DIMACS Technical Report 95-30
August 1995

Algorithms for matrix groups and the Tits alternative

by

Robert Beals[1,2,3]

[1]DIMACS Postdoc
[2]School of Mathematics Institute For Advanced Study Olden Lane Princeton, New Jersey, 08540
[3]Research supported by: NSF Mathematical Sciences Posdoctoral Fellowship, Institute for Advanced Study, and the Sloan Foundation

---

DIMACS is a cooperative project of Rutgers University, Princeton University, AT&T Bell Laboratories and Bellcore.

DIMACS is an NSF Science and Technology Center, funded under contract STC–91–19999; and also receives support from the New Jersey Commission on Science and Technology.

ABSTRACT

Tits has shown that a finitely generated linear group either contains a nonabelian free group or has a solvable subgroup of finite index. We give a polynomial time algorithm for deciding which of these two conditions holds for a given finitely generated matrix group over an algebraic number field. Noting that many computational problems are undecidable for groups with nonabelian free subgroups, we investigate the complexity of problems relating to linear groups with solvable subgroups of finite index.

For such a group G, we are able in polynomial time to compute a homomorphism ϕ such that $\phi(G)$ is a finite matrix group and the kernel of ϕ is solvable. If in addition G has a nilpotent subgroup of finite index, we obtain much stronger results. These include an effective encoding of elements of G such that the encoding length of an element obtained as a product of length $\leq \ell$ over the generators is $O(\log \ell)$ times a polynomial in the input length. This result is the best possible; it has been shown by Tits and Wolf that if a finitely generated matrix group does not have a nilpotent subgroup of finite index, then the number of group elements expressible as words of length ℓ over the generators grows as c^ℓ for some constant $c > 1$ depending on G.

For groups with abelian subgroups of finite index, we obtain a Las Vegas algorithm for several basic computational tasks including membership testing and computing a presentation. This generalizes recent work of Beals and Babai, who give a Las Vegas algorithm for the case of finite groups, as well as recent work of Babai, Beals, Cai, Ivanyos, and Luks, who give a deterministic algorithm for the case of abelian groups.

1 Introduction

1.1 Groups and computation

For several decades, computers have been used for group theoretic investigations. The design and implementation of practical algorithms for group theory is the domain of computational group theory (cf. [Atk84]). Such algorithms have been developed for permutation groups [Si1] (cf. [Bu]), finitely presented groups (cf.[Si2]), and, more recently, finite matrix groups [NP, HR, H+94a, H+94b]. Many of these algorithms have been implemented in popular programming packages [Can, Sch+], which are used by hundreds of researchers in group theory, design theory, coding theory, and other areas. Methods for handling infinite matrix groups, particularly abelian, nilpotent, and solvable groups, are currently being developed by G. Ostheimer [Os].

The study of the asymptotic complexity of group theoretic algorithms is much newer than computational group theory, with the oldest papers dating back only 15 years [Ba79, FHL]. As is the case with computational group theory, the area of permutation groups is particularly well studied [Ba91a, KLu, Lu93]. Polynomial time algorithms for several classes of finite matrix groups have been obtained in the last few years [Lu92, BBR, BB]. (Some important complexity results for these groups were obtained ten years ago [BSz, Ba85].) For infinite matrix groups, polynomial time algorithms are known only for the class of abelian groups [KLi, Ge93a, Ge93b, CLZ, BBCIL].

This research has benefited theoretical computer science in several ways. One of the motivations for the invention of interactive proofs was Babai's investigation of complexity in finite matrix groups [Ba85]. Much of our understanding of the complexity of graph isomorphism comes from research in permutation group algorithms [Ba79, FHL, Lu82]. Numerous applications have been found in cryptography [BCY, BY, IY, M, MM89, MM92].

Asymptotic complexity theory has also benefited computational group theory. While "polynomial time" and "practical" are two different notions, often the structure used to design a polynomial time algorithm can be refined to yield a practical algorithm. For example, Luks's polynomial time algorithm for computing the composition factors of a permutation group [Lu87] introduced the basic approach used by several improved algorithms [BLS87, BLS88, BS]. An implementation of the [BS] algorithm is included with the group theory package GAP [Sch+].

This paper is an initial investigation into algorithms for infinite nonabelian matrix groups. We give a number of polynomial time algorithms for problems which were not previously known to be decidable. These algorithms are unlikely to be practical in their current form except for small dimension. However, the results are new, even for fixed dimension. We hope that the approach introduced here will serve as a guide for future research, both theoretical and practical, in this area.

1.2 Matrix groups: problems and obstacles

We consider the basic computational problem of testing membership in finitely generated matrix groups over number fields. Membership testing naturally reduces to two problems which are better suited for divide-and-conquer techniques: constructive membership testing, and computing a presentation. *Constructive membership testing* for a group G generated by a set S is the task of expressing a given g as a straight line program over S, or determining that $g \notin G$. A *presentation* for a group is a set of abstract generators and relations which define the group. These two problems are naturally occurring subproblems for many group computations, including (nonconstructive) membership testing (see, for example, [Lu92]).

In recent years, algorithms have been obtained for two important classes of groups: finite groups and abelian groups. For finite groups, Beals and Babai [BB] give Las Vegas algorithms for constructive membership and computing a presentation. Also, Babai, Beals, and Rockmore [BBR] have given a deterministic polynomial time algorithm to test finiteness. For abelian groups, Babai, Beals, Cai, Ivanyos, and Luks [BBCIL] give deterministic algorithms for both problems, building on work of Cai, Lipton, and Zalcstein [CLZ] and Ge [Ge93a, Ge93b] (G. Ostheimer [Os] has developed similar methods independently). We extend both of these lines of research, giving Las Vegas algorithms for groups with abelian subgroups of finite index. We obtain weaker results for two larger classes of groups: those with a nilpotent subgroup of finite index, and those with a solvable subgroup of finite index.

We briefly recall the definitions of solvable and nilpotent groups. A group G is *solvable* if the *derived series* $G = G^0 \geq G^1 \geq \ldots$ reaches the identity in a finite number of steps, where G^{i+1} is the subgroup of G^i generated by all elements of the form $x^{-1}y^{-1}xy$, with $x, y \in G^i$. A group G is *nilpotent* if the *lower central series* $G = G_0 \geq G_1 \geq \ldots$ reaches the identity in a finite number of steps, where G_{i+1} is the group generated by elements of the form $x^{-1}y^{-1}xy$, where $x \in G$ and $y \in G_i$.

