The Complexity of the A B C Problem Resolved

Jin-yi Cai*
Department of Computer Science
SUNY at Buffalo
Buffalo, NY 14260
cai@cs.buffalo.edu

Richard J. Lipton†
Department of Computer Science
Princeton University
Princeton, NJ 08544
rjl@cs.princeton.edu

Yechezkel Zalcstein
Division of Computer and Computation Research
National Science Foundation
4201 Wilson Boulevard
Arlington, VA 22230
zzalcste@nsf.gov

1 Introduction

The membership problem for groups is the following. Given a finite number of elements $a_1, \ldots, a_k$ of a group G and a target element g, does g belong to the subgroup generated by $a_1, \ldots, a_k$? The most useful and interesting representation of groups is by matrices over a field. Most groups occurring in physics and many finite simple groups arise as groups of matrices. Thus in this case the group G is the group of all invertible matrices (of a fixed dimension) over a field. However, in this generality the problem is undecidable [21], even for groups of 4 by 4 matrices. Thus work on this problem has focused on *finite* groups. Babai and Szemeredi [4] have shown that, for finite groups of matrices over a finite field the problem is in NP. This problem has also motivated the introduction of the complexity class AM [1]. Since, as has been observed in [4], the membership problem for finite matrix groups over finite fields is at least as difficult as the discrete logarithm problem, most research has concentrated on groups represented by *permutations* rather than matrices. Over the years, this research has culminated in efficient sequential and parallel algorithms [2, 3, 9, 11]. The case of a matrix representation turned out to be much harder and has been solved only very recently by Luks [19] for solvable groups over restricted finite fields and by Beals and Babai [5] for arbitrary finite groups. The latter algorithm is Las Vegas. For semigroups the problem is even harder. It is known to be undecidable even if the target element g is the zero matrix and the matrices are 3 by 3 [22] . For finite semigroups the problem is PSPACE-complete [13] and even for finite *commutative* semigroups it is NP-hard [6].

*Research supported in part by NSF grants CCR 9057486 and CCR 9319093, and an Alfred P. Sloan Fellowship.
†Research supported, in part, by NSF grant CCR-9304718.

In this paper, we extend previous results in two directions. First, we allow infinite systems, where even decidability is not clear, and second, we consider semigroups as well. On the other hand, since the general problem is undecidable, some restrictions are needed. In 1980, Kannan and Lipton [15] solved the following *orbit problem*, the case of a *cyclic* semigroup, by giving a polynomial time algorithm:

> Given two commuting matrices A and B over the rational numbers, does there exist a nonnegative integer i, such that $A^i = B$?

The following *generalized orbit problem*, or *the A B C problem* has remained open since 1980:

> Given commuting matrices A, B and C over the rational numbers, does there exist nonnegative integers i and j, such that $A^i B^j = C$?

A host of other problems are reducible to the orbit problem [15]. In this paper, we resolve the A B C problem, by giving a polynomial time algorithm which not only decides the solvability of a given instance A, B and C, but also finds all solutions (in a polynomial sized encoding) if any solution exists. We also obtain the corresponding result for the group case, as a corollary. We would like to point out that the semigroup problem is more complicated, even when the generating matrices are invertible, since we do not allow the use of the inverse operation. The results are new, even in the group case. Prior to this work, the best upper bound for the problem in the group case is that it is decidable, a result by Kopytov [12]. Nothing was previously known for the semigroup case. The proof in [12], which constructs "yes" and "no" lists, does not give any complexity bounds. Our result also generalizes to apply to the case where the underlying field is an algebraic number field. We hope that our techniques can be generalized to the case with k matrices. However, the number of generators must be fixed, otherwise, as noted above, the problem is already NP-hard in the finite case.

