Determination Problems for Orbit Closures and Matrix Groups

RIDA AIT EL MANSSOUR, University of Oxford, United Kingdom
GEORGE KENISON, Liverpool John Moores University, United Kingdom
MAHSA SHIRMOHAMMADI, CNRS - IRIF, France
ANTON VARONKA, TU Wien, Austria
JAMES WORRELL, University of Oxford, United Kingdom

Computational problems concerning the orbit of a point under the action of a matrix group occur throughout computer science, including in program analysis, complexity theory, quantum computation, and automata theory. In many cases the focus extends beyond orbits proper to orbit closures under a suitable topology. Typically one starts from a group and a set of points and asks questions about the orbit closure of the set under the action of the group, e.g., whether two given orbit closures intersect.

In this paper we consider a collection of what we call determination problems concerning matrix groups and orbit closures. These problems begin with a given variety and seek to understand whether and how it arises either as an algebraic matrix group or as an orbit closure. The *how* question asks whether the underlying group is s-generated, meaning it is topologically generated by s matrices for a given number s. Among other applications, problems of this type have recently been studied in the context of synthesising loops subject to certain specified invariants on program variables.

Our main result is a polynomial-space procedure that inputs a variety and a number s and determines whether the given variety arises as an orbit closure of a point under an s-generated commutative algebraic matrix group. The main tools in our approach are structural properties of commutative algebraic matrix groups and module theory. We leave open the question of determining whether a variety is an orbit closure of a point under an s-generated algebraic matrix group (without the requirement of commutativity).

CCS Concepts: • **Theory of computation → Invariants**; **Logic and verification**; • **Computing methodologies → Algebraic algorithms**.

Additional Key Words and Phrases: Algebraic Loop Invariant, Zariski Closure, Polynomial Space, Algebraic Reasoning, Program Synthesis.

1 Introduction

1.1 Orbits and their Closures

The general linear group $\mathrm{GL}_d(\mathbb{F})$ is the group of all $d \times d$ invertible matrices with entries in a field $\mathbb{F}$. Let G be a subgroup of $\mathrm{GL}_d(\mathbb{F})$, acting on $\mathbb{F}^d$ via left multiplication, that is, $g \in G$ maps $\boldsymbol{v} \in \mathbb{F}^d$ to $g \cdot \boldsymbol{v}$. The *orbit* of $\boldsymbol{v} \in \mathbb{F}^d$ under G is the set $G \cdot \boldsymbol{v} \coloneqq \{g \cdot \boldsymbol{v} \mid g \in G\}$. The computational study of orbits of subgroups of $\mathrm{GL}_d(\mathbb{F})$ stretches back many decades. One of the most fundamental problems is determining whether a given pair of vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{F}^d$ lie in the same orbit under the action of a finitely generated subgroup G of $\mathrm{GL}_d(\mathbb{F})$, that is, whether some element $g \in G$ maps $\boldsymbol{v}$ to $\boldsymbol{u}$. Over the field $\mathbb{Q}$ this problem is undecidable in general, but it is decidable if G is commutative, and even in polynomial time if G is generated by a single matrix [2, 27].

For many applications it makes sense to study *orbit closures* in place of orbits.

Orbit closures can be seen as an instance of abstract semantics in program analysis, where one over-approximates the set of reachable program states by a set from a particular abstract domain (such as intervals, octogons, or polyhedra). In the case at hand, the abstract domain is the collection

Authors' Contact Information: Rida Ait El Manssour, University of Oxford, Oxford, United Kingdom, rida.aitelmanssour@cs.ox.ac.uk; George Kenison, Liverpool John Moores University, Liverpool, United Kingdom, g.j.kenison@ljmu.ac.uk; Mahsa Shirmohammadi, CNRS - IRIF, Paris, France, mahsa@irif.fr; Anton Varonka, TU Wien, Vienna, Austria, anton.varonka@tuwien.ac.at; James Worrell, University of Oxford, Oxford, United Kingdom, jbw@cs.ox.ac.uk.

of *algebraic subsets* of $\mathbb{F}^d$, where $S \subseteq \mathbb{F}^d$ is algebraic if it arises as the set of common zeros of a collection of polynomials in $\mathbb{F}[x_1, \ldots, x_d]$. One can thus think of algebraic sets as surfaces in $\mathbb{F}^d$.

The algebraic subsets of $\mathbb{F}^d$ form the closed sets of a topology, called the *Zariski topology*. Formally, the *orbit closure* of $\boldsymbol{v} \in \mathbb{F}^d$ under the action of a subgroup G of $\mathrm{GL}_d(\mathbb{F})$ is the closure $\overline{G \cdot \boldsymbol{v}}$ of the orbit with respect to the Zariski topology, that is, $\overline{G \cdot \boldsymbol{v}}$ is the smallest algebraic superset of the orbit $G \cdot \boldsymbol{v}$. When $\mathbb{F}$ is the field of complex numbers, the Zariski topology is coarser than the familiar Euclidean topology, and thus the Zariski closure contains the Euclidean closure. However the Zariski and Euclidean closures of an orbit coincide when G is a linear algebraic group (that is, when G is both a group and algebraic set, see Section 2).

The orbit closure of a point $\boldsymbol{w} \in \mathbb{F}^d$ is determined by the set of polynomial relationships that hold on every point in the orbit of $\boldsymbol{w}$. For instance, consider the following matrix M and the vector $\boldsymbol{w}$:

$$M := \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \qquad \boldsymbol{w} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

The orbit of $\boldsymbol{w}$ under the group generated by M consists of the vectors $M^n \cdot \boldsymbol{w} = (F_{n+1}, F_n)^\mathsf{T}$, where F_n is the nth Fibonacci number for $n \in \mathbb{Z}$. Based on the identity $(F_{n+1}^2 - F_{n+1}F_n - F_n^2)^2 = 1$, which holds for all $n \in \mathbb{Z}$, the orbit closure is the curve $\{(x, y) : (x^2 - xy - y^2)^2 = 1\}$.

Orbit closure problems arise across many areas of computer science, including complexity theory, program analysis, quantum computation and automata theory [8, 10, 14, 15, 22, 23, 30, 32, 37, 38]. Among many computational questions studied in this context are orbit-closure containment and intersection: the former asks whether a vector $\boldsymbol{u} \in \mathbb{F}^d$ is contained in the orbit closure $\overline{G \cdot \boldsymbol{v}}$ of another vector $\boldsymbol{v}$, and the latter asks whether two such closures intersect. A striking application arises in geometric complexity theory, where the **VP** = **VNP** problem[1] has been studied in terms of orbit-closure containment with respect to the action of general linear group $\mathrm{GL}_d(\overline{\mathbb{Q}})$ on polynomial rings [8, 10].

In certain applications, such as non-convex optimisation problems, non-commutative rational identity testing, and graph isomorphism [5, 9, 16, 17], one considers the orbit closure of a vector $\boldsymbol{v}$ under the action of an explicitly given algebraic group G, which is specified as the set of common zeros of a finite collection of polynomials. In other contexts, such as quantum computing and program analysis [1, 15, 22], the goal is to compute the orbit closure of a group G that is *implicitly* presented via a finite set of topological generators. More precisely, the group is expressed as $G = \overline{\langle M_1, \ldots, M_s \rangle}$ with the matrices M_i given as input. We refer to these two settings as *explicit* and *implicit* presentations of the orbit-closure problems, respectively.

The explicit orbit-closure containment and intersection problems[2] can be directly formulated as existential formulas of first-order logic over the base field $\mathbb{F}$ [5]. In the other direction, it was shown in [5] that orbit containment over $\mathbb{R}$ is polynomial-time equivalent to the existential theory of the reals, and it is **NP**-hard over $\overline{\mathbb{Q}}$. Further applications of explicit orbit-closure problems have been identified in [9], particularly in combinatorial optimisation and dynamical systems, under the assumption that the underlying group is commutative. The complexity of these problems is fully resolved in the case of a key subclass of commutative groups (namely, tori) but remains open for general commutative groups [9].

In the implicit orbit-closure problem, the main challenge in computing an orbit closure lies in computing a set of polynomials whose zero set is the Zariski closure of the group in question. For matrices over $\mathbb{Q}$ an algorithm for this task was given in [15]. This algorithm is not known to be elementary [35, Appendix C]; an elementary procedure albeit of severalfold exponential time was

[1] The **VP**=**VNP** problem is the algebraic variant of **P**=**NP** problem.

[2] In [5], the problems are simply called orbit-closure containment and intersection.

recently provided in [35]. Recent progress on computing multiplicative relations among algebraic numbers [12] shows that for the special case of cyclic groups, the approach in [1, Theorem 1.1], which originally results in a polynomial-space algorithm, gives a polynomial-time upper bound for the implicit orbit-closure computation of cyclic groups (and cyclic semigroups). It remains a challenging open problem to close the complexity gap for the explicit orbit-closure problems in the general setting.

The study of implicit orbit-closure problems in quantum computation and automata theory have led to the resolution of the equivalence problem for deterministic top-down tree-to-string transducers [38] and the threshold problem for quantum automata [15]. Orbit closures feature prominently in program analysis when one wants to automatically compute polynomial invariants of certain classes of loop programs [14, 23, 30, 32, 37].

1.2 Determination Problems

In this paper, we investigate a series of determination problems related to groups and their orbit closures. These problems start with a given algebraic set (commonly referred to as a variety) and examine whether it can be realized as a linear algebraic group or as an orbit closure, with the constraint that the underlying group be topologically s-generated. Here we define an algebraic group $G \leq \mathrm{GL}_d(\overline{\mathbb{Q}})$ to be *topologically s-generated* if there is a set $S \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ of matrices with cardinality s such that $G = \overline{\langle S \rangle}$. By Proposition 6, algebraic groups are always topologically generated by a finite set.

In this context, determining whether a variety $Z \subseteq \overline{\mathbb{Q}}^d$ arises as an orbit closure under the action of G is, in principle, straightforward. Let $\mathrm{Sym}(Z)$ denote the subgroup of $\mathrm{GL}_d(\overline{\mathbb{Q}})$ consisting of all matrices that leave Z invariant, defined as $\mathrm{Sym}(Z) := \{A \in \mathrm{GL}_d(\overline{\mathbb{Q}}) : A(Z) = Z\}$. Any group G having Z as an orbit-closure is necessarily a subgroup of $\mathrm{Sym}(Z)$, and so we may assume without loss of generality that G is $\mathrm{Sym}(Z)$. But $\mathrm{Sym}(Z)$ is definable in first-order logic over $\overline{\mathbb{Q}}$ and hence the orbit closure of any point of Z under $\mathrm{Sym}(Z)$ is definable over the real closed field $\overline{\mathbb{Q}} \cap \mathbb{R}$.[3] Hence the question of whether Z arises as the orbit closure of a point under $\mathrm{Sym}(Z)$ reduces to the decision problem for the theory of real closed fields. It further holds by Proposition 6 that $\mathrm{Sym}(Z)$ has a finitely generated subgroup whose Zariski closure is $\mathrm{Sym}(Z)$ itself. Whence Z is the orbit closure of a point under some finitely generated matrix group if and only if it is the orbit closure of a point under $\mathrm{Sym}(Z)$. However, it is more challenging to determine whether a given variety is the orbit closure of a topologically s-generated group than simply determining whether it is an orbit closure.

Our main determination problems are as follows:

- The **Group Determination** problem asks, given $s \in \mathbb{N}$ and a family of m polynomials in $\mathbb{Q}[\{x_{i,j}\}_{1 \leq i,j \leq d}]$, each of total degree at most b, to determine whether their zero locus $Z \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ is an s-generated algebraic matrix group.
- The **Orbit-closure Determination** problem asks, given $s \in \mathbb{N}$ and a family of m polynomials in $\mathbb{Q}[\{x_i\}_{1 \leq i \leq d}]$, each of total degree at most b, determine whether their zero locus $Z \subseteq \overline{\mathbb{Q}}^d$ is the orbit closure of some point $v \in \overline{\mathbb{Q}}^d$ under the action of an s-generated algebraic matrix group.

In our complexity analysis we refer to the tuple (s, d, m, b) as the parameters of the problem instances. By Propositions 9 and 10 and Remark 13, the minimum number s of topological generators for the groups we study, commutative algebraic groups, is upper bounded by d.

[3] The Zariski closure of every first-order definable set over $\overline{\mathbb{Q}}$ is equal to its Euclidean closure and is thus first-order definable over $\overline{\mathbb{Q}} \cap \mathbb{R}$ after identifying each element of $\overline{\mathbb{Q}}$ as the pair of its real and imaginary parts.

This paper focuses on addressing the complexity of the determination problems for commutative algebraic matrix groups. Our main result, Theorem 2, is a polynomial-space procedure for both problems, based on reductions to the decision problem for formulas in a fragment of the first-order theory of algebraically closed or real closed fields of characteristic zero.

Example 1. Let $Z \subseteq \overline{\mathbb{Q}}^4$ be the set of common zeros of the two polynomials $Q_1 := x_2^2 - x_1 - x_4$ and $Q_2 := -2x_4x_2 - 2x_3^2 - \frac{1}{5}x_2x_3$. The task underlying this instance of the orbit-closure determination problem (as above) is to determine whether Z is the orbit closure of some point $v \in \overline{\mathbb{Q}}^4$ under the action of an algebraic matrix group with a prescribed number of generators.

Our nondeterministic procedure in Theorem 2 shows that Z is the orbit closure of a 1-generated algebraic matrix group, i.e., $Z = \overline{\langle M \rangle \cdot \boldsymbol{v}}$. The following matrix M and vector $\boldsymbol{v}$ pair witness that Z is the orbit-closure of a one-generated (cyclic) matrix group:

$$M = \begin{pmatrix} 25 & 0 & -1 & 20 \\ 0 & 5 & 0 & 0 \\ 0 & -\frac{1}{2} & 5 & 0 \\ 0 & 0 & 1 & 5 \end{pmatrix} \quad \text{and} \quad \boldsymbol{v} = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}.$$

An account of the steps taken to produce M and $\boldsymbol{v}$ is given in Example 21. ◀

The extension of our results from the case of commuting matrices to the case of general algebraic matrix groups appears to be challenging. To approach the above version of the orbit determination problem, we rely on the observation that with respect to a convenient basis an orbit closure itself carries the structure of an algebraic matrix group. We then examine structural properties of semisimple and unipotent linear algebraic groups to identify when the orbit, seen as a group, arises as the closure of a commutative group.

1.3 Application in Programming Languages

A research direction closely related to orbit-closure determination is the synthesis of linear loops. While computing the Zariski closure of an orbit is used to compute polynomial invariants of loops, the loop synthesis problem reverses this direction: given an input polynomial invariant, one seeks to construct update rules for loop variables that maintain the invariant.

Suppose a deterministic polynomial loop involves integer variables x, y, z, w, and in each iteration, both w and y are incremented by one. The goal is to determine suitable updates to x and z so that a specified invariant, such as $x^2 - y^2z^2 + z^3 = 0$, remains valid after each iteration, assuming that it holds initially. In other words, we have a partially defined loop

Assert: $x^2 - y^2z^2 + z^3 = 0$
  while $(*)$ **do**
    $x \leftarrow P(x, y, z, w)$;
    $y \leftarrow y + 1$;
    $z \leftarrow Q(x, y, z, w)$;
    $w \leftarrow w + 1$;
  end while

where $*$ denotes a wildcard that nondeterministically evaluates to false or true, and the task is to determine assignments $x \leftarrow P(x, y, z, w)$ and $z \leftarrow Q(x, y, z, w)$ such that $x^2 - y^2z^2 + z^3 = 0$ is a loop invariant. A valid solution in this case is to set $P(x, y, z, w) := y(y^2 - w^2)$ and $Q(x, y, z, w) := y^2 - w^2$. This kind of synthesis task is closely connected to the classical algebraic geometry problem of parametrising varieties defined by polynomial equations [1].

Recent developments have focused on techniques for synthesising deterministic linear loops, often by encoding the synthesis task as a constraint-solving problem over algebraic structures [1,

20, 24, 25, 29]. Geometrically, this amounts to finding an infinite orbit of a cyclic matrix group that remains entirely within the variety defined by the invariant. Some formulations restrict attention to infinite orbits in order to avoid trivial cases where the synthesis reduces to solving a finite system of polynomial equations [20].