For general matrix groups, the membership problem is *undecidable*, even for 4×4 integer matrix groups [Mih]. Also, the problem of testing membership in a normal subgroup given by normal generators is undecidable for 2×2 matrix groups. These undecidability results stem from the presence of nonabelian free subgroups of $GL(n, \mathbf{Z})$. A group generated by a set S is *free* if no nonempty reduced word in the generators and their inverses represents the identity. (A word is reduced if it contains no substrings of the form gg^{-1} or $g^{-1}g$.) An important question, then, from the point of view of decidability, is which matrix groups contain nonabelian free subgroups? In a celebrated result, Tits [Ti] has shown that a finitely generated matrix group either has a nonabelian free subgroup or a solvable subgroup of finite index (the two conditions are mutually exclusive). This dichotomy is known as the "Tits alternative".

1.3 Main results

We give an algorithm to decide the Tits alternative. While we solve it in polynomial time, we should point out that even the decidability of this problem is not immediate and it does not appear to have been known.

Recall that $GL(n, F)$ denotes the group of invertible $(n \times n)$ matrices over F, and $\leq$ stands for "subgroup".

Theorem 1.1 *Let F be a symbolically given algebraic number field, and let $G \leq GL(n, F)$ be a finitely generated group given by a list of generators. Then it is decidable in polynomial time whether or not G has a solvable subgroup of finite index.*

Combining this with the Tits alternative, we have:

Corollary 1.2 *It is decidable in polynomial time whether or not $G \leq GL(n, F)$ has a non-abelian free subgroup.*

A group with a solvable subgroup of finite index necessarily has a solvable normal subgroup of finite index. Such a group is called *solvable-by-finite*. More generally, if X and Y are classes of groups, then a group is X-by-Y if it has a normal subgroup from the class X such that the quotient belongs to Y. In this paper we shall be especially concerned with the following three classes of groups: solvable-by-finite, nilpotent-by-finite, and abelian-by-finite.

For solvable-by-finite groups, we obtain the following:

Theorem 1.3 *Let F be a symbolically given algebraic number field, and let $G \leq GL(n, F)$ be a finitely generated group given by a list of generators. Then in polynomial time, for some $m \leq n^2$, we can either compute a homomorphism $\phi : G \to GL(m, F)$ such that $\phi(G)$ finite and $\ker(\phi)$ is solvable, or determine that G is not solvable-by-finite.*

By considering the image of ϕ, this reduces the problem of determining whether or not G is solvable to the finite case. Luks [Lu92] has given deterministic polynomial time algorithms for a host of problems relating to finite solvable matrix groups, including solvability testing. Thus, we have:

Corollary 1.4 *Solvability of finitely generated matrix groups over number fields is decidable in polynomial time.*

For groups which are nilpotent-by-finite, we obtain stronger results, reducing most problems to the unipotent case. A matrix is *unipotent* if its minimal polynomial is a power of $(x-1)$. A matrix group is unipotent if it consists of unipotent matrices. A unipotent matrix group is necessarily nilpotent.

Theorem 1.5 *Let F be a symbolically given algebraic number field, and let $G \leq GL(n, F)$ be a finitely generated group given by a list of generators. We can decide in polynomial time whether or not G is nilpotent-by-finite, and if so we can compute a homomorphism ϕ of G with the following properties:*

1. *$\ker(\phi)$ is unipotent.*

2. *$\phi(G)$ is abelian-by-finite.*

3. *Constructive membership testing in $\phi(G)$ can be performed in Las Vegas polynomial time.*

In addition, normal generators for $\ker(\phi)$ *can be computed in Las Vegas polynomial time.*

For finitely generated abelian unipotent groups, [BBCIL] show how to compute an isomorphism to a lattice. This, combined with known polynomial time algorithms for lattices (cf. [Schr]), gives:

Corollary 1.6 *Constructive membership testing for abelian-by-finite matrix groups over number fields is in Las Vegas polynomial time.*

In working with matrix groups in characteristic 0, a problem even more basic than membership testing is estimating bit lengths of elements. This is nontrivial even for finite groups [Ba92, BBR]. In general, it is often desirable to represent elements of a group using an encoding other than the standard encoding (in which an element of $GL(n, F)$ is represented as an $n \times n$ array of elements of F). Let size(x) denote the bit length of x. For $G \leq GL(n, F)$ given by a generating set S, for technical reasons we assume that the encoding of S includes a description of F, so size(S) is the input length.

Let G be a finitely generated group, generated by the finite set S. The *growth function* of G with respect to S is the function g_S, where $g_S(m)$ is the number of distinct elements of G expressible as words of length m in G. The asymptotic behavior of g_S does not depend on the choice of generating set S. We say that G has *polynomial growth* if $g_S(m) = O(m^c)$ for some constant c, and we say that G has *exponential growth* if $g_S(m) = \Omega(c^m)$ for some constant $c > 1$. Milnor and Wolf (cf. [Mil, Wo]) have shown that a solvable-by-finite group G has polynomial growth if G is nilpotent-by-finite, and has exponential growth otherwise. Combined with the Tits alternative, this shows that a finitely generated linear group G has polynomial growth if G is nilpotent-by-finite, and has exponential growth otherwise. For a group with polynomial growth rate, it seems conceivable that there is an effective way of encoding elements of the group so that group multiplication does not cause a rapid increase in encoding length. We show that this is the case for matrix groups over number fields:

Theorem 1.7 *Let F be a symbolically given algebraic number field, and let $G \leq GL(n, F)$ be a finitely generated group given by a list S of generators. We can determine in polynomial time whether the growth rate of G is polynomial or exponential. If G has polynomial growth, then there is a group $\mathcal{G} \geq G$ and an encoding ζ of elements of $\mathcal{G}$ with the following properties:*

1. *Membership of g in $\mathcal{G}$ is decidable in polynomial time.*

2. *For $g \in \mathcal{G}$, we can compute $\zeta(g)$ in polynomial time.*

3. *Given $\zeta(g_1)$ and $\zeta(g_2)$, we can compute $\zeta(g_1 {g_2}^{-1})$ in polynomial time.*

4. *For any $g_1, \ldots, g_\ell \in \mathcal{G}$,*

$$\text{size}(\zeta(\prod_{i=1}^{\ell} g_i)) = O(\max_i\{\text{size}(\zeta(g_i))\} \ (\text{size}(S))^c \ \log \ell)$$

for some constant c.

In items 1–3 above, polynomial time *refers to time that is polynomial in the input length, where S is considered to be part of the input. In particular, the degree of the polynomial does not depend on S.*

Previously, it was known how to compute such a ζ for finite matrix groups [BBR] and for abelian matrix groups [BBCIL].

