Foundations of computing with infinite linear groups

Dane Flannery (joint with Alla Detinko)

Motivation: dearth of and need for *practical* methods to compute with linear groups over infinite fields.

Contrasts with the situation for matrix groups over finite fields; see the 'Matrix Group Recognition Project'.

Our approach uses traditional linear group theory and computational group theory to design algorithms for problems of basic, strategic interest (such as 'recognizing' an input linear group) and from further afield, in, e.g., topology, geometry, etc.

We emphasize practicality; software is a compulsory outcome: efficient algorithms, implemented in the MAGMA and GAP computer algebra systems.

Outline

- Computing with solvable-by-finite linear groups (lecture 1);
- Computing with groups in the other class of the Tits alternative: applications of the strong approximation property for linear groups (lectures 2, 3).

The lectures are based on joint work with

- Alla Detinko
- Alexander Hulpke;

and earlier work with Alla and Eamonn O'Brien.

Throughout $\mathbb{F}$ is a field.

For a commutative ring R with 1, as usual $\mathrm{GL}(n, R)$ denotes the group of invertible $n \times n$ matrices with entries in R.

That is, $\mathrm{GL}(n, R)$ is the unit group of $\mathrm{Mat}(n, R)$, the ring of all $n \times n$ matrices with entries in R.

$\mathrm{SL}(n, R)$ is the (normal) subgroup of $\mathrm{GL}(n, R)$ comprising all matrices with determinant 1. Note $\mathrm{GL}(n, R)/\mathrm{SL}(n, R) \cong R^\times$, the unit group of R.

A *linear group* (interchangeably, *matrix group*) *of degree* n *over* $\mathbb{F}$ is a subgroup of $\mathrm{GL}(n, \mathbb{F})$.

Setting up finite approximation for linear groups

We restrict to finitely generated linear groups; a natural restriction for computing.

Other formats are possible. A linear algebraic group may not be finitely generated, but can be input as a finite set (of defining polynomials).

Computing with linear algebraic groups is a well-developed area, tied to computation with Lie algebras. See the book by de Graaf (2017).

The method of *finite approximation* in linear group theory is classical. It was introduced by Mal'cev (1940s) and developed by Platonov, Wehrfritz et al. (1960s/70s).

Let R be a commutative ring with 1, and $\rho \subset R$ an ideal. Natural surjection $R \to R/\rho$ extends entrywise to a ring epimorphism

$$\mathrm{Mat}(n, R) \to \mathrm{Mat}(n, R/\rho),$$

which restricts to a group homomorphism

$$\mathrm{GL}(n, R) \to \mathrm{GL}(n, R/\rho).$$

We denote each of these maps by φ_ρ, and call them *congruence homomorphisms* modulo ρ.

Let $G \leq \mathrm{GL}(n, R)$. The kernel of $\varphi_\rho \colon G \to \mathrm{GL}(n, R/\rho)$ is a *congruence subgroup* of G.

Theorem (Mal'cev)

Let R be a finitely generated integral domain. Let $g_1, \ldots, g_r$ be distinct elements of $\mathrm{Mat}(n, R)$. Then there exists a maximal ideal ρ of R such that $\varphi_\rho(g_1), \ldots, \varphi_\rho(g_r)$ are distinct elements of $\mathrm{Mat}(n, R/\rho)$.

Note: a finitely generated integral domain is a homomorphic image of a polynomial ring $\mathbb{Z}[\mathrm{x}_1, \ldots, \mathrm{x}_m]$ for some $m \geq 0$.

