at least can be guessed with significant probability. Hence, we will assume $\inf u$ and the canonical length of u to be known, too.

3.1. A First Observation. We start with:

PROPOSITION 3.1. *Say that $u \in B_n$ has Δ-normal form*

$$\Delta^r U_1 \cdots U_s,$$

and that $w := u^m$ has Δ-normal form

$$\Delta^{r'} W_1 \cdots W_{s'}$$

such that $\ell := s' - (m-1)s > 0$. Then the last ℓ canonical factors $W_{(m-1)s+1} \cdots W_{s'}$ of the Δ-normal form of w form a tail of $U_1 \cdots U_s$, i.e.,

$$U_1 \cdots U_s = u' W_{(m-1)s+1} \cdots W_{s'}$$

for some positive $u' \in B_n^+$.

PROOF. Juxtaposing m copies of u and moving powers of Δ to the left, we can write w as

$$(2) \qquad w = \Delta^{mr} W_1^{(1)} \cdots W_{ms}^{(1)},$$

such that for $j = 0, \ldots, m-1$ and $k = 1, \ldots, s$ we have $W_{js+k}^{(1)} = \tau^{r(m-1-j)}(U_k)$, where τ denotes conjugation by Δ. As conjugation by Δ maps canonical factors to canonical factors, all $W_i^{(1)}$ are canonical factors, but the decomposition (2) is of course not necessarily the Δ-normal form of w.

The Δ-normal form of w can then be computed from (2) by successively making decompositions $W_i^{(1)} W_{i+1}^{(1)}$ left-weighted, e.g., using the normal form algorithm in [**CKL$^+$01**]. During this process, a number of decompositions

$$(3) \qquad w = \Delta^{r_l} W_1^{(l)} \cdots W_{s_l}^{(l)} \qquad (l > 1)$$

are formulated such that the l^{th} decomposition of w is derived from the $(l-1)^{\text{st}}$ by making exactly one product $W_i^{(l-1)}W_{i+1}^{(l-1)}$ of canonical factors left-weighted. Denote by LR (with L, R canonical factors) the new left-weighted decomposition of $W_i^{(l-1)}W_{i+1}^{(l-1)}$. If $L = \Delta$, then L can be moved to the left (by conjugating all canonical factors $W_j^{(l-1)}$ for $1 \le j < i$ with Δ and incrementing r_{l-1}). If that happens, or if $R = \varepsilon$, the new decomposition (3) gets shorter in the sense that $s_l < s_{l-1}$.

For some $l = l_0$, this procedure eventually terminates and for that l, the decomposition (3) is the Δ-normal form of w. We claim that for all l with $1 \le l \le l_0$, it holds that

$$(4) \qquad W_{(m-1)s+1}^{(l)} \cdots W_{s_l}^{(l)} \text{ is a tail of } U_1 \ldots U_s.$$

(Note here that by the assumption $s_{l_0} = s' > (m-1)s$ on the canonical length of w and the fact that $s_l \le s_{l-1}$ for all l, we are guaranteed $s_l \ge (m-1)s + 1$.)

Property (4) can be proven for all l by induction:

- For $l = 1$, statement (4) is clear by construction of $W_{(m-1)s+k}^{(1)} = U_k$ ($1 \le k \le s$).
- Assume (4) holds for $l - 1 \ge 1$. Only making a product $W_i^{l-1}W_{i+1}^{l-1}$ left-weighted obviously does not affect property (4). (Here, in the only interesting case $i = (m-1)s$, the very definition of left-weightedness guarantees that (4) still holds.)

  However, by making $W_i^{l-1}W_{i+1}^{l-1}$ left-weighted, the new decomposition LR might collapse in the sense that $L = \Delta$ or $R = \varepsilon$. Then the overall length decreases (i.e., $s_l < s_{l-1}$), canonical factors $W_j^{(l-1)}$ ($i < j \le s_{l-1}$) get shifted left, and possibly a Δ is "moved through" $W_j^{(l-1)}$ ($1 \le j < i$).

  For $i \le (m-1)s$, this does not affect (4), as with $W_{(m-1)s+1}^{(l-1)} \cdots W_{s_{l-1}}^{(l-1)}$ also $W_{(m-1)s+1}^{(l)} \cdots W_{s_l}^{(l)} = W_{(m-1)s+1+k}^{(l-1)} \cdots W_{s_{l-1}}^{(l-1)}$ is a tail of $U_1 \cdots U_s$ (for suitable $k \in \{1, 2\}$). But for $i > (m-1)s$, the case $L = \Delta$ is not possible (because then it would follow that $\Delta \ge U_1 \cdots U_s$, which contradicts the maximality of $r = \inf u$), and the case $R = \varepsilon$ does not affect (4). Taking it all together, (4) thus holds for l, too.