1.4 Methods

A matrix group $G \leq GL(V)$ is *primitive* if there is no direct sum decomposition $V_1 + \ldots + V_r$ of V such that G permutes the V_i. If such direct sum decompositions could be found for imprimitive groups, many computational problems would be reduced to the primitive case. Much is known about the structure of primitive solvable matrix groups [Su63, Su76]. However, no general imprimitivity algorithm is known (although Holt et al [H+94b] have developed practical methods for matrix groups over finite fields).

A natural combinatorial approach to matrix group imprimitivity is suggested by Atkinson's algorithm for permutation group imprimitivity [Atk75]. For permutation groups, a system of imprimitivity can be found by considering the orbits of a representation of degree n^2. The analogous approach for matrix groups would consider the invariant subspaces of a representation of dimension n^2, but it is not clear that this works in all cases. Nevertheless, by using a novel application of finiteness testing [BBR], we obtain a method which is sufficiently general for our purposes.

We find enough of the imprimitivity structure to decide the Tits alternative, and with some further refinements, we obtain our results for nilpotent-by-finite groups as well. This allows a reduction of the abelian-by-finite case to the abelian and the finite cases separately, for which algorithms exist (cf. [BBCIL, Ge93a, Os, BBR, BB]).

2 Preliminaries and background

2.1 Building blocks

We say that a group $\mathcal{G}$, together with an encoding ζ of elements of $\mathcal{G}$ satisfying conditions 1–4 of Theorem 1.7 is an *effective slow-growth encoding.* These conditions guarantee that any computation involving polynomially many group operations runs in polynomial time. We give some simple examples of such encodings, as well as some methods by which such encodings may be combined to create new ones. Babai, Beals, and Rockmore [BBR] have shown:

Theorem 2.1 *Let $G \leq GL(n, \mathbf{Q})$ be given by a finite list of generators. Then it is decidable in deterministic polynomial time whether or not G is finite. If G is finite, then in polynomial time we can compute a G-invariant n-dimensional lattice L and a G-invariant positive* *definite quadratic form B such that the group $\mathcal{G}$ of all elements of $GL(n, \mathbf{Q})$ which fix L and B, together with the standard encoding of rational matrices, is an effective slow-growth encoding.*♠

The generalization to finite subgroups of $GL(n, F)$ is obtained by considering F as a vector space over $\mathbf{Q}$.

The standard encoding is also sufficient for the additive group of a vector space:

Proposition 2.2 *Let F be a symbolically given algebraic number field. Then the standard encoding of elements of F^n is a slow-growth encoding.*♠

Unipotent groups are yet another example.

Lemma 2.3 *Let A be an n-dimensional algebra with unit element 1 and radical R. Let $\mathcal{G}$ be the multiplicative group of elements of the form $1 + r$ with $r \in R$. Then the standard encoding of elements of $\mathcal{G}$ is slow-growth.*

Proof: By Friedl–Rónyai, the radical R is polynomial time computable, so $\mathcal{G}$ is polynomial time recognizable (condition 1). Conditions 2–3 are trivial. It remains to show that condition 4 holds. Consider a product $g = g_1 g_2 \ldots g_\ell$. Let $g_i = 1 + r_i$, for $1 \leq i \leq \ell$. Then g can be expressed as a sum of products of the r_i. Since any product of n or more of the r_i is zero, we only need to consider products of length less than n. The number of summands is thus $O(\ell^n)$, each of which has bit length at most $O(n^c \max_i \mathrm{size}(r_i))$. Applying Proposition 2.2 gives the desired bound on the bit length of g. ♠

We now describe several methods of combining encodings.

Proposition 2.4 *Let $\mathcal{G}_1, \zeta_1$ and $\mathcal{G}_2, \zeta_2$ be groups with slow-growth encodings. For (g_1, g_2) in the direct product $\mathcal{G}_1 \times \mathcal{G}_2$, let $\zeta(g_1, g_2) = (\zeta_1(g_1), \zeta_2(g_2))$. Then $\mathcal{G}_1 \times \mathcal{G}_2, \zeta$ is a slow-growth encoding.*♠

Let G and H be groups, with G acting (on the left) on H as automorphisms. then the *semidirect product* $H \rtimes G$ is defined as the set of all ordered pairs $\{(h, g) \mid h \in H,\ g \in G\}$, with multiplication defined by:

$$(h_1, g_1)(h_2, g_2) = (h_1 \cdot g_1(h_2), g_1 g_2).$$

With this operation, $H \rtimes G$ is a group. It is not the case that given slow-growth encodings for G and H, we can construct a slow-growth encoding for $H \rtimes G$, since $H \rtimes G$ might have exponential growth. However, we do have:

Proposition 2.5 *Let $\mathcal{G}_1, \zeta_1$ and $\mathcal{G}_2, \zeta_2$ be groups with slow-growth encodings. Suppose that $\mathcal{G}_2$ acts on $\mathcal{G}_1$ as automorphisms such that the bit length of $\zeta_1(g_2(g_1))$ is polynomially bounded by $\zeta_1(g_1)$ and the bit length of the generating set for $\mathcal{G}_2$. Then the semidirect product $\mathcal{G}_1 \rtimes \mathcal{G}_2$ with the encoding $\zeta(g_1, g_2) = (\zeta_1(g_1), \zeta_2(g_2))$ is a slow-growth encoding.*♠

This condition is very restrictive, but we give three important examples:

Example 1 *Suppose that $\mathcal{G}_1$ is a linear group, and $\mathcal{G}_2$ acts linearly as a finite group on* $\mathrm{Env}(\mathcal{G}_1)$. *Then slow-growth encodings for $\mathcal{G}_1$ and $\mathcal{G}_2$ may be combined to obtain a slow-growth encoding for $G_1 \rtimes G_2$.*

Example 2 *Suppose that $\mathcal{G}_1$ is a lattice, and $\mathcal{G}_2$ acts linearly as a finite group on $\mathcal{G}_1$. Then slow-growth encodings for $\mathcal{G}_1$ and $\mathcal{G}_2$ may be combined to obtain a slow-growth encoding for $G_1 \rtimes G_2$.*

Let S_r denote the symmetric group of degree r.