Several concrete approaches have been proposed to instantiate these synthesis procedures. One line of work develops constraint-solving methods that generate loops preserving a given polynomial invariant, based on templates supplied by the user [24, 25]. Other approaches have focused on restricted classes of invariants, such as those defined by a single quadratic equation [20] or by ideals generated by pure binomials [29]. In the latter case, the invariant variety is a union of toric varieties, and the synthesis relies on a construction by [18, Proposition 14], showing that for every toric variety V one can construct a rational matrix M such that the Zariski closure of $\{M^n : n \in \mathbb{N}\}$ equals V. Using our terminology, toric varieties are topologically 1-generated.

A recent direction studies a strong variant of loop synthesis [1], where the objective is not merely to keep the orbit within the target variety, but that the orbit closure of the loop variables coincides exactly with the input variety. This corresponds computationally to an alternative formulation of the orbit-closure determination problem studied in this paper, where the goal is to decide whether a given variety arises as the orbit closure of some vector under a matrix group.

The approach in [1] is restricted to cyclic groups generated by a single rational matrix (modelling single-path loops) and assumes that the input includes an explicit bound on the bit-size of the generator. In contrast, the present work solves a broader version of the strong loop synthesis problem, allowing for an arbitrary number s of commuting generators, without requiring any a priori bit-size bound. To relax the bit-bound restriction, we formulate the problem over the field $\overline{\mathbb{Q}}$ of algebraic numbers. This widens the scope from deterministic to nondeterministic loops. For instance, given a variety $Z \subseteq \overline{\mathbb{Q}}^d$, our methods can determine whether Z is the orbit closure of a vector under a 2-generated matrix semigroup. In other words, whether there exist matrices $M_1, M_2 \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and a vector $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$ for which $Z = \overline{\langle M_1, M_2 \rangle \cdot \boldsymbol{v}}$. This problem corresponds to that of synthesising a nondeterministic loop of the form

$\boldsymbol{x} \leftarrow \boldsymbol{v}$;
while $(*)$ **do**
  if $(*)$ **then**
    $\boldsymbol{x} \leftarrow M_1\boldsymbol{x}$;
  else
    $\boldsymbol{x} \leftarrow M_2\boldsymbol{x}$;
  end if
end while

such that the orbit closure of the loop coincides exactly with the input variety Z.

Example 2. Let us revisit Example 1 from the viewpoint of loop synthesis. Recall that $Q_1 = x_2^2 - x_1 - x_4$ and $Q_2 = -2x_4x_2 - 2x_3^2 - \frac{1}{5}x_2x_3$. Consider the following task, *construct a deterministic linear loop for which* $Q_1(\boldsymbol{x}) = 0 \wedge Q_2(\boldsymbol{x}) = 0$ *is an invariant*. This task amounts to determining an initial assignment $\boldsymbol{x} \leftarrow \boldsymbol{v}$ and linear update $\boldsymbol{x} \leftarrow M\boldsymbol{x}$ in the following loop:

$\boldsymbol{x} \leftarrow \boldsymbol{v}$;
while $(*)$ **do**
  $\boldsymbol{x} \leftarrow M\boldsymbol{x}$;
end while

such that $Q_1(\boldsymbol{x}) = 0 \wedge Q_2(\boldsymbol{x}) = 0$ is a loop invariant. In other words, we require that the following Hoare triples are satisfied:

(1) $\{\text{true}\}\ \boldsymbol{x} \leftarrow \boldsymbol{v}\ \{Q_1(\boldsymbol{x}) = 0 \wedge Q_2(\boldsymbol{x}) = 0\}$;
(2) $\{Q_1(\boldsymbol{x}) = 0 \wedge Q_2(\boldsymbol{x}) = 0\}\ \boldsymbol{x} \leftarrow M\boldsymbol{x}\ \{Q_1(\boldsymbol{x}) = 0 \wedge Q_2(\boldsymbol{x}) = 0\}$.

Note that the above synthesis task may admit multiple solutions, including trivial solutions in which M is the identity matrix. However, there is a natural notion of a maximal (or most permissive) solution, which is a loop for which $Q_1(\boldsymbol{x}) = 0 \wedge Q_2(\boldsymbol{x}) = 0$ is the strongest invariant. Geometrically, a solution is maximal if the set Z of common zeros of Q_1 and Q_2 is the Zariski closure of the set $\langle M\rangle \cdot \boldsymbol{v}$ of reachable values of the program variables. A motivation to look for such a solution is that it enables one to establish liveness properties, guaranteeing that certain program configurations can be reached. (A sufficient condition for the orbit $\langle M\rangle \cdot \boldsymbol{v}$ to intersect a target set Y is that the Zariski closure Z meet the interior of Y.) The algorithm we present in Theorem 2 finds a most permissive solution.

The requirement that the linear loop be deterministic (have a single-path) translates, in the algebraic setting, to the requirement that Z be the orbit closure of a 1-generated algebraic matrix group. In this example, we instantiate the loop above with e.g.

$$M = \begin{pmatrix} 25 & 0 & -1 & 20 \\ 0 & 5 & 0 & 0 \\ 0 & -\frac{1}{2} & 5 & 0 \\ 0 & 0 & 1 & 5 \end{pmatrix} \quad \text{and} \quad \boldsymbol{v} = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}$$

and so deduce that $Z = \overline{\langle M\rangle \cdot \boldsymbol{v}}$.

1.4 Overview of the Main Result

We reduce our determination problems to satisfiability problems in fragments of the first-order theory of the algebraically closed or real-closed fields. These theories are formulated in the first-order language of rings, which includes constant symbols 0 and 1, as well as binary function symbols for addition and multiplication. The theory of algebraically closed field is the set of all sentences in this language that are true over $\overline{\mathbb{Q}}$. Similarly, the theory of the real numbers consists of those sentences that are true over $\mathbb{R}$. The following theorem gives a complexity bound on the decision problem for this theory.

Theorem 1 ([3, 4, 11]). *Consider a first-order sentence in the language of rings that mentions m polynomials in d variables, with total degree at most b, and with k quantifier alternations. The truth of such a sentence over $\overline{\mathbb{Q}}$ and $\overline{\mathbb{Q}} \cap \mathbb{R}$ can both be decided in time $(mb)^{d^{2k+2}}$.*

Following [3, Remark 13.10], the truth of first-order sentences, over both $\overline{\mathbb{Q}}$ and $\overline{\mathbb{Q}} \cap \mathbb{R}$, with a fixed number of alternations can be decided in space $(d \log b)^{O(1)}$.

The following theorem is our main contribution:

Theorem 2. *The orbit-closure determination problem for commutative matrices with parameters (s, d, m, b) can be decided in time $(mb)^{\mathrm{poly}(d)}$, and in space bounded by $(d \log b)^{O(1)}$.*

The proof of Theorem 2 shows how, given a variety $Z \subseteq \overline{\mathbb{Q}}^d$, to recover an s-generated group $G \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ and vector $\boldsymbol{v} \in Z$ such that $Z = \overline{G \cdot \boldsymbol{v}}$. The following is a summary of the key elements. We will explain technical terms and expand on each point immediately below:

(1) With respect to a suitable basis, the vector $\boldsymbol{v}$ has all its entries 0 or 1 and hence can be guessed in nondeterministic polynomial time. Crucially, given this form for $\boldsymbol{v}$, the orbit closure Z itself carries a group structure under component-wise multiplication.
(2) The semisimple part G_s of G is characterised up to isomorphism by an additive subgroup Λ of $\mathbb{Z}^k$ for some positive integer k such that the quotient group $\mathbb{Z}^k/\Lambda$ has s generators.

Moreover, the generators of Λ have bit-size bounded polynomially in the description of Z; hence Λ can also be guessed in nondeterministic polynomial time.

(3) The existence of an s-generated unipotent group G_u such that $Z = \overline{G_u \cdot G_s \cdot \boldsymbol{v}}$ can be expressed in first-order logic given descriptions of Z, $\boldsymbol{v}$, and Λ. This relies on the fact that for a unipotent matrix U, the matrix exponential U^n is a matrix of polynomials in n.

We write Id_d for the identity matrix of dimension d. A matrix $M \in \overline{\mathbb{Q}}^{d\times d}$ is *unipotent* if $(M - \mathrm{Id}_d)^d$ is the zero matrix, and *semisimple* if it is diagonalisable over $\overline{\mathbb{Q}}$. Let G be a commutative algebraic group. By Fact 4, the subset G_s of semisimple matrices in G forms an algebraic subgroup; similarly, the set G_u of unipotent matrices in G forms an algebraic subgroup. Moreover, the decomposition $G = G_u \cdot G_s$ holds in the commutative setting.

For a given instance of the orbit-closure determination problem, let Z be the zero locus of the input polynomials. If Z is the orbit closure of a vector $\boldsymbol{v}$ under the action of some group G, by the above, it can be written as $\overline{G_u \cdot G_s \cdot \boldsymbol{v}}$. If both G_s and G_u are s-generated, then by Remark 13, the group G is also s-generated.

A key technical result, Lemma 15, states that for given a commutative matrix group G, generated by matrices $M_1, \ldots, M_s$ and a vector $\boldsymbol{v}$, we can choose a change of basis (given by an invertible matrix P) that brings both the group G and the vector $\boldsymbol{v}$ into a simplified and canonical form. In this new basis:

- the semisimple subgroup becomes diagonal;
- the unipotent subgroup becomes upper unitriangular; and
- $\boldsymbol{v}$ is mapped to a binary vector $T\mathbf{1}$, where $T \in \{0,1\}^{d\times k}$ satisfies $T^{\mathsf{T}}T = \mathrm{Id}_k$.

Our algorithm in Theorem 2 guesses the matrix T. The algorithm proceeds by assuming that $P\boldsymbol{v} := T\mathbf{1}$. Write

$$PG_sP^{-1} = \overline{\langle D_i : 1 \le i \le s\rangle}\,.$$

In this new basis, the orbit-closure of $P\boldsymbol{v}$ under $\langle D_i : 1 \le i \le s\rangle$ forms a linear algebraic group, and becomes the zero set of a pure binomial ideal I. Recall that a pure binomial ideal in the variables $\boldsymbol{x} = (x_1, \ldots, x_d)$ is an ideal generated by polynomials of the from $\boldsymbol{x}^{\boldsymbol{\lambda}} - \boldsymbol{x}^{\boldsymbol{\lambda'}}$, where $\boldsymbol{\lambda}, \boldsymbol{\lambda'} \in \mathbb{N}^d$. Associated with any such ideal is its exponent lattice Λ, a subgroup of $\mathbb{Z}^d$ collecting the vectors $\boldsymbol{\lambda} - \boldsymbol{\lambda'}$. Studying this lattice provides key structural insight into the underlying variety defined by I.

By the requirement on the number of generators, the orbit closure $\overline{\langle D_i : 1 \le i \le s\rangle \cdot P\boldsymbol{v}}$ seen as a linear algebraic group must also have s topological generators. By Proposition 9, the torsion subgroup of $\mathbb{Z}^d/\Lambda$ has at most s generators. A careful analysis in Claim 20, combined with Proposition 8, shows that the degree bound b on the defining polynomials of Z also applies to the generators of the binomial ideal I. The degree bound b allows the algorithm to guess a lattice Λ, generated by vectors with entries bounded in absolute value by b, and such that $\mathbb{Z}^d/\Lambda$ has at most s generators, in line with Proposition 9. For this lattice $\Lambda \subseteq \mathbb{Z}^d$, we define H_Λ as a subgroup of the d-dimensional multiplicative variety, where $\boldsymbol{a} \in H_\Lambda$ if $\boldsymbol{a}^{\boldsymbol{\lambda}} = 1$ for all $\boldsymbol{\lambda} \in \Lambda$.

By Remark 13, there exist unipotent matrices $U_1, \ldots, U_s$ that topologically generate the unipotent subgroup G_u of G. Furthermore, by Proposition 10 the following equality holds:

$$G_u = \left\{ \exp\left(\sum_{i=1}^{s} t_i \log U_i \right) : t_1, \ldots, t_s \in \overline{\mathbb{Q}} \right\}.$$

Since the power series of exp has at most d terms in this case (see Section 2), it follows that the group G_u is definable by a first-order formula with parameters $U_1, \ldots, U_s$.

We are now in a position to encode the given instance of the problem as a first-order sentence in the theory of real-closed fields. An important observation enabling this reduction is that, for linear algebraic groups, the Zariski and Euclidean closures of an orbit coincide (see Fact 5). In our encoding, we rely on this fact to describe the algebraic closure of a constructible set using Euclidean conditions: by requiring that for all $\varepsilon > 0$, there exists a group element constructed by our guesses mapping the image within ε of some vector in the input Z.

More in detail, the orbit-closure verification task is expressed as a first-order sentence with quantifier prefix of the form $\exists^*\forall^*\exists^*$, that is, a block of existential quantifiers followed by a block of universal quantifiers and a final existential quantifier. The outermost existential quantifiers encode the possible choices of the matrices P and $U_1, \ldots, U_s$, while the equality of PZ and

$$\overline{\langle U_i : 1 \leq i \leq s \rangle} \cdot TH_\Lambda$$

is encoded by a $\forall^*\exists$-sentence with parameters P and $U_1, \ldots, U_s$. The algorithm returns "yes", meaning that Z is an orbit closure of a vector $\boldsymbol{v}$ under the action of the group G, if the above sentence is satisfiable over $\mathbb{R}$. By Theorem 1, the truth of such a sentence can be decided in time $(mb)^{\mathrm{poly}(d)}$. Then the overall complexity bound follows from the fact that the number of choices of the lattice Λ and vector $P\boldsymbol{v}$ is at most $(2b)^{d^2+1}$. This concludes our informal overview of the proof of Theorem 2; the detailed proof can be found in Section 4.

Orbit-Closure vs. Group Determination. En route to proving Theorem 2 on orbit-closure determination, we consider a variant—namely group determination.

For group determination we first consider a simpler setting where the input polynomial ideal $I \subseteq \mathbb{Z}[\boldsymbol{x}]$ is a pure binomial ideal. Recall that we can associate a lattice $\Lambda = \{\boldsymbol{\lambda} - \boldsymbol{\lambda'} \in \mathbb{Z}^d : \boldsymbol{x}^{\boldsymbol{\lambda}} - \boldsymbol{x}^{\boldsymbol{\lambda'}} \in I\}$ with I. It is well-known that, if the input ideal I defines a group G then it is necessarily topologically generated by diagonal matrices. If the torsion subgroup of $\mathbb{Z}^d/\Lambda$ is s-generated, then, by Proposition 9, the minimal number of topological generators of G is either s or $s+1$ (depending on the rank of Λ). However, in the setting of orbit-closure determination, this lower bound on the number of generators may no longer hold, as shown by the following example.

Example 3. Let $\Lambda \subseteq \mathbb{Z}^3$ be the associated lattice of the pure binomial ideal $I := \langle x_1^2 - x_3^2, x_2^2 - x_3^2 \rangle$. That is, $\Lambda = \{c_1(2,0,-2)^\mathsf{T} + c_2(0,2,-2)^\mathsf{T} : c_1, c_2 \in \mathbb{Z}\}$. By the fundamental theory of finitely generated abelian groups, the quotient group $\mathbb{Z}^3/\Lambda$ is isomorphic to $(\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z}) \times \mathbb{Z}$. In particular, the torsion subgroup of $\mathbb{Z}^3/\Lambda$ is $(\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})$, a product of two cyclic groups, which is 2-generated.

Here and in the following, let $\Delta(x_1, x_2, x_3) \in \mathrm{GL}_3(\overline{\mathbb{Q}})$ denote a diagonal matrix with x_1, x_2, x_3 (in this order) on the diagonal. Let $G \leq \mathrm{GL}_3(\overline{\mathbb{Q}})$ be the algebraic subgroup of diagonal matrices induced by I, that is,

$$G = \{\Delta(\boldsymbol{a}) : P(\boldsymbol{a}) = 0 \text{ for all } P \in I\}.$$

Observe that G is isomorphic to H_Λ (recall that the elements of the latter are $\boldsymbol{a}$ such that $P(\boldsymbol{a}) = 0$). By Proposition 9, the minimal number of topological generators for G is 2. Indeed, G is topologically generated by the diagonal matrices $\Delta(2,-2,2)$ and $\Delta(-2,2,2)$.

