ISRAEL JOURNAL OF MATHEMATICS 109 (1999), 93–116

DECIDING FINITENESS FOR MATRIX GROUPS OVER FUNCTION FIELDS

BY

DANIEL N. ROCKMORE*

Department of Mathematics, Dartmouth College
Hanover, NH 03755, USA
e-mail: rockmore@cs.dartmouth.edu

AND

KI-SENG TAN

Department of Mathematics, National Taiwan University
Taipei 106, Taiwan
e-mail: tan@math.ntu.edu.tw

AND

ROBERT BEALS

Department of Mathematics, University of Arizona
Tuscon, AZ 85721, USA
e-mail: beals@math.arizona.edu

ABSTRACT

Let $\mathbf{F}$ be a field and t an indeterminate. In this paper we consider aspects of the problem of deciding if a finitely generated subgroup of $\mathrm{GL}(n,\mathbf{F}(t))$ is finite. When $\mathbf{F}$ is a number field, the analysis may be easily reduced to deciding finiteness for subgroups of $\mathrm{GL}(n,\mathbf{F})$, for which the results of [1] can be applied. When $\mathbf{F}$ is a finite field, the situation is more subtle. In this case our main results are a structure theorem generalizing a theorem of Weil and upper bounds on the size of a finite subgroup generated by a fixed number of generators with examples of constructions almost achieving the bounds. We use these results to then give exponential deterministic algorithms for deciding finiteness as well as some preliminary results towards more efficient randomized algorithms.

* Supported in part by NSF DMS Awards 9404275 and Presidential Faculty Fellowship.

Received April 15, 1997 and in revised form February 2, 1998

1. Introduction

Recently, computational group theory has directed increased attention to the development of algorithms for studying matrix groups. In particular, various *recognition* algorithms are of importance. These algorithms efficiently identify a group given by a set of generators by deciding various properties of the group from the generators (cf. [1, 2, 3, 13, 14, 15] as well as the volumes [8, 9] and the many references therein).

For a potentially infinite group, possibly the most fundamental property to be determined of a set of generators is that of *finiteness*. That is, given $S \subset \Gamma$, for an infinite group Γ, decide if $\langle S \rangle$, the subgroup generated by S, is finite.

In this paper we consider the problem of finiteness for matrix groups with entries in function fields. That is, $\Gamma = \mathrm{GL}(n, \mathbf{F}(t))$, for t an indeterminate and $\mathbf{F}$ either a number field or the finite field of q elements. In the former case, the analysis is easily reduced to that of deciding finiteness for $\mathrm{GL}(n, \mathbf{F})$, for which efficient algorithms are known [1]. On the other hand, for the finite field case, no such quick reduction is apparent as some of the key tools used in the number field case (the semisimplicity of the enveloping algebra and the bounded order of any finite subgroup) no longer obtain. New ideas seem necessary here.

In Section 2 we consider the number field case. We show that in polynomial time we can either recognize the group as infinite or construct an isomorphic subgroup of $\mathrm{GL}(m, \mathbf{F})$ (for m not necssarily equal to n). In the latter case the techniques of [1] can be applied. Thus, for number fields we are able to give a polynomial-time algorithm for deciding finiteness (Theorem 2.1).

The main new results of this paper are in Section 3 where we take up the case of $\mathbf{F} = \mathbf{F}_q$, the field of q elements. Positive characteristic makes for a quite different situation and our main results treat this case. The techniques used in characteristic zero depend quite heavily on two conditions which do not extend to positive characteristic:

(1) The enveloping algebra of any finite matrix group over a field of characteristic zero will be semisimple.

(2) Finite subgroups of $\mathrm{GL}(n, \mathbf{F}(t))$ are necessarily conjugate to finite subgroups of $\mathrm{GL}(n, \mathbf{F})$.

Section 3.1 gives our first main result which is a structure theorem (in the spirit of (2)) for finite subgroups of $\mathrm{GL}_n(\mathbf{F}_q(t))$.

THEOREM 3.2: *Let $G \leq \mathrm{GL}(n, \mathbf{F}_q(t))$. Then G is finite if and only if G is* *conjugate to a subgroup of* $\mathrm{GL}(n, \mathbf{F}_q(t))$ *of the form*

$$\begin{pmatrix} A_1 & * & \cdots & * \\ 0 & A_2 & \cdots & * \\ 0 & 0 & \vdots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_r \end{pmatrix} \tag{1}$$

where the $A_i \in \mathrm{GL}(n_i, \mathbf{F}_q)$ *and* $*$ *indicates blocks of elements in the upper triangle which are all in* $\mathbf{F}_q[t]$ *and of bounded degree.*

In fact, Theorem 3.2 is a special case of a general decomposition theorem for matrix groups over local fields (Theorem 3.1). This extends an earlier result of Weil ([19], Theorem 1).

Notice that Theorem 3.2 implies that finite subgroups of $\mathrm{GL}(n, \mathbf{F}_q(t))$ can be arbitrarily large. This too differs from the number field case. Nevertheless, it is still possible to obtain bounds on the size of these groups in terms of the number of generators and various other parameters. For example (as the referee points out), a group generated by a finite set of elements of the form (1) is an extension of a finite group by a group which is both locally finite of exponent at most p^{n-1} and nilpotent of class at most $n-1$. Thus, in this case it is not too difficult to see that a group generated by a finite number of such elements will be finite (Corollary 3.6) and have order bounded in terms of these parameters. In Section 3.2 we give a more detailed investigation of this situation and arrive at a more precise estimate (cf. Theorems 3.3 and 3.9). In Section 3.3 we give examples indicating that our upper bounds may be tight.

Section 3.4 derives algorithms to determine finiteness using the results of the previous sections. Theorem 3.2 essentially reduces the problem of deciding finiteness in positive characteristic to finding invariant subspaces of $V \cong (\mathbf{F}_q(t))^n$. Once an invariant subspace is found it is easy to check if the associated restricted representation is defined only over $\mathbf{F}_q$. Direct application of this line of approach gives deterministic algorithms for deciding finiteness for which we can only derive an exponential upper bound.

Of a more speculative nature is a randomized approach. Parker's "Meat-Axe" [16] is a randomized algorithm used to decompose modular representations of finite groups. We outline how an adaptation of this method could be used in our situation.

ACKNOWLEDGEMENT: We thank the referee for a careful reading of the paper and pointing out the connection with the theory of buildings.

2. Number fields

In this section we consider the finiteness problem for $\mathrm{GL}(n, \mathbf{F}(t))$ for $\mathbf{F}$ a number field. By the results of [1], it is enough to give a polynomial-time reduction to the case of $\mathrm{GL}(n, \mathbf{F})$.

THEOREM 2.1: *Let $\mathbf{F}$ be a number field and let $\langle S\rangle = G \leq \mathrm{GL}(n, \mathbf{F}(t))$. Then in polynomial time either G can be transformed into an equivalent subgroup of $\mathrm{GL}(n, \mathbf{F})$, or shown to be infinite.*

We begin by showing that any finite subgroup of $\mathrm{GL}(n, \mathbf{F}(t))$ is isomorphic to a finite subgroup of $\mathrm{GL}(n, \mathbf{F})$.

LEMMA 2.2: *Let $\mathbf{F}$ be an infinite field and let $G \leq \mathrm{GL}(n, \mathbf{F}(t))$ be finite. Then there exists a finite extension F'/F and a matrix $Y \in \mathrm{GL}(n, F'(t))$ such that $YAY^{-1} \in \mathrm{GL}(n, \mathbf{F})$ for all $A \in G$.*

Proof: If G is finite, then there exists a finite extension $\mathbf{F}'$ of $\mathbf{F}$ such that the defining representation of G in $\mathrm{GL}(n, \mathbf{F}(t))$ is equivalent to a representation of G in $\mathrm{GL}(n, \mathbf{F}')$. In this case, let $X \in \mathrm{GL}(n, \mathbf{F}'(t))$ be such that

$$XAX^{-1} \in \mathrm{GL}(n, \mathbf{F}') \tag{2}$$

for all $A \in G \leq \mathrm{GL}(n, \mathbf{F}(t))$.