Example 3 *Suppose that $\mathcal{G}_1$ is an iterated direct product $\mathcal{G}^r$, and $\mathcal{G}_2 = S_r$ acts by permuting the factors in the direct product. Suppose that ζ_1 is obtained from a slow-growth encoding of $\mathcal{G}$ by repeated applications of Proposition 2.4. Then slow-growth encodings for $\mathcal{G}_1$ and $\mathcal{G}_2$ may be combined to obtain a slow-growth encoding for $\mathcal{G}_1 \rtimes \mathcal{G}_2$.*

In this last example, the group $\mathcal{G}_1 \rtimes \mathcal{G}_2$ is also called the *wreath product* of $\mathcal{G}$ by S_r, denoted $\mathcal{G} \wr S_r$.

2.2 Algebra decomposition

Let F be a field, and let $G \leq GL(n, F)$. Then we denote by $\mathrm{Env}(G)$ the *enveloping algebra* of G, i.e., the set of F-linear combinations of elements of G. We denote by $M_n(F)$ the algebra of $n \times n$ matrices over F.

Let $\mathcal{A}$ be an algebra over the field F. The *radical* $\mathrm{Rad}(\mathcal{A})$ is the largest nilpotent ideal of $\mathcal{A}$. The quotient $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$ is *semisimple*, i.e. its radical is 0. If $\mathcal{A}$ and $\mathcal{B}$ are subalgebras of an algebra, then the *centralizer* of $\mathcal{A}$ in $\mathcal{B}$, denoted $C_{\mathcal{A}}(\mathcal{B})$, is the set $\{a \in \mathcal{A} \mid \forall b \in \mathcal{B} : ab = ba\}$. If $\mathcal{A} = \mathcal{B}$ this is the *center* of $\mathcal{A}$, denoted $Z(\mathcal{A})$. Centralizers can be found in polynomial time by solving linear equations.

There is a nice correspondence between the structure of $\mathrm{Env}(G)$ and the action of G on F^n. If $\mathrm{Env}(G)$ is the direct sum of ideals I_i and I_2, then F^n is the direct sum of G-invariant subspaces I_1F^n and I_2F^n. In addition, the action of G is *completely reducible* (i.e., F^n is the direct sum of G-invariant subspaces on which G acts irreducibly) iff $\mathrm{Env}(G)$ is semisimple.

We require the following result of Friedl and Rónyai [FR]:

Theorem 2.6 *Let $\mathcal{A}$ be a subalgebra of $M_n(F)$, given by a list of generators. Then* $\mathrm{Rad}(\mathcal{A})$ *can be calculated in polynomial time, along with the minimal ideals of* $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$.♠

We indicate a way in which the Friedl–Ronyai result can be used as a divide-and-conquer tool. Once the decomposition of $\mathrm{Env}(G)/\mathrm{Rad}(\mathrm{Env}(G))$ is known, a basis can easily be chosen for F^n such that the matrices $g \in G$ have the form

$$\begin{pmatrix} \eta_1(g) & * & \cdots & * \\ 0 & \eta_2(g) & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \eta_r(g) \end{pmatrix},$$

where each $\mathrm{Env}(\eta_i(G))$ is a simple algebra. Let η denote the representation:

$$\eta(g) = \begin{pmatrix} \eta_1(g) & 0 & \cdots & 0 \\ 0 & \eta_2(g) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \eta_r(g) \end{pmatrix}$$

Then the kernel of η is solvable (it consists of unipotent matrices), so G is solvable-by-finite iff $\eta(G)$ is solvable-by-finite. We will make use of the homomorphism η throughout the paper.

An element e of an algebra is an *idempotent* if $e^2 = e$. The Friedl–Rónyai algorithm actually calculates, for each ideal I of $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$, an idempotent e in the center of $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$ such that e is the multiplicative identity of I. For two distinct such ideals I_1, I_2, we have $e_1 e_2 \in I_1 \cap I_2 = \{0\}$, so $e_1 e_2 = 0$.

The classical theorem of Wedderburn shows that $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$ is the direct sum of simple algebras. These simple components are precisely the minimal ideals found by the Friedl–Rónyai algorithm. If $\mathcal{A}$ is the enveloping algebra of a nilpotent group, then something more can be said. We first quote the following result of Suprunenko [Su76, p. 223, Theorem 2]:

Theorem 2.7 *Let F be an arbitrary field. The irreducible components of a nilpotent indecomposable subgroup of $GL(n, F)$ are pairwise equivalent.*♠

Corollary 2.8 *Let $G \leq GL(n, F)$ be nilpotent, and let $\mathcal{A} = \mathrm{Env}(G)$. Then there exist ideals $\mathcal{A}_1, \ldots, \mathcal{A}_r$ such that $\mathcal{A}$ is the direct sum of the $\mathcal{A}_i$ and each $\mathcal{A}_i/\mathrm{Rad}(\mathcal{A}_i)$ is simple.*♠

We quote two results of [BBCIL], which we prove in Appendix A.

Lemma 2.9 *Given a matrix $a \in M_n(F)$, we can in polynomial time find matrices a_d and a_n in the algebra generated by a, such that a_d is diagonalizable, (over an extension field), a_n is nilpotent, and $a = a_d + a_n$.*

Corollary 2.10 *Let $\mathcal{A}$ be a commutative algebra over the number field F, such that $K = \mathcal{A}/\mathrm{Rad}(A)$ is simple. Then in polynomial time, we can find a subalgebra of $\mathcal{A}$ isomorphic to K.* ♠

Note that K is an extension field of F. This Corollary will allow us to view F^n as a vector space (of smaller dimension) over K.

We quote below some well-known results on group algebras (cf. [We, Chapter 1]):

Theorem 2.11 *Suppose that $H \leq G \leq GL(n, F)$, and $G : H$ is finite. Then if $\mathrm{Env}(H)$ is semisimple, so is $\mathrm{Env}(G)$.*

Theorem 2.12 (Clifford) *Suppose that $G \leq GL(n, F)$ with $\mathrm{Env}(G)$ semisimple. Then for any normal subgroup N of G, $\mathrm{Env}(N)$ is semisimple.*

3 Structure of solvable-by-finite groups

In this Section we consider in some detail what it means for a group to have a solvable subgroup of finite index. The results presented here are nonconstructive,

Suppose that the group $G \leq GL(n, F)$ has a solvable subgroup H of finite index. Then the intersection of the conjugates of H is a normal subgroup N of finite index, so G is solvable-by-finite.

A linear group is *triangularizable* if, over some extension field K of F, all elements of the group may be simultaneously put in upper triangular form. A theorem of Mal'cev (cf. [We, Theorem 3.6]) states that a solvable linear group has a triangularizable normal subgroup of finite index. Let T be a triangularizable normal subgroup of N of finite index. Then the intersection of the G-conjugates of T is a triangularizable normal subgroup M of G of finite index in G. Without loss of generality we suppose that $M = N$.