Proof. For each pair k, l, $1 \leq k < l \leq r$, choose i, j such that $\delta_{k,l} := (g_k)_{i,j} - (g_l)_{i,j} \neq 0$. Let $\delta = \prod_{k,l} \delta_{k,l}$; so $\delta \neq 0$. Since the Jacobson radical of R is equal to its nilradical (by finite generation of R; fact from commutative algebra) and hence zero, $\varphi_\rho(\delta) = \prod_{k,l} \varphi_\rho(\delta_{k,l}) \neq 0$ for some maximal ideal ρ. Thus, $\forall\, k, l$, $\varphi_\rho(g_k), \varphi_\rho(g_l)$ differ in some position, i.e., $\varphi_\rho(g_k) \neq \varphi_\rho(g_l)$. □

Main context: if $G = \langle X \rangle$, the group generated by finite $X \subseteq \mathrm{GL}(n, \mathbb{F})$, then let $R \subseteq \mathbb{F}$ be the subring generated by the entries of g, g^{-1} $\forall\, g \in X$. Hence R is a finitely generated integral domain and $G \leq \mathrm{GL}(n, R)$.

R/ρ in Malcev's theorem is a *finite field* (another fact from commutative ring theory; e.g., $\mathbb{Z}[\mathrm{x}]/\rho$ is a finite field for maximal ideals ρ), so

Corollary

If R a finitely generated integral domain then $\mathrm{GL}(n, R)$ is residually finite.

A group is *residually finite* if the intersection of its finite-index subgroups is trivial; each non-identity element survives in some finite quotient.

In particular, *finitely generated linear groups are residually finite*. These statements are the essence of finite approximation.

Fields for computing

Lemma

If $\mathbb{F}/\mathbb{E}$ is a finitely generated field extension, then $\mathbb{F}$ is a finite-degree extension of $\mathbb{E}(\xi_1, \ldots, \xi_m)$ for some $\mathbb{E}$-algebraically independent $\xi_1, \ldots, \xi_m \in \mathbb{F}$, $m \geq 0$.

'$\mathbb{E}$-algebraically independent' means that $f(\xi_1, \ldots, \xi_m) \neq 0$ for all non-zero polynomials $f(\mathrm{x}_1, \ldots, \mathrm{x}_m) \in \mathbb{E}[\mathrm{x}_1, \ldots, \mathrm{x}_m]$.

$\{\xi_1, \ldots, \xi_m\}$ is a *transcendence basis* for $\mathbb{F}$ over $\mathbb{E}$;

$\mathbb{F}$ as in the lemma is an *algebraic function field*: a finite-degree extension of a field of rational functions.

By the lemma, in algorithms we restrict input to finitely generated $G \leq \mathrm{GL}(n, \mathbb{F})$ where $\mathbb{F}$ is one of

(i) $\mathbb{Q}$;

(ii) an algebraic number field;

(iii) a function field $\mathbb{P}(\mathrm{x}_1, \ldots, \mathrm{x}_m)$ where $\mathbb{P}$ is finite or a number field;

(iv) an algebraic function field: finite-degree extension of a field as in (iii).

Computation with fields of types (i)–(iv) is supported in Magma.

Applying finite approximation I: deciding finiteness

In computing with a potentially infinite group G, naturally we ask at the outset: is G finite? Is the question even decidable (does an algorithm to answer it exist)?

Decidability of finiteness of residually finite groups is unknown.

We show that finiteness is decidable for finitely generated linear groups, over 'any' field, i.e., after reducing to one of the four main field types.

Previous work on finiteness testing: Babai, Beals, & Rockmore, over $\mathbb{Q}$ (1993); Ivanyos, over function fields of positive characteristic (2001).

At the heart of our finiteness-testing algorithm is the following.

Theorem (Selberg–Wehrfritz)

Let $G \leq \mathrm{GL}(n, R)$, where R is a finitely generated integral domain. Then G has a finite-index normal subgroup N such that the finite-order elements of N are unipotent.

Unipotent: all eigenvalues are 1. A unipotent subgroup of $\mathrm{GL}(n, \mathbb{F})$ can be conjugated to a group of (upper) unitriangular matrices.

Also, if $\mathrm{char}\, \mathbb{F} = 0$ then a unipotent element $\neq 1_n$ of $\mathrm{GL}(n, \mathbb{F})$ has infinite order; if $\mathrm{char}\, \mathbb{F} = p > 0$ then a unipotent element has p-power order.