Since G is finite and $\mathbf{F}$ is infinite, there exists $\alpha \in \mathbf{F}$ such that α is not a root of any denominator occurring among X, X^{-1} and the matrices of G. For any matrix $B \in \mathrm{GL}(n, \mathbf{F}'(t))$ let $B|_\alpha$ denote the matrix obtained by evaluating all entries of B at α (assuming they are defined). Note that for $B, C \in \mathrm{GL}(n, \mathbf{F}'(t))$,

$$B|_\alpha C|_\alpha = (BC)|_\alpha. \tag{3}$$

Thus, $B \in \mathrm{GL}(n, \mathbf{F}'(t))$ implies that $B|_\alpha \in \mathrm{GL}(n, \mathbf{F}')$ and $(B^{-1})|_\alpha = (B|_\alpha)^{-1}$. These observations imply

$$\begin{aligned} A|_\alpha &= (X^{-1}XAX^{-1}X)|_\alpha \\ &= X^{-1}|_\alpha \, (XAX^{-1})|_\alpha \, X|_\alpha \\ &= X^{-1}|_\alpha \, (XAX^{-1}) \, X|_\alpha \\ &= (X^{-1}|_\alpha \, X) \, A \, (X^{-1} \, X|_\alpha) \end{aligned}$$

where the next to last equality follows from (2). Finally, since $\alpha \in \mathbf{F}$, $A|_\alpha \in \mathrm{GL}(n, \mathbf{F})$ for all $A \in G < \mathrm{GL}(n, \mathbf{F}(t))$. ∎

COROLLARY 2.3: *Let $G \leq \mathrm{GL}(n, \mathbf{Q}(t))$ be finite. Then there exists a finite extension $K/\mathbf{Q}$ and a matrix $Y \in \mathrm{GL}(n, K(t))$ such that $YAY^{-1} \in \mathrm{GL}(n, \mathbf{Z})$ for all $A \in G$.*

Proof: By Theorem 2.4 of [1], any finite subgroup of $\mathrm{GL}(n, \mathbf{Q})$ is conjugate to a finite subgroup of $\mathrm{GL}(n, \mathbf{Z})$. The result then follows using Lemma 2.2. ∎

COROLLARY 2.4:

(1) *Let $\mathbf{F}$ be of characterstic zero. If $G \leq \mathrm{GL}(n, \mathbf{F}(t))$ is finite, then* $\mathrm{trace}(A) \in \mathbf{F}$ *for all $A \in G$.*

(2) *Let $G \leq \mathrm{GL}(n, \mathbf{Q}(t))$ be finite, then* $\mathrm{trace}(A) \in \mathbf{Z}$ *for all $A \in G$.*

For any set of matrices $S \subset \mathrm{GL}(n, \mathbf{F}(t))$, let $\mathrm{env}_{\mathbf{F}}(S)$ denote the $\mathbf{F}$-linear span of the group generated by S. Lemma 2.2 shows that if G is finite, then $\dim_{\mathbf{F}}(\mathrm{env}_{\mathbf{F}}(G)) \leq n^2$. This is the key to our efficient finiteness test. The first step is to generate $\mathbf{F}$-linearly independent elements in $\mathrm{env}_{\mathbf{F}}(G)$ until either more than n^2 such elements are obtained, or a basis over $\mathbf{F}$ of dimension less than n^2 is found. In the former case the group is infinite. In the latter case, the basis over $\mathbf{F}$ can be used to find a faithful representation of G in $\mathrm{GL}(d, \mathbf{F})$ for some $d \leq n^2$, which can be tested to be finite using the results of [1].

By successive matrix multiplications and Gaussian elimination we have the following result.

LEMMA 2.5: *Let $\langle S \rangle = G \leq \mathrm{GL}(n, \mathbf{F}(t))$. Then in polynomial time, either a basis $A_1, \ldots, A_d$ for $\mathrm{env}_{\mathbf{F}}(G)$ $(d \leq n^2)$ can be constructed, or we may deduce that G is infinite. In the former case, the A_i can be taken to be in S^{i-1}.*

When G is finite, using Lemma 2.5 we can now construct an isomorphism of G with a subgroup of $\mathrm{GL}(d, \mathbf{F})$ by considering the faithful representation of G on $\mathrm{env}_{\mathbf{F}}(G)$. Theorem 2.1 is now proved.

H. Bass points out that it is easy to see that this argument works for any finite number of indeterminates and thus,

THEOREM 2.6: *Let $\langle S \rangle = G \leq \mathrm{GL}(n, \mathbf{F}(t_1, \ldots, t_m))$ for independent indeterminates $t_1, \ldots, t_m$. Then in polynomial time either G can be transformed into an equivalent subgroup of $\mathrm{GL}(n, \mathbf{F})$, or shown to be infinite.*

3. Positive characteristic

It appears that the arguments used in the characteristic zero case are not applicable to positive characteristic. For example, in general, if the group is finite, the associated enveloping algebra will not be semisimple. Also, subgroups of $\mathrm{GL}(n, \mathbf{F}_q(t))$ can be arbitrarily large. (When $n = 2$, consider the upper triangular subgroups generated by elements with monomials of differing degrees in the upper corner.) Consequently, new ideas seem to be necessary. As a first step we prove a structure theorem for finite subgroups of $\mathrm{GL}(n, \mathbf{F}_q(t))$. This will follow from a much more general decomposition theorem, which is a mild generalization of a classical result of Weil [19]. A particular case of this was also considered in the context of computing L-series for modular functions over function fields [18].

Notice that a subgroup of $\mathrm{GL}(n, \mathbf{F}_q(t))$ is finite if and only if its enveloping algebra over $\mathbf{F}_q$ is finite, or equivalently, finite dimensional. The structure theorem (Theorem 3.2) allows us to give some bounds on the size of the enveloping algebra over $\mathbf{F}_q$, and consequently some coarse bounds on the size of finite subgroups of $\mathrm{GL}(n, \mathbf{F}_q(t))$ generated by a fixed number of elements (Corollary 3.12, Corollary 3.12 and Corollary 3.15).

With these results in hand we may then give some naive algorithms for determining finiteness. We anticipate that there is much room for improvement.

3.1 STRUCTURE THEOREMS FOR $\mathrm{GL}(n, \mathbf{F}_q(t))$. We first introduce some notation which will be in use throughout this section.

Let $K = \mathbf{F}_q((\frac{1}{t}))$, $\mathcal{O} = \mathbf{F}_q[[\frac{1}{t}]]$ and $\pi = 1/t$. Let $\Gamma = \mathrm{GL}(n, \mathbf{F}_q[t]) \subset \mathrm{GL}(n, K)$ and let $\mathcal{Z} \simeq K^*$ be the center of $\mathrm{GL}(n, K)$. Define

$$\mathcal{R} = \{\mathbf{r} = (r_1, \dots, r_n) | \, r_i \in \mathbf{Z}, 0 = r_1 \le r_2 \le \cdots \le r_n\}$$

and, for each $\mathbf{r} \in \mathcal{R}$, define

$$\rho_{\mathbf{r}} = \begin{pmatrix} t^{r_1} & 0 & \cdots & 0 \\ 0 & t^{r_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t^{r_n} \end{pmatrix}.$$

The following theorem is basically from [17].

THEOREM 3.1: *With the notation above, for each $g \in \mathrm{GL}(n, K)$ there exist $\gamma \in \Gamma$, $\underline{s} \in \mathcal{R}$, $\xi \in \mathrm{GL}(n, \mathcal{O})$ and $\zeta \in \mathcal{Z}$ such that*

$$g = \gamma \cdot \rho_{\underline{s}} \cdot \xi \cdot \zeta.$$

Proof: Let $\mathbf{T}$ be the Bruhat-Tits building associated to $\mathrm{PGL}(n,K)$ (or $\mathrm{SL}(n,K)$). The vertices of $\mathbf{T}$ can be identified with the cosets

$$\mathrm{GL}(n,K)/\mathcal{Z}\cdot\mathrm{GL}(n,\mathcal{O}).$$

Let $\mathbf{Y}$ be the sector ("quartier") of $\mathbf{T}$, whose vertices are the cosets $y_{\mathbf{r}}$ (for $\mathbf{r}\in\mathcal{R}$) where

$$y_{\mathbf{r}}=\begin{pmatrix}\pi^{r_1} & 0 & \cdots & 0\\ 0 & \pi^{r_2} & \cdots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \cdots & \pi^{r_n}\end{pmatrix}\cdot\mathcal{Z}\cdot\mathrm{GL}(n,\mathcal{O}).$$

