Algorithms for matrix groups

Eamonn O'Brien

University of Auckland

June 2015

REPUBLIKA SLOVENIJA
MINISTRSTVO ZA IZOBRAŽEVANJE, ZNANOST IN ŠPORT

Introduction

$G = \langle X \rangle \leq \mathrm{GL}(d, q)$

Can we answer the following?

- $|G|$
- Composition series or chief series for G
- Sylow p-subgroups
- Conjugacy classes of elements or subgroups of G
- Normaliser of $H \leq G$
- Maximal subgroups of G
- Automorphism group of G

Challenge problems

Problem

Find the order of $H \leq \mathrm{GL}(6, 5^2)$.

Problem

Given $g \in \mathrm{GL}(6, 5^2)$ *find its order.*

Problem

Find the normaliser in $\mathrm{GL}(8, 5^2)$ *of a subgroup of moderate index.*

If $G = \mathrm{Sym}(10^6)$, we can answer readily most questions about G, using "efficient" algorithms.

The "matrix recognition" project

Goal: efficient algorithms, both theoretically and practically.

One measure of algorithm performance:
in time polynomial in the size of the input

If f and g are real-valued functions, defined on all sufficiently large integers, then $f(n) = O(g(n))$ means $|f(n)| < C|g(n)|$ for some positive constant C and all sufficiently large n.

Matrix group complexity

$|\mathrm{GL}(d, q)| = O(q^{d^2})$

For $G = \langle X \rangle \leq \mathrm{GL}(d, q)$, $\log |G| < d^2 \log q$.

Observe $\log q$ bits are required for each of the d^2 entries in matrix

Input size is $|X| d^2 \log q$.

Desire: algorithms whose complexity involves $\log q$, not q.

Another measure: practical, implemented in GAP or Magma.

Complexity: cross characteristic representations

$G \simeq \mathrm{SX}(n, q)$ but $G \leq \mathrm{GL}(d, F)$ where $\chi(F) \neq \chi(\mathrm{GF}(q))$.

Landazuri & Seitz (1974), Seitz & Zalesskii (1993): faithful projective representations in cross characteristic have degree d that is **polynomial** in q.

Complexity measured in terms of d and $\log q$, but $d = O(q)$.

Since input complexity contains q, algorithms can have runtime complexity involving q: often **easy** to obtain algorithm which is polynomial in q.

So focus is often defining characteristic representation.

Outline of lecture series

- Randomness and classes of algorithms
- Basic tasks: multiplication, powering
- Permutation group analogues
- Recognition strategies
- Geometry after Aschbacher
- Identifying simple groups
- Simple groups: the tasks
- Constructive Recognition
- CompositionTree
- The Souble Radical model

References

For further details, consult these and their reference lists, available from www.math.auckland.ac.nz/~obrien

- "Algorithms for matrix groups", Groups St Andrews 2009 in Bath, LMS Lecture Notes **388**, 297–323, 2011.
- "Towards effective algorithms for linear groups", *Finite Geometries, Groups and Computation*, (Colorado), pp. 163-190. De Gruyter, Berlin, 2006.
- (with Dietrich and Leedham-Green) "Effective black-box constructive recognition of classical groups", *J. Algebra* 2015.
- Holt, Eick and O'B: "Handbook of Computational Group Theory", 2005.

Randomness

$|\mathrm{GL}(d, q)| = O(q^{d^2})$

Many algorithms are **randomised**: use random search in G to find elements having prescribed property $\mathcal{P}$.

Example

- Characteristic polynomial having factor of degree $> d/2$.
- Order divisible by prescribed prime.

Common feature: algorithms depend on detailed analysis of **proportion** of elements of finite simple groups satisfying $\mathcal{P}$.

Classes of algorithms

Definition

A Monte Carlo algorithm is a randomised algorithm which may return an incorrect answer to a decision question, the probability of this event being less than some ϵ.

If one of the answers is always correct, then it is **one-sided**.

Definition

A Las Vegas algorithm is one which never returns an incorrect answer, but may report failure with probability less than ϵ.

Assume we determine a lower bound, say $1/k$, for proportion of elements in G satisfying Property $\mathcal{P}$.

To find element satisfying $\mathcal{P}$ by random search with a probability of failure less than given $\epsilon \in (0, 1)$: choose a sample of uniformly distributed random elements in G of size at least $\lceil -\log_e(\epsilon) \rceil k$.

Black-box groups

Babai & Szemerédi (1984)

Group elements represented by bit-strings of uniform length.

Operations: multiplication, inversion, and checking for equality with the identity element.

Representation-independent: model includes permutation groups and matrix groups defined over $\mathrm{GF}(q)$.

Definition

Black-box algorithm does not use specific features of the group representation, nor particulars of how group operations are performed; it uses only these operations.

Oracles

Common assumption is that *oracles* are available to perform certain tasks, often those not known to be solvable in polynomial time.

Examples of oracles include:

- ξ to construct a (nearly) uniformly distributed random element of G as an SLP in X.
- ρ to compute the order of a given $g \in G$.
- Π to compute a given power of $g \in G$.

Discrete Log Oracle

Given $G \leq \mathrm{GL}(d, \mathbb{Z})$, and $x \in \mathrm{GL}(d, \mathbb{Z})$: is $x \in G$?

Mihailova (1958): membership problem is undecidable for $d \geq 4$.

$\mathrm{GF}(q)$: $|\mathrm{GL}(d, q)| = O(q^{d^2})$

Membership decidable from exhaustive search.

Difficult even for $\ldots 1 \times 1$ matrices over $\mathrm{GF}(\mathrm{q})$:

Example

$H := \langle [561], [520], [320] \rangle \leq \mathrm{GL}(1, 593)$.

Membership related to Discrete log problem

Problem

$F = \mathrm{GF}(q)$, $\omega \in F$ primitive.
Given $\alpha \in F^{\times}$, determine k so that $\alpha = \omega^k$.

No polynomial-time algorithm known.

Realising oracles

Desire: polynomial-time black box algorithm to realise oracle.

Complexity often depends on input representation.

Example

Order of $g \in G$.

If $G \leq S_n$, trivial.

If $G \leq \mathrm{GL}(d, q)$, hard.

If G is black-box, then only available algorithm is: enumerate $g, g^2, g^3, \ldots$ until $g^n = 1$.

Oracle: Generate random elements of finite group

Babai (1991): Vertex-transitive graph approach

Independent nearly uniformly random distributed elements of finite group $G = \langle X \rangle$ can be found after a preprocessing stage consisting of $O(\log^5 |G|)$ group operations.

Preprocessing proceeds in $O(\log |G|)$ phases.

In each phase, random walk of random length between 1 and $O((\log |G|)^4)$ performed on Cayley graph of G.

Element found when walk finished is added to generators of G.

Walk is repeated $O(\log |G|)$ times.

Final list S of $O(\log |G|)$ elements input to construction phase.

Random element is *random subproduct* of S:

$$g_1^{\epsilon_1} \cdots g_m^{\epsilon_m}$$

where $S = \{g_1, \ldots, g_m\}$ and $\epsilon_i \in \{0, 1\}$ (chosen independently).

For $G \leq \mathrm{GL}(d, q)$, $\log |G| < d^2 \log q$.

Initialisation phase $O(d^{10} \log^5 q)$.

Cost per random element is $O(\log |G|)$.

CLMNO (1995): Product replacement algorithm

Input: ordered list of generators $[g_1, \ldots, g_m]$ for G.

Accumulator: r initialised to be identity of G.

Basic step:

- Select at random i, j where $1 \leq i, j \leq m$.
- Replace g_i by either $g_i g_j$ or $g_j g_i$.
- Multiply r by g_i.

Basic step repeated a number, say t, of times.

Now to obtain random element: execute basic operation once, and return r as random element.

Cost: after initialisation, two matrix multiplications.

Theorem

Let T be set of all m-tuples of generators of G. Then the algorithm constructs a Markov chain over state space T, and if m is at least twice the size of a minimal generating set of generators for G, this Markov chain is connected and aperiodic.

The random walk approaches a limiting distribution at exponential rate $O((1-\delta)^t)$ where t is number of steps taken.

Mixing time

What can we say about the "mixing time", t?

Variety of statistical tests applied to test outcome of algorithm. Practical: excellent.

- Diaconis & Saloff-Coste (1997, 1998):
  $t = O(\delta^2(G, S) \cdot m)$, where $\delta(G, S)$ is the maximal diameter for the Cayley graph of G wrt generating set S.
  Comparison of two Markov chains on different but related state spaces and combinatorics of random paths.
- Pak (2001): Mixing time is polynomial. Multi-commodity flow technique.
- Lubotzky & Pak (2002):
  Does the group of automorphisms of a free group of rank > 3 have Kazhdan's property (T)? If so, then "graph of states" is well-behaved, giving excellent mixing time.

Basic operations

Multiplication of two $d \times d$ matrices A and B
Cost of $A \times B$ using conventional algorithm is $O(d^3)$.

Strassen: $O(d^{\log_2(7)})$

Coppersmith & Winograd (1990): $O(d^{2.37})$

Where do we notice improvements? Perhaps for $d \geq 100$.

Null space

The *nullspace* of a $d \times c$ matrix A is the subspace of F^d consisting of those $v \in F^d$ for which $v \cdot A = 0$.

The standard method to compute the nullspace involves performing elementary column operations on A, which do not change the nullspace of A.

To use row operations, replace A by the $c \times d$ matrix $B := A^T$, and then calculate the space of $v \in K^d$ such that $B \cdot v^T = 0$.

Characteristic and minimal polynomials

The *minimal polynomial* $m(x)$ of $A \in \mathrm{GL}(d, F)$ is the monic polynomial of least degree such that $m(A)$ is the zero matrix.

Lemma

The minimal polynomial is unique. If $c(x)$ is the characteristic polynomial of A then $m(x)$ divides $c(x)$.

How do we calculate the minimal polynomial of A?

Let v be nonzero vector in $V = F^d$. Key idea: generate images $v, vA, vA^2, \ldots$ until we find a linear relation among them.

This describes the minimal polynomial of A restricted to the subspace W spanned by the vectors vA^j.

If $W < V$ repeat calculation for $v \notin W$.

Iterating, obtain polynomials f_i and subspaces W_i, where f_i is the minimal polynomial of A restricted to W_i and $\sum W_i = V$.