We will explain briefly why the problem is challenging. Let us assume first (this is the easiest case) that the matrices A, B and C are all diagonalizable. Since they commute they must be simultaneously diagonalizable. In this case the problem is reduced to n instances of the following problem (n is the size of the matrices): Given algebraic numbers α, β and γ, are there non-negative integers i, j such that $\alpha^i \beta^j = \gamma$? Or more generally, for a fixed k, given algebraic numbers $\alpha_1, \ldots, \alpha_k$, and η, are there non-negative integers $i_1, \ldots, i_k$ such that

$$\alpha_1^{i_1} \ldots \alpha_k^{i_k} = \eta \tag{1}$$

holds? Even the much easier problem of *checking*, given $\alpha_1, \ldots, \alpha_k$, $i_1, \ldots, i_k$ and η, if (1) holds is nontrivial and has been solved only recently [10]. In [15] for the case $k = 1$, a polynomial time algorithm is given. In case some of the α's are not units, one can use Kummer's unique factorization of ideals, and apply an argument based on the exponential growth of the norms $[\mathbf{N}(\alpha_i)]^i$. However, the hard case is when all the α's are units. In [15], this was also the hard case, and a deep result of Blanksby and Montgomery [7] was needed to handle it. The result of [7] is inadequate, however, for the case $k \geq 2$. Another natural attempt is to use Dirichlet's Unit Theorem to decompose the α's in terms of fundamental units. However, the complexity of computing such decompositions has not, to our knowledge, been analyzed and it may not be computationally tractable. Furthermore, Dirichlet's theorem handles only the group case. Finally, in reducing the general case to the diagonalizable case, Dirichlet's theorem is not enough. We need a description of all solutions of (1). Such a description has been obtained only recently by Masser [20].

If at least one of the matrices is not diagonalizable, it is tempting to believe that, since the matrices commute, they can be brought simultaneously to Jordan canonical form (this misconception is held even by some algebraists!). If this were true, an extension of the argument in the diagonalizable case could possibly be constructed. However, this is not true, as the example in the Appendix shows. The above discussion indicates some of the difficulties inherent in this problem.

Our results are:

Theorem 1.1 *There is a polynomial-time algorithm for the following problem: Given k algebraic numbers $\alpha_1, \ldots, \alpha_k$, and η, are there non-negative integers $i_1, \ldots, i_k$ such that Equation (1) holds? If so, find all solutions (in a polynomial sized encoding).*

Theorem 1.2 *There is a polynomial-time algorithm to decide the solvability of the A B C problem; if it is solvable find all solutions (in a polynomial sized encoding).*

Theorem 1.3 *There is a polynomial-time algorithm to solve the A B C problem for the group case, i.e., the exponents are allowed to be arbitrary integers.*

The plan for this paper is as follows. The failure of simultaneous Jordanization is one of the main difficulties in this problem. We will introduce a method of successive restriction to handle this. This technique, coupled with Theorem 3.2 proved using Masser's Theorem, form the backbone of our proof. In Section 2 we will deal with some preliminary issues. Proofs in Section 2 are omitted. In Section 3 we prove Theorem 3.2. In Section 4, we first give an illustration of the power of our new techniques in a simple setting, by giving a transparent proof of the Kannan-Lipton Theorem. Then we use our method of successive restriction to complete our proof of the A B C problem.

2 Preliminaries

Let A, B and $C \in \mathbf{Q}^{n \times n}$ be $n \times n$ rational matrices. The input size to the problem is n plus the sum of binary lengths of all entries. By an extension of a technique from [15], we may assume that C is given as a polynomial of A and B with rational coefficients. This is certainly a necessary condition, and whether C is such a polynomial can be decided in polynomial time, as it is a question of linear dependence over the rationals [14]. Here we use the Cayley-Hamilton theorem $\chi_A(A) = \chi_B(B) = 0$, to limit the power of A and B to at most $n-1$, where χ_A and χ_B are characteristic polynomials of A and B respectively.

Thus, we are given $A, B \in \mathbf{Q}^{n \times n}$ such that A, B commute, and a polynomial p in A and B. The question is to decide if p can be expressed as nonnegative powers of A and B. In [15], the computation of eigenvalues could be circumvented, but this cannot be done here. We will compute in various number fields, mostly symbolically, i.e., we will have an irreducible polynomial f (proven so by the L^3–algorithm), and each element in the field will be represented by a polynomial of degree $< \deg(f)$ [17, 16]. This can be done over any number fields as well, not just $\mathbf{Q}$. However, an important issue is that we can not allow the degree of the extension over $\mathbf{Q}$ to be too large; in general, to stay within polynomial time, we cannot operate in the splitting field of an irreducible polynomial of degree say, n, since the Galois group structure of this extension field might be too large. It is not known how to compute the Galois group in polynomial time in general, and in fact it is believed to be not possible [18].