Changing our viewpoint, let us consider I as defining a subvariety V of $\overline{\mathbb{Q}}^3$ rather than a set of diagonal matrices in $\mathrm{GL}_3(\overline{\mathbb{Q}})$. The variety V can be written as the orbit closure of a vector $\boldsymbol{v}$ under the action of the 1-generated (cyclic) group $\overline{\langle M \rangle}$, where

$$M = \begin{pmatrix} 0 & -2 & 0 \\ 2 & 0 & 0 \\ 0 & 0 & 2 \end{pmatrix} \quad \text{and} \quad \boldsymbol{v} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

◀

1.5 Outline and Structure

The remainder of the paper is structured as follows. In Section 2 we give useful notations and definitions (we refer the reader to the Appendix for extended preliminaries). In Section 3 we present procedures for the group determination problem (Propositions 12 and 14). In Section 4 we present procedures for the orbit-closure determination problem (Proposition 18 and Theorem 2). Finally, in Section 5 we present two algorithms that compute a generator in the special case that the input variety is the Zariski closure of a cyclic group.

2 Primer on Algebraic Geometry

In this section we introduce basic notions from algebraic geometry that are needed in the paper. We refer the reader to [13] for further details and examples.

Ideals and Varieties. Recall that an algebraic number is a complex number that is the root of a univariate polynomial with integer coefficients. The collection of all algebraic numbers forms a sub-field of $\mathbb{C}$, which is denoted $\overline{\mathbb{Q}}$. We write $\overline{\mathbb{Q}}[\boldsymbol{x}]$ for the ring of polynomials in the variables $\boldsymbol{x} = (x_1, \ldots, x_d)$ with coefficients in $\overline{\mathbb{Q}}$. A *polynomial ideal* I is an additive subgroup of $\overline{\mathbb{Q}}[\boldsymbol{x}]$ that is closed under multiplication in $\overline{\mathbb{Q}}[\boldsymbol{x}]$. Given a finite collection of polynomials $S = \{P_1, \ldots, P_k\} \subseteq \overline{\mathbb{Q}}[\boldsymbol{x}]$, we denote by $\langle S \rangle$ the ideal *generated by* S:

$$\langle S \rangle \coloneqq \{P_1Q_1 + \cdots + P_kQ_k : Q_1, \ldots, Q_k \in \overline{\mathbb{Q}}[\boldsymbol{x}]\}\,.$$

For example, $\langle x_1 \rangle$ is the ideal of all polynomials that are multiplies of x_1.

An *algebraic set* (or *variety*) is the set of common zeroes of a finite collection of polynomials. We may also refer to an algebraic set as the *zero locus* of these polynomials. By Hilbert's basis theorem every polynomial ideal $I \subseteq \overline{\mathbb{Q}}[\boldsymbol{x}]$ is finitely generated. Thus the zero set of I

$$V(I) \coloneqq \{\boldsymbol{a} \in \overline{\mathbb{Q}}^d : P(\boldsymbol{a}) = 0 \text{ for all } P \in I\}$$

is a variety. As an example, consider an ideal I of the polynomial ring $\overline{\mathbb{Q}}[x_1, x_2]$ generated by polynomials $\{x_1^2, x_2^2\}$. The ideal

$$I = \langle x_1^2, x_2^2 \rangle = \{x_1^2Q_1(x_1, x_2) + x_2^2Q_2(x_1, x_2) : Q_1, Q_2 \in \overline{\mathbb{Q}}[x_1, x_2]\}$$

contains polynomials such as x_1^2, $x_1^2 - \sqrt{2}x_2^2$, and $x_1^3 + x_1^2x_2 + x_1x_2^2 - x_2^3$. The zero set of I is the variety

$$V(I) \coloneqq \{(a_1, a_2) \in \overline{\mathbb{Q}}^2 : P(a_1, a_2) = 0 \text{ for all } P \in I\} = \{(a_1, a_2) \in \overline{\mathbb{Q}}^2 : a_1^2 = 0 \wedge a_2^2 = 0\}.$$

Dually, the set of polynomials that vanish on a set $E \subseteq \overline{\mathbb{Q}}^d$, denoted

$$I(E) \coloneqq \{P \in \overline{\mathbb{Q}}[\boldsymbol{x}] : P(a) = 0 \text{ for all } a \in E\},$$

is an ideal in $\overline{\mathbb{Q}}[\boldsymbol{x}]$ and is referred to as the *vanishing ideal* of E. Resuming our example, we can consider $E = V(I)$ for $I = \langle x_1^2, x_2^2 \rangle$. The vanishing ideal of this variety is

$$\begin{aligned} I(V(I)) &\coloneqq \{P \in \overline{\mathbb{Q}}[x_1, x_2] : P(a_1, a_2) = 0 \text{ for all } (a_1, a_2) \in V(I)\} \\ &= \{P \in \overline{\mathbb{Q}}[x_1, x_2] : P(a_1, a_2) = 0 \text{ for all } (a_1, a_2) \text{ such that } a_1^2 = a_2^2 = 0\} = \langle x_1, x_2 \rangle. \end{aligned}$$

The ideal $I(V(I))$ may, in general, be different from I.

Zariski Topology. A *topology* on a set X is a collection τ of subsets of X, called *closed sets*, that satisfy the following axioms:

(1) $\varnothing, X \in \tau$;
(2) An arbitrary intersection of a collection of sets in τ belongs to τ;
(3) Any finite union of sets in τ belongs to τ.

The *Zariski topology* on $\overline{\mathbb{Q}}^d$ has as its closed sets the varieties in $\overline{\mathbb{Q}}^d$. Given a set $E \subseteq \overline{\mathbb{Q}}^d$, we denote by $\overline{E}$ the *closure* of E in the Zariski topology, that is, the smallest algebraic set that contains E. If the closure $\overline{E}$ coincides with a variety V, we say that E is *dense* in V.

A closed set $A \subseteq \overline{\mathbb{Q}}^d$ is *irreducible* if it cannot be written as the union of two proper closed subsets. A maximal irreducible closed subset of A is called an *irreducible component* of A. By Hilbert's basis theorem every closed set A can be written as a finite union of its irreducible components. For example, the set $A = \{\boldsymbol{a} \in \overline{\mathbb{Q}}^2 : a_1 a_2 = 0\}$ has irreducible components $A_1 = \{\boldsymbol{a} \in \overline{\mathbb{Q}}^2 : a_1 = 0\}$ and $A_2 = \{\boldsymbol{a} \in \overline{\mathbb{Q}}^2 : a_2 = 0\}$. Given a closed set E, denote by $\dim E$ the dimension of the variety E, that is, the maximal length of a strictly decreasing chain of nonempty irreducible subvarieties of E. We define the dimension of an arbitrary set to be the dimension of its closure.

Given an ideal $I \subseteq \overline{\mathbb{Q}}[X]$ over the variables $X = \{x_{i,j}\}_{1 \le i,j \le d}$, and matrix $M \in \overline{\mathbb{Q}}^{d \times d}$, we write $M \cdot I$ for the ideal $\{P(MX) \in \overline{\mathbb{Q}}[X] : P \in I\}$. Clearly, $V(M \cdot I) = \{A \in \overline{\mathbb{Q}}^{d \times d} : MA \in V(I)\}$.

Linear Algebraic Groups. A matrix $M \in \overline{\mathbb{Q}}^{d \times d}$ is *nilpotent* if $M^n = 0$ for some $n \in \mathbb{N}$. It is *unipotent* if $M - \mathrm{Id}_d$ is nilpotent, and *semisimple* if it is diagonalisable over $\overline{\mathbb{Q}}$. The matrix $M \in \overline{\mathbb{Q}}^{d \times d}$ is called *upper triangular* if all entries below the main diagonal are zero. We use the term *upper unitriangular* to refer to an upper triangular matrix whose entries along the main diagonal are all ones.

Write $\mathrm{GL}_d(\overline{\mathbb{Q}})$ for the group of invertible $d \times d$ matrices with entries in $\overline{\mathbb{Q}}$. We identify $\mathrm{GL}_d(\overline{\mathbb{Q}})$ with the variety

$$\left\{(M, y) \in \overline{\mathbb{Q}}^{d \times d} \times \overline{\mathbb{Q}} : \det(M) \cdot y = 1\right\}.$$

Under this identification, matrix multiplication is a polynomial map $\mathrm{GL}_d(\overline{\mathbb{Q}}) \times \mathrm{GL}_d(\overline{\mathbb{Q}}) \to \mathrm{GL}_d(\overline{\mathbb{Q}})$, and, by Cramer's rule, matrix inversion is also a polynomial map $\mathrm{GL}_d(\overline{\mathbb{Q}}) \to \mathrm{GL}_d(\overline{\mathbb{Q}})$. A *linear algebraic group* (or algebraic matrix group) G is a Zariski-closed subgroup of $\mathrm{GL}_d(\overline{\mathbb{Q}})$. In other words, a subgroup $G \le \mathrm{GL}_d(\overline{\mathbb{Q}})$ is algebraic if it can be defined by polynomial equalities. As an example take the subgroup of diagonal matrices in $\mathrm{GL}_d(\overline{\mathbb{Q}})$, defined by

$$\{M = (m_{ij})_{1 \le i,j \le d} \in \mathrm{GL}_d(\overline{\mathbb{Q}}) : \quad m_{ij} = 0 \text{ for all } i \ne j\}.$$

Denote by G_s the subset of semisimple matrices in G, and by G_u the subset of unipotent matrices. The following fact is well-known, see for example [26, Chapter 6].

Fact 4. For a commutative algebraic group $G \le \mathrm{GL}_d(\overline{\mathbb{Q}})$, the algebraic subgroups G_s and G_u form algebraic subgroups of G; moreover, we have the decomposition of G into $G_u \cdot G_s$.

The next fact follows from Chevalley's Theorem; see [7, Chapter AG, Corollary 10.2] and Appendix B for details.

Fact 5. Let $G \le \mathrm{GL}_d(\overline{\mathbb{Q}})$ be an algebraic group and $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$ a vector, the Zariski and Euclidean closures of the orbit $G \cdot \boldsymbol{v}$ coincide.

We say that G is *topologically generated by* $S \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ if G is the smallest Zariski-closed subgroup of $\mathrm{GL}_d(\overline{\mathbb{Q}})$ that contains S, that is, $G = \overline{\langle S \rangle}$. If G is topologically generated by a set with s elements then we say that is G is *s-generated.*

Proposition 6. *Let $G \le \mathrm{GL}_d(\overline{\mathbb{Q}})$ be an algebraic group, then G is topologically generated by a finite set of matrices.*

The proof of Proposition 6 is given in Appendix B.

Lattices and the Multiplicative Group $\mathbb{G}_m^d$. The *rank* of an abelian group Λ is the size of a maximal linearly independent subset [31]. A subgroup $\Lambda \subseteq \mathbb{Z}^d$ is called a *lattice* (and has rank at most d). The set $\mathbb{Z}^d/\Lambda$ contains all cosets of Λ in $\mathbb{Z}^d$, that is, sets of the form $\boldsymbol{g} + \Lambda = \{\boldsymbol{g} + \boldsymbol{v} : \boldsymbol{v} \in \Lambda\}$

where $\boldsymbol{g} \in \mathbb{Z}^d$. Together with an additive operation, the set of cosets defines the *quotient group* $\mathbb{Z}^d/\Lambda$. The *torsion subgroup* of $\mathbb{Z}^d/\Lambda$ is the subgroup of $\mathbb{Z}^d/\Lambda$ consisting of all its elements of finite order.

The d-dimensional multiplicative group over $\overline{\mathbb{Q}}$ is defined as

$$\mathbb{G}_m^d = \mathbb{G}_m^d(\overline{\mathbb{Q}}) \coloneqq \left\{ \boldsymbol{a} \in \overline{\mathbb{Q}}^d : a_1 \cdots a_d \neq 0 \right\}. \tag{1}$$

Here the subscript m stands for *multiplicative*. Evidently, this is a commutative group with respect to pointwise multiplication. We identify $\mathbb{G}_m^d$ with the subgroup of diagonal matrices in $\mathrm{GL}_d(\overline{\mathbb{Q}})$ via the map Δ that sends $(a_1, \ldots, a_d) \in \mathbb{G}_m^d$ to the diagonal matrix $\Delta(a_1, \ldots, a_d) \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ which has $a_1, \ldots, a_d$ (in this order) on the diagonal and zeros elsewhere.

We can now explain two important ingredients of our results. We have so far identified the invertible diagonal matrices with the algebraic group $\mathbb{G}_m^d(\overline{\mathbb{Q}})$. Now, we define the necessary vocabulary to discuss the correspondence of algebraic subgroups of $\mathbb{G}_m^d(\overline{\mathbb{Q}})$ and lattices in $\mathbb{Z}^d$. This comes in handy when we need to decide whether a certain subgroup of diagonal invertible matrices is s-generated, in which we reduce it to a lattice problem.

Given a lattice $\Lambda \subseteq \mathbb{Z}^d$, define

$$H_\Lambda \coloneqq \{\boldsymbol{a} \in \mathbb{G}_m^d : \ \forall \boldsymbol{v} \in \Lambda.\ a_1^{v_1} \cdots a_d^{v_d} = 1\}.$$

The map $\Lambda \mapsto H_\Lambda$ is an isomorphism between lattices and algebraic subgroups of $\mathbb{G}_m^d$. This implies that $\mathbb{G}_m^d$ is topologically generated by any d-tuple $(g_1, \ldots, g_d)$ of multiplicatively independent elements of $\overline{\mathbb{Q}}$.

For variables $\boldsymbol{x} = (x_1, \ldots, x_d)$, write $\boldsymbol{x}^{\boldsymbol{\lambda}}$ for the monomial $\prod_{i=1}^d x_i^{\lambda_i}$. A *pure binomial ideal* is an ideal generated by polynomials of the form $\boldsymbol{x}^{\boldsymbol{\lambda}} - \boldsymbol{x}^{\boldsymbol{\lambda'}}$, where $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda'}$ are non-negative integer vectors. More generally, a *binomial ideal* is one that is generated by polynomials of the form $\boldsymbol{x}^{\boldsymbol{\lambda}} - \theta \boldsymbol{x}^{\boldsymbol{\lambda'}}$, where $\theta \in \overline{\mathbb{Q}}$. It is known that the vanishing ideal $I \subseteq \overline{\mathbb{Q}}[\boldsymbol{x}]$ of an algebraic subgroup of $\mathbb{G}_m^d$ is a pure binomial ideal.

Example 7. Let $\Lambda \coloneqq \{(2n, -n) : n \in \mathbb{Z}\} = (2, -1) \cdot \mathbb{Z}$ be a lattice in $\mathbb{Z}^2$ of rank 1.

A vector $(a, b) \in \mathbb{Z}^2$ is equivalent to $(a + 2b, 0)$ modulo Λ. Hence the quotient group $\mathbb{Z}^2/\Lambda$ comprises a set of cosets $\{(c, 0) + \Lambda : c \in \mathbb{Z}\}$ with group operation defined by componentwise addition. Thus $\mathbb{Z}^2/\Lambda$ is isomorphic to the additive group $\mathbb{Z}$, which has trivial torsion subgroup since 0 is the only element of finite order.

We have

$$\begin{aligned} H_\Lambda &= \{(x, y) \in \mathbb{G}_m^2 : \ \forall n \in \mathbb{Z}.\ x^{2n} y^{-n} = 1\} \\ &= \{(x, y) \in \mathbb{G}_m^2 : \ x^2 = y\}. \end{aligned}$$

The vanishing ideal of H_Λ is pure binomial, generated by $x^2 - y$. ◀

The following proposition shows how to recover the generators of the lattice and hence the pure binomial ideal that vanishes on an algebraic subgroup of $\mathbb{G}_m^d$ defined by its equations.

Proposition 8 ([6, Proposition 3.2.14]). *Let G be a subgroup of $\mathbb{G}_m^d$ defined by polynomial equations $\sum_{\boldsymbol{\lambda} \in \mathcal{L}_i} a_{i,\boldsymbol{\lambda}} \boldsymbol{x}^{\boldsymbol{\lambda}} = 0$ for $i = 1, \ldots, m$, where $\mathcal{L}_i \subseteq \mathbb{Z}^d$. Then $G = H_\Lambda$, where $\Lambda \subseteq \mathbb{Z}^d$ is generated by vectors of the form $\boldsymbol{\lambda}_i - \boldsymbol{\lambda}_i'$ with $\boldsymbol{\lambda}_i, \boldsymbol{\lambda}_i' \in \mathcal{L}_i$, for $i \in \{1, \ldots, m\}$.*

The next proposition follows from a standard result in Diophantine geometry concerning the number of generators of a subgroup of $\mathbb{G}_m^d$ (cf. [6, Chapter 3]); we give a proof in Appendix B.