Now suppose that $\mathrm{Env}(G)$ is semisimple. Then $\mathrm{Env}(N)$ is semisimple as well by Clifford's Theorem (Theorem 2.12). Therefore N is abelian (it is diagonalizable over an extension field). We suppose that for any positive integer ℓ, the groups N and $N^\ell = \{x^\ell \mid x \in N\}$ have the same enveloping algebra (otherwise replace N by N^ℓ; this can only happen finitely many times before we arrive at a subgroup with the desired property).

Let $\theta : G \to GL(M_n(F))$ denote the conjugation action of G on $M_n(F)$. It clear that $\mathrm{Env}(\theta(N))$ is semisimple, so by Theorem 2.11, $\mathcal{A} = \mathrm{Env}(\theta(G))$ is semisimple as well. Let $\mathcal{A}_1, \ldots, \mathcal{A}_s$ be the decomposition of $\mathcal{A}$ into a direct sum of minimal ideals, and let $W_i = \mathcal{A}_i M_n(F)$ (the W_i are subspaces of $M_n(F)$, but not necessarily subalgebras). Let $\mathcal{B}$ be the direct sum of those W_i such that the restriction of $\theta(G)$ to W_i is finite. Note that $\mathcal{B}$ is a subalgebra of $M_n(F)$. Let $\theta_{\mathcal{B}}$ be the conjugation action of G on $\mathcal{B}$, and note that $\theta_{\mathcal{B}}(G)$ is finite.

Lemma 3.1 *Let $\mathcal{B}$ be as above. Then $\mathcal{B} = C_{M_n(F)}(\mathrm{Env}(N))$, and $\mathrm{Env}(N) = C_{M_n(F)}(\mathcal{B})$.*

Proof: We first show that $C_{M_n(F)}(\mathrm{Env}(N)) \subseteq \mathcal{B}$. Suppose a centralizes $\mathrm{Env}(N)$. We wish to show that a has finitely many G-conjugates. Since N has finite index in G, this is equivalent to showing that a has finitely many N-conjugates. This is immediate, since $a^N = \{a\}$.

Next, suppose that a does not centralize N. We wish to show that a has infinitely many N-conjugates. Suppose to the contrary that $|a^N| = m$ for some finite m. Let $\ell = m!$, and note that a centralizes N^ℓ. Therefore $\mathrm{Env}(N) \neq \mathrm{Env}(N^\ell)$, contradicting our choice of N.

Now we need to show that $C_{M_n(F)}(\mathcal{B}) = \mathrm{Env}(N)$. Let $\mathcal{C} = C_{M_n(F)}(\mathcal{B})$. Since $\mathrm{Env}(N) \subseteq \mathcal{B}$, $\mathcal{C} \subseteq C_{M_n(F)}(\mathrm{Env}(N)) = \mathcal{B}$. Therefore it suffices to show that $\mathrm{Env}(N) = Z(\mathcal{B})$. Let $E_1, \ldots, E_r$ denote the minimal ideals of $\mathrm{Env}(N)$. Let $V_i = E_i F^n$. Then $V_1 + \ldots + V_r$ is a direct sum decomposition of F^n into N-invariant subspaces. These subspaces are $\mathcal{B}$-invariant as well, but even more can be said. Each V_i may be viewed as a vector space over the field E_i. For any $a \in \mathcal{B}$, the action of a on V_i is linear over E_i, since a centralizes E_i. Conversely, suppose $a \in M_n(F)$ leaves the V_i invariant and acts as E_i-linear transformations on each V_i. Then a centralizes N and so belongs to $\mathcal{B}$. So $Z(\mathcal{B})$ commutes with all E_i-linear transformations on V_i. Since the center of a full matrix algebra is the set of scalar matrices, we have $Z(\mathcal{B}) = \mathrm{Env}(N)$ as desired. ♠

So far we have assumed that $\mathrm{Env}(G)$ is semisimple. We now replace that with the weaker assumption that N is nilpotent. As in the semisimple case, by Theorem 2.7, the minimals ideals of $\mathrm{Env}(N)/\mathrm{Rad}(\mathrm{Env}(N))$ yield a direct sum decomposition $V_1 + \dots V_r$ of F^n into N-invariant subspaces. The group G permutes the V_i. We consider the case that this action is transitive. Let G_1 denote the subgroup $\{g \in G \mid gV_1 = V_1\}$. Then G is a subgroup of the wreath product $G_1 \wr S_r$. We now consider the $r = 1$ case (i.e., we focus our attention on G_1). In this case $\mathrm{Env}(N)/\mathrm{Rad}(\mathrm{Env}(N))$ is a field E, of finite degree d over F. By [Su76, page 228, Theorem 3], there is an embedding of E in $M_n(F)$ such that N acts as E-linear transformations (the embedding of E in $M_n(F)$ gives a natural way of interpreting F^n as an n/d-dimensional vector space over E). Indeed, the same reference shows that N acts as a subgroup of the direct product of the multiplicative group of E and a unipotent subgroup of $End_E(F^n)$. Since N is normal in G, we have that G acts by conjugation as automorphisms of E which fix F. The kernel H of this action has index at most d in G. Note that H acts as E-linear transformations, so we may consider H to be a subgroup of $GL(n/d, E)$.

We summarize this in our main structure theorem for nilpotent-by-finite groups:

Theorem 3.2 *Suppose $G \leq GL(n, F)$ is nilpotent-by-finite. Then G has a normal subgroup N of finite index such that:*

1. *$\mathrm{Env}(N)$ is a direct sum of ideals $\mathcal{B}_1, \dots, \mathcal{B}_r$, such that each $\mathcal{B}_i/\mathrm{Rad}(\mathcal{B}_i)$ is a field E_i.*
2. *F^n is a direct sum of N-invariant subspaces $V_1, \dots, V_r$, where $V_i = \mathcal{B}_i F^n$. These are permuted by G.*
3. *E_i can be embedded in $End_F(V_i)$ such that N acts as E_i-linear transformations on V_i.*
4. *For any $1 \leq i \leq r$, the subgroup $G_i = \{g \in G \mid gV_i = V_i\}$ has a subgroup H_i of index dividing the degree of E_i over F such that H_i acts as E_i-linear transformations on V_i.*

The subgroup N is a subgroup of the direct product of a unipotent normal subgroup U of G and the multiplicative groups of the fields E_i. ♠

In subsequent sections, we will give a polynomial time algorithm to find the V_i. This will reduce many problems for nilpotent-by-finite G to the case $r = d = 1$ (so $G = H$, and N is isomorphic to a subgroup of a direct product of F^* and a unipotent group).