Now $m(A)$ is $\mathrm{lcm}(f_i)$.

Let $0 \neq v \in F^d$. If $v \cdot A^j$ are linearly independent for $0 \leq j < r$, but $v \cdot A^r = \sum_{j=0}^{r-1} a_j v \cdot A^j$ for some $a_j \in F$, then

$$\rho_v(x) := x^r - a_{r-1}x^{r-1} - \cdots - a_1 x - a_0$$

is the unique monic polynomial of minimal degree over F such that $v \cdot \rho_v(A) = 0$.

Let $W = \langle v \cdot A^j \mid 0 \leq j < r \rangle$. The action of A on W with respect to this basis is

$$\begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ & & & \cdots & \\ 0 & 0 & 0 & \cdots & 1 \\ a_0 & a_1 & a_2 & \cdots & a_{r-1} \end{pmatrix}.$$

This is the *companion matrix* of the polynomial ρ_v.

Example

$$A := \begin{pmatrix} 0 & 2 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 2 & 2 & 0 & 2 \end{pmatrix}$$

with entries in $\mathrm{GF}(3)$.

Let $v = e_1 = (1, 0, 0, 0)$. Now $vA = (0, 2, 0, 1)$ and $vA^2 = (2, 1, 0, 2)$. Observe that $vA^2 \in \langle v, vA \rangle$ and $vA^2 = 2v + 2vA$. Let $W_1 = \langle v, vA \rangle$.

The matrix of A in its action on W_1 is

$$\begin{pmatrix} 0 & 1 \\ 2 & 2 \end{pmatrix}$$

and the minimal polynomial is $f_1(x) = x^2 + x + 1$.

Now choose $v = e_3 = (0, 0, 1, 0)$. Now $vA = (1, 2, 1, 0)$ and $vA^2 = (1, 0, 1, 1)$ and since $A^3 = 1$, $vA^3 = v$.

The matrix of A in its action on W_1 is

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

and the minimal polynomial is $f_2(x) = x^3 + 2$.

Now $V = W_1 + W_2$ and hence $m(x) = \mathrm{lcm}(f_1, f_2) = x^3 + 2$ is the minimal polynomial.

Costs for characteristic and minimal polynomials

Lemma

(i) *Multiplication and division operations for polynomials of degree d defined over $\mathrm{GF}(q)$ can be performed deterministically in $O(d \log d \log\log d)$ field operations. Using a Las Vegas algorithm, such a polynomial can be factored into its irreducible factors in $O(d^2 \log d \log\log d \log(qd))$ field operations.*

(ii) *Using Las Vegas algorithms, both the characteristic and minimal polynomial of $g \in \mathrm{GL}(d, q)$ can be computed in $O(d^3 \log d)$ field operations.*

Determine the order of a matrix

Let $g \in \mathrm{GL}(d, q)$.
Find $n \geq 1$ such that $g^n = 1$.

$\mathrm{GL}(d, q)$ has elements of order $q^d - 1$, Singer cycles, ...

so not practical to compute powers of g until $g^n = 1$.

To find $|g|$: probably requires factorisation of numbers of form $q^i - 1$, a hard problem.

Babai & Beals (1999):

Theorem

If the set of primes dividing a multiplicative upper-bound B for $|g|$ is known, then the precise value of $|g|$ can be determined in polynomial time.

Celler & Leedham-Green (1995): compute $|g|$ in time $O(d^4 \log q)$ subject to factorisation of $q^i - 1$ for $1 \leq i \leq d$.

- First compute a "good" multiplicative upper bound B for $|g|$.

Determine and factorise minimal polynomial for g as

$$m(x) = \prod_{i=1}^{t} f_i(x)^{m_i}$$

where $\deg(f_i) = d_i$ and $\beta = \lceil \log_p \max m_i \rceil$.

$B := \prod_{i=1}^{t} \operatorname{lcm}(q^{d_i} - 1) \times p^{\beta}$

Lemma

Let $B = \prod_{i=1}^{t} \mathrm{lcm}(q^{d_i} - 1) \times p^{\beta}$. *Then* $|g|$ *divides* B.

To see this, reduce g to Jordan normal form over the algebraic closure of $\mathrm{GF}(q)$.

Each eigenvalue lies in an extension field of $\mathrm{GF}(q)$ of dimension d_i and so has multiplicative order dividing $q^{d_i} - 1$.

If a block has size $\gamma_i > 1$, then the p-part of the order of the block is p^{δ} where $\delta = \lceil \log_p \gamma_i \rceil$.

Hence $o(g) | \mathrm{lcm}(q_i^d - 1) \times p^{\beta}$.

Can we use B to learn $|g|$?

1. Factorise $B = \prod_{i=1}^{m} p_i^{\alpha_i}$ where the primes p_i are distinct.
2. If $m = 1$, then calculate $g^{p_1^j}$ for $j = 1, 2, \ldots, \alpha_1 - 1$ until the identity is constructed.
3. If $m > 1$ then express $B = uv$, where u, v are coprime and have approximately the same number of distinct prime factors. Now g^u has order k dividing v and g^k has order ℓ say dividing u, and $|g|$ is $k\ell$. Hence the algorithm proceeds by recursion on m.

Celler & Leedham-Green prove the following:

Theorem

If we **can compute** *a factorisation of B, then the cost of the algorithm is $O(d^4 \log q \log\log q^d)$ field operations.*

If we don't complete the factorisation, then obtain *pseudo-order* of g – the order $\times$ some large primes.

Suffices for most theoretical and practical purposes.

Implementations in both GAP and Magma use databases of factorisations of numbers of the form $q^i - 1$, prepared as part of the Cunningham Project.

Example

$$A = \begin{pmatrix} 2 & 5 & 1 & 2 \\ 0 & 1 & 6 & 1 \\ 4 & 0 & 2 & 2 \\ 3 & 3 & 6 & 6 \end{pmatrix}$$

with entries in $\mathrm{GF}(7)$. A has minimal polynomial

$$m(x) = x^4 + 3x^3 + 6x^2 + 6x + 4 = (x+4)^2(x^2 + 2x + 2)$$

Hence $e_1 = 1, e_2 = 2$ and $\beta = \lceil \log_7 2 \rceil = 1$. Hence

$B = (7^1 - 1)(7^2 - 1)7^1 = 336.$

Now $336 = 2^4 \cdot 3 \cdot 7 = uv$ where $u = 2^4$ and $v = 3 \cdot 7$.

A^u has order dividing v. Reapply: $|A^u| = 21$.

A^v has order dividing u. Reapply: $|A^v| = 8$.

Conclude $|A| = 168$.

Element has even order?

Task: Determine if g has *even* order.

Can compute a multiplicative upper bound for $|g|$ in polynomial time.

To obtain $|g|$ requires knowledge of factorisation of B.

However, if we know B, then we can learn in polynomial time the *exact* power of 2 (or of any specified prime) which divides $|g|$.

By repeated division by 2, we write $B = 2^m b$ where b is odd.

Now we compute $h = g^b$, and determine (by powering) its order which divides 2^m.

Oracle: Compute power of element of group

We can compute large powers n of $g \in G$ in at most $2 \lfloor \log_2 n \rfloor$ multiplications by the standard doubling algorithm:

- $g^n = g^{n-1} g$ if n is odd
- $g^n = g^{(n/2)2}$ if n is even.

Black-box algorithm, complexity cost of $O(\log_2 n)$ multiplications.

Alternative for *matrices* by Leedham-Green and O'B (2009).

Frobenius normal form or . . .

Rational canonical form of a square matrix A is a canonical form that reflects the structure of the minimal polynomial of A. Can be constructed over given field, no need to extend field.

Definition

A is equivalent to $\begin{pmatrix} C_1 & 0 & \dots & 0 \\ 0 & C_2 & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & C_r \end{pmatrix}$.

Each block C_i is the companion matrix of monic $f_i \in F[x]$ and $f_i | f_{i+1}$ for $1 \leq i \leq r$.

The minimal polynomial of A is f_l and char poly is $f_1 \cdot f_2 \dots f_l$.

Frobenius normal form N of A is sparse.

Hence multiplication by N costs just $O(d^2)$ field operations.

A faster power algorithm

1. Construct the Frobenius normal form of g and record change-of-basis matrix C.
2. From the Frobenius normal form, read off the minimal polynomial $m(x)$ of g, and factorise $m(x)$ as a product of irreducible polynomials.
3. Compute multiplicative upper bound, B, to the order of g.
4. If $n > B$, then replace n by $n \bmod B$. By repeated squaring, calculate $x^n \bmod m(x)$ as a polynomial of degree $k-1$, where k is the degree of $m(x)$.
5. Evaluate this polynomial in g to give g^n.
6. Now compute $C^{-1}g^nC$ to return to the original basis.

Complexity of this task

Theorem

Let $g \in \mathrm{GL}(d, q)$ and let $0 \leq n < q^d$. This is a Las Vegas algorithm that computes g^n in $O(d^3 \log d + d^2 \log d \log\log d \log q)$ field operations.

Permutation groups: Base and strong generating set

G acts faithfully on $\Omega = \{1, \ldots, n\}$

Base: $B = [\epsilon_1, \epsilon_2, \ldots, \epsilon_m] \subset \Omega$ where $G_{\epsilon_1,\epsilon_2,\ldots,\epsilon_m} = 1$.

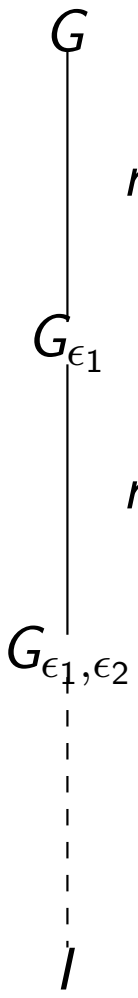

$r_1 = |\epsilon_1^G|$, orbit of ϵ_1 under G

$r_2 = |\epsilon_2^{G_{\epsilon_1}}|$, orbit of ϵ_2 under G_{ϵ_1}

$|G| = r_1 \ldots r_m$

Sims (1970, 1971): base and strong generating set (BSGS).

Base: sequence of points $B = [\epsilon_1, \epsilon_2, \ldots, \epsilon_k]$ where $G_{\epsilon_1,\epsilon_2,\ldots,\epsilon_k} = 1$.