Proposition 9. *Let $\Lambda \subseteq \mathbb{Z}^d$ be a lattice of rank r. Then*

(1) the torsion subgroup of $\mathbb{Z}^d/\Lambda$ is s_0-generated, for some $s_0 \leq r$, and

(2) *H_Λ is s-generated, where $s := s_0$ if Λ has full rank and otherwise $s := \max(s_0, 1)$. Furthermore, s is the minimal number of topological generators for H_Λ.*

Unipotent Matrices. For a $d \times d$ unipotent matrix A and a $d \times d$ nilpotent matrix B, define

$$\log(A) := \sum_{k=1}^{d-1} (-1)^{k+1} \frac{(A - \mathrm{Id}_d)^k}{k} \quad \text{and} \quad \exp(B) := \sum_{k=0}^{d-1} \frac{B^k}{k!}.$$

Let $G \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ be a commutative subgroup of unipotent matrices. The set $L := \{\log(A) : A \in G\}$ is a linear subspace of $\overline{\mathbb{Q}}^{d^2}$ consisting of nilpotent matrices [7, Chapter II, Section 7.3]. Moreover, $\exp\colon L \to G$ and $\log\colon G \to L$ yield polynomial isomorphisms between L and G as algebraic groups, with additive and multiplicative group structures respectively. Taken together, these observations lead to the following proposition.

Proposition 10. *Let G be a commutative group of unipotent matrices and $L := \{\log(A) : A \in G\}$ the associated linear subspace of nilpotent matrices. Then G has a topological generator of cardinality s if and only if L is spanned by a set of s matrices as a $\overline{\mathbb{Q}}$-vector space.*

Proof. For all $A_1, \ldots, A_s \in G$ we have the following equivalences:

$$\begin{aligned}
&\{A_1, \ldots, A_s\} \text{ topologically generates } G,\\
\Leftrightarrow\ &\{A_1^{n_1} A_2^{n_2} \cdots A_s^{n_s} : n_1, \ldots, n_s \in \mathbb{Z}\} \text{ is dense in } G,\\
\Leftrightarrow\ &\{\textstyle\sum_{i=1}^{s} n_i \log(A_i) : n_1, \ldots, n_s \in \mathbb{Z}\} \text{ is dense in } L,\\
\Leftrightarrow\ &\{\textstyle\sum_{i=1}^{s} t_i \log(A_i) : t_1, \ldots, t_s \in \overline{\mathbb{Q}}\} = L,
\end{aligned}$$

as desired. □

Example 11. We consider 2×2 matrices of the form

$$A_t := \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix},$$

where $t \in \overline{\mathbb{Q}}$. Since $(A_t - \mathrm{Id}_2)^2 = 0$, the matrix A_t is unipotent. Moreover, $G := \{A_t : t \in \overline{\mathbb{Q}}\}$ is a (multiplicative) commutative subgroup of $\mathrm{GL}_2(\overline{\mathbb{Q}})$. The matrix

$$\log(A_t) = A_t - \mathrm{Id}_2 = \begin{pmatrix} 0 & t \\ 0 & 0 \end{pmatrix}$$

is nilpotent for any $t \in \overline{\mathbb{Q}}$. It further holds that

$$L := \{\log(A_t) : t \in \overline{\mathbb{Q}}\} = \left\{ \begin{pmatrix} 0 & t \\ 0 & 0 \end{pmatrix} : t \in \overline{\mathbb{Q}} \right\}$$

is a linear subspace of $\overline{\mathbb{Q}}^4$. In fact, L is spanned by a single matrix over $\overline{\mathbb{Q}}$ and so, by Proposition 10, G is topologically 1-generated. Indeed, G is the Zariski closure of, say,

$$\langle A_1 \rangle = \left\langle \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \right\rangle = \left\{ \begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix} : n \in \mathbb{Z} \right\}.$$

The map exp is an isomorphism of L and G:

$$\exp : \begin{pmatrix} 0 & t \\ 0 & 0 \end{pmatrix} \mapsto \mathrm{Id}_2 + \begin{pmatrix} 0 & t \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}.$$

◀

3 Commutative Group Determination

Recall that the group determination problem with parameters (s, d, m, b) asks, given $s \in \mathbb{N}$ and a family of m polynomials in $\mathbb{Q}[\{x_{i,j}\}_{1\le i,j\le d}]$, each of total degree at most b, to determine whether their zero locus $Z \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ is an s-generated algebraic matrix group. In this section, we first demonstrate a procedure for this problem subject to the constraint that the underlying group is semisimple and commutative (Proposition 12). Next, we generalise this result by lifting the requirement that the matrices are semisimple (Proposition 14).

Proposition 12. *The group determination problem for commutative semisimple matrices with parameters (s, d, m, b) can be decided in time $(mb)^{\mathrm{poly}(d)}$, and in space bounded by $(d \log b)^{O(1)}$.*

Proof. Recall that the input to the problem consists of m polynomials, each of total degree at most b, together with a natural number s. Let Z be the subvariety of $\mathrm{GL}_d(\overline{\mathbb{Q}})$ defined by the input polynomials. The task is to determine whether Z is a group that is topologically generated by at most s commutative semisimple matrices.

Suppose that the input is a positive instance of the problem, meaning that Z is an algebraic group that is topologically generated by commutative semisimple matrices $M_1, \dots, M_s \in \mathrm{GL}_d(\overline{\mathbb{Q}})$. Recall that commutative semisimple matrices are simultaneously diagonalisable [21, Theorem 1.3.21]. Thus there exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that $D_i \coloneqq P^{-1}M_iP$ are diagonal. Let G be the subgroup of $\mathbb{G}_m^d$ defined by

$$G \coloneqq \{g \in \mathbb{G}_m^d : \Delta(g) \in \overline{\langle D_i : 1 \le i \le s\rangle}\}.$$

Since $Z = \overline{\langle M_i : 1 \le i \le s\rangle}$, we have $P^{-1}ZP = \{\Delta(g) : g \in G\}$. Moreover, as Z is the zero set of polynomials of total degree at most b and conjugation by the linear transformation P does not increase the total degrees of the transformed polynomials, it holds that $P^{-1}ZP$ is the zero set of polynomials of degree at most b. Hence, by Proposition 8, the group G has the form H_Λ for some lattice $\Lambda \subseteq \mathbb{Z}^d$ whose generators have norm at most b and such that H_Λ is s-generated.

Conversely, for every such lattice Λ, the variety $P\{\Delta(g) : g \in H_\Lambda\}P^{-1}$ is topologically s-generated by commutative semisimple matrices.

The above reasoning shows the correctness of the following nondeterministic decision procedure:

(1) Guess a lattice $\Lambda \subseteq \mathbb{Z}^d$ whose generators have norm at most b such that H_Λ is s-generated.
(2) Output "yes" if there exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that $P^{-1}ZP = \{\Delta(g) : g \in H_\Lambda\}$.

By Proposition 9, Step 1 amounts to guessing Λ subject to the condition that the torsion subgroup of $\mathbb{Z}^d/\Lambda$ is generated by at most s elements. We can understand this observation in terms of the number of non-unit elementary divisors of the desired lattice Λ; see Appendix B for more details.

Step 2 amounts to checking the truth in $\overline{\mathbb{Q}}$ of the sentence

$$\exists P \in \mathrm{GL}_d(\overline{\mathbb{Q}})\ \forall A = (a_{ij})_{1\le i,j\le d} \in \mathrm{GL}_d(\overline{\mathbb{Q}})$$
$$\left(P^{-1}AP \in Z \Leftrightarrow \textstyle\bigwedge_{i\neq j} a_{ij} = 0 \wedge (a_{11}, a_{22}, \dots, a_{dd}) \in H_\Lambda\right),$$

with respect to the theory of algebraically closed fields. By Theorem 1, this can be done in time $(mb)^{\mathrm{poly}(d)}$. The claimed running time for the overall procedure follows from the fact that the number of possibilities for the lattice Λ is at most $(2b)^{d^2}$. □

By the following remark, a commutative algebraic group G is s-generated if its unipotent subgroup G_u and semisimple subgroup G_s are both s-generated.

Remark 13. Let $U = \overline{\langle U_i : 1 \le i \le s\rangle}$ be a commutative unipotent s-generated algebraic group and $S = \overline{\langle S_i : 1 \le i \le s\rangle}$ a semisimple commutative s-generated algebraic group. Suppose moreover that the matrices $\{U_i, S_i : 1 \le i \le s\}$ are commutative. Then we have $\overline{\langle U_i, S_i : 1 \le i \le s\rangle} =$

$\overline{\langle S_i U_i : 1 \leq i \leq s \rangle}$. This relies on the fact that if $A \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ is semisimple and $B \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ is unipotent, then both A and B lie in the Zariski closure of the subgroup generated by their product AB; see [26, Section 15.3].

The next proposition generalises the procedure described in Proposition 12. In Proposition 14 we consider the group determination problem for s-generated commutative algebraic groups. Key to our generalisation is the determination of a matrix $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and properties associated with the semisimple and unipotent subgroups of the group $P^{-1}ZP$.

Proposition 14. *The group determination problem for commutative matrices with parameters (s, d, m, b) can be decided in time $(mb)^{\mathrm{poly}(d)}$, and in space bounded by $(d \log b)^{O(1)}$.*

Proof. Recall that the input to the problem consists of m polynomials, each of total degree at most b, together with a natural number s. Let Z be the subvariety of $\mathrm{GL}_d(\overline{\mathbb{Q}})$ defined by the input polynomials. The task is to determine whether Z is a group that is topologically generated by at most s commutative matrices.

Suppose that the input is a positive instance of the problem, that is, Z is an algebraic group that is topologically generated by commutative matrices $M_1, \ldots, M_s \in \mathrm{GL}_d(\overline{\mathbb{Q}})$. Recall that commutative matrices are simultaneously triangularisable [21, Theorem 2.3.3]. Thus there exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that there exist diagonal matrices D_i and upper unitriangular matrices U_i, $1 \leq i \leq s$, where $P^{-1}M_iP = D_iU_i$, and moreover D_i and U_i commute. Then we can recover $\overline{\langle D_i : 1 \leq i \leq s \rangle}$ as the set of diagonal matrices in $P^{-1}\,\overline{\langle M_i : 1 \leq i \leq s \rangle}\,P$. As in Proposition 12, it follows that $\overline{\langle D_i : 1 \leq i \leq s \rangle}$ is the zero locus of a system of polynomials of degree at most b.

By Proposition 8, $\overline{\langle D_i : 1 \leq i \leq s \rangle} = \{\Delta(g) : g \in H_\Lambda\}$ for some lattice $\Lambda \subseteq \mathbb{Z}^d$ that is generated by vectors having supremum norm at most b and H_Λ is s-generated. Note that $\overline{\langle U_i : 1 \leq i \leq s \rangle}$ is the set of upper unitriangular matrices in $P^{-1}\,\overline{\langle M_i : 1 \leq i \leq s \rangle}\,P$.

The decision procedure is as follows:

(1) Guess a lattice $\Lambda \subseteq \mathbb{Z}^d$ whose generators have norm at most b and such that H_Λ is s-generated.
(2) Output "yes" if there exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that
   (a) $G \coloneqq P^{-1}ZP$ is a commutative group of upper triangular matrices;
   (b) $\{A \in G : A \text{ diagonal}\} = \{\Delta(g) : g \in H_\Lambda\}$; and
   (c) $\{\log(A) : A \in G,\ A \text{ unipotent}\}$ is a linear variety of dimension at most s.

As in Proposition 12, guessing Λ that satisfies the condition of Item 1 is via Proposition 9. We claim that the procedure outputs "yes" if and only if the input is a positive instance of the problem. Indeed, Item 2a checks that G is a commutative algebraic matrix group. In this case, by Fact 4, both the set G_s of semisimple matrices in G and the set G_u of unipotent matrices in G form subgroups of G. Next, Items 2b and 2c respectively check that G_s and G_u are s-generated (relying on Propositions 9 and 10). This in turn implies that G is itself s-generated, as noted in Remark 13.

The existence of P satisfying Items 2a to 2c reduces to checking the truth in $\overline{\mathbb{Q}}$ of an $\exists^*\forall^*$-sentence in the language of fields. The existential quantifiers correspond to the possible choices of P, while the universal quantifiers range over entries of the group G defined in Item 2a. For a fixed choice of Λ, the truth of such a formula can be decided in time $(mb)^{\mathrm{poly}(d)}$ by Theorem 1. Given that the number of possible choices of the lattice Λ is at most $(2b)^{d^2}$ the claimed complexity bound immediately follows. □

4 Orbit-Closure Determination

Recall the aforementioned orbit-closure determination problem. The problem asks, given $s \in \mathbb{N}$ and a family of m polynomials in $\mathbb{Q}[\{x_i\}_{1\le i\le d}]$, each of total degree at most b, to determine whether their zero locus $Z \subseteq \overline{\mathbb{Q}}^d$ is the orbit closure of some point $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$ under the action of an s-generated algebraic matrix group, i.e., determine whether there exists an algebraic group topologically generated by matrices $M_1, \dots, M_s \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and vector $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$ such that Z is the Zariski closure of the orbit $\{M_1^{\ell_1} \cdots M_s^{\ell_s}\boldsymbol{v} : \ell_1, \dots, \ell_s \in \mathbb{Z}\}$.

Key to the main result in this section (Theorem 2) is Lemma 15. Briefly, the lemma describes a change of basis on a given commutative matrix group G and vector $\boldsymbol{v}$ such that in this new basis the elements of the semisimple subgroup $G_s \le G$ are diagonal, the elements of the unipotent subgroup $G_u \le G$ are upper unitriangular; and moreover $\boldsymbol{v}$ is mapped to a binary vector with a prescribed structure. In fact, the lemma permits us to focus on the orbit closure of the vector $\mathbf{1}$ (that is, the all-one vector).

Lemma 15. *Let $G = \overline{\langle M_i : 1 \le i \le s\rangle} \le \mathrm{GL}_d(\overline{\mathbb{Q}})$ be a commutative algebraic group and $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$. There exist $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$, and $T \in \{0,1\}^{d\times k}$ with the following properties:*

1. *the elements of PGP^{-1} consist of block diagonal matrices, where each block is a scalar multiple of a unitriangular matrix;*
2. *PG_sP^{-1} consists of diagonal matrices and PG_uP^{-1} consists of upper unitriangular matrices,*
3. *$T^{\mathsf{T}}T = \mathrm{Id}_k$,*
4. *$P\boldsymbol{v} = T\mathbf{1}$ is a 0-1 vector with at most one 1 per block.*

Proof. The proof of Lemma 15 relies on Theorem 3 and Corollary 17 in Section 4.1, immediately below. These results show how to find P satisfying Item (1) such that $P\boldsymbol{v} = \boldsymbol{e}_{i_1} + \cdots + \boldsymbol{e}_{i_k}$ is a sum of standard unit vectors of $\overline{\mathbb{Q}}^d$. Defining T to be the $d \times k$ matrix with columns $\boldsymbol{e}_{i_1}, \dots, \boldsymbol{e}_{i_k}$, it follows that $T^{\mathsf{T}}T = \mathrm{Id}_k$ and $P\boldsymbol{v} = T\mathbf{1}$.

□

4.1 A Convenient Base Change for Commuting Matrices

This section contains technical details concerning the proof of Lemma 15.

Theorem 3. *Let $A_1, \dots, A_m$ be pairwise commuting matrices in $\overline{\mathbb{Q}}^{d\times d}$, and $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$ a non-zero vector. There exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that*

- *PA_iP^{-1} is block diagonal;*
- *each block is a scalar multiple of an upper unitriangular matrix;*
- *$P\boldsymbol{v}$ is a 0-1 vector with at most one 1 per block.*

Proof. By [34, Theorem 1], there exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that PA_iP^{-1} is upper triangular for $i \in \{1, \dots, m\}$. We have the following easy refinement of this well-known fact.

Claim 16. *There exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that PA_iP^{-1} is upper triangular for $i \in \{1, \dots, m\}$ and $P\boldsymbol{v} = \boldsymbol{e}_j$ for some $j \in \{1, \dots, d\}$.*

Proof of Claim 16. As mentioned above, by [34, Theorem 1], there exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that PA_iP^{-1} is upper triangular for $i \in \{1, \dots, m\}$. This yields a basis $\boldsymbol{v}_1, \dots, \boldsymbol{v}_d$ of $\overline{\mathbb{Q}}^d$ such that $A_i\boldsymbol{v}_j \in \mathrm{span}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_j)$ for all i and j. Let ℓ be the minimum index such that $\boldsymbol{v} \in \mathrm{span}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_\ell)$. Then $\boldsymbol{v} = \sum_{i=0}^{\ell} \alpha_i\boldsymbol{v}_i$, where $\alpha_\ell \neq 0$ by minimality of ℓ. Hence $\boldsymbol{v}_\ell \in \mathrm{span}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_{\ell-1}, \boldsymbol{v})$. Consider a new basis in which $\boldsymbol{v}_\ell$ is replaced by $\boldsymbol{v}$. Note that $\mathrm{span}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_\ell) = \mathrm{span}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_{\ell-1}, \boldsymbol{v})$.