4 Algorithm for the Tits alternative

We are now ready to prove Theorems 1.1 and 1.3. By using the Friedl–Rónyai algorithm [FR], it suffices to consider the case that $\mathrm{Env}(G)$ is semisimple (otherwise, let η be the homomorphism described after Theorem 2.6, and replace G by $\eta(G)$). In other words, we are trying to determine if a completely reducible group is abelian-by-finite.

The algorithm that we present here will attempt to find the structure described in Section 3. If at any point this fails, then we know that the group G is not solvable-by-finite.

Let $\theta : G \to GL(M_n(F))$ denote the conjugation action of G on $M_n(F)$ (we are viewing $M_n(F)$ as a vector space). We compute $\mathcal{A} = \mathrm{Env}(\theta(G))$, and verify that $\mathcal{A}$ is semisimple (if not, then G is not solvable-by-finite). If $\mathcal{A}$ is semisimple, then we find its minimal ideals $\mathcal{A}_1, \ldots, \mathcal{A}_s$. Let $W_i = \mathcal{A}_i M_n(F)$, and let θ_i be the conjugation action of G on W_i. Using the finiteness algorithm of Babai, Beals, and Rockmore [BBR], we determine the set of i such that $\theta_i(G)$ is a finite group. Let $\mathcal{B}$ be the sum of the W_i such that $\theta_i(G)$ is finite. Note that $\mathcal{B}$ is a subalgebra of $M_n(F)$ (the W_i are merely subspaces). By Lemma 3.1, if G is solvable-by-finite then $\mathcal{B}$ is the centralizer in $M_n(F)$ of $\mathrm{Env}(N)$, for some abelian normal subgroup N of finite index in G. If this is the case, we verify it as follows:

By solving linear equations, we find the center $\mathcal{C}$ of $\mathcal{B}$, and verify that $\mathcal{B}$ and $\mathcal{C}$ are the centralizers in $M_n(F)$ of each other. This is the case if G is solvable-by-finite (cf. Lemma 3.1). Conversely, suppose that $\mathcal{C}$ is the full centralizer in $M_n(F)$ of $\mathcal{B}$. G acts by conjugation as a finite group on $\mathcal{B}$ so $G \cap C_{M_n(F)}(\mathcal{B})$ has finite index in G. But $C_{M_n(F)}(\mathcal{B}) = \mathcal{C}$ is commutative, so G is abelian-by-finite. This completes the proof of Theorem 1.1.♠

To prove Theorem 1.3, we let ϕ be the conjugation action of G on $\mathcal{B}$.♠

5 The nilpotent-by-finite case

We sketch the proofs of Theorems 1.5 and 1.7. A general divide-and-conquer strategy of the algorithm is to reduce computational problems to "simpler" groups using direct products and semidirect products.

We begin by applying the algorithm of the previous Section to $\eta(G)$. Let $f_1, \ldots, f_r$ be the central idempotents corresponding to the decomposition of $\mathcal{C}$. We wish to "lift" these central idempotents of $\mathrm{Env}(N)/\mathrm{Rad}(\mathrm{Env}(N))$ to central idempotents of $\mathrm{Env}(N)$. Let $G_i \leq G$ consist of all elements g such that $\eta(g)$ fixes V_i. Schreier generators (cf. [Si1, FHL]) for G_i can be found in polynomial time. Note that η extends naturally to an algebra homomorphism from $\mathrm{Env}(G)$ to $\mathrm{Env}(\eta(G))$, which we will also denote η. To lift f_i, find an $e_i \in Z(\mathrm{Env}(G_i))$ such that $\eta(e_i) = f_i$ (this reduces to solving a system of linear equations). By Lemma 2.9, we may assume that e_i is diagonalizable. By Corollary 2.8, this will succeed unless the group is not nilpotent-by-finite.

Assume that G is nilpotent-by-finite, and that $e_1, \ldots, e_r$ have been lifted. Let $V_i = e_i F^n$. The V_i are permuted by G, and G is easily represented as a subdirect product of subgroups of wreath products. Standard techniques [Si1, FHL], together with Propositions 2.4 and 2.5, now allow us to reduce our computational problems to the case $r = 1$.

We assume now that $r = 1$, so $\mathcal{C}$ is a field E of degree d over F for some d dividing n. The action of G on $\mathcal{C}$ corresponds to F-automorphisms of $\mathcal{C}$ (i.e. automorphisms which fix F). Therefore the centralizer in G of $\mathcal{C}$ has index at most d. Again, techniques of [Si1, FHL] allow us to reduce to the case that G acts trivially on $\mathcal{C}$.

We now assume that G acts trivially on $\mathcal{C}$, and we wish to embed $E = \mathcal{C}$ in the center of $\mathrm{Env}(G)$. Such an embedding exists iff G is nilpotent-by-finite. By Corollary 2.10, we can find such an embedding or show that none exists in polynomial time.

Once the field E has been embedded in the center of $\mathrm{Env}(G)$, we may effectively view G as a subgroup of $GL(n/d, E)$. Thus, we have reduced to the case that $E = F$.

We are now in the case that $G \leq GL(n, F)$ has a normal subgroup N of finite index such that N is a subdirect product of F^* and a unipotent group U. By applying the Friedl–Rónyai algorithm to $\mathrm{Env}(G)$, we are able to factor elements of G as products of the form $g = \eta(g)\mu(g)$, where $\eta(G)$ acts by conjugation as a finite group on $M_n(F)$ and the group $\overline{U}$ generated by the $\eta(G)$-conjugates of $\mu(G)$ is unipotent (this group may be bigger than U). Let θ denote the conjugation action of $\eta(G)$ on $M_n(F)$, and let π denote a homomorphism of $\eta(G)$ to $GL(kn, p)$ obtained by converting $\eta(G)$ to a rational matrix group of degree kn (k is the degree of F) and reducing the entries modulo p for some odd prime p. By [New, Theorem IX.7], π will act faithfully on any subgroup H of G such that $\eta(H)$ is finite.

We claim that G is isomorphic to the semidirect product $(\mu(G) \rtimes \theta(G)) \times \pi(G) \times \det(G)$. It is sufficient to show that for $g \in \ker(\pi) \cap \ker(\theta) \cap \ker(\det)$, $\eta(g) = 1$. If $\theta(g) = 1$, then $\eta(g)$ is a scalar matrix λI for some $\lambda \in F$. For such a g, if $\det(g) = 1$, then $\lambda^n = 1$, so π acts faithfully on $\langle g \rangle$. If in addition $\pi(g) = 1$, then $\eta(g) = 1$ as desired.