Hence, if $\mathrm{char}\, R = 0$ then N in the theorem is torsion-free (has no non-trivial finite-order elements).

DFO (2013): given finitely generated $G \leq \mathrm{GL}(n, R)$, we can construct a congruence homomorphism φ_ρ on $\mathrm{GL}(n, R)$ such that N as in the Selberg–Wehrfritz theorem is the congruence subgroup $G_\rho := \ker \varphi_\rho$ on G.

Here ρ is a certain maximal ideal of R with R obtained as usual from a generating set of $G \leq \mathrm{GL}(n, \mathbb{F})$, $\mathbb{F}$ one of the four main types of field.

Now periodic finitely generated linear groups are finite (Schur). And a finite-index subgroup of G is finitely generated (Schreier).

Thus, by the S–W theorem, finiteness of G can be decided by testing whether

- $G_\rho = \{1_n\}$ if $\mathrm{char}\, \mathbb{F} = 0$, or
- G_ρ is unipotent (a p-group) if $\mathrm{char}\, \mathbb{F} = p > 0$.

Call φ_ρ as above an *SW-homomorphism*, and its kernel G_ρ on G an *SW-subgroup* of G. To repeat: we can construct SW-homomorphisms for each of the four main field types.

Example. Let G be a finitely generated subgroup of $\mathrm{GL}(n, \mathbb{Q})$. Then $G \leq \mathrm{GL}(n, \mathbb{Z}[1/\mu])$ for some positive integer μ. Assume that $\mu = 1$ (if G is finite then it can be conjugated into $\mathrm{GL}(n, \mathbb{Z})$ by a result of Burnside).

Simple matrix arithmetic shows that the reduction modulo m map φ_m on $\mathrm{GL}(n, \mathbb{Z})$ for $m > 2$ has torsion-free kernel (Minkowski); so φ_m is an SW-homomorphism, and $G_m := \ker \varphi_m$ on G is an SW-subgroup.

G is finite $\Leftrightarrow G_m = \{1_n\}$.

How to test whether G_m is trivial?

Although an SW-subgroup G_ρ is finitely generated, we *don't* need a full generating set of G_ρ.

Instead we use the following standard technique. For $G = \langle g_1, \ldots, g_r \rangle$, take a presentation of the matrix group $\varphi_\rho(G)$ over a finite field.

The relators in this presentation need to be words $w_j(\varphi_\rho(g_1), \ldots, \varphi_\rho(g_r))$ in the $\varphi_\rho(g_i)$. Rewrite each w_j as $\widetilde{w_j} := w_j(g_1, \ldots, g_r) \in G$.

Lemma

The $\widetilde{w_j}$ are 'normal generators' of G_ρ: they generate a subgroup of G whose normal closure is G_ρ.

In particular, $G_\rho = \{1_n\} \Leftrightarrow$ each of its normal generators is 1_n.

Now suppose that we have tested $G \leq \mathrm{GL}(n, \mathbb{F})$ and found that G is infinite. A natural next step is to consider the Tits alternative for G.

Theorem (J. Tits, 1972)

Let G be a finitely generated linear group. Then either G is virtually solvable (solvable-by-finite, SF), or G contains a non-abelian free subgroup.

H SF means H has a (normal) solvable subgroup of finite index.

Tits' theorem partitions finitely generated linear groups into two very different classes. It is a fundamental problem to decide which of these classes contains a given $G \leq \mathrm{GL}(n, \mathbb{F})$.

Applying finite approximation II: deciding virtual solvability

In this application of finite approximation, our algorithm is suggested by the following.

Theorem (Kolchin–Lie–Mal'cev)

Let $G \leq \mathrm{GL}(n, \mathbb{F})$ be solvable. Then G has a UA (unipotent-by-abelian) normal subgroup of finite index. If furthermore G is connected in the Zariski topology on $\mathrm{GL}(n, \mathbb{F})$, then G is UA.

$H \leq \mathrm{GL}(n, \mathbb{F})$ being UA means that H has a normal unipotent subgroup N such that H/N is abelian.