We will need to compute the Jordan Normal Form (JNF) for either matrices A and B. As noted in the Appendix, even though A and B commute, it does not follow that they can be simultaneously put in JNF. What we will show in the paper is a technique of successive restrictions, which is of independent interest, and is almost as good for our purposes, as simultaneous JNF. But before we get to that, let's first ask can we actually compute a basis change for the JNF of a matrix A in polynomial time? The difficulty is that if we use the standard textbook algorithm, we will be dealing with the splitting field of χ_A over $\mathbf{Q}$. Thus the standard algorithm will not do. However there is a polynomial time algorithm [8] that finds an invertible matrix T, such that $T^{-1}AT = J_A$, the JNF for A. The matrix T has entries from the splitting field of χ_A, but the trick is that in [8] we can find a T, where every entry t in fact belongs to a much smaller extension field, which varies with t, and can be computed in polynomial time within the smaller field. The matrix J_A is *not* computed from the *product* $T^{-1}AT$; in fact, computing T^{-1} from T, say using Gaussian elimination, involves the splitting field again. However a different technique is used in [8] to compute J_A, and to compute T^{-1} from T, without ever going to the splitting field, all accomplished in polynomial time.

Next there is the issue of how to apply this basis transformation to another matrix M. The fact is, for a general matrix M, we cannot do it in polynomial time. The same difficulty with the splitting field gets in the way. *But we can do it for the matrix B and C*, using the fact that they commute with A. It is an interesting situation where we have the transformation matrix in hand, and we can apply to some matrix M, but not to all. It is only in this limited sense that we have obtained this basis transformation computationally. Fortunately, we *can* apply to those which we care about.

There is another technical issue concerning conjugates of the same irreducible polynomial. When it is necessary to distinguish one root from another, we will use a "good enough" rational approximation, which gives a polynomial number of bits, and uniquely identifies the root on the complex plane. This uses Liouville's theorem. The details of all these matters would take up too much space, and are omitted from this extended abstract.

In the following we will speak freely of computing JNF, and invariant subspaces, etc, all in polynomial time, without further comment. After some preparation, we can arrive at the following situation: We have a decomposition of $\mathbf{C}^n$ as a direct sum of subspaces $V_{\lambda,\mu}$, invariant for both A and B, where λ and μ range over $Spec(A)$ and $Spec(B)$, i.e., all eigenvalues of A and B, such that there is just one eigenvalue λ (and μ, respectively) for A (and B, respectively) on $V_{\lambda,\mu}$:

$$\mathbf{C}^n = \bigoplus_{\lambda \in Spec(A), \mu \in Spec(B)} V_{\lambda,\mu}.$$

Lemma 2.1 *Suppose A, B commute and C is a polynomial in A and B, then for any i, j,*

$$A^i B^j = C \iff \quad \forall \lambda \in Spec(A), \mu \in Spec(B), \quad (A\mid_{V_{\lambda,\mu}})^i (B\mid_{V_{\lambda,\mu}})^j = C\mid_{V_{\lambda,\mu}} .$$

For what follows, we will fix any pair of eigenvalues $\lambda \in Spec(A)$ and $\mu \in Spec(B)$, and consider $A\mid_{V_{\lambda,\mu}}$ and $B\mid_{V_{\lambda,\mu}}$. We will call them A and B respectively, when no confusion arises.

3 The affine lattice of solutions

We will need the following special case of a very general theorem of Masser [20].