By Theorem 1 of [17], under the action of $\mathrm{SL}(n,\mathbf{F}_q[t])$, the sector $\mathbf{Y}$ is a simplicial fundamental domain, namely, any simplex of $\mathbf{T}$ is equivalent by $\mathrm{SL}(n,\mathbf{F}_q[t])$ to a unique simplex of $\mathbf{Y}$. In particular, every vertex of $\mathbf{T}$ is equivalent by $\mathrm{SL}(n,\mathbf{F}_q[t])$ to some $y_{\mathbf{r}}$. Theorem 3.1 then follows, because

$$\begin{pmatrix}0 & \cdots & 0 & 1\\ 0 & \cdots & 1 & 0\\ \vdots & \vdots & \ddots & \vdots\\ 1 & 0 & \cdots & 0\end{pmatrix}\cdot y_{\mathbf{r}}\cdot\begin{pmatrix}0 & \cdots & 0 & 1\\ 0 & \cdots & 1 & 0\\ \vdots & \vdots & \ddots & \vdots\\ 1 & 0 & \cdots & 0\end{pmatrix}=\pi^{r_n}\cdot\rho_{\underline{s}},$$

where

$$\underline{s}=(0,\ldots,r_n-r_i,\ldots,r_n-r_2,r_n).\qquad\blacksquare$$

For $\mathbf{r}=(r_1,\ldots,r_n)\in\mathcal{R}$, define subgroup $\Gamma_{\mathbf{r}}\leq\Gamma$ by

$$\Gamma_{\mathbf{r}}=\{\gamma=(\gamma_{ij})\in\Gamma\mid t^{r_i-r_j}\gamma_{ij}\in\mathcal{O}\}.$$

Note that $\Gamma_{\mathbf{r}}$ is a finite subgroup of Γ. In particular, let $\gamma\in\Gamma_{\mathbf{r}}$, then (1) if $r_i=r_j$, then $\gamma_{ij}\in\mathbf{F}_q$; (2) if $r_i>r_j$, then $\gamma_{ij}=0$; (3) if $r_i<r_j$ then $\deg(\gamma_{ij})\leq r_j-r_i$.

THEOREM 3.2: *Suppose $G\leq\mathrm{GL}(n,\mathbf{F}_q(t))$ is finite. Then there exists $\Delta\in\mathrm{GL}(n,\mathbf{F}_q(t))$ and $\mathbf{r}\in\mathcal{R}$ such that*

$$\Delta\cdot G\cdot\Delta^{-1}\leq\Gamma_{\mathbf{r}}.$$

Proof: Since G is finite, under its natural action on $\mathbf{F}_q(t)^n$, G stabilizes a free $\mathbf{F}_q[t]$-submodule that is of rank n. This shows that G is conjugate to a subgroup of Γ. Without loss of generality, we assume that G is a subgroup of Γ. As a subgroup of $\mathrm{GL}(n,K)$, G also acts on K^n. Again, the finiteness of G implies that there exists $g \in \mathrm{GL}(n, K)$ such that $G \cdot g \subset g \cdot \mathrm{GL}(n, \mathcal{O})$. By Theorem 3.1 we can write

$$g = \gamma \cdot \rho_{\mathrm{r}} \cdot \xi \cdot \zeta.$$

Then for $x \in G$,

$$\rho_{\mathrm{r}}^{-1} \cdot \gamma^{-1} \cdot x \cdot \gamma \cdot \rho_{\mathrm{r}} \in \mathrm{GL}(n, \mathcal{O}) \cdot \mathcal{Z}.$$

Let $y = \gamma^{-1} \cdot x \cdot \gamma$. Then

$$\left(t^{r_j - r_i} y_{ij}\right) = \begin{pmatrix} t^{-r_1} & 0 & \cdots & 0 \\ 0 & t^{-r_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t^{-r_n} \end{pmatrix} \cdot \left(y_{ij}\right) \cdot \begin{pmatrix} t^{r_1} & 0 & \cdots & 0 \\ 0 & t^{r_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t^{r_n} \end{pmatrix} = \xi' \cdot \zeta'$$

for some $\zeta' \in \mathcal{Z}$ and for some $\xi' \in \mathrm{GL}(n, \mathcal{O})$.

Since $y \in \Gamma$, $\det(\rho_{\mathrm{r}}^{-1} \cdot y \cdot \rho_{\mathrm{r}}) = \det(y) \in \mathbf{F}_q{}^*$. Also, $\det(\xi') \in \mathbf{F}_q[[\frac{1}{t}]]^*$, and $\zeta'^n = \det(\zeta') \in \mathbf{F}_q[[\frac{1}{t}]]^*$. Hence, $\zeta' \in \mathbf{F}_q[[\frac{1}{t}]]^* \subset \mathrm{GL}(n, \mathcal{O})$. Thus, $t^{r_j - r_i} y_{ij} \in \mathcal{O}$ and $y \in \Gamma_r$. ∎

3.2 Bounds for finite subgroups of $\mathrm{GL}(n, \mathbf{F}_q(t))$. In characteristic zero finite subgroups of $\mathrm{GL}(n, \mathbf{F}(t))$ are bounded in size. Upper bounds can be derived from bounds for the size of subgroups of $\mathrm{GL}(n, \mathbf{Z})$. For this situation, W. Feit has shown [7], by using some unpublished results of B. Weisfeiler and the classification of finite simple groups (cf. [11]), that except for $n = 2, 4, 6, 7, 8, 9$, and 10, finite subgroups of $\mathrm{GL}(n, \mathbf{Z})$ have size at most $2^n(n!)$, for which the signed permutation matrices provide an example. In the other cases, the Weyl groups of some of the exceptional groups provide larger bounds. Beyond this, the paper [7] also contains a careful analysis of the finite subgroups of $\mathrm{GL}(n, \mathbf{F})$ for $\mathbf{F}$ a cyclotomic field. Other interesting results on finite subgroups of $\mathrm{GL}(n, \mathbf{C})$ and a wealth of references can be found in Friedland's recent paper [10].

In positive characteristic the situation is very different. Here, finite subgroups can be arbitrarily large. However, for a fixed number of generators a bound can be given. As the referee points out (see Section 1) a simple observation of the structure of these groups as a certain group extension give very coarse bounds on the size.

Our method is to instead consider the algebra over $\mathbf{F}_q$ generated by the generators of the group and to bound its dimension (over $\mathbf{F}_q$). This quickly translates into a bound on the size of the subgroup which they generate. We present two different approaches towards such results. The former uses a simple analysis of the matrix multiplications which can occur (Section 3.2.1) while the latter (Section 3.2.2) introduces the notion of the diameter of a finitely generated algebra and is able to obtain bounds in terms of the diameters of certain subalgebras of the full matrix algebras $M(n_i, \mathbf{F}_q)$ for a given composition $n_1 + \cdots + n_d = n$. We also discuss the issue of the "tightness" of these bounds.

3.2.1 General upper bounds. To begin, let us introduce some notation. For a commutative ring R we will use the notation $M(r \times s, R)$ to denote the R-module of $r \times s$ matrices with entries in R and $M(n, R)$ for $M(n \times n, R)$.

Let $\mathbf{n} = (n_1, \ldots, n_d)$ be a composition of n, denoted as $\mathbf{n} \models n$. Let $T(\mathbf{n}, \mathbf{F}_q(t)) \leq \mathrm{GL}(n, \mathbf{F}_q(t))$ be defined as the block upper triangular matrices $X = (X_{ij})$ such that $X_{ij} \in M(n_i \times n_j, \mathbf{F}_q(t))$ and

$$\left\{ \begin{array}{lcl} X_{ii} & \in & \mathrm{GL}(n_i, \mathbf{F}_q), \\ X_{ij} & = & 0 \text{ if } i > j. \end{array} \right. \tag{4}$$

Note that for $0 \leq r_1 \leq \cdots \leq r_n$, we have $\Gamma_{\mathbf{r}} \leq T(\mathbf{n}, \mathbf{F}_q(t))$ where

$$r_1 = \cdots = r_{n_1} < r_{n_1+1} = \cdots = r_{n_1+n_2} < \cdots < r_{n_1+\cdots+n_{d-1}+1} = \cdots = r_{n_1+\cdots+n_d}.$$

THEOREM 3.3: *Let $\langle S \rangle = G \leq \mathrm{GL}(n, \mathbf{F}_q(t))$ such that $|S| = r$ and G is finite. Then*

$$\dim_{\mathbf{F}_q}(\mathrm{env}_{\mathbf{F}_q}(G)) \leq \left\{ \begin{array}{cl} n & \text{if } r = 1, \\ \frac{1}{r}(r+1)^n & \text{if } r > 1. \end{array} \right.$$

Proof: First consider the case $r = 1$. Then $G = \langle A \rangle$ for some $A \in \mathrm{GL}(n, \mathbf{F}_q(t))$ of finite order. By Theorem 3.2, we know that A is conjugate to some $A' \in T(\mathbf{n}, \mathbf{F}_q(t))$ for $\mathbf{n}$ a composition of n. In particular, the characteristic polynomial of A is of degree n with coefficients in $\mathbf{F}_q$. Since A is a root of its characteristic polynomial, its powers will span a space of at most dimension n over $\mathbf{F}_q$.