This determines chain of stabilisers

$$G = G^{(0)} \geq G^{(1)} \geq \cdots \geq G^{(k-1)} \geq G^{(k)} = 1,$$

where $G^{(i)} = G_{\epsilon_1,\epsilon_2,\ldots,\epsilon_i}$.
S strong generating set: $G^{(i)} = \left\langle S \cap G^{(i)} \right\rangle$

Example

$G = \langle (1,5,2,6), (1,2)(3,4)(5,6) \rangle$
$B = [1,3]$
$G > G_1 > G_{1,3} = 1$
$S = \{(1,5,2,6), (1,2)(3,4)(5,6), (3,4)\}$

Central task: construct *basic orbits* – orbit B_i of the base point ϵ_{i+1} under $G^{(i)}$.

$$|G^{(i)} : G^{(i+1)}| = \#B_i$$

Schreier's Lemma gives generating set for each $G^{(i)}$.

Base image $B^g = [\epsilon_1^g, \ldots \epsilon_k^g]$ *uniquely* determines g:

if $B^g = B^h$ then $B^{gh^{-1}} = B$, so $gh^{-1} = 1$. Hence g can be represented as $|B|$-tuple.

Let U_i be transversal of $G^{(i+1)}$ in $G^{(i)}$.

Transversal provide normal form: every $g \in G$ has **unique** representation $g = u_k u_{k-1} \dots u_1$ where $u_i \in U_i$.

Sifting algorithm provides membership test for G: write $g \in G$ uniquely as $g = u_k u_{k-1} \dots u_1$ where $u_i \in U_i$.

For many interesting $G \leq S_n$, $|B|$ is small compared to n: *short base* groups.

Luks et al. (1980), Seress (2003): polynomial time.

Variations underpin both theoretical and practical approaches to permutation group algorithms.

Schreier-Sims for matrix groups

G acts faithfully on $V = F^d$: $v \cdot g$, for $v \in V$

Compute BSGS for G, viewed as permutation group on the vectors.

Base points: standard basis vectors for V.

Central problem: basic orbits B_i large. Usually $|B_1|$ is $|G|$.

Butler (1979): action of G on one-dimensional subspaces of V.

Murray & O'Brien (1995): heuristic algorithm to select base points: certain common eigenvectors/spaces for collection of elements of G.

Example

$J_4 \leq \mathrm{GL}(112, 2)$: choose base points of dimension $1, 11, 10, 1$ to get optimal orbits.

Critical for success: **index of one stabiliser in its predecessor**.

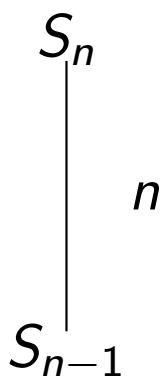

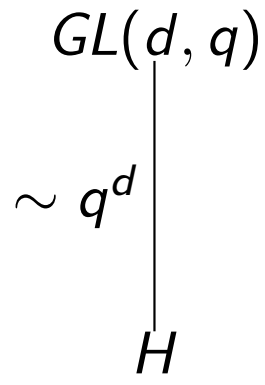

$|S_n : S_{n-1}| = n$

$\mathrm{GL}(d, q) \geq q^{d-1}.\mathrm{GL}(d-1, q) \geq \mathrm{GL}(d-1, q) \geq \ldots$

O'B and Wilson (2003): good base points for all repns of sporadics in Atlas. But ...

Example

Largest maximal subgroup $2^{11} : M_{24} \leq J_4$ index 173 067 389.

Neunhöffer et al. (2000s): use "helper subgroups" to construct large orbits; requires detailed knowledge of specific group. Most applicable for computations with specific group.

The basic strategies

- Geometry following Aschbacher
- Characteristic structure

Both provide composition series (and more) for G.

Geometry following Aschbacher

Aschbacher (1984)

G maximal subgroup of $\mathrm{GL}(d, q)$, let V be underlying vector space

- G preserves some natural linear structure associated with the action of G on V, and has normal subgroup related to this structure,
- or G is almost simple modulo scalars: $T \leq G/Z \leq Aut(T)$ where T is simple.

Geometry following Aschbacher: general strategy

1. Determine (at least one of) its Aschbacher categories.
2. If $N \lhd G$ exists, recognise N and G/N recursively, ultimately obtaining a composition series for the group.

7 categories giving normal subgroup

Prototype: G acts imprimitively on V

G preserves decomposition of V as direct sum

$$V_1 \oplus V_2 \oplus \cdots \oplus V_r$$

of $r > 1$ subspaces of dimension s, which are permuted transitively by G.

Then $\phi : G \to S_r$ where $r \leq d$ and $N = \ker \phi$.

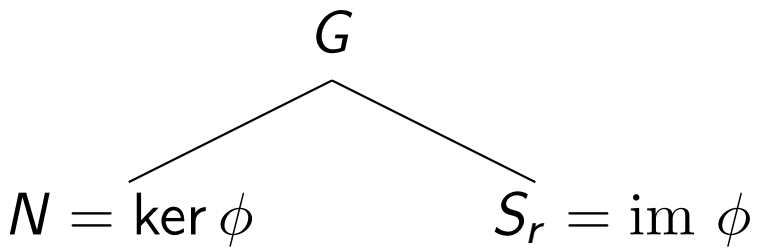

Holt, Leedham-Green, O'B & Rees (1996)

Geometry following Aschbacher: general strategy

$G = \langle X \rangle \leq \mathrm{GL}(d, q)$.

1. Determine (at least one of) its Aschbacher categories.
2. If $N \lhd G$ exists, recognise N and G/N recursively, ultimately obtaining a composition series for the group.
3. Otherwise G is either classical group in natural representation e.g. $G = \mathrm{SL}(d, q)$, invertible matrices of determinant 1; or $T \leq G/Z \leq Aut(T)$ where T is simple.
   - "Reduce" from G to (quasi)simple group L.
   - Name L.
   - Set up "constructive isomorphisms" between L and its *standard copy*.

CompositionTree: exploits geometry to produce composition series for G, factors are **leaves** of tree.

Aschbacher's classification

Aschbacher (1984)
G maximal subgroup of $\mathrm{GL}(d, q)$, let V be underlying vector space

- G preserves some natural linear structure associated with the action of G on V, and has normal subgroup related to this structure,
- or G is almost simple modulo scalars: $T \leq G/Z \leq Aut(T)$ where T is simple.

Theorem

$G \leq \mathrm{GL}(d, q)$ acts on $V := \mathbb{F}_q^d$, and Z is the subgroup of scalar matrices of G. If G is a maximal subgroup of $\mathrm{GL}(d, q)$, then one of the following is true:

C1. *G acts reducibly.*

C2. *G acts imprimitively.*

C3. *G acts on V as a group of semilinear automorphisms of a d/e-dimensional space over $GF(q^e)$, for some $e > 1$.*

C4. *G preserves a decomposition of V as a tensor product.*

C5. *G is definable modulo scalars over a subfield.*

C6. *$d = r^m$, prime r, and $G \leq$ normaliser of $ES(r^{2m+1})$, or a symplectic type group of order 2^{2m+2}.*

C7. *G preserves a decomposition of V as $V_1 \otimes V_2 \otimes \cdots \otimes V_m$, all of dimension $r > 1$, where $d = r^m$.*

C8. *G is classical group in its natural representation.*

C9. *$T \leq G/Z \leq \mathrm{Aut}\ T$, for non-abelian simple group T.*

A constructive version of Aschbacher's theorem?

Given $G = \langle X \rangle \leq \mathrm{GL}(d, F)$ acting on V.

Constructively decide (at least one of) its Aschbacher categories.

If $\ker \phi = N \lhd G$ exists, then construct both N and $\mathrm{im}\ \phi$.

Desirable: Polynomial-time decision.

Clifford's theorem (1937)

Describes the relation between representations of a group G and those of a normal subgroup N of finite index.

Theorem

Let $\pi : G \to GL(d, F)$ be an irreducible representation of G for a field F. Then the restriction of π to N breaks up into a direct sum of inequivalent irreducible representations of N of equal dimensions. These irreducible representations of N lie in one orbit of the action of G by conjugation on the equivalence classes of irreducible representations of N. The number of distinct summands is bounded by the index of N in G.

A constructive version of Clifford's theorem

Let non-scalar $N \lhd G$. Consider the restriction of V to N.

For some $t \geq 1$, V decomposes as direct sum $W_1 \oplus W_2 \oplus \cdots \oplus W_t$ of irreducible FN-modules, all same dimension.

For some $r, s \geq 1$, with $rs = t$, the W_is partition into r sets, each containing s pairwise-isomorphic FN-modules.

If $V_1, V_2, \ldots, V_r$ are each the sum of s pairwise isomorphic W_is, so that $V = V_1 \oplus V_2 \oplus \cdots \oplus V_r$, then G permutes the V_is transitively.

- If $r > 1$ then G acts imprimitively on V (C2).
- If $r = 1$ and $t > 1$ and the W_i are absolutely irreducible as FN-modules, then V is a tensor product preserved by G (C4).
- If $r = 1$ and the W_i are not absolutely irreducible as FN-modules, then G is semilinear (C3).

Otherwise, $r = s = 1$ and N acts absolutely irreducibly on V.

$N/Z(N) \cong N_0 \times N_0 \times \cdots \times N_0$ of m copies of simple group N_0.

If N_0 is cyclic, then G normalises a symplectic-type group (C6).

Otherwise N_0 is non-abelian simple.

- If $m > 1$ G is tensor-induced (C7).
- If $m = 1$, G is almost simple.

The SMASH algorithm

Holt, Leedham-Green, O'B, Rees (1996): constructive realisation.

Assume G acts absolutely irreducibly on V and $S \subseteq G$ contains at least one non-scalar element.

SMASH investigates whether G preserves decomposition with respect to $\langle S \rangle^G$.

Problem

How can we construct elements of relevant N?

Many heuristics apply.

Example

Assume $\phi : G \longmapsto S_r$. If $|g|$ not valid for S_r, then $g \in \ker \phi$.

Further developed and analysed by Neunhöffer (2008).

C1: Reducible groups

G acts reducibly on V if there exists $0 \neq U < V$ fixed by G.

Image of homomorphism ϕ is action of G on U, kernel of ϕ centralises U.

Maximal subgroups of $\mathrm{GL}(d, q)$ in $C1$ are maximal parabolics.