Theorem 3.1 *Fix any $k \geq 1$. Let $\alpha_1, \ldots, \alpha_k$ be non-zero algebraic numbers over $\mathbf{Q}$. Let $D = [\mathbf{Q}(\alpha_1, \ldots, \alpha_k) : \mathbf{Q}]$ be the degree of the algebraic extension, and let h be the maximum height* [1] *of $\alpha_1, \ldots, \alpha_k$ over $\mathbf{Q}$. Then, the* relation group

$$L = \{\ (e_1, \ldots, e_k) \in \mathbf{Z}^k \mid \alpha_1^{e_1} \cdots \alpha_k^{e_k} = 1\ \}$$

is a lattice with a small basis. *More specifically, L has a generating set $v_1, \ldots, v_\ell \in \mathbf{Z}^k$, $1 \leq \ell \leq k$, such that the maximum entry in these vectors satisfy*

$$\max_{1 \leq i \leq \ell\ 1 \leq j \leq k} |v_{i,j}| \leq (ckh)^{k-1} D^{k-1} \frac{(\log(D+2))^{3k-3}}{(\log\log(D+2))^{3k-4}},$$

where c is some absolute constant.

In order to appreciate this remarkable theorem of Masser, let's look at its implication on the computational complexity of finding such a basis. Note that in the inequality, the left hand side refers to the *quantity* $v_{i,j}$, while the right hand side essentially refers to the binary length of the input data. (Recall that k is fixed.) Thus, an exhaustive search can be done in polynomial time to find a small basis.

We call a shift of a lattice $L + v$, for some $v \in \mathbf{Z}^k$, an *affine lattice*. We say that an affine lattice has *small description*, if it is $L + v$ for some small v and L has a small basis, here small means all entries are polynomially bounded in *quantity* (not merely in binary length).

Theorem 3.2 *Given non-zero algebraic numbers $\alpha, \beta_1, \ldots, \beta_k$, the set*

$$\{\ (j_1, \ldots, j_k) \mid \alpha\beta_1^{j_1} \cdots \beta_k^{j_k} = 1\ \}$$

is either empty, or is an affine lattice with rank at most k and a small description. Moreover, it is decidable in polynomial time whether it is empty, and if not, to compute a small description in terms of an off-set vector v and a small basis. Finally, if the rank of the affine lattice is k, then all $\alpha, \beta_1, \ldots, \beta_k$ are roots of unity.

Proof: Consider the lattice

$$L = \{\ (i, j_1, \ldots, j_k) \mid \alpha^i \beta^{j_1} \cdots \beta^{j_k} = 1\ \}.$$

By Masser's Theorem, and an exhaustive search, we get a small basis $v_1, \ldots, v_\ell$, $\ell \leq k+1$. Now we wish to intersect this lattice with the affine lattice $\{(1, j_1, \ldots, j_k) \mid j_1, \ldots, j_k \in \mathbf{Z}\}$.

[1] The height function of an algebraic number is essentially the sum of the degree, and the binary length of all coefficients in the defining equation over $\mathbf{Q}$.

Write

$$\begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_\ell \end{pmatrix} = \begin{pmatrix} v_{10} & v_{11} & \dots & v_{1k} \\ v_{20} & v_{21} & \dots & v_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ v_{\ell 0} & v_{\ell 1} & \dots & v_{\ell k} \end{pmatrix},$$

where $v_{ij} \in \mathbf{Z}$ and $|v_{ij}| \leq n^{O(1)}$.

We can first perform a basis reduction to transform the basis to a so-called canonical basis in the sense of Smith and Hermite. In other words, we can compute the gcd of the first column of the matrix, and by elementary row transformations to arrive at a matrix whose first row has at most one (the first) entry to be non-zero. Continue on for the second column etc, we can arrive at the following form

$$\begin{pmatrix} v_{10} & v_{11} & \dots & v_{1k} \\ 0 & v_{21} & \dots & v_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \cdot \end{pmatrix},$$

where $0 \leq v_{10}$, $0 \leq v_{11} \leq v_{21}$, etc. Since k is fixed, there is no problem with respect to the size of the entries; they will remain polynomially bounded in quantity. The row vectors still form a basis for L. We will still call them $v_1, \ldots, v_\ell$.

Now it is clear that this lattice intersects with $\{(1, j_1, \ldots, j_k) \mid j_1, \ldots, j_k \in \mathbf{Z}\}$ iff $v_{10} = 1$, and if so, then the intersection is $L' + v_1$, where L' is the lattice spanned by $v_2, \ldots, v_\ell$.