As just shown, (4) holds in particular for $l = l_0$, and it follows that

$$W_{(m-1)s+1}^{(l_0)} \cdots W_{s_l}^{(l_0)} = W_{(m-1)s+1} \cdots W_{s'}$$

is a tail of $U_1 \ldots U_s$. $\qquad \square$

In a nutshell, Proposition 3.1 shows that, given only $w = u^m$, the exponent $m \in \mathbb{N}$ and the canonical length s of u, one can derive a tail of the "canonical part" of u by simply "reading off" canonical factors from the end of the Δ-normal form of w.

3.2. Uncovering the Remaining Part. Assume we have, e.g., using Proposition 3.1, obtained a tail u_R of the "canonical part" $U_1 \cdots U_s$ of u. Then, let $u_L \in B_n^+$ be the remaining canonical part of u, such that

$$(5) \qquad u = \Delta^r u_L u_R$$

with $r = \inf u$ as before assumed public. Suppose further that u_L itself is "only" a canonical factor. This latter assumption might seem a little odd, but in our experiments, this turns out to be actually the case for suggested instances of the REP. Note that this implies $\pi^{-1}(\pi(u_L)) = u_L$ for the canonical homomorphism π from Section 1.2.

Consider the projection

$$\overline{u} = \overline{\Delta}^r \ \overline{u_L} \ \overline{u_R} \tag{6}$$

of (5) into the symmetric group S_n, where we write $\overline{x}$ for $\pi(x)$. Obviously, $\overline{u_L}$, and thus by assumption about u_L also $u_L = \pi^{-1}(\overline{u_L})$, is uniquely determined by $\overline{u}$, $\overline{u_R}$ and r. Since r and u_R (and therefore also $\overline{u_R}$) are known, we only need to find $\overline{u}$ to obtain u_L and hence, using (5), also u (which is our final goal).

We can get $\overline{u}$ by solving the equation

$$\overline{w} = x^m \tag{7}$$

in the symmetric group S_n for unknown x. (Note that w and thus $\overline{w}$ is known.) Unfortunately, a solution to this root extraction problem in the symmetric group need not be unique, so there may be many candidates x for $\overline{u}$. However, for every such candidate, (5) can be checked, where $u_L = \pi^{-1}(\overline{u_L})$ is obtained in turn from (6) by setting $\overline{u} = x$. Of course, any solution to (5) solves already the REP in the braid group. Note also that the "right" $\overline{u}$ is among the candidates x for $\overline{u}$, so eventually a solution will be found.

3.3. Extracting Roots of Permutations. It remains to show how to get *all* solutions $x \in S_n$ to (7) for given $y := \overline{w} \in S_n$ and $m \in \mathbb{N}$. Suppose x consists of disjoint cycles $C_1, \ldots, C_\mu$. Then

$$y = x^m = C_1^m \cdots C_\mu^m,$$

where C_i^m consists of $\gcd(|C_i|, m)$ cycles of length $|C_i|/\gcd(|C_i|, m)$ each.

This allows to treat cycles of y with different lengths separately: let $D_1, \ldots, D_\nu$ be the disjoint cycles of y, and let $\mathcal{L} := \bigcup_{i=1}^{\nu} \{|D_i|\}$ be the set of cycle lengths occurring in y. Then, any m^{th} root x of y can be expressed as

$$x = \prod_{\ell \in \mathcal{L}} x_\ell, \tag{8}$$

where x_ℓ is an m^{th} root of the product

$$y_\ell := \prod_{|D_i|=\ell} D_i$$

of all cycles of y of length ℓ. Conversely, any product as in (8) with $x_\ell^m = y_\ell$ for all ℓ satisfies $x^m = y$. So without loss of generality, we may assume that y contains only cycles of length ℓ, i.e., that $y = y_\ell$ for some ℓ.

An m^{th} root x of y can then be constructed as follows: choose a divisor $a \geq 1$ of m that is coprime to ℓ such that $g := m/a \leq \nu$. Then $g = \gcd(g\ell, m)$, and thus any g cycles $D_{i_1}, \ldots, D_{i_g}$ can be combined into one larger cycle C of length $|C| = g\ell$ such that $C^m = D_{i_1} \cdots D_{i_g}$. This combination can be done (in ℓ^{g-1} different ways) by writing the D_{i_j} suitably "interleaved". (We omit details here.) The process can be repeated until there are no more cycles D_i left. That way, using backtracking, all possible solutions x can be obtained (however, the number of solutions may of course be large).