MeatAxe: Las Vegas algorithm to decide whether or not G acts irreducibly on V.

Original: Parker (1984).

Generalised by Holt & Rees (1994), analysis completed by Ivanyos & Lux (2000).

The MeatAxe

$G = \langle X \rangle$, M is FG-module, A is F-algebra spanned by X.

1. Select random $\theta \in A$, determine its characteristic polynomial $c(x)$ of θ, and factorise it.
2. Let $\chi = p(\theta)$ where $p(x)$ is an irreducible factor of $c(x)$. Hence χ has non-trivial nullspace N.
3. Now compute FG-submodule of M generated by non-zero vector in N. If we obtain proper submodule, G acts **reducibly** on V. Otherwise repeat Steps 2 and 3 for M^T.
4. If $p(x)$ has multiplicity one, then $\theta|_N$ acts with minimal polynomial $p(x)$ on N and $\dim(N) = \deg(p)$. So N is irreducible as $F\langle\theta\rangle$-module. Conclude G acts **irreducibly** on V.
5. Otherwise repeat the random selection.

f-cyclic matrices

$A \in \mathrm{Mat}(d, F)$ is *cyclic*: its characteristic polynomial coincides with its minimal polynomial. The vector space of $1 \times n$ matrices over F is cyclic as an $F\langle A \rangle$-module.

Definition

Let f a monic irreducible polynomial over F. A is f-**cyclic** if f divides the minimal polynomial $m(t)$ of A, but f does not divide $c(t)/m(t)$, where $c(t)$ is characteristic polynomial.

Family of f-cyclic matrices contains all cyclic matrices and also all matrices where f divides $c(t)$ with multiplicity one.

MeatAxe uses last case: proportion at least 0.08. Las Vegas algorithm, complexity $O(d^4 \log q)$.

Neumann & Praeger (1995): determine proportion of cyclic matrices in $\mathrm{Mat}(d, q)$.

Fulman (1999), Wall (1999): generating functions to study the proportion of cyclic matrices in $\mathrm{GL}(d, q)$.

Fulman, Neumann, and Praeger (2005): proportion in classical groups.

Neumann & Praeger (2001), Glasby (2006), Glasby & Praeger (2009): analysis of MeatAxe using such matrices.

C3: Semilinear groups

Maximal subgroups in $C3$ are $\mathrm{GL}(d/e, q^e).e$ where prime $e|d$.

FG-module is *absolutely irreducible* if it remains irreducible under any extension of F. Equivalently $C_{\mathrm{GL}(n,q)}(G)$ just scalars.

G not absolutely irreducible: there is an extension field $E = \mathrm{GF}(q^e)$ of F, where $e|d$, and V is a vector space of dimension d/e over E, with G acting linearly over E.
So $G \cong H \leq \mathrm{GL}(d/e, q^e)$.

Holt & Rees (1994): polynomial-time extension of the MeatAxe to determine centralising field of FG-module.

Semilinear: G acts semilinearly on V regarded as an E-space, where field automorphisms fix F.

So homomorphism $\alpha : G \longmapsto \mathrm{Gal}(E : F)$ where

$$(\lambda v)^g = \lambda^{g\alpha} v^g$$

for all $v \in V$, all $g \in G$, and all $\lambda \in E$.

Image is cyclic group, kernel is absolutely reducible and so conjugate to subgroup of $\mathrm{GL}(d/e, q^e)$.

Lemma

If G is semilinear, then V has a direct sum decomposition as isomorphic irreducible FG'-modules V_i, and G' does not act absolutely irreducibly on the V_i.

Holt *et al.* (1996): apply SMASH to normal generating set for G' to decide if absolutely irreducible group G acts semilinearly.

C2: Imprimitive groups

Maximal subgroups in $C2$ are stabilisers of direct sum decompositions $V = \oplus_{i=1}^{r} V_i$ where $\dim(V_i) = d/r = s$.

Homomorphism $\phi : G \longmapsto \mathrm{Sym}(r)$, its kernel is a normal subgroup of G.

Holt *et al.* (1996): algorithm to decide if absolutely irreducible group G acts imprimitively on V.

MinBlocks: given a non-trivial subspace of a block of imprimitivity, find the block system with minimal block dimension that contains this subspace.

Smash applies when G does not act faithfully on the system of blocks: use element orders and characteristic polynomials of random elements to find non-scalar $g \in G$ which must lie in the kernel of the homomorphism from G to S_r.

Lemma

Let absolutely irreducible G act imprimitively on V and let H be the stabiliser of one such block W. $\mathrm{Hom}_{FH}(W, V)$ has dimension 1 over F.

Proof.

V is isomorphic to the induced module W^G, where W is regarded as an FH-module.

Thus, W must be irreducible as an FH-module, since otherwise V would not be irreducible as an FG-module.

Frobenius Reciprocity: $\mathrm{Hom}_{FG}(W^G, V) = \mathrm{Hom}_{FH}(W, V)$.

Since V is absolutely irreducible FG-module, $\mathrm{Hom}_{FH}(W, V)$ has dimension 1 over F. □

We apply MinBlocks to images of composition factors of appropriate dimension to find W.

So if we can construct the stabiliser H in G of a block W, then we can find W!

Assume G preserves a maximal system of imprimitivity on r blocks of size s, so action is primitive and H must be a maximal subgroup of G of index r.

We construct H by working up a chain of subgroups, starting with a cyclic subgroup and then adjoining new generators.

Polynomial-time? Difficulty of analysis of SMASH imprimitive case.

C5. Smaller field modulo scalars

Maximal subgroups of $\mathrm{GL}(d, K)$: conjugates of $\mathrm{GL}(d, F).Z$ where $F < K$ and Z is centre of $\mathrm{GL}(d, K)$.

$G = \langle X \rangle$ absolutely irreducible subgroup of $\mathrm{GL}(d, K)$, and $F < K$.

Glasby, Leedham-Green & O'B (2005): algorithm to decide if $G \cong H \leq \mathrm{GL}(\mathrm{d}, \mathrm{F}).\mathrm{Z}$.

Consider *base case*: is $G \cong H \leq \mathrm{GL}(d, F)$?

$F[G]$ denotes the set of F-linear combinations of the elements of G as an F-subalgebra of $\mathrm{M}(\mathrm{d}, \mathrm{K})$.

Lemma

If G can be written over the smaller field F, then so can the F-algebra $F[G]$.

1. Repeatedly select random $a \in \mathrm{M}(\mathrm{d}, \mathrm{F})$ until either char poly $c_a(t)$ does not lie in $F[t]$, or until $c_a(t) \in F[t]$ and a has an eigenvalue $\lambda \in F$ with multiplicity 1. In the former case return *false*.

2. Find a non-zero λ-eigenvector v for a.

3. Construct sufficient images of v under the action of G to obtain a basis B of V.

4. Write generators of G with respect to the basis B, and return *false* if one does not lie in $\mathrm{M}(d, F)$. Otherwise return matrix with rows B.

Theorem

Let G be absolutely irreducible subgroup of $\mathrm{GL}(d, K)$, and let $F < K$. There is a subfield L of K containing F such that $F[G]$ is conjugate in $\mathrm{GL}(d, K)$ to the full matrix algebra $\mathrm{M}(d, L)$. Hence G can be written over L, but not over any proper subfield of L containing F.

Theorem

Let F be a proper subfield of a finite field L, and let a be a uniformly random element of $\mathrm{M}(\mathrm{d}, \mathrm{L})$. The probability, π, that $c(t) := \det(tI - a)$ does not lie in $F[t]$ satisfies $\pi > \frac{2}{3}(1 - (|F|/|L|)^d) \geq 1/2$.

Theorem

There is a constructive polynomial-time Las Vegas algorithm that takes as input F, and an absolutely irreducible group $G := \langle X \rangle \leq \mathrm{GL}(\mathrm{d}, \mathrm{K})$, and decides whether or not G is conjugate to a subgroup of $\mathrm{GL}(\mathrm{d}, \mathrm{F})$.

Smaller field modulo scalars: General case

Lemma

If G can be written over F modulo scalars in K and G' is absolutely irreducible, then G' can be written over F and the F-space spanned by such a basis for G' is unique up to multiplication by a scalar in $K^\times$.

Proof.

Multiplying each of $g, h \in G$ by a fixed scalar does not change the value of $[g, h]$. Uniqueness follows by applying Schur's Lemma to V as an absolutely irreducible KG'-module. □

G' acts absolutely irreducibly on V: apply the base case algorithm to G'.

Clifford theory: if G is primitive, tensor-indecomposable, and not semilinear, then G' satisfies this condition.

Alternative: analyse structure of KG-module

Carlson, Neunhöffer, Roney-Dougal (2009): polynomial-time Las Vegas algorithm to find a non-trivial reduction of irreducible groups that either:

- lie in C5;
- are semilinear;
- or have non-absolutely irreducible derived group.

A constructive version of Aschbacher?

- Reducible. Polynomial time? Yes.
- Imprimitive? No.
- Semilinear? Yes in certain cases.
- Tensor product? No.
- Defined mod scalars over subfield? Yes in certain cases.
- Normaliser of p-group? Yes in certain cases.
- Tensor induced? No.
- Classical group in natural representation? Yes.
- Almost simple modulo scalars?

Practical algorithms to decide membership available in MAGMA.

Work of many authors including: Brooksbank; Carlson, Neunhöffer and Roney-Dougal; Glasby; Leedham-Green and O'B; Niemeyer and Praeger.

Base cases for Aschbacher strategy

Classical group in natural representation or other **almost simple modulo scalars**.

Liebeck (1985): almost all maximal non-classical subgroups of $\mathrm{GL}(d, q)$ have order at most q^{3d}.

Landazuri & Seitz (1974), Seitz & Zalesskii (1993): lower bounds for degrees of nonlinear irreducible projective representations of finite Chevalley groups. Faithful projective representations in cross characteristic have degree that is polynomial in the size of the defining characteristic.

Principal focus: *matrix representations in defining characteristic*.

Hiss & Malle (2001), Lübeck (2001): absolutely irreducible representations of degree ≤ 250 of quasisimple groups.

Can we name the group?

A prime r dividing $b^e - 1$ is a *primitive prime divisor* of $b^e - 1$ if r does not divide $b^i - 1$ for $1 \leq i < e$.