Now assume $r > 1$. We will repeatedly make use of the following very pessimistic upper bound.

CLAIM: *Let $n_1, \ldots, n_d$ be positive integers. For $i = 1, \ldots, d-1$ let $X_{i,i+1} \in M(n_i \times n_{i+1}, \mathbf{F}_q(t))$ and*

$$V = \mathrm{span}_{\mathbf{F}_q}\left(\{ g_1 \cdot X_{1,2} \cdot g_2 \cdot X_{2,3} \cdot g_3 \cdots g_{d-1} \cdot X_{d-1,d} \cdot g_d | \; g_i \in \mathrm{GL}(n_i, \mathbf{F}_q) \} \right).$$

Then

$$\dim_{\mathbf{F}_q}(V) \leq n_1^2 n_2^2 \cdots n_d^2.$$

Proof: Notice that the entries of $X_{1,2}$ are contained in an $\mathbf{F}_q$-vector space of dimension at most $n_1 n_2$ (a set of $n_1 n_2$ rational functions over any field will span a vector space of at most dimension n_1n_2 over that field). Pre- and postmultiplication by any g_1 and g_2 only effects a linear combination of these. Further postmultiplication by $X_{2,3}$ gives linear combinations of now at most $n_1n_2n_2n_3 = n_1n_2^2n_3$ rational functions. Continuing in this fashion we see that the entries of any element in V can be linear combinations of at most $n_1n_2^2\cdots n_{d-1}^2n_d$ rational functions. As there are n_1n_d entries in any element of V, we see that

$$\dim_{\mathbf{F}_q}(V) \le n_1n_d\cdot n_1n_2^2\cdots n_{d-1}^2n_d = n_1^2n_2^2\cdots n_d^2. \qquad \blacksquare$$

Using the Claim we may now prove the bound for $r > 1$. Note that we may assume that G is generated by $S = \{A_1, A_2, ..., A_r\}$ such that each $A_k = (X_{ij}^{(k)}) \in T(\mathbf{n}, \mathbf{F}_q(t))$ and so is of a fixed upper-triangular block-form. (As indicated, we assume that $X_{ij}^{(k)}$ refers to the i,j block of size $n_i \times n_j$ in A_k.)

Let $E = \mathrm{env}_{\mathbf{F}_q}(G)$. In order to bound $\dim_{\mathbf{F}_q}(E)$ we use the simplification,

$$\dim_{\mathbf{F}_q} E \le \sum_i n_i^2 + \sum_{i<j} \dim_{\mathbf{F}_q} V_{i,j} \tag{5}$$

where

$$V_{i,j} = \mathrm{span}_{\mathbf{F}_q}\{X_{i,j}\colon X \in E\}. \tag{6}$$

That is, $V_{i,j}$ is the subspace of $M(n_i \times n_j, \mathbf{F}_q(t))$ spanned by the i,j blocks of all products of the form $A_{a_1}\cdots A_{a_k}$. For any such product we have that

$$(A_{a_1}\cdots A_{a_k})_{i,j} = \left(\sum_{i\le l_1\le\cdots\le l_{k-1}\le j} X_{il_1}^{(a_1)}\cdot X_{l_1l_2}^{(a_2)}\cdot\ldots\cdot X_{l_{k-2}l_{k-1}}^{(a_{k-1})}\cdot X_{l_{k-1}j}^{(a_k)}\right). \tag{7}$$

The claim shows that

$$\dim_{\mathbf{F}_q} V_{i,j} \le \sum_{i\le i_1<\cdots<i_m\le j} n_{i_1}^2\cdots n_{i_m}^2\cdot r^{m-1}. \tag{8}$$

Using (5) and (8) over all sequences $1 \le i \le i_1 < ... < i_m \le j \le d$ and reindexing we get

$$\begin{aligned}\dim_{\mathbf{F}_q}(E) &\le \sum_i n_i^2 + \sum_{1\le i<j\le d}\frac{1}{r}\sum_{i<i_1<\ldots<i_m<j}(rn_{i_1}^2)\cdots(rn_{i_m}^2)\\ &= \sum_i n_i^2 + \frac{1}{r}(rn_1^2+1)\cdot\ldots\cdot(rn_d^2+1) - \frac{1}{r}(1+rn_1^2+rn_2^2+\ldots+rn_d^2). \qquad (9)\end{aligned}$$

This shows that

$$\dim_{F_q} X < \frac{1}{r}(rn_1^2+1)\cdots(rn_d^2+1).$$

Finally, notice that for $r \geq 2$ and $m \geq 0$,

$$(rm^2 + 1) \leq (r+1)^m$$

and thus (as $\sum_i n_i = n$)

$$\dim_{\mathbf{F}_q}(X) \leq \frac{1}{r}(r+1)^n. \qquad ∎$$

Corollary 3.4: *Let $\langle S \rangle = G \leq \mathrm{GL}(n, \mathbf{F}_q(t))$ have finite order. Then*

$$|G| \leq \begin{cases} q^n - 1 & \textit{if } r = 1, \\ q^{\frac{1}{r}(r+1)^n} - 1 & \textit{if } r > 1. \end{cases}$$

In particular, every element of finite order in $\mathrm{GL}(n, \mathbf{F}_q(t))$ has order at most $q^n - 1$.

Proof: The size of the group is at most the size of the nonzero part of the enveloping algebra over $\mathbf{F}_q$. ∎

Another easy corollary also follows.

Corollary 3.5: *An element $A \in \mathrm{GL}(n, \mathbf{F}_q(t))$ has finite order if and only if the characteristic polynomial of A is defined over $\mathbf{F}_q$.*

Theorem 3.2 shows that any finite subgroup is conjugate to a subgroup of some $T(\mathbf{n}, \mathbf{F}_q(t))$. Either a slight modification of the above proof of Theorem 3.3, or the observation that this group is an extension of a finite group by a locally finite group of bounded exponent and finite nilpotency class implies that any finitely generated subgroup of $T(\mathbf{n}, \mathbf{F}_q(t))$ is finite. We record this fact next.

Corollary 3.6: *Any finitely generated subgroup of $T(\mathbf{n}, \mathbf{F}_q(t))$ is finite.*

The strength of the bounds of Theorem 3.3 and Corollary 3.4 is investigated in Section 3.3.

3.2.2 Upper bounds and diameters for algebras. Here we introduce the notion of **diameter** for a finitely generated algebra. In so doing we are able to obtain a different upper bound on the size of finite subgroups of $\mathrm{GL}(n, \mathbf{F}_q(t))$ generated by r elements in terms of "natural" combinatorial data derived from the generators.

Let A be an algebra over a field $\mathbf{F}$ and $S = \{X^{(1)}, \ldots, X^{(r)}\} \subset A$. As usual, let $\mathrm{env}_{\mathbf{F}}(S)$ denote the $\mathbf{F}$-subalgebra of A generated by S. Furthermore, for any integer $j \geq 0$, let S^j denote the subset of elements of A which can be written as products of j elements of S.

Definition 3.7: The subalgebra $\mathrm{env}_{\mathbf{F}}(S)$ is said to have **diameter** δ, written $\delta = \mathrm{diam}(\mathrm{env}_{\mathbf{F}}(S))$, if all elements of $\mathrm{env}_{\mathbf{F}}(S)$ can be written as $\mathbf{F}$-linear combinations of $S^0 \cup \cdots \cup S^\delta$ and δ is the least integer such that this is true. If $X \in \mathrm{env}_{\mathbf{F}}(S)$, define the **length** of X (with respect to S), denoted $\mathrm{len}_S(X)$, to be the smallest integer j (necessarily less than $\mathrm{diam}(\mathrm{env}_{\mathbf{F}}(S))$) such that $X \in \mathrm{span}_{\mathbf{F}}(S^0, \ldots, S^j)$.