To prove Theorem 1.7, by Proposition 2.5 it suffices to give slow-growth encodings separately for (supergroups of) $\mu(G), \theta(G), \pi(G)$, and $\det(G)$. For the first three, the standard encoding suffices (for $\theta(G)$ this is Theorem 2.1, for $\mu(G)$ this is Lemma 2.3). For $\det(G) \leq F^*$, we use Ge's multiplicative encoding [Ge93a].♠

To prove Theorem 1.5, let $\phi(g) = (\theta(g), \pi(g), \det(g))$. The group $\theta(G)$ is managed by [BB], and the group $\det(G)$ is handled with [Ge93a]. The group $\pi(\ker(\theta) \cap \ker(\det))$ has order $\leq n$, and poses no problem. The [Ge93a, BB] algorithms give a polynomial length straight-line program to compute normal generators for $\ker(\phi)$. By using the slow growth encoding of Theorem 1.7, such a straight-line program can be evaluated in polynomial time. ♠

References

[Atk75] M.D. Atkinson: An algorithm for finding the blocks of a permutation group, *Math. Comp.* **29** (1975), 911–913.

[Atk84] M. D. Atkinson, ed.: *Computational Group Theory* (Proceedings of the London Mathematical Society Symposium on Computational Group Theory), Academic Press, London 1984.

[Ba79] L. Babai: Monte-Carlo Algorithms in Graph Isomorphism Testing, unpublished manuscript (1979).

[Ba85] L. Babai: Trading group theory for randomness, *Proc. 17th ACM STOC*, Providence RI, 1985, 421-429.

[Ba91a] L. Babai: Computational complexity in finite groups, *Proc. International Congress of Mathematicians*, Kyoto 1990, Springer-Verlag, Hong Kong 1991, 1479–1489.

[Ba92] L. Babai: Deciding finiteness of matrix groups in Las Vegas polynomial time, *Proc. 3rd ACM-SIAM SODA (Symp. on Discrete Algorithms)*, 1992, 33–40.

[BBCIL] L. Babai, R. Beals, J-y. Cai, G. Ivanyos, E. M. Luks: Multiplicative equations over commuting matrices, manuscript, 1994.

[BBR] L. Babai, R. Beals, D. Rockmore: Deciding finiteness of matrix groups in deterministic polynomial time, *Israel J. Math.*, to appear.

[BLS87] L. Babai, E. M. Luks, Á. Seress, Permutation groups in *NC*, *in: Proc. 19th ACM STOC*, 1987, 409–420.

[BLS88] L. Babai, E. Luks, Á. Seress: Fast management of permutation groups, *Proc. 28th IEEE FOCS*, 1988, 272–282.

[BSz] L. Babai, E. Szemerédi: On the complexity of matrix group problems I, *Proc. 25th IEEE FOCS*, 1984, 229–240.

[BB] R. Beals, L. Babai: Las Vegas algorithms for matrix groups, *Proc. 34th IEEE FOCS* (1993), 427–436.

[BS] R. Beals, Á. Seress: Structure forest and composition factors in nearly linear time for small base groups, *Proc. 24th ACM STOC*, 1992, 116–125.

[BCY] G. Brassard, C. Crépeau, M. Yung: Everything in NP can be argued in *perfect* zero-knowledge in a *bounded* number of rounds, *Proc. of the 16th ICALP*, Springer-Verlag, Berlin 1989, 123–136.

[BY] G. Brassard, M. Yung: One-Way Group Actions, *Advances in Cryptology: CRYPTO '90 proceedings*, Lecture Notes in Computer Science **537**, Springer-Verlag, Berlin 1991, 94–107.

[Bu] G. Butler: *Fundamental Algorithms for Permutation Groups*, Lecture Notes in Computer Science **559**, Springer-Verlag, Berlin, 1991.

[CLZ] J-y. Cai, R. J. Lipton, Y. Zalcstein: The complexity of the membership problem for 2-generated commutative semigroups of rational matrices, *Proc. 35th IEEE FOCS* (1994), 135–142.

[Can] J. J. Cannon: An Introduction to the Group Theory Language, Cayley, *Computational Group Theory* (Proceedings of the London Mathematical Society Symposium on Computational Group Theory), M. D. Atkinson, ed., Academic Press, London 1984, 145–184.

[FHL] M. L. Furst, J. Hopcroft, E. M. Luks: Polynomial-time algorithms for permutation groups, *in: 21st IEEE FOCS*, 1980, 36–41.

[FR] K. Friedl, L. Rónyai: Polynomial time solutions of some problems in abstract algebra, *Proc. 17th ACM STOC* (1985), 153–162.

[Ga] F. R. Gantmacher: *Matrix Theory*, Chelsea Publishing, New York, 1959.

[Ge93a] G. Ge: Algorithms related to multiplicative representations of algebraic numbers, PhD thesis, Math Dept, U. C. Berkeley, 1993.

[Ge93b] G. Ge: Testing equalities of multiplicative representations in polynomial time, *Proc. 34th IEEE FOCS* (1993), 422–426.

[HR] D. F. Holt, S. Rees: Testing modules for irreducibility, J. Austral. Math. Soc. Ser. A **57** (1994).

[H+94a] D. F. Holt, C. R. Leedham-Green, E. A. O'Brien, S. Rees: Computing matrix group decompositions with respect to a normal subgroup, submitted to London Math. Soc., 1994.

[H+94b] D. F. Holt, C. R. Leedham-Green, E. A. O'Brien, S. Rees: Primitivity testing for matrix groups, submitted to London Math. Soc., 1994.

[IY] R. Impagliazzo, M. Yung: Direct minimum-knowledge computations, *Advances in Cryptology: CRYPTO '87 proceedings*, Lecture Notes in Computer Science **293**, Springer-Verlag, Berlin 1988, 40–51.

[KLi] R. Kannan, R. Lipton: Polynomial-time algorithms for the orbit problem, *J. ACM* **33**, No. 4 (1986), 808–821.

[KLu] W. M. Kantor, E. M. Luks: Computing in quotient groups, *in: Proc. 22nd ACM STOC*, 1990, pp. 524–534.

[Lu82] E. M. Luks: Isomorphism of graphs of bounded valence can be tested in polynomial time, *J. Comp. Syst. Sci.* **25** (1982), 42–65.

[Lu87] E. M. Luks: Computing the composition factors of a permutation group in polynomial time, *Combinatorica* **7** (1987), 87–99.