Zsigmondy (1892): $b^e - 1$ has ppd unless $(b, e) = (2, 6)$ or $e = 2, b = 2^n - 1$.

$$|\mathrm{GL}(d, q)| = q^{\binom{d}{2}} \prod_{i=1}^{d} (q^i - 1)$$

Hence ppds of $q^e - 1$ for various values of $e \leq d$ divide $|\mathrm{GL}(d, q)|$ and also orders of the various classical groups.

ppd-element: order a multiple of some ppd

Naming groups: Classical groups

Problem

Given $G = \langle X \rangle \leq \mathrm{GL}(d, q)$, *does* G *contain* $\mathrm{SX}(d, q)$*?*

Praeger & Neumann (1992), P & Niemeyer (1998): Monte Carlo polynomial-time algorithms to name classical group in natural repn.

Search for certain kinds of ppd-elements that occur with high probability in $\mathrm{SX}(d, q)$ and are in only a "small" number of other subgroups of $\mathrm{GL}(d, q)$.

Original motivation: Joachim Neubüser (1988) asked for analogue of algorithm to decide if $G \leq S_n$ contains A_n.

Naming groups of Lie type: Black box

Theorem (Babai, Kantor, Palfy, Seress, 2002)

Given a group G isomorphic to a simple group of Lie type of known characteristic, its standard name can be computed using a polynomial time Monte-Carlo algorithm.

Choose sample $\mathcal{L}$ of independent (nearly) uniformly distributed random elements of G.

Find the three largest integers $v_1 > v_2 > v_3$ such that a member of $\mathcal{L}$ has order divisible by a primitive prime divisor of one of $p^{v_i} - 1$.

Usually $\{v_1, v_2, v_3\}$ determines $|G|$ and name of G.

Altseimer & Borovik (2002): distinguish between $\mathrm{PSp}(2m, q)$ and $\Omega(2m+1, q)$, q odd and $m \geq 3$.

Finding the characteristic

BKPS and other algorithms assume that input G is a simple group of Lie type of known characteristic.

Problem

Given $G \leq \mathrm{GL}(d, q)$ where G is a group of Lie type in **unknown** *defining characteristic r. Can we determine r?*

Liebeck & O'B (2007):
Monte Carlo algorithm which proceeds recursively through centralisers of involutions to find $\mathrm{SL}(2, F_r)$. Now read off r.

Kantor & Seress (2009):
The three largest element orders determine the characteristic of Lie-type simple groups of odd characteristic.

Result: extremely powerful Monte Carlo algorithms to name group.

Constructive recognition

$C = \langle X \rangle \leq \mathrm{GL}(d, q)$ where C is (quasi)simple. e.g. $\mathrm{SL}(d, q)$, invertible matrices of determinant 1.

C is standard copy, sometimes known as "gold copy".

$G = \langle Y \rangle \cong C$.

Want to construct "effective" isomorphisms
$\phi : C \longmapsto G$ and $\tau : G \longmapsto C$.

Key idea: use standard generators.

Using standard generators

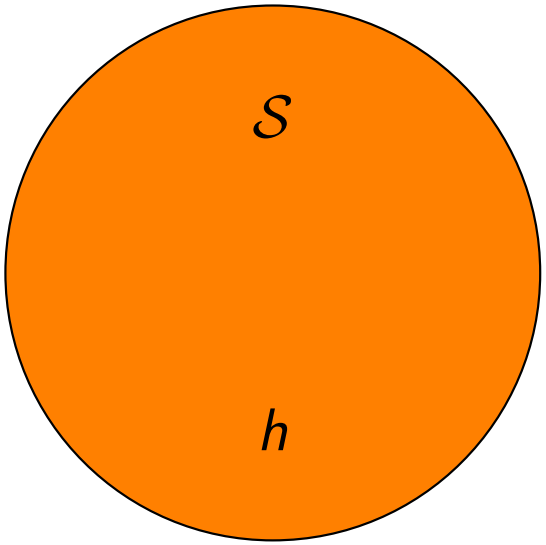

$C = \langle X \rangle$

Find $\mathcal{S} = w(X)$

Find $\bar{\mathcal{S}} = w(Y)$

Define $\phi : C \mapsto G : \mathcal{S} \mapsto \bar{\mathcal{S}}$

$h = w(\mathcal{S})$

Thus $\bar{h} = w(\bar{\mathcal{S}})$

$\langle Y \rangle = G$

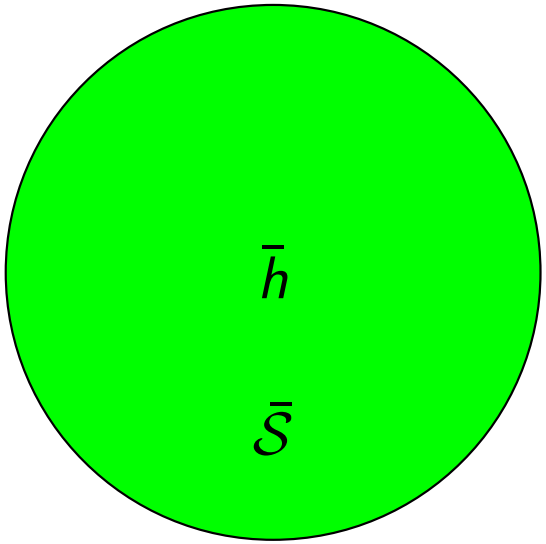

Motivation

Example

$C = \langle X \rangle = \mathrm{SL}(d, q)$
$G = \langle Y \rangle$ is symmetric square repn.

C is our "gold" copy in which we know information.

Examples include

- Conjugacy classes of elements.
- Maximal subgroups.

We know or can obtain these readily as words w in S.

If we know $\bar{S} \subset G$, we can evaluate w in $\bar{S}$.

So we now know this information in our arbitrary copy G.

Application I: Conjugacy classes of classical groups

Example: $C = \langle X \rangle = \mathrm{SX}(d, q)$
$G = \langle Y \rangle$ is symmetric cube.

Wall (1963): description of conjugacy classes and centralisers of elements of classical groups.

Liebeck, O'Brien (ongoing): algorithm, which given d and q, constructs classes for $\mathrm{SX}(d, q)$.

$\phi : C \longmapsto G$ now maps class reps and centralisers to G.

Example

Higman's (1961) count of p-groups of p-class 2.
Eick and O'B (1999): algorithm which, given d and p, counts precisely the number of d-generator p-groups of class 2.
Critical task: for each conjugacy class rep r in $G := \Lambda^2(\mathrm{GL(d, p)})$ use Cauchy-Frobenius theorem to count fixed points for r.

Application II: Maximal subgroups of classical groups

Kleidmann & Liebeck (1990): describe some maximal subgroups of classical groups where $d \geq 13$.

Bray, Holt & Roney-Dougal (2013): generating sets for geometric maximal subgroups, and all maximals for $d \leq 12$.

So obtain $M \leq C := \mathrm{SX}(d, q)$, classical group in natural representation.

Use $\phi : C \longmapsto G$ to construct image of M in arbitrary representation G.

Main tasks

- Define *standard generators* $\mathcal{S}$ for $C = \langle X \rangle$.
- Need algorithms to:
  - Construct $\mathcal{S}$ as *words* in X.
  - For $h \in C$, express h as $w(\mathcal{S})$ and so as $w(X)$.
- If $\langle Y \rangle = G \simeq C$ then:
  - Find standard generators $\bar{\mathcal{S}}$ in G as words in Y.
  - For $g \in G$, express g as $w(\bar{\mathcal{S}})$ and so as $w(Y)$.

Choose $\mathcal{S}$ so that solving for word in $\mathcal{S}$ is easy.

Now define isomorphism $\phi : C \longmapsto G$ from $\mathcal{S}$ to $\bar{\mathcal{S}}$

Effective: if $h = w(\mathcal{S})$ then $\phi(h) = w(\bar{\mathcal{S}})$.

Similarly $\tau : G \longmapsto C$.

Standard generators for $\mathrm{SL}(d, q)$

Leedham-Green & O'B (2008).
Natural module V for $C = \mathrm{SL}(d, q)$ with basis $\{e_1, \ldots, e_d\}$.
Define standard generators s, δ, u, v for C:

s, δ, u lie in copy of $\mathrm{SL}(2, q)$ and act on $\langle e_1, e_2 \rangle$ as:

$$s = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \delta = \begin{pmatrix} \omega & 0 \\ 0 & \omega^{-1} \end{pmatrix} \quad u = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

v maps

$$e_1 \mapsto e_d \mapsto -e_{d-1} \mapsto -e_{d-2} \mapsto -e_{d-3} \cdots \mapsto -e_1$$

Given $h \in C$, via echelonisation write $h = w(\mathcal{S})$.

Algorithm to construct standard generators

- Construct two subgroups H and K in G so

$$H = \begin{pmatrix} \mathrm{SX}_m & \\ & 1_{d-m} \end{pmatrix} \quad \text{and} \quad K = \begin{pmatrix} 1_m & \\ & \mathrm{SX}_{d-m} \end{pmatrix}$$

- Recursively construct standard generators $\mathcal{S}_H$ and $\mathcal{S}_K$ for H and K
- all but cycle from standard generators for G contained in $\mathcal{S}_H$
- cycle is constructed by glueing two cycles from $\mathcal{S}_H$ and $\mathcal{S}_K$. e.g. if $G = \mathrm{SL}(d, q)$ with even d and q, then

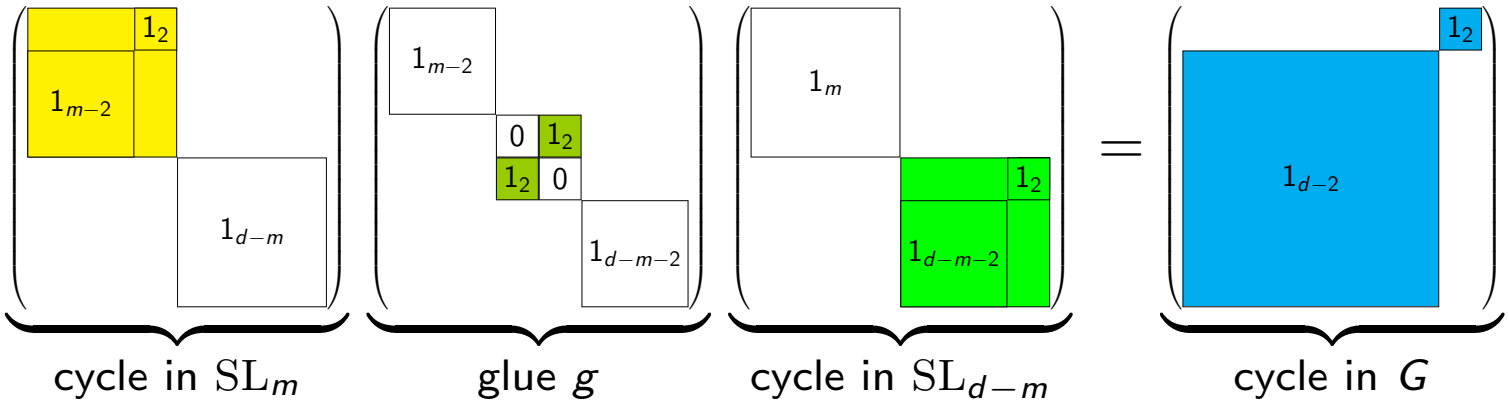

Leedham-Green and O'B, 2009; Dietrich, L-G, Lübeck, O'B, 2013; D, L-G, O'B, 2014

Theorem

There is a polynomial time Las Vegas algorithm that takes as input $G \cong \mathrm{SX}(d, q) = \langle X \rangle$ and returns standard generators $\mathcal{S}$ for G as words in X.