LEMMA 3.8: *Let all notation be as above and let $S \subset A$ and $\delta = \mathrm{diam}(\mathrm{env}_{\mathbf{F}}(S))$; then*

$$\dim_{\mathbf{F}}(\mathrm{env}_{\mathbf{F}}(S)) \leq 1 + r + \cdots + r^\delta.$$

Proof: There are at most r^j $\mathbf{F}$-linearly independent elements in S^j. ∎

Remark: Notice that diameter of an algebra is related, albeit in a seemingly loose fashion, to the concept of diameter of a finitely generated group. For example, let $X \in \mathrm{GL}(n, \mathbf{F}_q)$. Since X satisfies its characteristic polynomial, $\mathrm{diam}(\mathrm{env}_{\mathbf{q}}(X)) \leq n-1$. However, the order of X can be at most $q^n - 1$, in which case the diameter of the cyclic group generated by X is $\frac{1}{2}(q^n - 1)$.

At the very least, it is clear that for $S \subset \mathrm{GL}(n, \mathbf{F})$,

$$\mathrm{diam}(\mathrm{env}_{\mathbf{F}}(S)) \leq \mathrm{diam}(\langle S \rangle)$$

where the righthand side denotes the diameter of the subgroup of $\mathrm{GL}(n, \mathbf{F})$ generated by S.

Using the notion of diameter, another bound for the size of finite subgroups of $\mathrm{GL}(n, \mathbf{F}_q(t))$ can be obtained. To simplify the statement of the result, for any $X \in T(\mathbf{n}, \mathbf{F}_q(t))$ (for $\mathbf{n} = (n_1, \ldots, n_d) \models n$) let

$$X_{(i)} = \text{the } i,i \text{ block of } X. \tag{10}$$

Thus, $X_{(i)} \in \mathrm{GL}(n_i, \mathbf{F}_q)$ (cf. equation (4)).

THEOREM 3.9: *Let $\mathbf{n} \models n$ and $S = \{X^{(1)}, \ldots, X^{(r)}\} \subset T(\mathbf{n}, \mathbf{F}_q(t))$. Let*

$$S_i = \{X^{(1)}_{(i)}, \ldots, X^{(r)}_{(i)}\} \subset \mathrm{GL}(n_i, \mathbf{F}_q)$$

and

$$\delta_i = \mathrm{diam}(\mathrm{env}_{\mathbf{F}_q}(S_i)).$$

Then

$$\mathrm{diam}(\mathrm{env}_{\mathbf{F}_q}(S)) \leq \delta_1 + \cdots + \delta_d + d - 1.$$

Proof: It is enough to show that if $W = Y_1 Y_2 \cdots Y_m$ is a product of $m \geq \delta_1 + \cdots + \delta_d + d$ elements in S, then another expression for W can be found which is a linear combination of products of less than m elements in S.

Since $m \geq \delta_1 + \cdots + \delta_d + d$, W can be written as

$$W = W^{(1)} \cdots W^{(d)} \tag{11}$$

where for $i = 1, ..., d$, the element $W^{(i)}$ is in $S^{i'}$ for some $i' \geq \delta_i + 1$. Consider now the $1,1$ block of $W^{(1)}$ or in the notation of (10), $W^{(1)}_{(1)}$. Since $\operatorname{diam}(\operatorname{env}_{\mathbf{F}_q}(S_1)) = \delta_1$, then it must be the case that $W^{(1)}_{(1)}$ can be written as the $1,1$ block of a $\mathbf{F}_q$-linear combination of products of elements in S of length at most δ_1. Thus, let $Z^{(1)}$ be such an element, so that $\operatorname{len}_S(Z^{(1)}) \leq \delta_1$ and

$$W^{(1)}_{(1)} = Z^{(1)}_{(1)}.$$

Similarly, define $Z^{(j)}$ to be such that

- (1) $\operatorname{len}_S(Z^{(j)}) \leq \delta_j$ and
- (2) $W^{(j)}_{(j)} = Z^{(j)}_{(j)}$.

By condition (2)

$$(W^{(j)} - Z^{(j)})_{(j)} = 0. \tag{12}$$

Thus, using (12) it is easy to see

$$(W^{(1)} - Z^{(1)})(W^{(2)} - Z^{(2)}) \cdots (W^{(d)} - Z^{(d)}) = 0. \tag{13}$$

But (13) readily implies that

$$W = W^{(1)} \cdots W^{(d)} = Z$$

where Z is a linear combination of products of the form $A^{(1)} \cdots A^{(d)}$ where each for each i, $A^{(i)}$ is either equal to $W^{(i)}$ (and hence in $S^{i'}$ for some $i' > \delta_i$) or is a linear combination of elements which are in S^k for $k \leq \delta_i$ and furthermore **at least** one such i satisfies the latter condition. But this implies that Z is a linear combination of elements of length less than m and the theorem is proved. ∎

Suppose that $\mathbf{n} = (1^n)$. Then for every i, $n_i = 1$, $M(n_i, \mathbf{F}_q) \simeq \mathbf{F}_q$, and

$$\delta_i = 0.$$

COROLLARY 3.10: *Let all notation be as in Theorem 3.9. Let*

$$S = \{X_1, \ldots, X_r\} \subset T\left((1^n), \mathbf{F}_q(t)\right).$$

Then $\operatorname{diam}(S) \leq n - 1$.

COROLLARY 3.11: *Let all notation be as in Theorem 3.9. Let*

$$S = \{X_1, \dots, X_r\} \subset T\left((1^n), \mathbf{F}_q(t)\right).$$

Then $\dim_{\mathbf{F}_q}(\mathrm{env}_{\mathbf{F}_q}(S)) \leq 1 + r + \cdots + r^{n-1}$.

COROLLARY 3.12: *Let all notation be as in Theorem 3.9. Let*

$$S = \{X_1, \dots, X_r\} \subset T\left((1^n), \mathbf{F}_q(t)\right).$$

Let $G = \langle S \rangle$. *Then* $|G| \leq q^{1+r+\cdots+r^{n-1}} - 1$.

Corollary 3.10 can be generalized as follows. Let $U(\mathbf{n}, \mathbf{F}_q(t))$ denote the subgroup of $T(\mathbf{n}, \mathbf{F}_q(t))$ of block upper triangular matrices $((A_{ij}))$ with $A_{ii} = a_i \cdot I_{n_i}$, where $a_i \in {\mathbf{F}_q}^{\times}$. If $S \subset U(\mathbf{n}, \mathbf{F}_q(t))$, then $\delta_i = 0$ for every i.

COROLLARY 3.13: *Let all notation be as in Theorem 3.9. Let*

$$S = \{X_1, \dots, X_r\} \subset U\left(\mathbf{n}, \mathbf{F}_q(t)\right).$$

Then $\mathrm{diam}(S) \leq d - 1$.

COROLLARY 3.14: *Let all notation be as in Theorem 3.9. Let*

$$S = \{X_1, \dots, X_r\} \subset U\left(\mathbf{n}, \mathbf{F}_q(t)\right).$$

Then $\dim_{\mathbf{F}_q}(\mathrm{env}_{\mathbf{F}_q}(S)) \leq 1 + r + \cdots + r^{d-1}$.

COROLLARY 3.15: *Let all notation be as in Theorem 3.9. Let*

$$S = \{X_1, \dots, X_r\} \subset U\left(\mathbf{n}, \mathbf{F}_q(t)\right).$$

Let $G = \langle S \rangle$. *Then* $|G| \leq q^{1+r+\cdots+r^{d-1}} - 1$.

Remark: The above definitions and results suggest several natural questions. It would be of interest to better understand the diameters of various generating sets for subalgebras of $M(n, \mathbf{F}_q)$ and perhaps investigate their relationship to diameters of corresponding subgroups of $\mathrm{GL}(n, \mathbf{F}_q)$. Furthermore, it would be of interest to see under what conditions, if any, the bound of Corollary 3.14 is tight. The following section is a first step in this direction.

3.3 Explicit examples – bounds for $U(\mathbf{n}, \mathbf{F}_q(t))$. We now take up the problem of investigating if the bounds obtained in the previous section are tight.

Notice that if $r = 1$, the bound given by Corollary 3.4 is tight, since $\mathrm{GL}(n, \mathbf{F}_q)$ contains elements of order $q^n - 1$, so-called *Singer cycles*. They are realized as follows: Consider $\mathbf{F}_{q^n}$ as an n-dimensional vector space over $\mathbf{F}_q$. This gives a representation of $\mathbf{F}_{q^n}^\times$ as a subgroup of $\mathrm{GL}(n, \mathbf{F}_q)$, by considering the action of multiplication of $\mathbf{F}_{q^n}^\times$ on $\mathbf{F}_{q^n}$. Any generator of $\mathbf{F}_{q^n}^\times$ will then have order $q^n - 1$.