It can be shown that H is UA $\Leftrightarrow$ there exists a finite-degree field extension $\mathbb{E}/\mathbb{F}$ s.t. H is conjugate in $\mathrm{GL}(n, \mathbb{E})$ to a group of triangular matrices.

Lemma

A triangular matrix group is solvable.

Let $G \leq \mathrm{GL}(n, \mathbb{F})$ be finitely generated, $\mathbb{F}$ one of the four main types of field. Using vital results of Wehrfritz (2010), we are able to realize the Kolchin–Lie–Mal'cev theorem computationally: we construct a congruence homomorphism φ_ρ on G with $\varphi_\rho(G)$ over a finite field such that

$$G \text{ is SF } \implies G_\rho := \ker \varphi_\rho \text{ on } G \text{ is UA.}$$

Call such φ_ρ a *W-homomorphism*. Hence, by the above lemma:

G is SF $\Leftrightarrow$ the kernel of a W-homomorphism on G is UA.

Hence we can test virtual solvability and so decide the Tits alternative for finitely generated $G \leq \mathrm{GL}(n, \mathbb{F})$ in similar fashion to finiteness testing.

That is, we are testing whether a congruence subgroup G_ρ (for any W-homomorphism φ_ρ) has a particular property (in this problem, UA).

First we compute a normal generating set $\mathcal{N}$ for G_ρ.

The UA test works with the enveloping algebra of $\mathcal{N} \cup \mathcal{N}^{-1}$ (the smallest subalgebra of $\mathrm{Mat}(n, \mathbb{F})$ containing $\mathcal{N} \cup \mathcal{N}^{-1}$; a basis of it can readily be computed by a standard recursion). Again, we *don't* need a full generating set of G_ρ.

Cf. algorithms by Beals (1995, 2001) and Assmann & Eick (2007) to decide virtual solvability over $\mathbb{Q}$.

Computing over the infinite ground field

There are several sources of discomfort when computing with matrix groups over an infinite field.

'Entry explosion': too many multiplications can lead to huge matrix entries. Try to mitigate by replacement with matrix algebra computations, and transferring work to images over a finite field.

Decidability: problems may not have an algorithmic solution. E.g., by Mihailova (1958), the generalized word problem is undecidable in $\mathrm{SL}(4, \mathbb{Z})$.

On the positive side, our algorithms prove that other fundamental problems for finitely generated linear groups *are* decidable.

Software

Our finite approximation algorithms are available within MAGMA, implemented in collaboration with Eamonn O'Brien:

`http://magma.maths.usyd.edu.au/magma/handbook/matrix_groups_over_infinite_fields`

Implementations make heavy use of MGRP machinery in MAGMA (see Bäärnhielm et al., 2015) and associated advances in computational group theory, e.g., constructing 'short' presentations (for efficient determination of normal generators).

References

Classic texts on linear group theory:

- J. D. Dixon, *The structure of linear groups*, Van Nostrand Reinhold, 1971.
- D. A. Suprunenko, *Matrix groups*, American Math. Soc., 1976.
- B. A. F. Wehrfritz, *Infinite linear groups*, Springer-Verlag, 1971.

Algebraic groups:

- W. A. de Graaf, *Computation with linear algebraic groups*, CRC Press, 2017.

Finiteness testing:

- A. S. Detinko, D. L. Flannery, and E. A. O'Brien, Recognizing finite matrix groups over infinite fields, J. Symb. Comp. 50 (2013).

References (continued)

Deciding virtual properties of linear groups:

- B. A. F. Wehrfritz, Conditions for linear groups to have unipotent derived subgroups, J. Algebra 323 (2010).
- A. S. Detinko, D. L. Flannery, and E. A. O'Brien, Algorithms for the Tits alternative and related problems, J. Algebra 344 (2011).

Computing with matrix groups over finite fields:

- H. Bäärnhielm, D. F. Holt, C. R. Leedham-Green, and E. A. O'Brien, A practical model for computation with matrix groups, J. Symb. Comp. 68 (2015).