If the rank of the affine lattice is k, then $\ell = k + 1$. The basis matrix is a square matrix and all diagonal entries are non-zero. In particular, $v_{kk} \neq 0$, and $\beta_k^{v_{kk}} = 1$. Thus, β_k is a root of unity. Continuing, we have $v_{k-1,k-1} \neq 0$, and $\beta_{k-1}^{v_{k-1,k-1}} = \beta_k^{-v_{k-1,k}}$, which is a root of unity. Thus β_{k-1} is also a root of unity. The proof is completed by an easy induction. $\square$

If any of the algebraic numbers $\alpha, \beta_1, \ldots, \beta_k$ are 0, we will adopt the convention that $0^0 = 1$ and $0^i = 0$, for $i > 0$. Under this convention, the statement in Theorem 3.2 can be easily adapted to allow 0's.

4 The proof

We will first illustrate the new technique, in a simple setting, by demonstrating how swiftly the Kannan-Lipton Theorem [15] can be proved using this technique.

4.1 The Kannan-Lipton Theorem

Put A in JNF. Restrict to a subspace V_λ, where the only eigenvalue for A on V_λ is λ.

- Case 1. $\lambda = 0$. Then $A|_{V_\lambda}$ is nilpotent, $(A|_{V_\lambda})^n = 0$. Thus we need only check the cases $(A|_{V_\lambda})^i = (B|_{V_\lambda})$, for $0 \leq i \leq n$.

- Case 2. $\lambda \neq 0$. Consider the solutions to $\lambda^i = c$, which is obtained by looking at any diagonal entry in $(A|_{V_\lambda})^i = B|_{V_\lambda}$. If the solution set is empty, we are done. Suppose the solution set is an affine lattice of rank 0, i.e., there is only one small number

$i \leq n^{O(1)}$ possible, we can directly check it. The only real case is rank 1. Then by Theorem 3.2, λ is a root of unity . So must be c, hence $c \neq 0$. The order of λ must be polynomially bounded, since $\chi_A(\lambda) = 0$.

1. $A|_{V_\lambda}$ is diagonal. Then it is completely determined by Theorem 3.2.
2. $A|_{V_\lambda}$ is not diagonal. Then some Jordan block has size ≥ 2. We can symbolically compute the powers of $A|_{V_\lambda}$, up to $(A|_{V_\lambda})^i$, for $i \leq 2^{n^{O(1)}}$. This uses the fact that λ is a root of unity , and we know the close form formula for powers of JNF. Thus on the diagonal we get equation $\lambda^i = c$, and on the off-diagonal we get an equation $\binom{i}{1}\lambda^{i-1} = c'$. This gives a necessary condition that $i = \lambda c'/c$. If it is not a nonnegative integer, then there is no solution. Otherwise, we simply check directly.

The Kannan-Lipton Theorem is proved.

4.2 The method of successive restriction

Now the setting is $V_{\lambda,\mu}$. We will write A for $A|_{V_{\lambda,\mu}}$. Similarly for B. There will be four cases, depending on whether $\lambda = 0$ and/or $\mu = 0$. The cases when at least one of the eigenvalue is 0 are in fact simpler, since then the matrix is nilpotent, so that we need only check powers up to n. (Recall that on $V_{\lambda,\mu}$, each matrix has exactly one eigenvalue.) We will omit the details on these three cases, and assume $\lambda \neq 0$ and $\mu \neq 0$. If both A and B are diagonal matrices λI and μI, then this case is completely determined by Theorem 3.2. Suppose A is not diagonal. Put A in its JNF. Since $A \neq \lambda I$, at least one block is of dimension > 1.

Let
$$V' = \ker(\lambda I - A),$$
and
$$V'' = \ker(\lambda I - A)^2.$$
Since A and B commute, both V' and V'' are invariant subspaces of B as well as A. In terms of the Jordan form of A, V' corresponds to the first rows and columns of each λ-block. And V'' corresponds to the first and the second rows and columns (whenever a second row and column exists) of each λ-block. Since at least one block is of dimension > 1, $V' \neq V''$.