When $r > 1$, the situation is slightly more complicated. For example, we note that when $\mathbf{n} = (1,1)$ the bound in Corollary 3.11 is tight. In this case Theorem 3.3 yields

$$\dim_{\mathbf{F}_q}(\mathrm{env}_{\mathbf{F}_q}(G)) \leq \frac{1}{r}(r+1)^2 = r + 2 + \frac{1}{r}.$$

So for any $r > 1$, consider the generators,

$$\begin{pmatrix} a & 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & t \\ 0 & a \end{pmatrix}, \begin{pmatrix} 1 & t^2 \\ 0 & 1 \end{pmatrix}, \dots, \begin{pmatrix} 1 & t^{r-1} \\ 0 & 1 \end{pmatrix}$$

where a generates ${\mathbf{F}_q}^\times$.

The group

$$H = \left\{ \begin{pmatrix} \alpha & A + Bt \\ 0 & \beta \end{pmatrix} \mid \alpha, \beta \in {\mathbf{F}_q}^\times,\ A, B \in \mathbf{F}_q,\ A - B = \frac{\alpha - \beta}{a - 1} \right\}$$

is generated by $\begin{pmatrix} a & 1 \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & t \\ 0 & a \end{pmatrix}$. Also

$$\begin{pmatrix} 1 & t^2 \\ 0 & 1 \end{pmatrix}, \dots, \begin{pmatrix} 1 & t^{r-1} \\ 0 & 1 \end{pmatrix}$$

commute with each other. They generate a subgroup $G_{(r)}$ of $\mathrm{GL}(2, \mathbf{F}_q(t))$. It is not hard to check that H normalizes $G_{(r)}$ and

$$G = H \cdot G_{(r)} = \left\{ \begin{pmatrix} \alpha & A + Bt + a_2 t^2 + \ldots + a_{r-1} t^{r-1} \\ 0 & \beta \end{pmatrix} \mid \alpha, \beta \in {\mathbf{F}_q}^\times\ a_i \in \mathbf{F}_q, \right.$$
$$\left. A - B = \frac{\alpha - \beta}{a - 1} \right\}.$$

Notice that $\dim_{\mathbf{F}_q}(\mathrm{env}_{\mathbf{F}_q}(G)) = r + 1$.

In fact, this sort of construction, in which we fill the upper triangle with linearly independent polynomials, can be generalized to any $U(\mathbf{n}, \mathbf{F}_q(t))$ and $r \geq 0$, where as usual $\mathbf{n} = (n_1, \dots, n_d)$.

THEOREM 3.16: *There exist subsets $S = \{X_1, \dots, X_r\} \subset U(\mathbf{n}, \mathbf{F}_q(t))$ such that*

$$\dim_{\mathbf{F}_q}(\mathrm{env}_{\mathbf{F}_q}(S)) = 1 + r + \cdots + r^{d-1}.$$

Notice that any $X \in U(\mathbf{n}, \mathbf{F}_q(t))$ can be written as

$$X = I_n + N \tag{14}$$

where $N \in M(n, \mathbf{F}_q(t))$ is a strictly upper (block) triangular matrix. Call this subspace of strictly upper (block) triangular matrices $\mathcal{N}_{\mathbf{n}}$.

LEMMA 3.17: *Let $\mathbf{n} = (n_1, \dots, n_d) \vdash n$. Suppose that $N_1, \dots, N_d \in \mathcal{N}_{\mathbf{n}}$. Then*

$$N_1 \cdot N_2 \cdots N_d = 0.$$

The idea of the proof of Theorem 3.16 is to show that in fact, for a "generic" choice of $X_1, \dots, X_r$ the bound holds. This is accomplished by showing that an algebra generated by block matrices with entries given by distinct (and therefore algebraically independent) indeterminates has dimension $1+r+\cdots+r^{d-1}$ over $\mathbf{F}_q$ and then proving that for an appropriate specialization over $\mathbf{F}_q(t)$ the dimension remains the same.

Proof (Theorem 3.16): To begin, define a set of algebraically independent (over $\mathbf{F}_q$) variables, $\{\xi_{\alpha,\beta,i,j,k}\}$ for $k = 1, \dots, r$, $1 \leq j < i \leq d$ and $1 \leq \alpha \leq n_i$ and $1 \leq \beta \leq n_j$. For i, j, k with $1 \leq k \leq r$ and $1 \leq j < i \leq d$ define block matrices of these independent variables $\tilde{X}_k = (A_{i,j}^{(k)})$ according to

$$\begin{cases} A_{ij}^{(k)} = 0 & \text{for } i < j, \\ A_{ii}^{(k)} = I_{n_i}, & \\ A_{ij}^{(k)} = (\xi_{\alpha,\beta,i,j,k}) & \text{for } i > j,\ 1 \leq \alpha \leq n_i,\ 1 \leq \beta \leq n_j. \end{cases} \tag{15}$$

Thus, each $\tilde{X}_k$ has a decomposition

$$\tilde{X}_k = I_n + \tilde{N}_k. \tag{16}$$

The proof depends on the following two claims whose proofs we postpone.

CLAIM 1: *The products $\tilde{N}_{v_1} \cdots \tilde{N}_{v_m}$ for $1 \leq m \leq d-1, 1 \leq v_1, \dots, v_k \leq r$ and I_n are linearly independent over $\mathbf{F}_q$.*

CLAIM 2: *The enveloping algebra $\mathrm{env}_{\mathbf{F}_q}(\{\tilde{X}_1, \dots, \tilde{X}_r\})$ is precisely the vector space spanned by the matrices of Claim 1.*

Assuming Claims 1 and 2, we see that, from Lemma 3.17,

$$\dim_{\mathbf{q}}(\mathrm{env}_{\mathbf{F}_q}(\{\tilde{X}_1,\ldots,\tilde{X}_r\})) = 1 + r + \cdots + r^{d-1} := e(r).$$

The goal now is to find an appropriate specialization over $\mathbf{F}_q(t)$.

To this end, let $\{\tilde{Y}_1,\ldots,\tilde{Y}_{e(r)}\}$ be a basis of $\mathrm{env}_{\mathbf{F}_q}(\{\tilde{X}_1,\ldots,\tilde{X}_r\})$. We want to choose $f_{\alpha,\beta,i,j,k} \in \mathbf{F}_q(t)$ such that after substituting $f_{\alpha,\beta,i,j,k}$ for $\xi_{\alpha,\beta,i,j,k}$ in $\{\tilde{Y}_1,\ldots,\tilde{Y}_{e(r)}\}$ and $\{\tilde{X}_1,\ldots,\tilde{X}_r\}$ (thereby obtaining matrices $\{Y_1,\ldots,Y_{e(r)}\}$ and $\{X_1,\ldots,X_r\}$ in $M(n,\mathbf{F}_q(t))$), we are able to maintain the independence of the Y_i over $\mathbf{F}_q$. It is easy to see that the Y_i will span the algebra generated by the X_i over $\mathbf{F}_q$.

For this, let $\mathbf{a} = (a_1,\ldots,a_{e(r)}) \in {\mathbf{F}_q}^{e(r)}$ be any nonzero vector. Then the matrix

$$\tilde{Y}_{\mathbf{a}} = \sum_{i=1}^{e(r)} a_i\tilde{Y}_i \neq 0.$$

Consequently, at least one entry of $\tilde{Y}_{\mathbf{a}}$ is nonzero. Choose one and call it $F_{\mathbf{a}}$. Then $F_{\mathbf{a}}$ is a nonzero polynomial in the indeterminates $\xi_{\alpha,\beta,i,j,k}$.

Observation: If after substitution of some choice of $f_{\alpha,\beta,i,j,k}$ for the $\xi_{\alpha,\beta,i,j,k}$ the polynomial $F_{\mathbf{a}}$ is still nonzero, then it will be the case that

$$\sum_{i=1}^{e(r)} a_iY_i \neq 0.$$

By the Observation, finding a specialization of the $\xi_{\alpha,\beta,i,j,k}$ that leaves the Y_i linearly independent over $\mathbf{F}_q$ is then reduced to finding a set of polynomials $\{f_{\alpha,\beta,i,j,k}\}$ such that

$$F(f_{\alpha,\beta,i,j,k}) := \prod_{\substack{\mathbf{a}\in{\mathbf{F}_q}^{e(r)}\\ \mathbf{a}\neq 0}} F_{\mathbf{a}}(f_{\alpha,\beta,i,j,k}) \neq 0. \tag{17}$$

Since $\mathbf{F}_q[t]$ is infinite, such set of polynomials exists. ∎

To complete the proof of Theorem 3.16 we now prove Claims 1 and 2.