3.4. Summary of the Algorithm. In summary, our algorithm thus works as depicted in Figure 1. Some comments are in place: the algorithm takes additionally $\inf u$ and $\sup u$ as input. When not already knowing these values (e.g., because they might be fixed in a cryptographic system), they can simply be guessed: for example, $\inf u$ lies between 0 and $\inf w/m$, and the canonical length s of u lies between 0 and the canonical length s' of w. The algorithm can also be run in parallel for all possible guesses.

Steps 1 and 2 of the algorithm make use of Theorem 3.1 to read off a tail u_R of the canonical factors of u from the Δ-normal form of w. Steps 3 and 4 assume that the remaining canonical part u_L of u is only a canonical factor; then, the method described in Section 3.2 is used to reduce the remaining problem to a root extraction problem in the symmetric group S_n. Step 3 can be implemented by, e.g., using the method described in Section 3.3. Note that a potential solution u' to the REP can easily be checked for correctness, and thus the algorithm never generates wrong output (but it may well output "`fail`" to indicate that something went wrong).

Input: $w = u^m \in B_n$, $m \in \mathbb{N}$, $r = \inf u$, canonical length s of u
Output: either $u' \in B_n$ with $u'^m = w$ or "`fail`"

(1) Let $\ell := \max\{0, s' - (m-1)s\}$ for the canonical length s' of w.
(2) Let $u_R \in B_n$ be the product of the last ℓ canonical factors of w's Δ-normal form.
(3) Solve $\pi(w) = x^m$ in the symmetric group S_n for unknown $x \in S_n$; let $\mathcal{X}$ be the set of solutions.
(4) For all $x \in \mathcal{X}$ do
  - let $u_L := \pi^{-1}(\pi(\Delta)^{-r}\, x\, \pi(u_R)^{-1}) \in B_n$
  - let $u' := \Delta^r u_L u_R$
  - if $u'^m = w$, terminate with output w
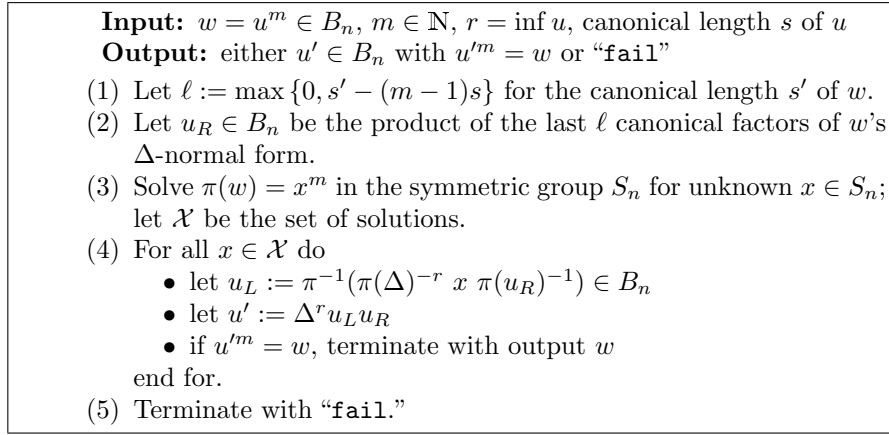end for.
(5) Terminate with "`fail`."

FIGURE 1. Summary of our algorithm

3.5. Improvements and Optimizations. If the algorithm fails, not all is lost. For example, there is a way to make the crucial assumption that the "remaining canonical part" u_L of u is a canonical factor potentially more likely. To see how this works, first observe that to find an m^{th} root of $w \in B_n$, it suffices to find an m^{th} root u' of $w' := \alpha^{-1}w\alpha$, where $\alpha \in B_n$ is arbitrary but known. Namely, then $u := \alpha u' \alpha^{-1}$ is an m^{th} root of w as desired, because

$$u^m = (\alpha u' \alpha^{-1})^m = \alpha u'^m \alpha^{-1} = \alpha\alpha^{-1}w\alpha^{-1}\alpha = w.$$