Taking Q to be the matrix with columns $\boldsymbol{v}_1, \dots, \boldsymbol{v}_{\ell-1}, \boldsymbol{v}, \boldsymbol{v}_{\ell+1}, \dots, \boldsymbol{v}_n$ we have that $Q\boldsymbol{e}_\ell = \boldsymbol{v}$. Let $P := Q^{-1}$ then we have that PA_iP^{-1} is upper triangular for $i \in \{1, \dots, m\}$. This concludes the proof of the claim. □

By [36, Theorem 2.4] we can write $\overline{\mathbb{Q}}^d = V_1 \oplus \cdots \oplus V_k$ such that each subspace V_i is invariant under $A_1, \ldots, A_m$ and for all $1 \leq i \leq m$ and $1 \leq j \leq k$ the restriction of A_i to V_j has a single eigenvalue. Let $\boldsymbol{v} = \boldsymbol{v}_1 + \cdots + \boldsymbol{v}_k$ be the corresponding decomposition of $\boldsymbol{v}$.

Fix a subspace V_j. We apply Claim 16 to the restrictions of $A_1, \ldots, A_m$ to V_j and the vector $\boldsymbol{v}_j$. We obtain a basis for V_j that includes $\boldsymbol{v}_j$ and such that the restriction of A_i to V_j is upper triangular. □

The next corollary follows from Theorem 3.

Corollary 17. *Let $A_1, \ldots, A_m$ be matrices in $\mathrm{GL}_d(\overline{\mathbb{Q}})$ and let $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ be such that PA_iP^{-1} is block diagonal where each block is the scalar multiple of an upper unitriangular matrix. Given $M \in \langle A_1, \ldots, A_m \rangle$, M is semi-simple if and only if PMP^{-1} is diagonal and M is unipotent if and only if PMP^{-1} is upper unitriangular.*

Proof. Clearly PMP^{-1} is block diagonal. Thus M is semisimple if and only if it is blockwise semisimple and it is unipotent if and only if it is blockwise unipotent. Moreover, since each block is upper triangular with constant diagonal, such a block is diagonalisable if and only if it is already diagonal and is unipotent if and only if it is has all ones along the diagonal. □

4.2 Decision Procedures for the Orbit-Closure Determination

The main contributions of this section are the procedures for certain cases of the orbit-closure determination problem (Proposition 18 and Theorem 2). The procedure in Proposition 18 supposes that the generators of the matrix group are semisimple and commutative. The procedure in Theorem 2 lifts the requirement that the generators are semisimple. We illustrate the procedures with a worked example (Example 21) at the close of this section.

Proposition 18. *The orbit-closure determination problem for commutative semisimple matrices with parameters (s, d, m, b) can be decided in time $(mb)^{\mathrm{poly}(d)}$, and in space bounded by $(d \log b)^{O(1)}$.*

Proof. Recall that the input to the problem consists of m polynomials, each of total degree at most b, together with a natural number s. Let Z be the subvariety of $\overline{\mathbb{Q}}^d$ defined by the input polynomials. The task is to determine whether Z is an orbit-closure under a group that is generated by at most s commutative semisimple matrices.

Suppose that the input is a positive instance of the problem, that is, $Z = \overline{\langle M_i : 1 \leq i \leq s \rangle \cdot \boldsymbol{v}}$ for semisimple commutative matrices $M_1, \ldots, M_s \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and a vector $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$. By Lemma 15, there exist matrices $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and $T \in \{0,1\}^{d \times k}$ such that $D_i \coloneqq PM_iP^{-1}$ are diagonal, and moreover $T^{\mathsf{T}}T = \mathrm{Id}_k$ and $P\boldsymbol{v} = T\mathbf{1}$ hold.

For all $i \in \{1, \ldots, s\}$, denote by $D'_i \in \mathrm{GL}_k(\overline{\mathbb{Q}})$ the diagonal matrix uniquely defined by the requirement that $D_iT = TD'_i$. Write $G \coloneqq \left\{ g \in \mathbb{G}_m^k : \Delta(g) \in \overline{\langle D'_i : 1 \leq i \leq s \rangle} \right\}$. Then we have

$$PZ = \overline{\langle D_i : 1 \leq i \leq s \rangle \cdot T\mathbf{1}} = \overline{T\langle D'_i : 1 \leq i \leq s \rangle \cdot \mathbf{1}} = \overline{\{Tg : g \in G\}}.$$

Note moreover that by the definition of G we have

$$G = \left\{ g \in \mathbb{G}_m^k : Tg \in PZ \right\}$$

and so G is defined by polynomials of total degree at most b. It follows from Proposition 8 that $G = H_\Lambda$ for some lattice $\Lambda \subseteq \mathbb{Z}^k$ that is generated by vectors whose entries have absolute value at most b.

Conversely, suppose that $PZ = \overline{\{Tg : g \in H_\Lambda\}}$ for some matrices $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and $T \in \{0,1\}^{d \times k}$ such that $T^{\mathsf{T}}T = \mathrm{Id}_k$, and lattice Λ as above. Then

$$Z = \overline{\langle M_i : 1 \leq i \leq s \rangle \cdot \boldsymbol{v}},$$

where $\boldsymbol{v} = P^{-1}T\mathbf{1}$ and $M_i := P^{-1}D_iP$ with $D_i \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ being any diagonal matrices such that $D_iT = T\Delta(g_i)$ for some topological generators $\{g_1, \ldots, g_s\}$ of H_Λ.

In summary, the decision procedure is as follows:

(1) Guess $k \subseteq \{0, \ldots, d\}$ and $T \in \{0,1\}^{d\times k}$ such that $T^\intercal T = \mathrm{Id}_k$.
(2) Guess a lattice $\Lambda \subseteq \mathbb{Z}^k$ whose generators have norm at most b and such that H_Λ is s-generated (see Proposition 9).
(3) Determine whether there exists $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that $PZ = \overline{\{Tg : g \in H_\Lambda\}}$.

Step 3 can be reduced in polynomial time to checking the truth of a $\exists^*\forall^*\exists^*$-sentence in the theory of real closed fields. The outermost existential quantifiers correspond to the choice of the matrix P. Then the right-to-left inclusion in the equation $PZ = \overline{\{Tg : g \in H_\Lambda\}}$ is expressed by the formula

$$\forall g \in H_\Lambda\, \exists \boldsymbol{z} \in Z\, (P\boldsymbol{z} = Tg),$$

while, by Fact 5, the left-to-right inclusion is expressed by the formula

$$\forall \boldsymbol{z} \in Z\, \forall \varepsilon > 0\, \exists g \in H_\Lambda\ (\|P\boldsymbol{z} - Tg\| < \varepsilon)\,.$$

By Theorem 1, the truth of such a sentence can be decided in time $(mb)^{\mathrm{poly}(d)}$. The claimed overall complexity bound now follows from the fact that there are at most $(2b)^{d^2+1}$ choices of the lattice Λ and matrix U. □

We motivate the constructive subroutines in Proposition 18 with loop synthesis.

Example 19. Let us determine whether there is a single-path linear loop with update matrix $M \in \mathrm{GL}_2(\overline{\mathbb{Q}})$ and initial vector $\boldsymbol{v} \in \overline{\mathbb{Q}}^2$ such that the zero set $Z \subseteq \overline{\mathbb{Q}}^2$ of the ideal $I := \langle 4x^2+y^2+4xy-x-y\rangle$ satisfies $Z = \overline{\langle M\rangle \cdot \boldsymbol{v}}$. In other words, we seek a matrix M and vector $\boldsymbol{v}$ in the loop

$\boldsymbol{x} \leftarrow \boldsymbol{v}$;
while $(*)$ **do**
  $\boldsymbol{x} \leftarrow M\boldsymbol{x}$;
end while

such that following two Hoare triples are satisfied:

$$\{\mathbf{true}\}\ \boldsymbol{x} \leftarrow \boldsymbol{v}\ \ \{4x^2 + y^2 + 4xy - x - y = 0\}$$
$$\{4x^2 + y^2 + 4xy - x - y = 0\}\ \boldsymbol{x} \leftarrow M\boldsymbol{x}\ \{4x^2 + y^2 + 4xy - x - y = 0\}$$

and such that $4x^2 + y^2 + 4xy - x - y = 0$ is moreover the strongest polynomial invariant satisfied by all reachable program state (i.e., all other invariants lie in the ideal I).

Suppose that in Steps 1 and 2 of the procedure in Proposition 18 we guess $H_\Lambda := \{(x, y) \in \mathbb{G}_m^2 : x^2 - y = 0\}$ and the matrix $T := \mathrm{Id}_2$. For Step 3, we want to find all invertible matrices $P = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{GL}_2(\overline{\mathbb{Q}})$ such that $V(P^{-1} \cdot I) = \overline{\{Th : h \in H_\Lambda\}}$. This is equivalent to the requirement that the two polynomials

$$4(ax + by)^2 + (cx + dy)^2 + 4(ax + by)(cx + dy) - ax - by - cx - dy$$

and $x^2 - y$ are multiples of one another. Therefore the polynomials defining P^{-1} comprise the following ideal:

$$J_{P^{-1}} := \langle 4a^2 + c^2 + 4ac - b - d, 4b^2 + d^2 + 4bd, 8ab + 2cd + 4ad + 4bc, a + c\rangle$$
$$= \langle a + c, c^2 - b - d, 2bc + cd, (2b + d)^2\rangle.$$

One choice of P^{-1} is $\left(\begin{smallmatrix} 1 & -1 \\ -1 & 2 \end{smallmatrix}\right)$. Thus we realise Z as the orbit closure $\overline{\langle M \rangle \cdot \boldsymbol{v}}$ where

$$M \coloneqq P^{-1} \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix} P = \begin{pmatrix} 0 & -2 \\ 4 & 6 \end{pmatrix} \quad \text{and} \quad \boldsymbol{v} \coloneqq P^{-1} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

◀

The following theorem is our main contribution, which provides a decision procedure for the orbit-closure determination problem for commutative groups. The generalization of this result to the case of general matrix groups appears to be challenging.

Theorem 2. *The orbit-closure determination problem for commutative matrices with parameters (s, d, m, b) can be decided in time $(mb)^{\mathrm{poly}(d)}$, and in space bounded by $(d \log b)^{O(1)}$.*

Proof. Recall that the input to the problem consists of m polynomials, each of total degree at most b, together with a natural number s. Let Z be the subvariety of $\overline{\mathbb{Q}}^d$ defined by the input polynomials. The task is to determine whether Z is an orbit closure under an algebraic group that is topologically generated by at most s commutative matrices.

Suppose that the input is a positive instance of the problem, that is, $Z = \overline{\langle M_i : 1 \le i \le s \rangle \cdot \boldsymbol{v}}$ for commutative matrices $M_1, \ldots, M_s \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and a vector $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$. Using Fact 4, let G_u be the subgroup of unipotent elements of G and G_s be the subgroup of semisimple elements of G. We have $Z = \overline{G_u \cdot G_s \cdot \boldsymbol{v}}$.

Applying Lemma 15 to the group G, we obtain $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and $T \in \{0,1\}^{d \times k}$ such that PG_sP^{-1} is a group of diagonal matrices, PG_uP^{-1} is a group of upper unitriangular matrices, $T^{\mathsf{T}}T = \mathrm{Id}_k$, and $P\boldsymbol{v} = T\mathbf{1}$. In particular, we have $PG_sP^{-1} = \overline{\langle D_i : 1 \le i \le s \rangle}$ for diagonal matrices $D_1, \ldots, D_s$ and $PG_uP^{-1} = \overline{\langle U_i : 1 \le i \le s \rangle}$ for upper unitriangular matrices $U_1, \ldots, U_s$, such that $PM_iP^{-1} = D_iU_i$ for $1 \le i \le s$.

For all $i \in \{1, \ldots, s\}$, denote by $D_i' \in \mathrm{GL}_k(\overline{\mathbb{Q}})$ the diagonal matrix uniquely defined by the requirement that $D_iT = TD_i'$. Furthermore, write

$$G' \coloneqq \{g \in \mathbb{G}_m^k : \Delta(g) \in \overline{\langle D_i' : 1 \le i \le s \rangle}\}.$$

Then we have

$$\begin{aligned} PZ &= \overline{P \cdot G_u \cdot G_s \cdot \boldsymbol{v}} \\ &= \overline{\langle U_i : 1 \le i \le s \rangle \cdot \langle D_i : 1 \le i \le s \rangle \cdot P\boldsymbol{v}} \\ &= \overline{\langle U_i : 1 \le i \le s \rangle \cdot \langle D_i : 1 \le i \le s \rangle \cdot T\mathbf{1}} \\ &= \overline{\langle U_i : 1 \le i \le s \rangle \cdot T\langle D_i' : 1 \le i \le s \rangle \cdot \mathbf{1}} \\ &= \overline{\langle U_i : 1 \le i \le s \rangle \cdot TG'} \end{aligned}$$

where the last equality follows from Proposition 10.

To obtain a degree bound on G', we prove that $G' = \{g \in \mathbb{G}_m^k : P^{-1}Tg \in Z\}$.

Claim 20. $G' = \{g \in \mathbb{G}_m^k : P^{-1}Tg \in Z\}$.

Proof of Claim 20. The left-to-right inclusion is obvious. Now let $g \in \mathbb{G}_m^k$ such that $P^{-1}Tg \in Z$. Since the Zariski closure of

$$\overline{\langle U_i : 1 \le i \le s \rangle \cdot TG'} \tag{2}$$

coincides with its Euclidean closure (by Fact 5), there exists a sequence of vectors $(h_n)_{n \ge 0}$ in (2) such that $\lim_{n \to \infty} h_n = Tg$. By construction, the elements of (2) are given by blocks, therefore we

can argue blockwise. Note that a block of Tg has only one non-zero entry, let us consider one of them:

$$(0, \dots, 0, g_i, 0, \dots, 0)^{\mathsf{T}}$$

where g_i corresponds to the ith entry in the block. The corresponding block of h_n has the form

$$(h_n^{(1)}, \dots, h_n^{(i)}, 0, \dots, 0)^{\mathsf{T}},$$

and moreover the corresponding block of $P\boldsymbol{v}$ has 1 in the ith entry and 0 elsewhere. Therefore, we have

$$\lim_{n\to\infty} \begin{pmatrix} h_n^{(i)} & * & * & * & * \\ & \ddots & \ddots & * & * \\ & & \ddots & * & * \\ & & & & h_n^{(i)} \end{pmatrix} \begin{pmatrix} 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ g_i \\ 0 \\ \vdots \end{pmatrix} \tag{3}$$

where the matrix

$$\begin{pmatrix} h_n^{(i)} & * & * & * & * \\ & \ddots & \ddots & * & * \\ & & \ddots & * & * \\ & & & & h_n^{(i)} \end{pmatrix}$$

is a block of an element in PGP^{-1} for every $n \in \mathbb{N}$, and its corresponding semisimple element has as block $h_n^{(i)}$ Id. From (3) we conclude that

$$\lim_{n\to\infty} (h_n^{(i)} \operatorname{Id}) \cdot \begin{pmatrix} 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ g_i \\ 0 \\ \vdots \end{pmatrix}$$

and moreover that the diagonal blocks $(h_n^{(i)} \operatorname{Id})$ correspond to the blocks of diagonal matrices that belong to PGP^{-1}. Thus, there exists a sequence of diagonal matrices $(D^{(n)})_{n\geq 0} \subseteq PGP^{-1}$ such that

$$\lim_{n\to\infty} D^{(n)} T\mathbf{1} = Tg.$$

Let $\widetilde{D}^{(n)}$ be the diagonal matrix uniquely defined by $D^{(n)}T = T\widetilde{D}^{(n)}$. Then we have $(\widetilde{D}^{(n)})_{n\geq 0} \subseteq \overline{\langle D_i' : 1 \leq i \leq s\rangle}$. Hence we have,

$$\lim_{n\to\infty} T\widetilde{D}^{(n)}\mathbf{1} = T\Delta(g)\mathbf{1},$$

and since multiplication by T defines an injective map we have

$$\lim_{n\to\infty} \widetilde{D}^{(n)} = \Delta(g),$$

which proves that $g \in G'$. This ends the proof of Claim 20. □

It follows that G' is defined by polynomials of total degree at most b and hence has the form H_Λ for some lattice $\Lambda \subseteq \mathbb{Z}^k$ whose generators have norm at most b and such that H_Λ is s-generated (see Proposition 9).