Let $n_1 = \dim V'$, and $n_1 + n_2 = \dim V''$, then $n_1 \geq n_2 \geq 1$. Let $v_{11}, \ldots, v_{1k_1}, v_{21}, \ldots, v_{2k_2}, \ldots, v_{\ell 1}, \ldots, v_{\ell k_\ell}$ be the basis vectors for which A has its Jordan form, where $v_{i1}, \ldots, v_{ik_i}$ is the basis vectors corresponding to the ith block. We also assume that the blocks are ordered so that $k_1 \geq \ldots \geq k_\ell$. Then $v_{11}, \ldots, v_{\ell 1}$ spans the eigenspace $V' = \ker(\lambda I - A)$, and consequently $\ell = n_1$. By assumption $k_1 \geq 2$, and let $\ell' = \max\{i \mid k_i \geq 2\}$, then $v_{11}, \ldots, v_{\ell,1}, v_{12}, \ldots, v_{\ell' 2}$ spans V'' and $\ell' = n_2$.

Consider $A \mid_{V''}$ and $B \mid_{V''}$. If we order the basis vectors as $v_{11}, \ldots, v_{\ell,1}, v_{12}, \ldots, v_{\ell' 2}$, then A has the form
$$A = \begin{pmatrix} \lambda I_{n_2} & 0 & I_{n_2} \\ 0 & \lambda I_{n_1 - n_2} & 0 \\ 0 & 0 & \lambda I_{n_2} \end{pmatrix}.$$

By decomposing B in blocks of the same dimensions, one can easily verify that

A and B commute $\iff$ B is of the following block upper triangular form $B = \begin{pmatrix} X & Y & Z \\ 0 & U & V \\ 0 & 0 & X \end{pmatrix}$, where X, Y, Z, U, V are of the appropriate dimensions. (If $n_1 = n_2$ then the middle blocks $\lambda I_{n_1-n_2}$ in A, and Y, U and V in B are understood not to appear.)

4.2.1 Restriction on V'

Choose a basis transformation of V' such that $B\mid_{V'}$ is in its Jordan form. Note that on V', $A\mid_{V'}$ is the scalar matrix λI_{n_1} which is unchanged under all basis transformations. Thus, under this basis, $A\mid_{V'}$ is still the scalar matrix λI_{n_1}, and $B\mid_{V'}$ is a direct sum of Jordan blocks all of which have μ on its diagonal, since $Spec(B\mid_{V'}) = \{\mu\}$. By comparing eigenvalues, we get $\lambda^i\mu^j = c$. By Theorem 3.2 we get either an empty set of solutions, or an affine lattice of small description. If it is empty, we are done; or the rank is 0, we can directly check it. If the rank is 1, then we can reduce it to Kannan-Lipton's Theorem. In fact, let $i = i_0 + as$ and $j = j_0 + bs$, for small i_0, j_0, a, b, then we have to solve for s in $(A^aB^b)^s = (CA^{-i_0}B^{-j_0})$.

Now suppose the rank is 2. Then both λ and μ are roots of unity. We note that in this case we can compute high powers of A and B, up to $2^{n^{O(1)}}$, in polynomial time. This is accomplished separately for A and B by first putting the target matrix in JNF, then taking its power, and finally revert it back.

Suppose $B\mid_{V'}$ is not a scalar matrix μI. Then some μ-block of $B\mid_{V'}$ has dimension > 1. Then consider

$$U' = \ker(\mu I - B\mid_{V'}),$$

and

$$U'' = \ker(\mu I - B\mid_{V'})^2.$$

As before, both U' and U'' are invariant subspaces of A as well as B. Since at least one block of $B\mid_{V'}$ is of dimension > 1, $U' \neq U''$.

Let $m_1 = \dim U'$ and $m_1 + m_2 = \dim U''$, then $m_1 \geq m_2 \geq 1$, and if we order the basis vectors appropriately, we have $A\mid_{U''} = \lambda I_{m_1+m_2}$ and

$$B\mid_{U''} = \begin{pmatrix} \mu I_{m_2} & 0 & I_{m_2} \\ 0 & \mu I_{m_1-m_2} & 0 \\ 0 & 0 & \mu I_{m_2} \end{pmatrix}.$$

(Again, if $m_1 = m_2$, then the middle blocks are understood not to appear.)