In other words, we may as well preprocess w by conjugation as long as we keep track of the conjugating braid α. Then it might be useful to *cycle* or *decycle* w (and thus u) so that both end up in their *super summit set* (*SSS*). (Here, cycling and decycling are simply special types of conjugation operations, and the super summit set of a braid w is the subset of w's conjugacy class of braids with minimal canonical length; cf., e.g., [**EM94**] for details. An upper bound on the number of needed cycling/decycling operations is provided in [**BKL01**].) The intuition here is that the SSS contains exactly the "dense" conjugates of u resp. w. One may hope

that for these conjugates (call them u' and w') the canonical length of $w' = u'^m$ is close to m times the canonical length of u'. Given that, one can hopefully read off larger tails of u' from the normal form of w' in Step 2 of the algorithm. Also, this makes $\inf u = \inf w/m$ and $\sup u = \sup w/m$ very likely, so that the algorithm can guess r and s with high probability on its own.

Another rather trivial optimization is to not only consider the Δ-normal form (or, left-normal form), but in addition also the right-normal form of w. An obvious analogue to Theorem 3.1 shows then how to read off a head $u_{L'}$ of the canonical part (in the right-normal form) of u from the right-normal form of w. Then it can be hoped that the remaining canonical part $u_{R'}$ of u is only a canonical factor and the problem can be projected into the symmetric group S_n just as in Section 3.2.

3.6. Experimental Results. We have implemented our algorithm (taking into account the optimizations from Section 3.5) in the computer algebra system MAGMA [**BCP97**] on a standard PC. Several different parameter choices have been tested, and in particular we tried to cover concrete suggestions made in the literature (cf. Section 2.2). In Table 1, we have summarized our success rates for extracting *square* roots for different choices of the braid index n and types of instance generation. The table considers the following types of instances:

Artin: Here the MAGMA command `u:=Random(B,r,r)` is used to choose u as a product of r randomly chosen Artin generators (either positive or negative), and our choices for r and n follow [**SDG02**, Section 4.4].

Canonical: Here the MAGMA command `u:=Random(B,0,0,r,r)` is used to choose u as a product of r randomly chosen canonical factors. The parameters $n = 30$, $r = 15$ are taken from [**SDG02**, Section 4.4].

[Sib03]: Here u, is chosen according to [**Sib03**, Algorithme 3.2.2] with $\ell = r$. Not being aware of a more concrete suggestion, we adopted the parameter choices from the "Canonical" instances.

n	r	Instance Type	Total samples	Success rate
30	15	Canonical	1000	99.5%
30	15	[**Sib03**]	1000	99.6%
30	1000	Artin	1000	96.7%
60	30	Canonical	500	93.6%
60	30	[**Sib03**]	500	94.0%
60	2000	Artin	500	41.0%

TABLE 1. Experimental results.

It should be noted that "success" indicates that the algorithm found a complete square root of w on input u only[1]. The reason for failure was—independently of parameter choice—almost always the fact that there were too many solutions to (7) to handle them efficiently. In other words, the problem was not the computation in the braid group B_n, but the task of extracting roots in the S_n. So our algorithm seems to succeed almost always in reducing a REP in the B_n to a REP in the S_n. However, note that even to find only one solution to a REP in the B_n, we may have to find *all* solutions to a REP in the S_n.

---

[1]$\inf u$ and $\sup u$ were—after (de)cycling w—guessed as $\inf u = \inf w/2$ and $\sup u = \sup w/2$

This also means that the success rate of the algorithm can be raised significantly (at least for $n > 30$) if it is allowed to run longer (and thus to consider more solutions to (7)). For the computations documented in the above table, the algorithm typically succeeded within seconds.

The results seem to generalize to exponents $m > 2$ and larger braid index n, only the number of solutions to (7) grows substantially with the number of divisors of m. So for large m with many divisors, the attack becomes impractical because not all solutions to (7) can be found efficiently. It should be stressed however, that also here the only problem is the number of solutions in the symmetric group; the reduction to the REP in the S_n itself experimentally works almost always.

3.7. Countermeasures. The heuristic algorithm from [**HS02**] for the conjugacy search problem $w = \alpha^{-1} u \alpha$ in braid groups is closely related to our attack. For their attack, it is crucial that the "distance" of canonically obtained representatives of u and w in their super summit set is small. Actually, the countermeasure from [**Sib03**] against that attack consists in generating u, w with large distance in their SSS. However, this approach of [**Sib03**] does not immediately carry over to our setting, since the braids u, w, u_R, and u_L are not conjugated in general. Additionally, they cannot be chosen separately: in fact, w, u_R, and u_L directly depend on u.