Conversely, suppose that there exist $k \in \{0,\dots,d\}$ and $T \in \{0,1\}^{d\times k}$ such that $T^{\mathsf{T}}T = \mathrm{Id}_k$, together with $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$, commuting matrices $U_1,\dots,U_s,D_1,\dots,D_s \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ with U_i upper unitriangular and D_i diagonal, and an s-generated lattice $\Lambda \subseteq \mathbb{Z}^k$ such that

$$TH_\Lambda = \overline{\langle D_1,\dots,D_s\rangle \cdot T\mathbf{1}}$$

satisfying

$$PZ = \overline{\langle U_i : 1 \le i \le s\rangle \cdot TH_\Lambda}. \tag{4}$$

Then $Z = \overline{\langle M_i : 1 \le i \le s\rangle \cdot \boldsymbol{v}}$, where $M_i \coloneqq P^{-1}U_iD_iP$ and $P\boldsymbol{v} = T\mathbf{1}$.

In summary, the decision procedure is as follows:

(1) Guess $k \in \{0,\dots,d\}$ and $T \in \{0,1\}^{d\times k}$ such that $T^{\mathsf{T}}T = \mathrm{Id}_k$.
(2) Guess a lattice $\Lambda \subseteq \mathbb{Z}^k$ whose generators have norm at most b and such that H_Λ is s-generated, say by $g_1,\dots,g_s \in H_\Lambda$ (see Proposition 9).
(3) Return "yes" if there exist $P \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and commuting matrices $U_1,\dots,U_s,D_1,\dots,D_s \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that the U_i are upper unitriangular, the D_i are diagonal, $D_iT = T\Delta(g_i)$ holds for all $i \in \{1,\dots,s\}$, and (4) holds.

Step 3 can be reduced to checking the truth of an $\exists^*\forall^*\exists$-sentence over the theory of real closed fields. The outer group of existential quantifiers range over the possible choices of the matrices P and $U_1,\dots,U_s$. The rest of the formula checks (4). The right-to-left inclusion is expressed by the formula

$$\forall t_1 \cdots \forall t_s\, \forall g \in H_\Lambda \cdot \exp(\textstyle\sum_{i=1}^s t_i \log U_i)Tg \in PZ.$$

By Fact 5, the left-to-right inclusion is expressed by the formula

$$\forall \boldsymbol{z} \in Z\, \forall \varepsilon > 0 \exists t_1 \cdots \exists t_s\, \exists g \in H_\Lambda \cdot \|P\boldsymbol{z} - \exp(\textstyle\sum_{i=1}^s t_i \log U_i)Tg\| < \varepsilon.$$

By Theorem 1, the truth of such a sentence can be decided in time $(mb)^{\mathrm{poly}(d)}$. Then the overall complexity bound follows from the fact that the number of choices of the lattice Λ and matrix T is at most $(2b)^{d^2+1}$. □

Example 21, below, applies the procedure in Theorem 2 to the variety we first saw in Examples 1 and 2 in the Introduction. The calculations involved in the preparation of Examples 21 and 22 were performed in Macaulay2 [19].

Example 21. Let $Z \subseteq \overline{\mathbb{Q}}^4$ be the zero set of the ideal $I = \langle Q_1, Q_2\rangle$, defined in Example 2, where

$$Q_1 = x_2^2 - x_1 - x_4 \quad \text{and} \quad Q_2 = -2x_4x_2 - 2x_3^2 - \frac{1}{5}x_2x_3.$$

Let us determine whether the zero set Z is equal to the orbit closure of a 1-generated algebraic matrix group. To answer this affirmatively, we construct a matrix $M \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ and vector $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$ such that Z is the orbit closure $\overline{\langle M\rangle \cdot \boldsymbol{v}}$.

Suppose that in Steps 1 and 2 of the procedure in Theorem 2 we nondeterministically guess

$$H_\Lambda \coloneqq \{(x,y) \in \mathbb{G}_m^2 : x^2 - y = 0\} \text{ and } T \coloneqq \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Let $g = (5, 25)$ be a topological generator of H_Λ, and let

$$D = \begin{pmatrix} 5 & 0 & 0 & 0 \\ 0 & 5 & 0 & 0 \\ 0 & 0 & 5 & 0 \\ 0 & 0 & 0 & 25 \end{pmatrix},$$

which satisfies $DT = T\Delta(g)$.

For Step 3, we would like to find a set of invertible matrices $P \subseteq \mathrm{GL}_4(\overline{\mathbb{Q}})$ (with $P^{-1} = (p_{ij})_{\{1\le i,j\le 4\}}$) and a matrix

$$U \coloneqq \begin{pmatrix} 1 & \lambda & 0 & 0 \\ 0 & 1 & \lambda & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

for some $\lambda \in \overline{\mathbb{Q}}$, such that

$$V(P^{-1} \cdot I) = \overline{\{\exp(t \log U)Th : h \in H_\Lambda, t \in \overline{\mathbb{Q}}\}}.$$

Note that

$$\exp(t \log U) = \begin{pmatrix} 1 & t\lambda & \frac{t(t-1)}{2}\lambda^2 & 0 \\ 0 & 1 & t\lambda & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}; \quad \text{thus} \quad \exp(t \log U)T \begin{pmatrix} x \\ x^2 \end{pmatrix} = \begin{pmatrix} \frac{t(t-1)}{2}\lambda^2 x \\ t\lambda x \\ x \\ x^2 \end{pmatrix}.$$

The ideal defining $\overline{\{\exp(t \log U)Th : h \in H_\Lambda, t \in \overline{\mathbb{Q}}\}}$ is $H \coloneqq \langle x_3^2 - x_4, x_1x_3 - \frac{1}{2}x_2^2 + \frac{\lambda}{2}x_2x_3\rangle$. Consider the ideal

$$I_{P^{-1}} = \langle F_1(P^{-1}X), F_2(P^{-1}X)\rangle \subseteq \mathbb{Q}[P^{-1} = (p_{ij}), \lambda, y]/\langle (\det P^{-1})y - 1\rangle[x_1, x_2, x_3, x_4].$$

By applying Algorithm ContainmentIso[4] and eliminating λ we obtain the following ideal, defining the set of admissible choices of P^{-1}:

$$\begin{aligned} J_{P^{-1}} \coloneqq \langle & p_{34}, p_{31}, p_{24}, p_{22}, p_{21}, p_{13} + p_{43}, p_{12} + p_{42}, \\ & p_{11} + p_{41}, p_{33}p_{44}, p_{32}p_{44}, p_{23}p_{44}, p_{14}p_{44} + p_{44}^2, \\ & p_{23}p_{33} + 10p_{33}^2 + 10p_{23}p_{43}, 2p_{32}^2 + p_{23}p_{41}, p_{23}^2 - p_{14} - p_{44}\rangle. \end{aligned}$$

One may choose, for example,

$$P^{-1} = \begin{pmatrix} \frac{1}{2} & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & \frac{1}{2} & 0 & 0 \\ -\frac{1}{2} & 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad \lambda = -\frac{1}{5}.$$

Thus we realise Z as the orbit closure $\overline{\langle M\rangle \cdot \boldsymbol{v}}$, with $M = P^{-1}UDP$ and $\boldsymbol{v} = P^{-1}T\mathbf{1}$, that is

$$M = \begin{pmatrix} 25 & 0 & -1 & 20 \\ 0 & 5 & 0 & 0 \\ 0 & -\frac{1}{2} & 5 & 0 \\ 0 & 0 & 1 & 5 \end{pmatrix} \quad \text{and} \quad \boldsymbol{v} = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}.$$

[4]The algorithm ContainmentIso inputs two ideals I_1 and I_2 and outputs the locus of points P^{-1} for which $P^{-1} \cdot I_1 \subseteq P^{-1} \cdot I_2$. This algorithm thereby solves a generalisation of the ideal membership algorithm since it determines the containment of an ideal into another after a change of variables. Clearly ContainmentIso can also be applied to determine equality after a change of variables, since $P^{-1} \cdot I_1 = P^{-1} \cdot I_2$ if and only if $P^{-1} \cdot I_1 \subseteq P^{-1} \cdot I_2$ and $P^{-1} \cdot I_2 \subseteq P^{-1} \cdot I_1$. See [28, Algorithm 2.9] for more details.

◀

5 Algorithms to Compute Generators

In Section 3 we gave algorithms to determine whether a given variety is the Zariski closure of a commutative matrix group. These procedures can also be used to find a set of generators of such a group, by quantifier elimination in the theory of algebraically closed fields. In this section, we describe specialised algorithms that can be directly implemented using standard computer-algebra software. The cases we consider compute a generator when the input variety is the Zariski closure of a cyclic group. These algorithms rely on Gröbner-basis techniques. The first algorithm finds a semisimple generator, if one exists, while the second algorithm finds a generator in the general case.

Let $I \subseteq \overline{\mathbb{Q}}[\boldsymbol{x}]$ be an ideal and denote by $\sqrt{I}$ the *radical* of I, defined as

$$\sqrt{I} \coloneqq \{f \in \overline{\mathbb{Q}}[\boldsymbol{x}] \mid f^n \in I \text{ for some } n \in \mathbb{N}\}.$$

By Hilbert's Nullstellensatz, the ideal of all polynomials that vanish on $V(I) \in \overline{\mathbb{Q}}^d$ is $\sqrt{I}$. The ideal I is *primary* if $fg \in I$ with $f, g \in \overline{\mathbb{Q}}[\boldsymbol{x}]$ implies that $f \in I$ or $g^n \in I$ for some $n \in \mathbb{N}$, and it is *prime* if it satisfies the stronger condition that $fg \in I$ only if $f \in I$ or $g \in I$. Recall that the radical of a primary ideal is necessarily prime.

A polynomial ideal I can be written as the intersection of primary ideals, giving the so-called *primary decomposition* of I. It is known that there exists a unique irredundant primary decomposition $I = \bigcap_{i=1}^{\ell} Q_i$, that is, a finite set $\{Q_1, \ldots, Q_\ell\}$ of primary ideals such that

- the prime ideals $\sqrt{Q_i}$ are all distinct; and
- $\bigcap_{i \neq j} Q_i \nsubseteq Q_j$ holds for all $j \in \{1, \ldots, \ell\}$.

The prime ideals, the $\sqrt{Q_i}$'s, are called the *associated primes* of I. An associated prime $\sqrt{Q}$ of the ideal I is called *minimal* if it does not contain any other associated primes of I.

Both algorithms take as input a variety Z, given as the zero set of an ideal $I \subseteq \mathbb{Q}[X]$, $X = \{x_{i,j}\}_{1 \leq i,j \leq d}$. We will assume *a priori* that Z is a commutative subgroup of $\mathrm{GL}_d(\overline{\mathbb{Q}})$. Verifying that Z is commutative entails first checking that Z is closed under matrix multiplication (which implies closure under matrix inversion), which amounts to showing that

$$F(XY) \in \sqrt{I(X) + I(Y)}$$

for all polynomials F in I, where $X = \{x_{i,j}\}_{1 \leq i,j \leq d}$, $Y = \{y_{i,j}\}_{1 \leq i,j \leq d}$. Commutativity is captured by showing that

$$XY - YX \in \sqrt{I(X) + I(Y)}\,.$$

5.1 Semisimple Generator

In the following we describe a procedure that, given an ideal $I \subseteq \mathbb{Q}[X], X = \{x_{i,j}\}_{1 \leq i,j \leq d}$, determines whether there exists a semisimple matrix $M \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that I is the vanishing ideal of the group $\langle M \rangle$ and which moreover outputs such an M in case the answer is "yes". We show that if such an M exists then it can be chosen such that its eigenvalues lie in the number field $\mathbb{Q}(\zeta_q)$, where q is the number of minimal associated primes of I and ζ_q is a primitive qth root of unity.

Let the input ideal I be generated by a finite collection of polynomials $F_1, \ldots, F_k \in \mathbb{Q}[X]$ with q minimal associated primes. Write $Z \coloneqq V(I)$ for the zero locus of I, assumed to be a commutative linear algebraic group. The general procedure of the algorithm is depicted in Figure 1. The ideal J_0 defined in **Line 1** is an ideal of the ring $\mathbb{Q}[P, X, Q]$, where the relations $\{x_{i,j}\}_{i \neq j}$ and $PQ - \mathrm{Id}_d$ ensure that every point $(\widetilde{P}, \widetilde{X}, \widetilde{Q}) \in V(J_0)$ comprises a diagonal matrix $\widetilde{X}$ and an invertible matrix $\widetilde{P}$ with

Fig. 1. A procedure for the group determination problem of cyclic groups, specific to semisimple generators.

Cyclic Groups: Semisimple Generator	
Input:	An ideal $I \subseteq \mathbb{Q}[X]$ generated by $F_1, \ldots, F_k \in \mathbb{Q}[X]$, $X = \{x_{i,j}\}_{1 \le i,j \le d}$, with q minimal associated primes, such that $V(I)$ is a commutative linear algebraic group.
Output:	Determine whether there exists a semisimple matrix $M \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that $V(I) = \overline{\langle M \rangle}$. If "yes", output such a matrix M.
Line 1:	Define the ideal $J_0 \subseteq \mathbb{Q}[P, X, Q]$ as follows $J_0 \coloneqq \langle F_1(PXQ), \ldots, F_k(PXQ), PQ - \mathrm{Id}_d, \{x_{i,j}\}_{i \neq j} \rangle$ where $P = \{p_{i,j}\}_{1 \le i,j \le d}$, $X = \{x_{i,j}\}_{1 \le i,j \le d}$, and $Q = \{q_{i,j}\}_{1 \le i,j \le d}$.
Line 2:	Write $J \coloneqq \sqrt{J_0 \cap \mathbb{Q}[X]}$. Compute the unique irredundant primary decomposition $J = \bigcap_{s \in S} \mathcal{P}_s$.
Line 3:	Check whether all primary components $\mathcal{P}_s$ of J are binomials using Gröbner basis computation; **return** "no" if this test fails.
Line 4:	Let $\mathcal{P}_0$ be one of the primary components of J such that $\mathrm{Id}_d \in V(\mathcal{P}_0)$.
Line 5:	Following Proposition 9, construct a rational diagonal matrix D for $\mathcal{P}_0$, such that all the entries of $\sqrt[q]{D}$ lie in $\mathbb{Q}[\zeta_q]$. Write $D_q \coloneqq \sqrt[q]{D}$.
Line 6:	Check whether for all $i \in \{1, \ldots, q-1\}$, the ideal $D_q^i \cdot \mathcal{P}_0$ is a primary component of J; **return** "no" if this test fails.
Line 7:	Write $I_q \coloneqq \bigcap_{1 \le i \le q} D_q^i \cdot \mathcal{P}_0$.
Line 8:	Check whether $J = \bigcap_{\sigma \in S_d} M_\sigma I_q M_\sigma^{-1}$ where M_σ is the permutation matrix corresponding to $\sigma \in S_d$; **return** "no" if this test fails.
Line 9:	Define the ideal $J_1 \coloneqq \langle F_1(QD_qP), \ldots, F_k(QD_qP), PQ - I \rangle \cap \mathbb{Q}[P]$. Pick $\widetilde{P} \in V(J_1)$.
Line 10:	Check whether $I = \widetilde{P} I_q \widetilde{P}^{-1}$; **return** "no" if this test fails.
Return:	"yes" together with the matrix $\widetilde{P} D_q \widetilde{P}^{-1}$.

$\widetilde{P}^{-1} = \widetilde{Q}$ satisfying $\widetilde{P}\widetilde{X}\widetilde{P}^{-1} \in Z$. The aim is to find a single such point $(\widetilde{P}, \widetilde{X}, \widetilde{Q}) \in V(J_0)$ satisfying

$$Z = \overline{\langle \widetilde{P}\widetilde{X}\widetilde{P}^{-1} \rangle}\,.$$

Subsequently, the ideal J defined in **Line 2** contains all diagonal conjugates of each matrix in Z.

In particular, for each matrix $M \in Z$ not only does the diagonal matrix D satisfying $M = \widetilde{P} D \widetilde{P}^{-1}$ lie in $V(J)$, but also all the diagonal matrices D for which $M_\sigma D M_\sigma^{-1}$ and the permutation $\sigma \in S_d$ lie in $V(J)$.