Effective complexity: $O(d^4 \log q)$

Theorem (Liebeck & O'B; TAMS, 2014)

Similar statement for exceptional groups.

Odd characteristic

Key: *centralisers of involutions* and statistical group theory.

$G = \mathrm{SX}(d, q)$.

t is involution in G, with eigenspaces E_+ and E_-

$C_G(t)$ is $(\mathrm{GL}(E_+) \times \mathrm{GL}(E_-)) \cap \mathrm{SL}(d, q)$.

A *strong involution* in $\mathrm{SX}(d, q)$ has -1-eigenspace of dimension in range $(d/3, 2d/3]$.

$G = \mathrm{SX}(d, q)$ for q odd

1. Find and construct strong involution t having -1-eigenspace of dimension m.
2. Now construct $C_G(t)$. Construct the direct summands of the derived group to obtain $\mathrm{SX}(m, q)$ and $\mathrm{SX}(d - m, q)$ as *subgroups* of G.
3. Recursively construct standard generators for $\mathrm{SX}(m, q)$ and $\mathrm{SX}(d - m, q)$.
4. Construct centraliser C of involution

$$\begin{pmatrix} I_{m-2} & 0 & 0 \\ 0 & -I_4 & 0 \\ 0 & 0 & I_{d-m-2} \end{pmatrix}$$

5. Within C solve constructively for matrix g

$$\begin{pmatrix} I_{m-2} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & I_{d-m-2} \end{pmatrix}$$

6. Now m-cycle v_m and $(d-m)$-cycle v_{d-m} "glued" together by g to produce d-cycle $v_m g v_{d-m}$.

Cost of finding a strong involution

First step: search for an element of $\mathrm{SX}(d, q)$ of even order that has as a power a strong involution.

Theorem (Lübeck, Niemeyer, Praeger, 2009)

For an absolute constant c, the proportion of $g \in \mathrm{SX}(d, q)$ such that a power of g is a strong involution is $\geq c/\log d$.

Recursion to smaller cases requires additional results.

Theorem (Leedham-Green & O'B, 2009)

For some absolute constant c, the proportion of $g \in \mathrm{SX}(d, q)$ such that a power of g is a "suitable" involution is $\geq c/d$.

Constructing centralisers

Bray (2001): Monte Carlo algorithm to construct $C_G(t)$ for involution $t \in G$.

Algorithm exploits properties of dihedral group.

Construct random conjugate t^g of t.

1. If $[t, g]$ has odd order $2m + 1$, then $g[t, g]^m$ commutes with t.
2. If $[t, g]$ has even order $2m$, both $[t, g]^m$ and $[t, g^{-1}]^m$ commute with t.

So convert random elements of G into elements of $C_G(t)$.

Elements not, in general, uniformly-distributed, but:

Lemma

If g is uniformly distributed among the elements of G for which $[t, g]$ has odd order, say $2n + 1$, then $g[t, g]^n$ is uniformly distributed among the elements of $C_G(t)$.

If odd order case occurs *sufficiently often*, we can construct nearly-uniformly distributed random elements of $C_G(t)$ in polynomial time.

Proportion of odd order elements

Theorem (Parker & Wilson, 2009; Liebeck, 2015)

Let G be a simple group of Lie type, of Lie rank r, defined over field of odd characteristic. The probability that $[t, g]$ has odd order, where t is a fixed involution and g is a random element of G, is at least c/r for some absolute constant c.

Example: lower bound for $\mathrm{PSL}_d(q)$ is $\frac{1}{12d}$.

Method: for each class of involutions, find a dihedral group generated by two involutions of this class, and show that a significant proportion of pairs of involutions in this class generate such a dihedral group.

Cost of construction of centraliser

Bray (2001)
Parker & Wilson (2010)
Holmes, Linton, O'B, Ryba, Wilson (2008)
Let μ, ξ and ρ denote the costs of a group operation, constructing a random element of G, and an order oracle respectively.

Theorem

Let H be a simple group of Lie rank r defined over a field of odd characteristic. The centraliser in H of an involution can be computed in time $O(r(\xi + \rho)\log(1/\epsilon) + \mu r^2)$ with probability of success at least $1 - \epsilon$, for $\epsilon > 0$.

This is a black-box Monte Carlo algorithm. Similar statement for even char.

Even characteristic: Problems

- Involutions cannot be found efficiently by a random search Guralnick & Lübeck (2001): proportion of elements in G of even order is $< 5/q$;
- Groups for a recursion cannot be found in centraliser; Aschbacher & Seitz (1976): various types of involutions.

Theorem (Aschbacher & Seitz)

If $g \in G$ is a good involution, then, mod base change,

$$C_G(g) = \begin{pmatrix} \mathrm{GL_r} & * & * \\ & \mathrm{GL_{d-m}} & * \\ & & \mathrm{GL_r} \end{pmatrix} \cap G \quad \text{or} \quad C_G(g) = \begin{pmatrix} Sp_r & * & * \\ & \mathrm{SX}_{d-m} & * \\ & & Sp_r \end{pmatrix}$$

where $r = \mathrm{rank}(g-1)$, $m = 2r$, and SX_{d-m} same type as G.

Finding involutions is specific instance of:

Problem

Find element of order p in $G = \langle X \rangle$, a group of Lie type in characteristic p, as a word in X.

$\rho(G)$ is proportion of p-singular elements in G.
Kantor, Isaacs, Spaltenstein (1995); Guralnick & Lübeck (2003)

Theorem

$\frac{2}{5q} < \rho(G) < \frac{5}{q}$ *where G is a group of Lie type defined over* $\mathrm{GF}(q)$.

So random search requires $O(q)$ random selections.

The key case

$\mathrm{SL}(2, q) \simeq \langle X \rangle$.

Critical task: find transvection as word in X.

Proportion is $O(1/q)$, can't search randomly.

Equivalent task: constructive recognition of $\mathrm{SL}(2, q)$.

Constructive recognition for $\mathrm{SL}(2, q)$

Landazuri & Seitz (1974), Seitz & Zalesskii (1993): faithful projective representations in cross characteristic have degree that is **polynomial** in q, so critical focus is defining characteristic representation.

Let $\tau(d)$ denote the number of factors of d.

Theorem (Conder, Leedham-Green, O'B, 2006)

$G \leq \mathrm{GL}(d, F)$ for $d \geq 2$, where F has same characteristic as $\mathrm{GF}(q)$. Assume that G is isomorphic modulo scalars to $\mathrm{PSL}(2, q)$. Subject to a fixed number of calls to a Discrete Log Oracle, there exists a Las Vegas algorithm that constructs an epimorphism from G to $\mathrm{PSL}(2, q)$ at a cost of at most $O(d^5\tau(d))$ field operations.

Theorem (Brauer & Nesbitt, 1940)

Let F be an algebraically closed field of characteristic p, and let V be an irreducible $F[G]$-module for $G = \mathrm{SL}(2, q)$, where $q = p^e$. Then $V \simeq T_1 \otimes T_2 \otimes \cdots \otimes T_t \otimes_{\mathrm{GF}(q)} F$, where T_i is the s_i-fold symmetric power S_i of the natural $\mathrm{GF}(q)[G]$-module M twisted by the f_i th power of the Frobenius map, with $0 \leq f_1 < f_2 < \cdots < f_t < e$, and $1 \leq s_i < p$ for all i.

G absolutely irreducible representation of $\mathrm{SL}(2, q)$.
Three components to constructive recognition algorithm for G.

1. Decompose tensor product to obtain one symmetric power T_i.
2. Decompose T_i to obtain SL(2, q) in its natural representation.
3. Construct standard generators for SL(2, q).

Standard generators for H := SL(2, q) in natural repn

1. Find $A \in H$ of order $q - 1$ and B a random conjugate of A.
2. Compute eigenvectors u and v of A, with corresponding eigenvalues a and a^{-1}.
3. Find a random element C of H and an i such that $B^i C$ fixes $\langle u \rangle$, if such an i exists. If A and $B^i C$ lie in $\mathrm{SL}(2, q)$ and have common eigenvector u, then $S = [A, B^i C]$ is a transvection fixing u.
4. Similarly, find a random element D of H and a j such that $B^j D$ fixes $\langle v \rangle$ and $T = [A, B^j D]$ is not trivial. Now, T is a non-trivial transvection fixing v.
5. Write S, T, A with respect to the ordered basis (u, v) to obtain generating set for $\mathrm{SL}(2, q)$.

Step 3 is critical: Find a random element C of H and an i such that B^iC fixes $\langle u \rangle$, with corresponding eigenvalue a.

B^iC fixes $\langle u \rangle$ if and only if $a^{2i} = \mu$ where $\mu \in \mathrm{GF}(q)$.

Its solution relies on *discrete log*.

Easy to find elements of order $q-1$:
proportion is $\phi(q-1)/2(q-1) > 1/2 \log\log q$.

Now given $x \in \mathrm{SL}(2,q)$, use echelonisation to write x as word in S, T, A.