Similarly, choosing braids with a large *ultra summit set* (USS, cf. [**Geb03**]) can prove useful to thwart the algorithm of [**Geb03**] for the conjugacy search problem in braid groups. Again, this approach does not directly relate to our setting, since the braids u and u_R (whose "distance" u_L directly influences whether or not our algorithm succeeds) are neither conjugated in general, nor can they both be chosen separately.

Another, more obvious way to defeat the attack is to provoke a situation in which there is a large number of solutions to the REP (7) in the S_n, so that $\pi(u)$ cannot be found efficiently. As just mentioned, this can be done by choosing large m with a large number of divisors. However, this makes, e.g., the authentication system of [**SDG02**] rather inefficient, because the public key then becomes huge.

Note that choosing, e.g., *pure* braids u (i.e., braids u with $\pi(u) = \mathrm{id}$) in itself does not defeat the attack, as then $\pi(u)$, which is the whole goal of solving (7), is public. An ad hoc strategy which for $m = 2$ successfully counters our attack is the following: Select a pure braid u and insert at random positions in u distinct Artin generators such that $\pi(u)$ consists only of cycles of length 2. Then, $\pi(w) = \pi(u^2) = \mathrm{id}$, so the number of solutions to the REP in S_n is large, and guessing $\pi(u)$ becomes more difficult.

While this choice of instances defeats our attack, we do not endorse the security of this construction and think further investigation of the choice of parameters for the REP in B_n is required. (In particular, in face of our heuristic reduction, what role does the canonical length of u play?)

4. Conclusion

We have described a heuristic algorithm for the root problem in braid groups. This algorithm does not solve the root problem in the general case, yet it applies to most of the cases considered for cryptographic purposes. We have run various experiments with parameters proposed for braid group based cryptosystems to back this result.

Furthermore, we believe that our algorithm can be improved to succeed for some parameters not considered yet for cryptographic applications, and it seems an interesting question how to efficiently find cryptographically satisfying instances of the root problem in braid groups.

Acknowledgments

References

[Art25]   Emil Artin, *Theorie der Zöpfe*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **4** (1925), 47–72.

[BCP97]   Wieb Bosma, John Cannon, and Catherine Playoust, *The Magma algebra system I: The user language*, Journal of Symbolic Computation **24** (1997), 235–265.

[BKL01]   Joan S. Birman, Ki Hyoung Ko, and Sang Jin Lee, *The infimum, supremum, and geodesic length of a braid conjugacy class*, Advances in Mathematics **164** (2001), 41–56, Online available at `http://mail.konkuk.ac.kr/~sangjin/webfiles/articles/geolen.pdf`.

[CJ03]   Jung Hee Cheon and Byungheup Jun, *A polynomial time algorithm for the braid Diffie-Hellman conjugacy problem*, Advances in Cryptology, Proceedings of CRYPTO 2003 (Dan Boneh, ed.), Lecture Notes in Computer Science, no. 2729, Springer-Verlag, 2003, Online available at `http://eprint.iacr.org/2003/019.ps`, pp. 212–225.

[CKL+01] Jae Choon Cha, Ki Hyoung Ko, Sang Jin Lee, Jae Woo Han, and Jung Hee Cheon, *An efficient implementation of braid groups*, Advances in Cryptology, Proceedings of ASIACRYPT 2001 (Colin Boyd, ed.), Lecture Notes in Computer Science, no. 2248, Springer-Verlag, 2001, Online available at `http://crypt.kaist.ac.kr/pre_papers/braid-impl.ps`, pp. 144–156.

[Deh04]   Patrick Dehornoy, *Braid-based cryptography*, Group Theory, Statistics, and Cryptography (Alexei G. Myasnikov, ed.), Contemporary Mathematics, no. 360, American Mathematical Society, 2004, Online available at `http://www.math.unicaen.fr/~dehornoy/Surveys/Dgw.ps`, pp. 5–33.

[EM94]   Elsayed A. Elrifai and H. R. Morton, *Algorithms for positive braids*, Quarterly Journal of Mathematics **45** (1994), 479–497, Online available at `http://www.liv.ac.uk/~su14/papers/EM04.ps.gz`.

[FS87]   Amos Fiat and Adi Shamir, *How to prove yourself: Practical solutions to identification and signature problems*, Advances in Cryptology, Proceedings of CRYPTO '86 (Andrew M. Odlyzko, ed.), Lecture Notes in Computer Science, no. 263, Springer-Verlag, 1987, pp. 186–194.