Due to this fact, we cannot simply employ Proposition 9 to construct a generator for J. Instead, in **Line 4**, we isolate a primary component $\mathcal{P}_0$ of J containing Id_d. In the following line, we apply Proposition 9 to the binomial ideal $\mathcal{P}_0$ and construct a diagonal matrix D such that $V(\mathcal{P}_0) =$

$\overline{\langle D \rangle}$. Since $V(\mathcal{P}_0)$ is irreducible, the matrix D can be chosen to have rational entries, from which it follows that the entries of $\sqrt[q]{D}$ lie in $\mathbb{Q}[\zeta_q]$.

The assertion in **Line 6** verifies whether the orbit of D_q cycles between the irreducible components of $V(J)$; this ensures that $V(I_q) = \overline{\langle D_q \rangle}$ is included in $V(J)$, where I_q is defined in **Line 7**. Next, our procedure checks whether J is equal to the intersection of $M_\sigma I_q M_\sigma^{-1}$. The necessity of the latter test is due to the above-mentioned fact that $V(J)$ contains all diagonal conjugates of each matrix in Z; see Example 22. The rest of the algorithm is straightforward.

Example 22. Let $F_1 := 2z + w$, $F_2 := 2x - 2y + 3w$, and $F_3 := 4y^2 - 4yw + w^2 - 4y + 4w$. Consider the following ideal as an input to the procedure in Figure 1:

$$I := \langle F_1, F_2, F_3 \rangle \subseteq \mathbb{Q}\left[\begin{pmatrix} x & z \\ w & y \end{pmatrix}\right].$$

The ideal I is prime (meaning that $q = 1$) and $V(I)$ is a commutative linear algebraic group. The output of our procedure shows that there exists M such that $V(I) = \overline{\langle M \rangle}$, and such that the eigenvalues of M lie in $\mathbb{Q}$.

Following the algorithms, the ideal J defined in **Line 2** has two primary components

$$\mathcal{P} := \langle w, z, y^2 - x \rangle \qquad \text{and} \qquad \mathcal{P}' := \langle w, z, x^2 - y \rangle.$$

Since $I_2 \in V(\mathcal{P} \cap \mathcal{P}')$, we can pick any of these ideals as $\mathcal{P}_0$ in **Line 4**. Following Proposition 9 in **Line 5**, we may construct diagonal matrices $D = \Delta(4, 2)$ and $D' = \Delta(2, 4)$ such that

$$V(\mathcal{P}) = \overline{\langle \Delta(4, 2) \rangle} \qquad \text{and} \qquad V(\mathcal{P}') = \overline{\langle \Delta(2, 4) \rangle}.$$

Clearly, matrices D and D' are conjugates under permutation of diagonals, implying that the assertion in **Line 8** holds. (The above is an indication (1) that permutations of matrices arising from one choice of $\mathcal{P}_0$ under M_σ are suitable for other possible choices of $\mathcal{P}_0$, and (2) the necessity of the check in **Line 8**). Following **Line 9** for $D_q = \Delta(2, 4)$, defining the ideal

$$J_1 := \langle F_1(QD_qP),\ F_2(QD_qP),\ F_3(QD_qP) \rangle \cap \mathbb{Q}[P],$$

we have that

$$J_1 = \langle p_3 - 2p_4, p_1 - p_2 \rangle \subseteq \mathbb{Q}\left[\begin{pmatrix} p_1 & p_2 \\ p_3 & p_4 \end{pmatrix}\right].$$

Subsequently, one choice for a semisimple generator of $V(I)$ is the following matrix M:

$$M := \begin{pmatrix} 6 & 2 \\ -4 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ -2 & -1 \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ -2 & -1 \end{pmatrix}^{-1}.$$

◀

5.2 General Generator

We employ the algorithm from the previous subsection to provide a procedure that, given an algebraic set $Z \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ determines whether there exists a matrix $M \in \mathrm{GL}_d(\overline{\mathbb{Q}})$ such that $Z = \overline{\langle M \rangle}$ and which moreover outputs such an M in the affirmative case.

Let the input ideal I be generated by a finite collection of polynomials $F_1, \ldots, F_k \in \mathbb{Q}[X]$ with q minimal associated primes. Write $Z := V(I)$ for the zero locus of I that is, by assumption, a commutative linear algebraic group. Our algorithm first calls a (modified variant of) the procedure in Figure 1, with the input ideal I, to check whether the subgroup G_s of all semisimple matrices in Z is one-generated. The modification is as follows: (1) the assertion in **Line 10** is omitted (as this assertion requires that Z is generated with a single semisimple matrix), and (2) the algorithm outputs D_q and the ideal J_1 defining the locus point of a suitable P.

After the above subprocedure returns, our algorithm proceeds by verifying that the subgroup G_u of all unipotent matrices is one-generated. For this purpose, it checks

- whether $V(I + \langle (X - \mathrm{Id}_d)^n \rangle)$ is a commutative linear algebraic group; and
- whether $V(I + \langle (X - \mathrm{Id}_d)^n \rangle)$ is one-dimensional.

The algorithm returns "no" if either of the subgroups G_s or G_u is not one-generated. Otherwise, the procedure defines the ideal $H \subseteq \mathbb{Q}[P, X, Q]$ by

$$H \coloneqq \langle\, F_1(PXQ),\, \ldots,\, F_k(PXQ),\, PQ - \mathrm{Id}_d,\, \{x_{i,j}\}_{j \neq i, i+1} \,\rangle$$

where

$$P = (p_{i,j})_{1 \leq i,j \leq d}, \qquad X = \{x_{i,j}\}_{1 \leq i,j \leq d} \qquad \text{and} \qquad Q = (q_{i,j})_{1 \leq i,j \leq d}\,.$$

It returns "yes" together with the matrix $\widetilde{P} D_q \widetilde{X} \widetilde{P}^{-1}$ where $(\widetilde{P}, \widetilde{X}) \in J + H$.

Acknowledgments

Mahsa Shirmohammadi and Rida Ait El Manssour were supported by the ANR grant VeSyAM (ANR-22-CE48-0005). James Worrell was supported by EPSRC Fellowship EP/X033813/1.

This paper is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 10103444). Anton Varonka gratefully acknowledges the support of the ERC consolidator grant ARTIST 101002685.

References

[1] R. Ait El Manssour, G. Kenison, M. Shirmohammadi, and A. Varonka. 2025. Simple linear loops: algebraic invariants and applications. *Proc. ACM Program. Lang.*, 9, POPL, Article 26, 27 pages. DOI: 10.1145/3704862.

[2] L. Babai, R. Beals, J.-Y. Cai, G. Ivanyos, and E. M. Luks. 1996. Multiplicative equations over commuting matrices. In *Proceedings of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, 28-30 January 1996, Atlanta, Georgia*. 498–507.

[3] S. Basu, R. Pollack, and M.-F. Roy. 2006. *Algorithms in real algebraic geometry*. (Second ed.). Vol. 10. Springer-Verlag, Berlin, x+662.

[4] S. Basu, R. Pollack, and M.-F. Roy. 1996. On the combinatorial and algebraic complexity of quantifier elimination. *J. ACM*, 43, 6, 1002–1045. DOI: 10.1145/235809.235813.

[5] M. Bläser, C. Ikenmeyer, V. Lysikov, A. Pandey, and F.-O. Schreyer. 2021. On the orbit closure containment problem and slice rank of tensors. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2565–2584.

[6] E. Bombieri and W. Gubler. 2006. *Heights in Diophantine Geometry*. Cambridge University Press.

[7] A. Borel. 1991. *Linear Algebraic Groups*. Springer New York. DOI: 10.1007/978-1-4612-0941-6.

[8] P. Bürgisser. 2024. Completeness classes in algebraic complexity theory. *arXiv preprint arXiv:2406.06217*.

[9] P. Bürgisser, M. L. Doğan, V. Makam, M. Walter, and A. Wigderson. 2021. Polynomial Time Algorithms in Invariant Theory for Torus Actions. In *36th Computational Complexity Conference (CCC 2021)*. V. Kabanets, (Ed.) Vol. 200. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 32:1–32:30. DOI: 10.4230/LIPIcs.CCC.2021.32.

[10] P. Bürgisser, J. M. Landsberg, L. Manivel, and J. Weyman. 2011. An overview of mathematical issues arising in the geometric complexity theory approach to vp≠vnp. *SIAM J. Comput.*, 40, 4, 1179–1209. DOI: 10.1137/090765328.

[11] A. L. Chistov and D. Y. Grigor'Ev. 1984. Complexity of quantifier elimination in the theory of algebraically closed fields. In *International Symposium on Mathematical Foundations of Computer Science*. Springer, 17–31.

[12] T. Combot. 2025. Computing linear relations between polynomial roots. *Mathematics of Computation*.

[13] D. A. Cox, J. B. Little, and D. O'Shea. 2015. *Ideals, varieties, and algorithms*. (Fourth ed.). An introduction to computational algebraic geometry and commutative algebra. Springer, Cham, xvi+646. DOI: 10.1007/978-3-319-16721-3.

[14] J. Cyphert and Z. Kincaid. 2024. Solvable polynomial ideals: the ideal reflection for program analysis. *Proc. ACM Program. Lang.*, 8, POPL, 724–752. DOI: 10.1145/3632867.

[15] H. Derksen, E. Jeandel, and P. Koiran. 2005. Quantum automata and algebraic groups. *J. Symb. Comput.*, 39, 3-4, 357–371.

[16] H. Derksen and V. Makam. 2020. Algorithms for orbit closure separation for invariants and semi-invariants of matrices. *Algebra & Number Theory*, 14, 10, 2791–2813. DOI: 10.2140/ant.2020.14.2791.

[17] M. A. Forbes and A. Shpilka. 2013. Explicit noether normalization for simultaneous conjugation via polynomial identity testing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. P. Raghavendra, S. Raskhodnikova, K. Jansen, and J. D. P. Rolim, (Eds.) Springer Berlin Heidelberg, 527–542.

[18] F. Galuppi and M. Stanojkovski. 2021. Toric varieties from cyclic matrix semigroups. *Rend. Istit. Mat. Univ. Trieste*. DOI: 10.13137/2464-8728/33099.

[19] D. R. Grayson and M. E. Stillman. [n. d.] Macaulay2, a software system for research in algebraic geometry. Available at http://www2.macaulay2.com. ().

[20] S. Hitarth, G. Kenison, L. Kovács, and A. Varonka. 2024. Linear Loop Synthesis for Quadratic Invariants. In *41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024)*. O. Beyersdorff, M. M. Kanté, O. Kupferman, and D. Lokshtanov, (Eds.) Vol. 289. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 41:1–41:18. DOI: 10.4230/LIPIcs.STACS.2024.41.

[21] R. A. Horn and C. R. Johnson. 2012. *Matrix Analysis*. (2nd ed.). Cambridge University Press.

[22] E. Hrushovski, J. Ouaknine, A. Pouly, and J. Worrell. 2023. On strongest algebraic program invariants. *Journal of the ACM*, 70, 5, 1–22.

[23] E. Hrushovski, J. Ouaknine, A. Pouly, and J. Worrell. 2018. Polynomial invariants for affine programs. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. Association for Computing Machinery, 530–539. DOI: 10.1145/3209108.3209142.

[24] A. Humenberger, D. Amrollahi, N. Bjørner, and L. Kovács. 2022. Algebra-Based Reasoning for Loop Synthesis. *Form. Asp. Comput.*, 34, 1, Article 4, 31 pages. DOI: 10.1145/3527458.

[25] A. Humenberger, N. S. Bjørner, and L. Kovács. 2020. Algebra-Based Loop Synthesis. In *Integrated Formal Methods - 16th International Conference, IFM 2020, Lugano, Switzerland, November 16-20, 2020, Proceedings*. B. Dongol and E. Troubitsyna, (Eds.) Vol. 12546. Springer, 440–459. DOI: 10.1007/978-3-030-63461-2_24.

[26] J. E. Humphreys. 1975. *Linear Algebraic Groups*. Vol. 21. Springer. DOI: 10.1007/978-1-4684-9443-3.

[27] R. Kannan and R. J. Lipton. 1986. Polynomial-time algorithm for the orbit problem. *J. ACM*, 33, 4, 808–821. DOI: 10.1145/6490.6496.

[28] L. Katthän, M. Michałek, and E. Miller. 2017. When is a polynomial ideal binomial after an ambient automorphism? *Foundations of Computational Mathematics*, 19. DOI: 10.1007/s10208-018-9405-0.

[29] G. Kenison, L. Kovács, and A. Varonka. 2023. From Polynomial Invariants to Linear Loops. In *Proceedings of the 2023 International Symposium on Symbolic and Algebraic Computation, ISSAC 2023, Tromsø, Norway, July 24-27, 2023*. A. Dickenstein, E. P. Tsigaridas, and G. Jeronimo, (Eds.) ACM, 398–406. DOI: 10.1145/3597066.3597109.

[30] Z. Kincaid, J. Cyphert, J. Breck, and T. W. Reps. 2018. Non-linear reasoning for invariant synthesis. *Proc. ACM Program. Lang.*, 2, POPL, 54:1–54:33. DOI: 10.1145/3158142.

[31] S. Lang. 2002. *Algebra*. Springer New York. DOI: 10.1007/978-1-4613-0041-0.

[32] M. Müller-Olm and H. Seidl. 2004. Computing polynomial program invariants. *Inf. Process. Lett.*, 91, 5, 233–244. DOI: 10.1016/J.IPL.2004.05.004.

[33] D. Mumford. 1999. *The Red Book of Varieties and Schemes*. Springer Berlin Heidelberg. DOI: 10.1007/b62130.

[34] M. Newman. 1967. Two classical theorems on commuting matrices. *Journal of Research of the National Bureau of Standards Section B Mathematics and Mathematical Physics*, 69.

[35] K. Nosan, A. Pouly, S. Schmitz, M. Shirmohammadi, and J. Worrell. 2022. On the Computation of the Zariski Closure of Finitely Generated Groups of Matrices. In *Proceedings of the 2022 International Symposium on Symbolic and Algebraic Computation*, 129–138.

[36] J. Ouaknine, A. Pouly, J. S. Pinto, and J. Worrell. 2019. On the decidability of membership in matrix-exponential semigroups. *J. ACM*, 66, 3, 15:1–15:24. DOI: 10.1145/3286487.

[37] S. Sankaranarayanan, H. Sipma, and Z. Manna. 2004. Non-linear loop invariant generation using gröbner bases. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2004, Venice, Italy, January 14-16, 2004*. N. D. Jones and X. Leroy, (Eds.) ACM, 318–329. DOI: 10.1145/964001.964028.

[38] H. Seidl, S. Maneth, and G. Kemper. 2015. Equivalence of deterministic top-down tree-to-string transducers is decidable. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*. V. Guruswami, (Ed.) IEEE Computer Society, 943–962. DOI: 10.1109/FOCS.2015.62.

A Extended Background

Here we give an extended preliminary section, which collects together background material and definitions relevant to the paper.

A.1 Ideals and Varieties

Let $\overline{\mathbb{Q}}$ denote the field of algebraic numbers and write $\overline{\mathbb{Q}}[x_1, \ldots, x_d]$ for the ring of polynomials with coefficients in $\overline{\mathbb{Q}}$ over the variables $x_1, \ldots, x_d$. A polynomial ideal I is an additive subgroup of $\overline{\mathbb{Q}}[x_1, \ldots, x_d]$ that is closed under multiplication in $\overline{\mathbb{Q}}[x_1, \ldots, x_d]$. Given a finite collection of polynomials $S \subseteq \overline{\mathbb{Q}}[x_1, \ldots, x_d]$, we denote by $\langle S \rangle$ the ideal generated by S.

We briefly describe useful algebraic operations for ideals. The *sum* of I and J, which we denote by $I + J$, is the ideal

$$I + J \coloneqq \{f + g : f \in I \text{ and } g \in J\}.$$

The sum $I + J$ is the smallest ideal containing both I and J and, in addition, if $I = \langle f_1, \ldots, f_r \rangle$ and $J = \langle g_1, \ldots, g_s \rangle$, then $I + J = \langle f_1, \ldots, f_r, g_1, \ldots, g_s \rangle$. The *product* of I and J, which we denote by IJ, is the ideal given by

$$IJ = \left\{ \textstyle\sum_{\ell=1}^{r} f_\ell g_\ell : f_1, \ldots, f_r \in I, g_1, \ldots, g_r \in J, r \in \mathbb{N} \right\}.$$

The *intersection* $I \cap J$ is the set of polynomials common to both I and J. Given an ideal $I \subseteq \overline{\mathbb{Q}}[X]$, where $X = \{x_{i,j}\}_{1 \le i,j \le d}$, and matrix $M \in \overline{\mathbb{Q}}^{d \times d}$, we write $M \cdot I$ for the ideal $\{f(MX) \in \overline{\mathbb{Q}}[X] : f \in I\}$.