Implied by the forms of $A\mid_{U''}$ and $B\mid_{U''}$ is the fact that the first and the last m_2 basis vectors together generate an invariant subspace for both A and B. By focusing on this subspace, we get a necessary condition of $A^iB^j = C$ in the form of a pair of equations

$$\lambda^i\mu^j = c, \qquad j\lambda^i\mu^{j-1} = c'.$$

This gives us, upon substitution, $j = \mu c'/c$, which reduces the problem to the one-variable case.

4.2.2 Restriction on V''

Now we suppose on V', $B\mid_{V'} = \mu I_{n_1}$, thus on V'',

$$B\mid_{V''} = \begin{pmatrix} \mu I_{n_2} & 0 & Z \\ 0 & \mu I_{n_1-n_2} & V \\ 0 & 0 & \mu I_{n_2} \end{pmatrix}.$$

(Once again, the middle blocks disappear if $n_1 = n_2$.)

Suppose $n_1 > n_2$. $V^* = \text{span}\{v_{n_2+1\,1}, \ldots, v_{n_1 1}\}$, the subspace corresponding to the middle blocks, is an invariant subspace for both A and B. Consequently we may consider the quotient space V''/V^*, and the induced action of A and B on this quotient space, denoted by $\tilde{A}$ and $\tilde{B}$ respectively. (If $n_1 = n_2$, then $V^* = 0$ and $V''/V^* = V''$.) Under the induced basis from $\{v_{11}, \ldots, v_{n_2,1}, v_{1,2}, \ldots, v_{n_2,2}\}$ for V''/V^*,

$$\tilde{A} = \begin{pmatrix} \lambda I_{n_2} & I_{n_2} \\ 0 & \lambda I_{n_2} \end{pmatrix},$$

and

$$\tilde{B} = \begin{pmatrix} \mu I_{n_2} & Z \\ 0 & \mu I_{n_2} \end{pmatrix}.$$

If $Z = 0$ then we will get a necessary condition

$$\lambda^i \mu^j = c, \quad i\lambda^{i-1}\mu^j = c'.$$

This gives us, upon substitution, $i = \lambda c'/c$, which reduces the problem to the one-variable case.

If $Z \neq 0$ then we will get a necessary condition

$$\lambda^i \mu^j = c, \quad j\lambda^i \mu^{j-1} z + i\lambda^{i-1}\mu^j = c',$$

for some known algebraic numbers c, c' and $z \neq 0$. Upon substitution, we get

$$\frac{i}{\lambda} + z\frac{j}{\mu} = \frac{c'}{c}.$$

Solving i in terms of j,

$$i = \frac{\lambda c'}{c} - \frac{z\lambda}{\mu} j.$$

If $\frac{z\lambda}{\mu}$ is not rational, then we will have at most one pair of integral solution (i, j), and we can easily find it and test it. If $\frac{z\lambda}{\mu}$ is rational, then unless $\lambda c'/c$ is rational, there is no solution, and if $\lambda c'/c$ is rational, then there is an integral relationship between i and j,

$$ui + vj = w.$$

We can solve this equation in general form, $i = i_0 + v_1 t$ and $j = j_0 - u_1 t$, where $u_1 = u/\gcd(u, v)$ and $v_1 = v/\gcd(u, v)$. Substituting back in $A^i B^j = C$, we get $(A^{i_0} B^{j_0})(A^{v_1} B^{-u_1})^t = C$, which reduces it to the case with one variable t.

Acknowledgement

We would like to thank Ravi Kannan, Susan Landau, Laci Lovász and Andrew Odlyzko for many helpful discussions. We would also like to thank Mike Robinson for the reference to Masser's paper.

References

[1] L. Babai, Trading group theory for randomness, STOC 1985,421-429.

[2] L. Babai, E. Luks and A. Seress, Permutation groups in NC, STOC 1987, 272-282.

[3] L. Babai, G. Cooperman, L. Finkelstein, E. Luks and A. Seress, Fast monte carlo algorithms for permutation groups, STOC 1991, 90-100.

[4] L. Babai and E. Szemeredi, On the complexity of matrix group problems I, FOCS 1984,229-240.

[5] R. Beals and L. Babai, Las vegas algorithms for matrix groups, FOCS 1993, 427-436.

[6] M. Beaudry, Membership testing in commutative transformation semigroups, Information and Computation, Vol 79 (1988), 84–93.