Other recognition algorithms for $\mathrm{SL}_2(q)$

Kantor & Kassabov (2015); Borovik & Yalcinkaya (2015)

isomorphism between a black-box copy of $\mathrm{SL}_2(q)$ and the natural copy in time that is quadratic in the characteristic of $\mathrm{GF}(q)$.

Char 2: polynomial-time, no reliance on discrete log, implemented, practical.

Even characteristic – The general approach

- Find $H = \mathrm{SX}(m, q) \leq G$ where $m \in [d/3, 2d/3]$ is even or $4|m$; if G is linear or unitary, then so is H, otherwise Ω^+;

(via base change) $H =$ 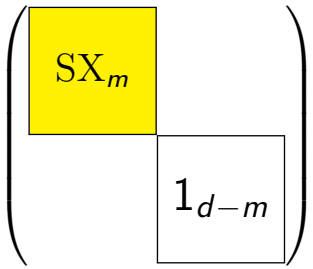

and $K =$ 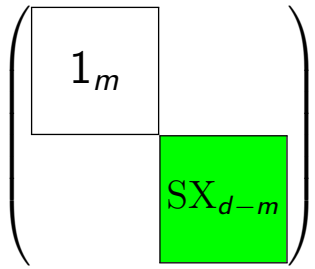

- Recursion: construct standard generators of SX_m in H and a *good* involution $g \in H$ with $r = \mathrm{rank}(g - 1) = m/2$
- in $C_G(g)$ find $K = \mathrm{SX}(d - m, q) \leq G$
- Recursion: construct standard generators of SX_{d-m} in K
- glue the cycles of SX_m and SX_{d-m}

That we can construct H and K is a consequence of the following.

Theorem

There is a black-box Las Vegas algorithm which takes as input $G \cong \mathrm{SX}_d(q)$, which is not a base case, and constructs $H \leq G$ with $H \cong \mathrm{SX}_m(q)$, admitting $K \leq C_G(H)$ with $K \cong \mathrm{SX}_{d-m}(q)$; in general, $m \in [d/3, 2d/3]$ is even. If G is linear or unitary, then so is H. In all other cases, H is of type Ω^+ and m is divisible by 4. If G has type Ω^-, then $m \in \{d-4, d-6\}$ is divisible by 4. The time required is $\mathrm{O}(\mathrm{d}(\xi+\rho)+\mu)$.

Correctness established by Praeger et al. (2015); complexity follows from Dietrich et al. (2015).

Other black-box algorithms for classical groups

Theorem (Kantor & Seress, 2001)

There is a Las Vegas algorithm which when given a perfect group $G = \langle X \rangle \leq GL(V)$ where $G/Z(G)$ is isomorphic to a classical simple group of known characteristic produces a constructive isomorphism $G/Z \longmapsto C$.

Algorithm not polynomial in size of input: factor of q.

Brooksbank & Kantor (2001): algorithms can be made polynomial in $\log q$ given an *oracle* for constructive membership testing in $\langle X \rangle \cong \mathrm{SL}(2, q)$.

B & K (2001-2008): Black-box algorithms for the classical families which run in polynomial time subject to existence of $\mathrm{SL}(2, q)$ *oracle*.

Base cases for recursion in our algorithms

$\mathrm{SL}(d, q)$ where $d = 2, 3$.

Conder, Leedham-Green, O'B (2006): $\mathrm{SL}_2(q)$.

Lübeck, Magaard and O'B (2008): $\mathrm{SL}_3(q)$.

$\mathrm{SU}(d, q)$ where $d = 3, 4$; $Sp(4, q)$; $\Omega^{\pm}(d, q)$ for $d \leq 8$.

Brooksbank, B. & Kantor (2001-2008): natural, black-box

Practical versions/mixtures developed and implemented by Clarkson (2014).

Constructive recognition for A_n

- A_n: Bratus & Pak (2000), Holt; Beals et al. (2001-05). Jambor, Leuner, Niemeyer & Plesken (2013). Black-box.

Exceptional groups

- Kantor & Magaard (2013): black-box algorithms with complexity $O(q)$.
- Liebeck & O'Brien (2014): Standard generators: those which satisfy reduced Curtis-Steinberg-Tits presentations.

  Polynomial time Las Vegas subject to oracles.
- Bäärnhielm (2006-2014): Algorithms for matrix representations of Suzuki, large and small Ree groups.
- Bäärnhielm and Bray (2010): black-box for Suzuki.

Last three available in MAGMA.

Constructive recognition for the sporadics

Wilson (1996): standard generators for sporadic $G = \langle Y \rangle$

Bray & Wilson: black-box algorithms to find these (as words) in Y.

Two methods to solve constructive membership problem for G.

- Random Schreier works well for many – with careful choice of base points (O'B & Wilson, 2002).
- REDUCTION algorithm of Holmes et al. (2008): reduces constructive membership problem in G to three instances of the same problem for involution centralisers in G.

Writing elements as SLPs in classical groups

Elliot Costi (2009): defining char repns

Csaba Schneider and Praeger: black box repns

Theorem (Costi, 2009; Praeger & Schneider, 2014)

$G \simeq \mathrm{SX}(d, q)$: algorithms to write element of G as word in $\mathcal{S}$.

- $G = \mathrm{SX}(d, q)$: algorithms to write element of G as SLP in our standard generators.

  Complexity: $O(d^3 \log q)$

- G is defining char (projective) irreducible representation of $\mathrm{SX}(d, q)$.

  Complexity: $O(d^4 n^3 \log^3 q + d^2 n^4 \log q)$.

Complete implementation available in Magma.

⏩ Skip details

Polynomial-time algorithms for unipotent groups

- Stabiliser of subspace algorithm.

  Input: unipotent $K \leq \mathrm{GL}(\mathrm{d}, \mathrm{q})$ and $U \leq V$.

  Output: a canonical element $\overline{U}$ of the orbit of U under K; and $k \in K$ such that $U^k = \overline{U}$, and generators for the stabiliser of U in K.
- Constructively decide membership in unipotent group.

Defining characteristic representations

$H = \mathrm{SL}(d, q)$, $G \leq \mathrm{GL}(n, F)$ is (projective) irreducible representation in defining char acting on V, $\phi : H \longmapsto G$.

Let K be maximal parabolic subgroup of $\mathrm{SL}(d, q)$ that fixes the space spanned by first basis element.

$$\begin{pmatrix} \det^{-1} & 0 & 0 & 0 \\ \star & & & \\ \vdots & & \mathrm{GL(d-1,q)} & \\ \star & & & \end{pmatrix}$$

Since $K\phi$ is p-local, it stabilises a proper $K\phi$-submodule U of V.

Consider elementary abelian $E \leq H$ generated by

$$\begin{pmatrix} 1 & \star & \dots & \star \\ 0 & & & \\ \vdots & & I_{d-1} & \\ 0 & & & \end{pmatrix}$$

Critical reduction

Construct $x \in E\phi$ as an SLP that maps $W := U^g$ to U.

Hence $U^{gx} = U$ and so preimage of gx is in K.

So we have "killed" the first row of the preimage of gx.

Dualise to kill first column, obtaining $g_1 := \begin{pmatrix} \alpha & 0 \\ 0 & A \end{pmatrix}$

$t\phi := g_1^{-1} \cdot T_{1,j}^{\phi} \cdot g_1 \in E\phi$ where $T_{1,j}$ is transvection with non-zero entry in $(1, j)$ position.

Use **membership test** for $t\phi$ in $E\phi$ to obtain preimage $t \in E$.

Read off from t (scalar multiple of) j-th row of preimage in $\mathrm{SL}(d, q)$ of g_1.

So reduce problem to **natural representation** in rank $d - 1$.

Geometry following Aschbacher: general strategy

$G = \langle X \rangle \leq \mathrm{GL}(d, q)$.

1. Determine (at least one of) its Aschbacher categories.
2. If $N \lhd G$ exists, process N and G/N recursively.
3. Otherwise G is either classical group in natural representation or $T \leq G/Z \leq Aut(T)$ where T is simple.
   - "Reduce" from G to quasisimple group L.
   - Name L.
   - Set up "effective" isomorphisms between L and its standard copy S.

The C9 case

$L \leq G/Z \leq \mathrm{Aut}(L)$ so $G \simeq Z.L.E$.

- Use determinant map to ensure that $|Z|$ is a divisor of $\gcd(d, q-1)$.
- Calculate the stable derivative $D = G^{(\infty)}$ of G.
- Construct $\phi : G \longmapsto E$ by letting G act on cosets of $H = \langle Z, D \rangle$.

  $Hx = Hy \iff xy^{-1} \in H$

  Use "order of element modulo normal subgroup" algorithm to determine to decide membership in H.

The composition tree for G

Bäärnhielm, Leedham-Green & O'B
Neunhöffer & Seress

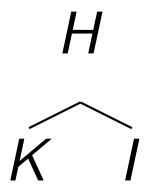

- Node: section H of G.
- Image I: image under homomorphism or isomorphism. Images usually correspond to Aschbacher category, but also others e.g determinant map.
- Kernel K.
- Leaf is "composition factor" of G: simple modulo scalars. Cyclic not necessarily of prime order.

Tree is constructed in right depth-first order.

If node H is not a leaf, construct recursively subtree rooted at I, then subtree rooted at K.

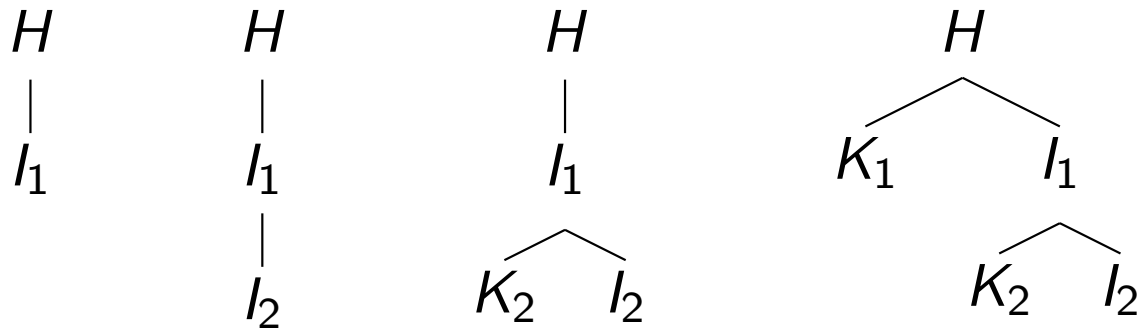

Constructing kernels

Assume $\phi : H \longmapsto I$ where $K = \ker \phi$.