An *algebraic set* (or *variety*) is the set of common zeroes of a finite collection of polynomials. By Hilbert's basis theorem every polynomial ideal $I \subseteq \overline{\mathbb{Q}}[x_1, \ldots, x_d]$ is finitely generated. Thus the set

$$V(I) \coloneqq \{\boldsymbol{a} \in \overline{\mathbb{Q}}^d : f(\boldsymbol{a}) = 0 \text{ for all } f \in I\}$$

is a variety. As an aside, the varieties $V(I+J)$ and $V(IJ)$ are easily shown to be equal to $V(I) \cap V(J)$ and $V(I) \cup V(J)$, respectively. Taking, as above, $M \in \overline{\mathbb{Q}}^{d \times d}$, we have

$$V(M \cdot I) = \{A \in \overline{\mathbb{Q}}^{d \times d} : MA \in V(I)\}.$$

Let $I \subseteq \overline{\mathbb{Q}}[x_1, \ldots, x_d]$ be an ideal and denote by $\sqrt{I}$ the *radical* of I, defined as

$$\sqrt{I} \coloneqq \{f \in \overline{\mathbb{Q}}[x_1, \ldots, x_d] \mid f^n \in I \text{ for some } n \in \mathbb{N}\}.$$

By Hilbert's Nullstellensatz, the ideal of all polynomials that vanish on $V(I) \in \overline{\mathbb{Q}}^d$ is $\sqrt{I}$.

A polynomial ideal I is *principal* if there exists a polynomial f for which $\langle f \rangle = I$. The ideal I is *primary* if for all $f, g \in \overline{\mathbb{Q}}[x_1, \ldots, x_d]$, if $fg \in I$ then $f \in I$ or $g^n \in I$ for some $n \in \mathbb{N}$, and it is *prime* if it satisfies the stronger condition that $fg \in I$ only if $f \in I$ or $g \in I$. Recall that the radical of a primary ideal is necessarily prime.

A polynomial ideal I can be written as the intersection of primary ideals, giving the so-called *primary decomposition* of I. It is known that there exists a unique irredundant primary decomposition $I = \bigcap_{i=1}^{\ell} Q_i$, that is, a finite set $\{Q_1, \ldots, Q_\ell\}$ of primary ideals such that (1) the prime ideals $\sqrt{Q_i}$ are all distinct; and (2) $\bigcap_{i \neq j} Q_i \not\subseteq Q_j$ holds for all $j \in \{1, \ldots, \ell\}$. The prime ideals, the $\sqrt{Q_i}$'s, are called the *associated primes* of I. An associated prime $\sqrt{Q}$ of the ideal I is called *minimal* if it does not contain any other associated primes of I.

Let k be a field. Consider a system of m equations in d variables $x_1, \ldots, x_d$ with coefficients in k:

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1d}x_d &= b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2d}x_d &= b_2 \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{md}x_d &= b_m. \end{aligned}$$

We call the solution set for such a system a *linear variety*. The *dimension* of a non-empty linear variety is the number of independent equations in its definition.

A.2 The Zariski Topology

The *Zariski topology* on $\overline{\mathbb{Q}}^d$ has as its closed sets the varieties in $\overline{\mathbb{Q}}^d$. Given a set $E \subseteq \overline{\mathbb{Q}}^d$, we denote by $\overline{E}$ the closure of E in the Zariski topology, i.e., the smallest algebraic set that contains E. A closed set $A \subseteq \overline{\mathbb{Q}}^d$ is *irreducible* if it cannot be written as the union of two closed proper subsets. A maximal irreducible closed subset of A is called an *irreducible component* of A. By Hilbert's basis theorem every closed set A can be written as a finite union of its irreducible components.

Given a closed set E, denote by $\dim E$ the dimension of the variety E, that is, the maximal length of a strictly decreasing chain of nonempty irreducible subvarieties of E. We define the dimension of an arbitrary set to be the dimension of its closure.

A.3 Linear Algebraic Groups

The general linear group $\mathrm{GL}_d(\mathbb{F})$ is the group of all $d \times d$ invertible matrices with entries in a given field $\mathbb{F}$. The orbit closure of $\boldsymbol{v} \in \mathbb{F}^d$, denoted by $\overline{G \cdot \boldsymbol{v}}$, is the closure of $G \cdot \boldsymbol{v}$ in the Zariski topology. Recall that the closed sets of the Zariski topology are algebraic sets–that is, sets of common zeros of a finite collection of polynomials.

Recall that a matrix $M \in \overline{\mathbb{Q}}^{d\times d}$ is *nilpotent* if $M^n = 0$ for some $n \in \mathbb{N}$. It is *unipotent* if $M - \mathrm{Id}_d$ is nilpotent, and *semisimple* if it is diagonalisable over $\overline{\mathbb{Q}}$. A matrix $M \in \overline{\mathbb{Q}}^{d\times d}$ is called *upper triangular* if all entries below the main diagonal are zero. We use the term *upper unitriangular* to refer to an upper triangular matrix whose entries along the main diagonal are all ones.

Write $\mathrm{GL}_d(\overline{\mathbb{Q}})$ for the group of $d\times d$ invertible matrices with entries in $\overline{\mathbb{Q}}$. We identify $\mathrm{GL}_d(\overline{\mathbb{Q}})$ with the variety $\left\{(M, y) \in \overline{\mathbb{Q}}^{d\times d} \times \overline{\mathbb{Q}} : \det(M) \cdot y = 1\right\}$. Under this identification, matrix multiplication is a polynomial map $\mathrm{GL}_d(\overline{\mathbb{Q}}) \times \mathrm{GL}_d(\overline{\mathbb{Q}}) \to \mathrm{GL}_d(\overline{\mathbb{Q}})$, and, by Cramer's rule, matrix inversion is also a polynomial map $\mathrm{GL}_d(\overline{\mathbb{Q}}) \to \mathrm{GL}_d(\overline{\mathbb{Q}})$. A *linear algebraic group* G is a Zariski-closed subgroup of $\mathrm{GL}_d(\overline{\mathbb{Q}})$. Given a linear algebraic group G, it is well-known that it has a unique irreducible component that contains the identity matrix, which we denote by G°. The cosets of G° are the irreducible components of G.

We say that G is *topologically generated by* $S \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ if G is the smallest Zariski closed subgroup of $\mathrm{GL}_d(\overline{\mathbb{Q}})$ that contains S, that is, $G = \overline{\langle S \rangle}$. If G is topologically generated by a set with s elements then we say that is G is *s-generated*. Denote by G_s the subset of semisimple matrices in G, and by G_u the subset of unipotent matrices. If G is a commutative algebraic group then G_s and G_u form algebraic subgroups; moreover we have $G = G_u \cdot G_s$.

The d-dimensional multiplicative group over $\overline{\mathbb{Q}}$ is defined by

$$\mathbb{G}_m^d = \mathbb{G}_m^d(\overline{\mathbb{Q}}) := \left\{\boldsymbol{a} \in \overline{\mathbb{Q}}^d : a_1 \cdots a_d \neq 0\right\}.$$

Here the subscript m stands for *multiplicative*. Evidently this is a commutative group with respect to pointwise multiplication.

We identify $\mathbb{G}_m^d$ with the subgroup of diagonal matrices in $\mathrm{GL}_d(\overline{\mathbb{Q}})$ via the map Δ that sends $(a_1, \ldots, a_d) \in \mathbb{G}_m^d$ to the diagonal matrix $\Delta(a_1, \ldots, a_d) \in \mathrm{GL}_d(\overline{\mathbb{Q}})$.

Given a subgroup $\Lambda \subseteq \mathbb{Z}^d$, define

$$H_\Lambda := \{\boldsymbol{a} \in \mathbb{G}_m^d : \forall \boldsymbol{v} \in \Lambda\, (a_1^{v_1} \cdots a_d^{v_d} = 1)\}.$$

The map $\Lambda \mapsto H_\Lambda$ is an isomorphism between subgroups of $\mathbb{Z}^d$ and algebraic subgroups of $\mathbb{G}_m^d$. This implies that $\mathbb{G}_m^d$ is topologically generated by any d-tuple $(g_1, \ldots, g_d)$ of multiplicatively independent elements of $\overline{\mathbb{Q}}$. It also follows that the vanishing ideal $I \subseteq \overline{\mathbb{Q}}[x_1, \ldots, x_d]$ of an algebraic

subgroup of $\mathbb{G}_m^d$ is a so-called *pure binomial ideal*; that is, an ideal generated by polynomials of the form $x_1^{\alpha_1} \cdots x_d^{\alpha_d} - x_1^{\beta_1} \cdots x_d^{\beta_d}$, where $\alpha_1, \ldots, \alpha_d$ and $\beta_1, \ldots, \beta_d$ are non-negative integers. A mere *binomial ideal* is one that is generated by polynomials of the form $x_1^{\alpha_1} \cdots x_d^{\alpha_d} - \lambda x_1^{\beta_1} \cdots x_d^{\beta_d}$, where $\lambda \in \overline{\mathbb{Q}}$.

For a $d \times d$ unipotent matrix A and nilpotent matrix B, define

$$\log(A) \coloneqq \sum_{k=1}^{d-1} (-1)^{k+1} \frac{(A-I)^k}{k} \quad \text{and} \quad \exp(B) \coloneqq \sum_{k=0}^{d-1} \frac{B^k}{k!}.$$

Let $G \subseteq \mathrm{GL}_d(\overline{\mathbb{Q}})$ be a commutative subgroup of unipotent matrices. Recall that $L \coloneqq \{\log(A) : A \in G\}$ is a linear subspace of $\overline{\mathbb{Q}}^{d^2}$ consisting of nilpotent matrices [7, Chapter II, Section 7.3].

A.4 Lattices

The *rank* of an abelian group Λ is the size of a maximal linearly independent subset [31]. A subgroup $\Lambda \subseteq \mathbb{Z}^d$ is called a *lattice* (and has rank at most d). The *torsion subgroup* of $\mathbb{Z}^d/\Lambda$ is the subgroup of $\mathbb{Z}^d/\Lambda$ consisting of all elements of finite order.

A.5 Modules over Principal Ideal Domains

Let R be a commutative ring. A *module* M is an additive abelian group together with an operation $R \times M \to M$ such that for all $r, s \in R$ and $x, y \in M$ we have

$$(r+s)x = rx + sx \quad \text{and} \quad r(x+y) = rx + ry.$$

By definition of an operation, we have $1x = x$ for all $x \in M$. A module over $\mathbb{Z}$ is abelian group and, vice versa, an abelian group is a module over $\mathbb{Z}$.

Let M be a module over a ring R and let $S \subseteq M$. The set S is a *basis* of M if S is non-empty, S generates M, and the elements of S are linearly independent. A *free module* is a module which admits a basis, or the zero module. If S is a basis of a non-zero module M, then every element of M can be uniquely expressed as a linear combination of elements of S.

An *integral domain* is a non-zero commutative ring where the product of two non-zero elements is itself non-zero. A *principal ideal domain* is an integral domain wherein every ideal is principal. Examples of principal ideal domains include $\mathbb{Z}$ and the rings $\mathbb{F}[x]$, the univariate polynomials with coefficients in the field $\mathbb{F}$.

Broadly speaking, a finitely generated module over a principal ideal domain is given by a direct sum of cyclic modules (i.e., modules with a single generator). One formulation of this structural result is the Elementary Divisors Theorem, see, for example, [31].

Theorem 4 (Elementary Divisors Theorem). *Let F be a free module over a principal ideal domain R and M a non-trivial finitely generated submodule. Then there exists a basis of F, elements $e_1, \ldots, e_m$ in this basis, and non-zero elements $a_1, \ldots, a_m \in R$ such that*

(1) $a_1 e_1, \ldots, a_m e_m$ form a basis of M over R, and
(2) $a_i \mid a_{i+1}$ for $i = 1, \ldots, m-1$.

B Proofs Omitted from Section 2

Recall that the family of *constructible sets* is the smallest class that contains the algebraic sets and is also closed under boolean operations [3, Chapter 1] and that a *morphism* is a map that is given locally by polynomials.

Fact 5. Let $G \leq \mathrm{GL}_d(\overline{\mathbb{Q}})$ be an algebraic group and $\boldsymbol{v} \in \overline{\mathbb{Q}}^d$ a vector, the Zariski and Euclidean closures of the orbit $G \cdot \boldsymbol{v}$ coincide.

Fact 5 follows from the observation that the orbit $G \cdot \boldsymbol{v}$ is the image of a variety under a morphism and is thus a constructible set by Chevalley's Theorem (see [7, Chapter AG, Corollary 10.2]). In our setting, the Zariski and Euclidean closures of a constructible set coincide [33, Chapter 1, §10, Corollary 1].

Proposition 6. *Let $G \leq \mathrm{GL}_d(\overline{\mathbb{Q}})$ be an algebraic group, then G is topologically generated by a finite set of matrices.*

Proof. Let U be the group $\overline{\langle G_u \rangle}$. It is shown in [35, Proof of Lemma 6] that U is a normal subgroup of G, and that it is topologically generated by $\dim U$ elements.

Recall that the quotient of a linear algebraic group by a normal subgroup is itself isomorphic to a linear algebraic group (possibly in higher dimension [26, Section 11.5]). By construction, the quotient G/U is a linear algebraic group that consists only of semisimple elements. Therefore, $(G/U)^\circ = G^\circ/U$ is a torus and by [18, Proposition 14] it is 1-generated.

Let $h \in G^\circ$ be such that $G^\circ/U = \overline{\langle h \rangle}$, and define $H \coloneqq \overline{\langle h, U \rangle} \subseteq G^\circ$. Then $G^\circ/U = H/U$ and hence $\dim H = \dim G^\circ$. Since G° is irreducible we have $G^\circ = H$; hence G° is topologically generated by at most $\dim U + 1$ elements.

In order to topologically generate G, it is sufficient to take the topological generators of G° and one element from every other irreducible component of G. Hence, G is topologically generated by $\dim U + |G/G^\circ|$ elements. □

Proposition 9. *Let $\Lambda \subseteq \mathbb{Z}^d$ be a lattice of rank r. Then*

1. *the torsion subgroup of $\mathbb{Z}^d/\Lambda$ is s_0-generated, for some $s_0 \leq r$, and*
2. *H_Λ is s-generated, where $s \coloneqq s_0$ if Λ has full rank and otherwise $s \coloneqq \max(s_0, 1)$. Furthermore, s is the minimal number of topological generators for H_Λ.*

Proof. Let $\Lambda \subseteq \mathbb{Z}^d$ have rank r and elementary divisors $d_1, \ldots, d_r$, where $d_i \mid d_{i+1}$ for all $i \in \{1, \ldots, r-1\}$. Write s_0 for the number of non-unit elementary divisors. Hence, there is a basis $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_d$ of $\mathbb{Z}^d$ such that Λ is generated by the vectors $d_1\boldsymbol{u}_1, \ldots, d_r\boldsymbol{u}_r$. Then the torsion subgroup of $\mathbb{Z}^d/\Lambda$ is $\mathbb{Z}/d_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d_r\mathbb{Z}$, which is s_0-generated. This shows Item 1.

The map $\varphi \colon \mathbb{G}_m^d \to \mathbb{G}_m^d$, defined by $\varphi(\boldsymbol{a}) = (\boldsymbol{a}^{\boldsymbol{u}_1}, \ldots, \boldsymbol{a}^{\boldsymbol{u}_d})$ is a Zariski-continuous group automorphism of $\mathbb{G}_m^d$ that maps H_Λ to the group $G \coloneqq \Omega_{d_1} \times \cdots \times \Omega_{d_r} \times \mathbb{G}_m^{d-r}$, where Ω_k denotes the group of all k-th roots of unity for k a positive integer. Clearly H_Λ is s-generated if and only if G is s-generated.

Write $F \coloneqq \Omega_{d_1} \times \cdots \times \Omega_{d_r}$, so $G = F \times \mathbb{G}_m^{d-r}$ and F is s_0-generated. Hence if $r = d$ then G is also s_0-generated, while if $r < d$ then G is s_0-generated unless $s_0 = 0$, in which case G is 1-generated. □