Proof (of Claim 1): Consider the i,j entry of an arbitrary product $\tilde{N}_{v_1}\cdots\tilde{N}_{v_m}$ with $1 \leq m \leq d-1$, $1 \leq v_1,\ldots,v_k \leq r$. If we let x^c_{ab} denote the a,b entry of $\tilde{N}_{v_c}$ then

$$\begin{aligned}(\tilde{N}_{v_1}\cdots\tilde{N}_{v_m})_{ab} &= \sum_{a_1}\sum_{a_2}\cdots\sum_{a_m} x^{v_1}_{aa_1}x^{v_2}_{a_1a_2}\cdots x^{v_m}_{a_{m-1}b}\\ &= \sum_{a<a_1<a_2<\cdots<a_m} x^{v_1}_{aa_1}x^{v_2}_{a_1a_2}\cdots x^{v_m}_{a_{m-1}b}\end{aligned} \tag{18}$$

where the second line follows from the fact that the $\tilde{N}_{v_i}$ are all strictly upper triangular.

Since a different choice of m would create monomials of different degree, it is clear that the various monomials for distinct m are linearly independent. Similarly, the expression (18) shows that for fixed m, distinct choices of vectors of indices $(v_1, \ldots, v_m)$ are also independent. ∎

Proof (of Claim 2): It is enough to show that $E = \mathrm{env}_{\mathbf{F}_q}(\{\tilde{X}_1, \ldots, \tilde{X}_r\})$ is contained in the vector space spanned by the monomials in Claim 1 as the reverse inclusion is clear. For this we need only show that each $\tilde{X}_k^{-1}$ is in the space.

Lemma 3.17 implies that $N_k^d = 0$. Since $\tilde{X}_k^{-1} = I - \tilde{N}_k + \tilde{N}_k^2 + \cdots + (-1)^i \tilde{N}_k^i + \cdots$, it is in that space. ∎

3.4 Algorithms.

3.4.1 Algorithm 1. Theorem 3.3 yields immediately a simple exponential (in n) algorithm for deciding finiteness for $\mathrm{GL}(n, \mathbf{F}_q(t))$.

That is, given the generators $\{S_1, \ldots, S_r\}$, attempt to construct a basis for $\mathrm{env}_{\mathbf{F}_q}(G)$. In a given round, we test if the products of the current independent set by the generators contain any new independent elements. Thus any given round takes at most $O(rm \cdot (n^2)^2)$ operations where m is the current number of independent elements. By Theorem 3.3 in at most $\frac{1}{r}(r+1)^n + 1$ rounds we will either have found a complete basis for $\mathrm{env}_{\mathbf{F}_q}(G)$ or generated too many independent elements. We record this simple algorithm here.

Theorem 3.18: *In at most $O(r \cdot (\frac{1}{r}(r+1)^n)^2 n^4) = O(n^4 \frac{1}{r}(r+1)^{2n})$ operations we can decide finiteness for a subgroup of $\mathrm{GL}(n, \mathbf{F}_q(t))$ generated by r elements.*

3.4.2 Algorithm 2. As a next step towards giving an efficient algorithm (deterministic or randomized) we outline another exponential algorithm. Although still exponential we believe that this different framework may be more amenable to improvement.

Theorem 3.6 indicates a possibly fruitful approach towards deciding finiteness. Letting V_0 denote the vector space $\mathbf{F}_q(t)^n$, Theorem 3.6 indicates that we should be searching for a flag of V_0,

$$V_0 > V_1 > \cdots > V_d > V_{d+1} = \{0\},$$

with subspaces V_i such that the successive quotients V_i/V_{i+1} are (1) G-invariant and (2) such that the induced action of G on V_i/V_{i+1} gives a representation of G

in $\mathrm{GL}(n_i, \mathbf{F}_q)$ where $n_i = \dim(V_i/V_{i+1})$. By Theorem 3.6 it is enough to check (1) and (2) on the generators.

Direct implementation of this idea still seems to give an exponential upper bound.

THEOREM 3.19: *Let $S = \{S_1, \ldots, S_r\} \subset \mathrm{GL}(n, \mathbf{F}_q(t))$. Then in at most $O(n^4 \cdot [(\frac{1}{r}(r+1)^n)^8])$ operations we can decide finiteness for $G = \langle S \rangle$.*

Proof: The idea of the algorithm is as follows: Suppose $\langle S \rangle = G \sim G' \leq T(\mathbf{n}, \mathbf{F}_q(t))$. Then we would have a homomorphism

$$\begin{array}{ccl} G & \longrightarrow & \mathrm{GL}(n_1, q) \times \cdots \times \mathrm{GL}(n_d, q) = D(\mathbf{n}), \\ A & \mapsto & \overline{A} \end{array}$$

given by projection onto the block-diagonal subgroup.

Suppose $\mathcal{A} = \{A_1, \ldots, A_m\} \subset G$ were such that it could be guaranteed that $\overline{\mathcal{A}} = \{\overline{A_1}, \ldots, \overline{A_m}\}$ span $\mathrm{env}_{\mathbf{F}_q}(\overline{G})$. Then, for every $A \in G$, either (1) $A \in \mathrm{span}_{\mathbf{F}_q}(\mathcal{A})$, or (2) $A - B$ is nilpotent for some $B \in \mathrm{span}_{\mathbf{F}_q}(\mathcal{A})$. Furthermore, if (2) holds, then the kernel of $A - B$ will contain some nonzero invariant subspace W, and V/W will also be G-invariant.

Thus, our method of attack is to attempt to successively apply the above idea until we arrive at a subspace W which is invariant and supports a representation over F_q. If G is finite this will be possible. This subspace W will then serve as V_d. Having done this, we then apply the algorithm to the quotient V/V_d and so on, ultimately arriving at a change of basis for G to a subgroup of $T(\mathbf{n}, \mathbf{F}_q(t))$. If G is infinite, at some point this algorithm will fail.

As usual, we make the inductive definition of $S^1 = S$ and for $k > 1$,

$$S^k = \bigcup_{A \in S} S^{k-1}A.$$

Since $D(\mathbf{n})$ can have at most dimension n^2 over $\mathbf{F}_q$ we have the following lemma.

LEMMA 3.20: *Let G be finite, with all notation as above. Then*

$$\mathrm{env}_{\mathbf{F}_q}(\overline{S^{n^2}}) = \mathrm{env}_{\mathbf{F}_q}(\overline{G}).$$

Now, suppose that $\mathcal{A} = \{A_1, \ldots, A_m\}$ is an orthogonal basis for

$$X = \mathrm{span}_{\mathbf{F}_q}(S^{n^2})$$

over $\mathbf{F}_q$. Notice that by Theorem 3.3 it will take at most $n^4 \cdot \frac{1}{r}(r+1)^n$ operations to compute $\mathcal{A}$. We can and do assume that $A_i \in G$ for all i. If each of the products A_iS_j is in the span of $\mathcal{A}$, then $\mathrm{env}_{\mathbf{F}_q}(G)$ is finite dimensional and G is finite. Otherwise, some product of this form is not in $\mathrm{span}_{\mathbf{F}_q}(\mathcal{A})$. Let A be such an element.

CLAIM: *Suppose G is finite. Let A and $\mathcal{A}$ be as above. If G is finite then we can compute elements $\alpha_i \in \mathbf{F}_q$ such that*

$$A' = A - \sum_{i=1}^{m} \alpha_i A_i$$

is nilpotent and nonzero.

Proof: Consider the $d \times d$ matrix T with i,j entry given by $\mathrm{trace}(A_iA_j)$. Note that if G is finite then T is defined over $\mathbf{F}_q$. At most $nm^2 < n\left(\frac{1}{r}(r+1)^n\right)^2$ operations are needed to form T. Also, $\mathrm{trace}(A_iA_j) = \mathrm{trace}(\overline{A_iA_j})$ and as a bilinear map from $\overline{X} \times \overline{X} \longrightarrow \mathbf{F}_q$ it is nondegenerate.

Consequently, we can assume that after reordering, the first $k \leq m$ rows of T are a basis for the span over $\mathbf{F}_q$ of all the rows of T. At most m^3 operations are required here.