[Lu92] E. M. Luks: Computing in solvable matrix groups, *Proc. 33rd IEEE FOCS* (1992), 111–120.

[Lu93] E. M. Luks: Permutation Groups and Polynomial-Time Computation, in *Groups and Computation*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science **11** (1993), 139–175.

[M] S. Magliveras: A cryptosystem from logarithmic signatures of finite groups, in *Proc. 29th Midwest Symposium on Circuits and Systems* (1986), 972–975.

[MM89] S. Magliveras, N. Memon: Properties of Cryptosystem PGM, in *Advences in Cryptology — CRYPTO '89*, Lecture Notes in Computer Science **435**, Springer-Verlag, Berlin (1989), 447–460.

[MM92] S. Magliveras, N. Memon: Algebraic Properties of Cryptosystem PGM, Journal of Cryptology **5** (1992), pp. 167–183.

[Mih] K. A. Mihaĭlova: The occurrence problem for a direct product of groups (in Russian), *Dokl. Akad. Nauk* **119** (1958), 1103–1105.

[Mil] J. Milnor: Growth of finitely generated solvable groups, *J. Differential Geometry* **2** (1968), 447–449.

[NP] P. M. Neumann, Cheryl E. Praeger: A recognition algorithm for the special linear groups, manuscript, 1990.

[New] M. Newman: *Integral Matrices*, Pure and Applied Mathematics, vol. 45, Academic Press, New York, 1972.

[Os] G. Ostheimer: PhD thesis, Dept. of Mathematics, Rutgers University, in preparation.

[Sch+] Martin Schönert et al.: *GAP – Groups, Algorithms, and Programming*, Lehrstuhl D für Mathematik, Rheinisch Westfälische Technische Hochschule, Aachen, Germany, first edition, 1992.

[Schr] A. Schrijver: *Theory of Linear and Integer Programming*, Wiley, New York, 1986.

[Si1] C. C. Sims: Computation with Permutation Groups, in: *Proc. Second Symposium on Symbolic and Algebraic Manipulation*, (S.R. Petrick, ed.), ACM, New York 1971, 23–28.

[Si2] C. C. Sims: *Computation with finitely presented groups*, Cambridge University Press, New York, 1994.

[Su63] D. Suprunenko: *Soluble and nilpotent linear groups*, Transl. of Math. Monographs, vol. 9, AMS, Rhode Island, 1963.

[Su76] D. Suprunenko: *Matrix Groups*, Transl. of Math. Monographs, vol. 45, AMS, Rhode Island, 1976.

[Ti] J. Tits: Free subgroups in linear groups, *J. Algebra* **20** (1972), 250–270.

[We] B. A. F. Wehrfritz: *Infinite linear groups*, Ergebnisse der Math. und ihrer Grenzgebiete, Band 76, Springer-Verlag, Berlin, 1973.

[Wo] J. A. Wolf: Growth of finitely generated solvable groups and curvature of Riemannian manifolds, *J. Differential Geometry* **2** (1968), 421–446.

A Finding the Jordan-Chevalley decomposition

We present the proof of Lemma 2.9. Corollary 2.10 follows immediately. Variants of this algorithm are well-known (see, e.g., [Ga, Chapter V]), but the complexity analysis seems to be new.

Let $a \in M_n(F)$ have minimal polynomial $m(x)$. Let $p(x)$ be the polynomial obtained by dividing $m(x)$ by $GCD(m(x), m'(x))$. Then $p(x)$ is square-free, and $p(x)^n \equiv 0 \bmod m(x)$. We wish to find a polynomial $q(x) \in F[x]$ such that $q(x) \equiv x \bmod p(x)$ and $p(q(x)) \equiv 0 \bmod p(x)^n$. These conditions remain satisfied if we add any multiple of $p(x)^n$ to $q(x)$, so we may require that $q(x)$ have degree less than dn, where d is the degree of p. With this degree restriction, the polynomial $q(x)$ exists and is unique. Let $\lambda_1, \ldots, \lambda_d$ be the roots of $p(x)$, possibly over an extension field of F. Then $q(x)$ is easily seen to be the unique polynomial of degree less than dn such that for all i, $q(x) \equiv \lambda_i \bmod (x - \lambda_i)^n$.

Once we have q, the element $a_d = q(a)$ is a diagonalizable element with $a_d - a$ nilpotent. We will find q by successively finding $q_0, q_1, \ldots$ such that for all ℓ, $q_\ell(x) \equiv x \bmod p(x)$ and $p(q_\ell(x)) \equiv 0 \bmod p(x)^{2^\ell}$. By working modulo $p(x)^{2^\ell}$, we insure that q_ℓ has degree less than $d2^\ell$.

We begin with $q_0 = x$. By the Euclidean algorithm, we find polynomials $s(x), t(x) \in F[x]$ such that $s(x)p(x) + t(x)p'(x) = 1$. Let $q_1(x) = x - t(x)p(x)$. Consider $q_1(x + y) \bmod (p(x), y^2)$:

$$q_1(x+y) \equiv q_1(x) + y q_1'(x) \equiv x + y(1 - t(x)p'(x)) \equiv x \bmod (p(x), y^2).$$

By induction, suppose that we have q_ℓ such that for all i, $q_\ell(x) \equiv \lambda_i \bmod (x - \lambda_i)^{2^\ell}$. We define $q_{\ell+1}(x)$ as the remainder of $q_1(q_\ell(x)) \bmod p(x)^{2^{\ell+1}}$. Then $q_{\ell+1}(x) \equiv \lambda_i \bmod (x - \lambda_i)^{2^{\ell+1}}$. After $\lceil \log_2 n \rceil$ iterations, we have found $q(x)$.

It remains to be shown that the coefficients of q have polynomially bounded bit lengths. To show this, we give another characterization of $q_\ell(x)$. It is not hard to see that $q_\ell(x)$ is the unique polynomial of degree less than $d2^\ell$ such that:

$$\begin{array}{rcl}
q(x) & \equiv & x \bmod p(x) \\
q'(x) & \equiv & 0 \bmod p(x) \\
q''(x) & \equiv & 0 \bmod p(x) \\
\vdots & \vdots & \vdots \\
q^{(2^\ell - 1)}(x) & \equiv & 0 \bmod p(x)
\end{array}$$

Note that the above conditions are equivalent to a system of (polynomial time computable) linear equalities in the coefficients of $q_\ell(x)$. By a result of Edmonds (cf. [Schr]), the bit length of a unique solution of a system of linear equalities is polynomially bounded. (Note that this yields a second algorithm to find q, but it is not as fast as the first one.)

This completes the proof of Lemma 2.9.♠