[Gar69]   Frank A. Garside, *The braid group and other groups*, Quarterly Journal of Mathematics **20** (1969), 235–254.

[Geb03]   Volker Gebhardt, *A new approach to the conjugacy problem in Garside groups*, Journal of Algebra (2003), To be published, online available at `http://arxiv.org/ps/math.GT/0306199`.

[GKT+05] David Garber, Shmuel Kaplan, Mina Teicher, Boaz Tsaban, and Uzi Vishne, *Probabilistic solutions of equations in the braid group*, Advances in Applied Mathematics **35** (2005), no. 3, 323–334, Online available at `http://arxiv.org/ps/math.GR/0404076`.

[GM03]   Juan González-Meneses, *The nth root of a braid is unique up to conjugacy*, Algebraic and Geometric Topology **3** (2003), 1103–1118, Online available at `http://www.maths.warwick.ac.uk/agt/ftp/main/2003/agt-3-39.ps`.

[GMR89]   Shafi Goldwasser, Silvio Micali, and Charles Rackoff, *The knowledge complexity of interactive proof systems*, SIAM Journal on Computing **18** (1989), no. 1, 186–208.

[HS02]   Dennis Hofheinz and Rainer Steinwandt, *A practical attack on some braid group based cryptographic primitives*, Public Key Cryptography, Proceedings of PKC 2003 (Yvo Desmedt, ed.), Lecture Notes in Computer Science, no. 2567, Springer-Verlag, 2002, pp. 187–198.

[Hug02]   Jim Hughes, *A linear algebraic attack on the AAFG1 braid group cryptosystem*, Information Security and Privacy, Proceedings of ACISP 2002 (Lynn Batten and Jennifer Seberry, eds.), Lecture Notes in Computer Science, no. 2384, Springer-Verlag, 2002, Online available at `http://www.network.com/hughes/ACISP02.pdf`, pp. 176–189.

[KLC+00] Ki Hyoung Ko, Sang Jin Lee, Jung Hee Cheon, Jae Woo Han, Ju-sung Kang, and Choonsik Park, *New public-key cryptosystem using braid groups*, Advances in Cryptology, Proceedings of CRYPTO 2000 (Mihir Bellare, ed.), Lecture Notes in Computer Science, no. 1880, Springer-Verlag, 2000, pp. 166–183.

[LC05]   Sunder Lal and Atul Chaturvedi, *Authentication schemes using braid groups*, lanl.arXiv.org ePrint Archive, July 2005, Online available at `http://arxiv.org/pdf/cs.CR/0507066`.

[Lee04]   Sang Jin Lee, *Growth of minimal word-length in Garside groups*, lanl.arXiv.org ePrint Archive, November 2004, Online available at `http://arxiv.org/ps/math.GT/0411470`.

[SDG02]   Hervé Sibert, Patrick Dehornoy, and Marc Girault, *Entity authentication schemes using braid word reduction*, Discrete Applied Mathematics (2002), To be published, online available at `http://eprint.iacr.org/2002/187.ps`.

[Sib03]   Hervé Sibert, *Algorithmique des groupes de tresses*, Ph.D. thesis, Université de Caen, 2003, Online available at `http://www.math.unicaen.fr/~sibert/These.pdf`.

[Sty79]   V.B. Styshnev, *The extraction of a root in a braid group*, Mathematical of the USSR, Izvestija **13** (1979), 405–416.

[SU04]   Vladimir Shpilrain and Alexander Ushakov, *The conjugacy search problem in public key cryptography: unnecessary and insufficient*, IACR ePrint Archive, November 2004, Online available at `http://eprint.iacr.org/2004/321.pdf`.

[Tsa05]   Boaz Tsaban, *On an authentication scheme based on the root problem in the braid group*, lanl.arXiv.org ePrint Archive, September 2005, Online available at `http://arxiv.org/ps/cs.CR/0509059`.

Institut für Algorithmen und Kognitive Systeme, Universität Karlsruhe, Am Fasanengarten 5, 76131 Karlsruhe, Germany
  E-mail address: `groch@ira.uka.de`

Centrum voor Wiskunde en Informatica, Cryptology and Information Security Group, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands
  E-mail address: `Dennis.Hofheinz@cwi.nl`

Department of Mathematical Sciences, Florida Atlantic University, 777 Glades Road, Boca Raton, FL 33431, USA
  E-mail address: `rsteinwa@fau.edu`