Consider the new row vector v_A with i^{th} entry given by $\mathrm{trace}(AA_i)$. Since $\overline{A} \in \mathrm{span}_{\mathbf{F}_q}(\mathcal{A})$, it must be that v_A is in the span of the first k rows of T so that there exist $\alpha_i \in \mathbf{F}_q$ such that

$$\mathrm{trace}(AA_j) = \sum_{i=1}^{k} \alpha_i \, \mathrm{trace}(A_iA_j).$$

But this then implies that $\mathrm{trace}((A - \sum_i \alpha_i A_i)A_j) = 0$ for all A_j. Since trace is nondegenerate, this can only mean that $\overline{A - \sum_i \alpha_i A_i} = 0$. But since the A_i span the subalgebra of the block-diagonal entries, this must mean that $A' = A - \sum_i \alpha_i A_i$ is equivalent to some matrix contained in the span of the strictly upper triangular block of $T(\mathbf{n}, \mathbf{F}_q(t))$, and thus, if nonzero, must be nilpotent.

Let $W' = \mathrm{kernel}(A') < \mathbf{F}_q(t)^n$. Then $W = W' \cap S_1W' \cap \cdots \cap S_rW'$ will be G-invariant and nonzero, assuming G is finite. We can then iterate the above on W.

If G is infinite, it will be detected at one of several places:

(1) More than $\frac{1}{r}(r+1)^n$ independent elements will be generated to span S^{n^2}, contradicting Theorem 3.3.

(2) For some i,j, $\mathrm{trace}(A_iA_j) \notin \mathbf{F}_q$.

(3) $W = 0$.

If G is finite then the above procedure will need to be executed at most n^2 times. This yields an upper bound of on the total number of operations required of

$$O\left(n^2 \cdot \left[n\left(\frac{1}{r}(r+1)^n\right)^3\right] \cdot \left[n\left(\frac{1}{r}(r+1)^n\right)^2\right] \cdot \left[\left(\frac{1}{r}(r+1)^n\right)^3\right]\right)$$
$$\leq O\left(n^4 \cdot \left[\left(\frac{1}{r}(r+1)^n\right)^8\right]\right). \blacksquare$$

3.4.3 A practical approach. The exponential upper bounds of the last section suggest that alternative techniques should be sought for implementation. In this section we suggest a simple randomized algorithm whose motivation owes much to the approach taken by Parker's "Meat-Axe", an algorithm which decomposes modular representations of finite groups. We are at present unable to give a proof of reliability here, appealing only to the success of the Meat-Axe as an indication that this idea may prove useful for implementation.

The main tool we use is a result of S. P. Norton, which is also a theoretical lynchpin in the Meat-Axe.

THEOREM (due to S. P. Norton, cf. [16], Section 5): *Let* $\mathbf{F}$ *denote any field and* $\mathcal{S} = \{S_1, \dots, S_r\} \subset M_n(F)$. *Then for any* $B \in \mathrm{env}_{\mathbf{F}}(\mathcal{S})$, *at least one of the following must hold:*

(1) *B is non-singular;*

(2) *At least one non-zero null vector of B lies in a proper subspace invariant under $\mathcal{S}$;*

(3) *Every non-zero null vector of B^T lies in a proper subspace invariant under $\mathcal{S}^T = \{S_1^T, \dots, S_r^T\}$;*

(4) *There is no proper subspace invariant under $\mathcal{S}$.*

Thus, let $\mathcal{S} = \{S_1, \dots, S_n\} \subset \mathrm{GL}(n, \mathbf{F}_q(t))$. Norton's Theorem indicates the following algorithm for deciding finiteness for $\mathcal{S}$.

Randomized Algorithm.

STEP 1: As in the description of Algorithm 1 (cf. Section 3.3), attempt to generate $n^2 + 1$ independent elements over $\mathbf{q}$. If this is not possible (this can be determined in polynomial time), then $\langle \mathcal{S} \rangle$ is finite. Otherwise, proceed to Step 2.

STEP 2: Generate a singular element $B \in \mathrm{env}_{\mathbf{q}(t)}(\mathcal{S})$. Check if the translates of one of its null-vectors generate an invariant subspace. If one does, then perform a change of basis, thereby simultaneously rewriting the generators in some block upper triangular form,

$$S_i \sim \begin{pmatrix} A^i_{1,1} & * \\ 0 & A^i_{2,2} \end{pmatrix},$$

and $A^i_{j,j}$ is $d_j \times d_j$. Now return to Step 1, successively using as input the sets $\{A^1_{j,j}, \ldots, A^r_{j,j}\}$.

If no null-vector generates a nontrivial invariant space, then proceed to Step 3.

STEP 3: Take any non-zero null vector of B^T. If this does not generate a nontrivial invariant subspace for $\mathcal{S}^T$, then G is infinite. Otherwise, we may now find a change of basis given a block upper triangularization of the a group isomorphic to the group generated by $\mathcal{S}^T$. Note that the "transpose group" is finite if and only if the original group is finite. Return now to Step 1 with the blocks for the transposed group and continue.

The difficulty with this algorithm lies in Step 2. If we could guarantee that B has rank $n-1$, then up to scalar multiples there would be a unique null vector to test. Otherwise there are an infinite number. Thus, it is here that we would have to apply randomization. We would construct B in some randomized fashion. Parker points out that almost immediately elements of rank $n-1$ are found. If in fact elements of rank $n-1$ are not constructed, further randomization could then be applied and a null vector could be chosen randomly. The hope again is that (assuming invariant subspaces exist) with high probability a vector is found generating an invariant subspace.

References

[1] L. Babai, R. Beals and D. Rockmore, *Deciding finiteness for matrix groups in deterministic polynomial time*, in *Proc. ISSAC '93*, ACM Press, New York, 1993, pp. 117–126.

[2] R. Beals, *Algorithms for matrix groups and the Tits alternative*, in *Proceedings of 36th Annual Symposium on Foundations of Computer Science*, ACM Press, New York, 1995, pp. 593–602.

[3] R. Beals and L. Babai, *Las Vegas algorithms for matrix groups*, in *Proceedings of 34th Annual Symposium on Foundations of Computer Science*, ACM Press, New York, 1993, pp. 427–436.

[4] K. Friedl and L. Ronyai, *Polynomial time solutions of some problems in computational algebra*, in *Proc.* 17^{th} *ACM STOC*, ACM Press, New York, 1985, pp. 153–162.

[5] C. W. Curtis and I. Reiner, *Representation Theory of Finite Groups and Associative Algebras*, Wiley, New York, 1966.

[6] L. E. Dickson, *Algebras and Their Arithmetics*, University of Chicago Press, Chicago, 1923.

[7] W. Feit, *The orders of finite general linear groups*, preprint, 1996.

[8] L. Finkelstein and W. Kantor (eds.), *Groups and Computation, I*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Volume 11, American Mathematical Society, Providence, RI, 1993.

[9] L. Finkelstein and W. Kantor (eds.), *Groups and Computation, II*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Volume 28, American Mathematical Society, Providence, RI, 1997.

[10] S. Friedland, *Discrete groups of unitary isometries and balls in hyperbolic manifolds*, in *Proceedings of the Fourth Conference of the International Linear Algebra Society (Rotterdam, 1994)*, Linear Algebra and its Applications **241/243** (1996), 305–341.

[11] D. Gorenstein, *The Classification of Finite Simple Groups.*, Vol. 1, Plenum Press, New York, 1982.

[12] N. Jacobson, *Basic Algebra I*, Freeman and Co., San Francisco, 1974.

[13] P. Neumann and C. Praeger, *A recognition algorithm for special linear groups*, Proceedings of the London Mathematical Society **65** (1993), 555–603.

[14] A. Niemayer and C. E. Praeger, *A recognition algorithms for classical groups over finite fields*, Proceedings of the London Mathematical Society, to appear.

[15] A. Niemayer and C. E. Praeger, *Implementing a recognition algorithm for classical groups, in Groups and Computation, II*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Volume 28, American Mathematical Society, Providence, RI, 1997, pp. 273–296.

[16] R. A. Parker, *The computer calculation of modular characters (the Meat-Axe)*, in *Computational Group Theory* (M. Atkinson, ed.), Academic Press, London, 1984, pp. 267–274.

[17] C. Soulé, *Chevalley groups over polynomial rings*, in *Homological Group Theory* (C. T. C. Wall, ed.), London Mathematical Society Lecture Note Series **36**, Cambridge University Press, Cambridge, 1979, pp. 359–367.

[18] K.-S. Tan and D. Rockmore, *Computation of L-series for elliptic curves over function fields*, Journal für die reine und angewandte Mathematik **424** (1992), 107–135.

[19] A. Weil, *On the analogue of the modular group in characteristic p*, in *Functional Analysis and Related Fields*, Proceedings of a Conference in Honor of M. Stone, University of Chicago, Springer-Verlag, Berlin, 1968.