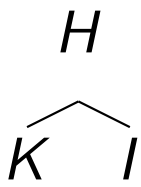

Sometime easy to obtain theoretically generating sets for $\ker \phi$.

Two approaches to construct kernel.

1. Random generation of the kernel

Assume $\phi : H \longmapsto I$ where $K = \ker \phi$.

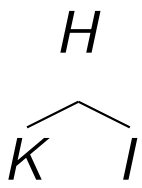

Let $x_1, \ldots, x_t$ be generating set for H.

Let $y_j = \phi(x_j)$ so $I = \langle y_1, \ldots, y_t \rangle$.

Let $h \in H$ and let $i = \phi(h)$.

Write $i = w(y_1, \ldots, y_t)$.

Let $\bar{h} = w(x_1, \ldots, x_t)$.

Now $k = h\bar{h}^{-1} \in K := \ker \phi$.

Choose random $h \in H$ to obtain random generator k of K.

Randomised algorithm to construct the kernel – but assumes that we can write $i = w(y_1, \ldots, y_t)$.

2. Normal generators for kernel

Construct normal generating set for K, by evaluating relators in presentation for I and take normal closure.

To obtain presentation for node: need only presentation for associated kernel and image.

So inductively need to know presentations **only for the leaves** – or composition factors.

Short presentations for finite groups

Babai and Szemerédi (1984): *length* of a presentation $P = \{X \mid R\}$ is number of symbols to write down the presentation.

Each generator is single symbol, relator is a string of symbols, exponents written in binary.

Example

S_n generated by $t_k = (k, k+1)$ for $1 \leq k < n$ with relations:

- $t_k^2 = 1$ for $1 \leq k < n$,
- $(t_{k-1} t_k)^3 = 1$ for $1 < k < n$,
- $(t_j t_k)^2 = 1$ for $1 \leq j < k-1 < n-1$.

Number of relations is $n(n-1)/2$, and presentation length is $O(n^2)$.

S_n acts on deleted permutation module: cost of evaluation of relations is $O(n^5)$.
Goal: short presentations on bounded number of relations.

Theorem (Guralnick, Kantor, Kassabov, Lubotzky, 2008)

Every non-abelian finite simple group of rank n over $\mathrm{GF}(q)$, with possible exception of Ree groups ${}^2G_2(q)$, has a presentation with a bounded number of generators and relations and total length $O(\log n + \log q)$.

Exploits results of:

- Campbell, Robertson and Williams (1990): $\mathrm{PSL}(2, p^n)$ has presentation on (at most) 3 generators and a bounded number of relations.
- Hulpke and Seress (2003): $\mathrm{PSU}(3, q)$

Previous best: Babai *et al.* (1997) presentation of length $O(\log^2 |G|)$. Modifications of Curtis-Steinberg-Tits presentations for groups of Lie rank at least 2.

Constructive version:

L-G and O'B (2013): explicit short presentations for the classical groups on our standard generators.

Liebeck and O'B (2014): explicit reduced Curtis-Steinberg-Tits presentations for exceptional groups.

Short presentations for S_n and A_n

Theorem (GKKL, 2006; Bray-Conder-LG-O'B, 2006)

A_n and S_n have presentations with a bounded number of generators and relations, and length $O(\log n)$.

Theorem (Bray-Conder-LG-O'B, 2006)

Let p be an odd prime, and let λ be a primitive element of $\mathrm{GF}(p)$, with inverse μ. Then

$$\{\, a, c, t \mid a^p, acacac^{-1}, (a^{(p+1)/2}ca^4c)^2, t^2, [t,a],$$

$$[t, ca^{\lambda}ca^{\mu}c], [t,c]^3, (tt^c tt^{ca})^2, (tt^c tt^{ca^{\lambda}})^2, (at^c)^{p+1} \,\}$$

is a 3-generator 10-relator presentation of length $O(\log p)$ for S_{p+2}, in which att^c stands for a $(p+2)$-cycle and t stands for a transposition.

Previous best results: length $O(n \log n)$ (Moore, 1897)

Theorem (GKKL, 2008)

A_n has presentation on 3 generators, 4 relations, length $O(\log n)$.

S_n: presentation of length $O(n^2)$ on $(1, 2)$ and $(1, 2, \ldots, n)$ and 78 relations.

Output of CompositionTree

Given $G = \langle X \rangle \leq \mathrm{GL}(d, q)$ as input.
Output:

- a composition series: $1 = G_0 \lhd G_1 \lhd G_2 \cdots \lhd G_m = G$.
- A representation $S_k = \langle X_k \rangle$ of G_k / G_{k-1}
- Effective maps $\tau_k : G_k \to S_k$, $\phi_k : S_k \to G_k$
  τ_k epimorphism with kernel G_{k-1}
- Map to write $g \in G$ as word in X.

Construct presentation for group defined by tree and verify that G satisfies the relations.

Hence construction of tree is Las Vegas algorithm.

Characteristic structure

Finite G has characteristic series $\mathcal{C}$ of subgroups:

$$1 \leq O_\infty(G) \leq S^*(G) \leq P(G) \leq G$$

$O_\infty(G) =$ largest soluble normal subgroup of G, soluble radical

$S^*(G)/O_\infty(G) =$ Socle $(G/O_\infty(G)) = T_1 \times \ldots \times T_k$ where T_i non-abelian simple

$\phi : G \longmapsto \mathrm{Sym}(k)$ is repn of G induced by conjugation on $\{T_1, \ldots, T_k\}$ and $P(G) = \ker \phi$

$P(G)/S^*(G) \leq \mathrm{Out}(T_1) \times \ldots \times \mathrm{Out}(T_k)$ and so is soluble

$G/P(G) \leq \mathrm{Sym}(k)$

Exploiting the characteristic series $\mathcal{C}$

Cannon, Holt et al. (2000s): use $\mathcal{C}$ in practical algorithms.

$$1 \leq L := O_\infty(G) \leq S^*(G) \leq P(G) \leq G$$

Also compute series

$$1 = N_0 \lhd N_1 \lhd \cdots \lhd N_r = L \lhd G$$

where $N_i \trianglelefteq G$ and N_i/N_{i-1} is elementary abelian.

The Soluble Radical model

$$1 = N_0 \lhd N_1 \lhd \cdots \lhd N_r = L \leq S^*(G) \leq P(G) \leq G$$

where $N_i \trianglelefteq G$ and N_i/N_{i-1} is elementary abelian.

Given a **problem**:

Solve problem first in $G/L = G/N_r$, and then, successively, solve it in G/N_i, for $i = r-1, \ldots, 0$.

$H := G/L$ has trivial Fitting subgroup.
So H has a socle S which is direct product of non-abelian simple groups T_i and these are permuted under conjugation by H.

Problem may have nice solution for H.

In many cases, easy to reduce the computation for TF-group H to almost simple groups.

Almost simple groups: Conjugacy classes

Wall (1963): description of conjugacy classes and centralisers of elements of classical groups.

Liebeck & O'B (ongoing): algorithms, which given d and q, constructs classes for $\mathrm{SX}(d,q) \leq K \leq \mathrm{CX}(d,q)$.

Embed TF-group $H = G/L$ in direct product W of $\mathrm{Aut}(T_i) \wr \mathrm{Sym}(d_i)$, where T_i occurs d_i times as socle factor.

Conjugacy class representatives in wreath products described theoretically (Hulpke 2004; Cannon & Holt, 2006).

Example: Automorphism group of G

Cannon & Holt, 2003

$H := G/L$ permutes the direct factors of its socle S by conjugation.

Embed H in direct product D of $\mathrm{Aut}(T_i) \wr \mathrm{Sym}(d_i)$, where T_i occurs d_i times as socle factor of S.

$\mathrm{Aut}(H)$ is normaliser of the image of H in D.

Now lift results through elementary abelian layers, computing $\mathrm{Aut}(G/N_i)$ successively.

Suppose $N \leq M \leq G$, where both M, N char in G and M/N is elementary abelian of order p^d.

Suppose $A_M = \mathrm{Aut}(G/M)$ is known.

All automorphisms of G fix both M and N.

$A_N = \mathrm{Aut}(G/N)$ has normal subgroups $C \leq B$
B induces identity on G/M
C induces identity on both G/M and M/N.

M/N is $\mathbb{F}_p(G/M)$-module.

- Elements of C correspond to derivations from G/M to M/N.
- Elements of B/C correspond to module automorphisms of M/N. Can choose M and N to ensure that these tasks "easy".
- Hardest task: determine $S \leq A_M$ which lifts to G/N. $S \leq A'$, subgroup of A_M whose elements preserve the isomorphism type of module M/N.

  G/N split extension of M/N by G/M?
  If so, all elements of A' lift.
  Otherwise, must test each element of A' for lifting.

Examples of algorithms using Soluble Radical model

- Determine conjugacy classes of elements of G; (Cannon & Souvignier, 1997)
- Determine maximal subgroups of G; (Cannon & Holt, 2004) and (Eick & Hulpke, 2001)
- Determine the automorphism group of G; (Cannon & Holt, 2003)
- Determine conjugacy classes of subgroups of G; (Cannon, Cox & Holt, 2001)

Constructing and using $\mathcal{C}$

Bäärnhielm, Holt, Leedham-Green & O'B (2014): refine composition series obtained from "geometric model" to obtain chief series reflecting characteristic structure.

Holt and others: developed Soluble Radical model algorithms using tree as infrastructure [replacement for BSGS].

Publicly available in MAGMA; parts available in GAP.

Black-box model to exploit this chain pioneered by Babai and Beals.

Babai, Beals, Seress (2009):

Theorem

$\mathcal{C}$ can be constructed directly in black-box groups in polynomial time (subject to Discrete Log solution and some other restrictions).

Challenge problems

Problem

Find the order of $H \leq \mathrm{GL}(6, 5^2)$.

Yes, in practice.

Problem

Given $g \in \mathrm{GL}(6, 5^2)$ *find its order.*

Yes, in practice.

Problem

Find the normaliser in $\mathrm{GL}(8, 5^2)$ *of a subgroup of moderate index.*

Some progress . . .: Hannah Coutts