[7] P. E. Blanksby and H. .L. Montgomery, Algebraic integers near the unit circle. Acta Arithmetica, (1971), 355–369.

[8] J. Cai, Computing Jordan Normal Forms exactly for commuting matrices in polynomial time. SUNY Buffalo CS Tech Report 94-16.

[9] M. Furst, J. Hopcroft and E. Luks, "Polynomial time algorithms for permutation groups", Proc. 21st IEEE FOCS (1980), pp60–78.

[10] G. Ge, Testing equalities of multiplicative representations in polynomial time, FOCS 1993, 422-426.

[11] D.E. Knuth, "Notes on Efficient Representation of permutation groups", (1981), unpublished manuscript.

[12] L. Kopytov, On the occurrence problem for groups, Algebra and Logic, 1968.

[13] D. Kozen, Lower bounds for natural proof systems, FOCS 1977,254-266.

[14] R. Kannan, The size of numbers in the analysis of certain algorithms, PH. D. Dissertation, Operations Research Dept., Cornell University, Ithaca, NY. 1980.

[15] R. Kannan and R. Lipton, The orbit problem is decidable, STOC 1980, 252-261. See also "Polynomial-time algorithms for the orbit problem", JACM, Vol 33, No. 4, 1986, 808–821.

[16] R. Kannan, A. K. Lenstra, L. Lovász, "Polynomial factorization and non-randomness of bits of algebraic numbers and certain transcendental numbers", Math. of Comp., Vol. 50, No. 181, Jan. 1988, 235–250.

[17] A. K. Lenstra, H. W. Lenstra, L. Lovász, Factoring polynomials with rational coefficients, Math. Ann. 261:515-534, 1982.

[18] S. Landau, personal communications.

[19] E. Luks, Computing in solvable matrix groups, FOCS 1992, 111-120.

[20] D. W. Masser, Linear relations on algebraic groups, in *New Advances in Transcendence Theory*, A. Baker,ed. *Cambridge University Press*, 1988, 248-262.

[21] K. A. Mihailova, The occurrence problem for a direct product of groups, Doklady Akads. Nauk (USSR), 119 (1958), 1103-1105.

[22] M. Paterson, Unsolvability in 3 by 3 matrices, J. of Math. and Physics, Vol 49 (1970), 105–107.

Appendix

Consider the 4 by 4 matrices

$$X = \begin{pmatrix} \lambda I + U & 0 \\ 0 & \lambda I + U \end{pmatrix} = \begin{pmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & 0 & 0 \\ 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & \lambda \end{pmatrix}$$

and

$$Y = \begin{pmatrix} \lambda I & I \\ 0 & \lambda I \end{pmatrix} = \begin{pmatrix} \lambda & 0 & 1 & 0 \\ 0 & \lambda & 0 & 1 \\ 0 & 0 & \lambda & 0 \\ 0 & 0 & 0 & \lambda \end{pmatrix},$$

where

$$I = I_{2\times 2} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

and

$$U = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$$

and λ is *any* complex number.

To verify that $XY = YX$, we can check directly, or we can go as follows: Since $X = \lambda I_4 + \begin{pmatrix} U & 0 \\ 0 & U \end{pmatrix}$ and $Y = \lambda I_4 + \begin{pmatrix} 0 & I \\ 0 & 0 \end{pmatrix}$, we only need to verify that

$$\begin{pmatrix} U & 0 \\ 0 & U \end{pmatrix} \cdot \begin{pmatrix} 0 & I \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & I \\ 0 & 0 \end{pmatrix} \cdot \begin{pmatrix} U & 0 \\ 0 & U \end{pmatrix},$$

which is just $\begin{pmatrix} 0 & U \\ 0 & 0 \end{pmatrix}$.

Now, X is already in its Jordan form, while the (unique) Jordan form for Y is X. This can be seen by the basis change of $\{e_1, e_2, e_3, e_4\} \rightarrow \{e_1, e_3, e_2, e_4\}$.

Suppose there exists a non-singular matrix Z, such that $Z^{-1}XZ = X$ and $Z^{-1}YZ = X$. Then clearly, we must have $X = Y$, a contradiction.