

Polynomial Invariants for Affine Programs

Ehud Hrushovski

Mathematical Institute

Oxford University, UK

ehud.hrushovski@maths.ox.ac.uk

Joël Ouaknine

Max Planck Institute for Software Systems

Saarland Informatics Campus, Germany

Department of Computer Science

Oxford University, UK

joel@mpi-sws.org

Amaury Pouly

Max Planck Institute for Software Systems

Saarland Informatics Campus, Germany

pamaury@mpi-sws.org

James Worrell

Department of Computer Science

Oxford University, UK

jbw@cs.ox.ac.uk

Abstract

We exhibit an algorithm to compute the strongest polynomial (or algebraic) invariants that hold at each location of a given affine program (i.e., a program having only non-deterministic (as opposed to conditional) branching and all of whose assignments are given by affine expressions). Our main tool is an algebraic result of independent interest: given a finite set of rational square matrices of the same dimension, we show how to compute the Zariski closure of the semigroup that they generate.

ACM Reference Format:

Ehud Hrushovski, Joël Ouaknine, Amaury Pouly, and James Worrell. 2018. Polynomial Invariants for Affine Programs. In *Proceedings of* . ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Acknowledgments

Joël Ouaknine was supported by ERC grant AVS-ISS (648701), and James Worrell was supported by EPSRC Fellowship EP/N008197/1.

1 Introduction

Invariants Invariants are one of the most fundamental and useful notions in the quantitative sciences, appearing in a wide range of contexts, from gauge theory, dynamical systems, and control theory in physics, mathematics, and engineering to program verification, static analysis, abstract interpretation, and programming language semantics (among others) in computer science. In spite of decades of scientific work and progress, automated invariant synthesis remains a topic of active research, particularly in the fields of theorem proving and program analysis, and plays a central role in methods and tools seeking to establish correctness properties of computer programs; see, e.g., [21], and particularly Sec. 8 therein.

Affine Programs Affine programs are a simple kind of non-deterministic imperative programs (which may contain arbitrarily nested loops) in which the only instructions are assignments whose right-hand sides are affine expressions, such as $x_3 := x_1 - 3x_2 + 7$. A

conventional imperative program can be abstracted to an affine program by replacing conditionals with nondeterminism and conservatively over-approximating non-affine assignments, see, e.g., [5]. In doing so, affine programs enable one to reason about more complex programs; a particularly striking example is the application of affine programs to several problems in inter-procedural analysis [5, 14, 30, 31].

Affine Invariants An affine invariant for an affine program with n variables assigns to each program location an affine subspace of $\mathbb{R}^n$ such that the resulting family of subspaces is preserved under the transition relation of the program. Such an invariant is specified by giving a finite set of affine relations at each location. The strongest (i.e., smallest) affine invariant is obtained by taking the affine hull of the set of reachable configurations (i.e., values of the program variables) at each program location. Equivalently, the strongest affine invariant is determined by giving, for each program location, the set of all affine relations holding at that location.

An algorithm due to Michael Karr in 1976 [20] computes the strongest affine invariant of an affine program. This allows to compute for any given affine relation and program location whether or not that relation holds among program variables whenever control reaches the given location. A more efficient reformulation of Karr's algorithm was given by Müller-Olm and Seidel [30], who moreover showed that if the class of affine programs is augmented with equality guards then it becomes undecidable whether or not a given affine relation holds at a particular program location. A randomised algorithm for discovering affine relations was proposed by Gulwani and Necula [14].

Polynomial Invariants A natural and more expressive generalisation of affine invariants are polynomial invariants. A polynomial invariant assigns to each program location an algebraic set (or variety, i.e., positive Boolean combination of polynomial equalities) such that the resulting family is preserved under the transition relation of the program. A polynomial invariant is specified by giving a set of polynomial relations that hold at each program location. The strongest (i.e., smallest) polynomial invariant is obtained by taking the Zariski closure of the set of reachable configurations in each location.

The problem of computing polynomial invariants for affine programs and related formalisms has been extensively studied over the past fifteen years, see, e.g., [6, 7, 11, 12, 16, 19, 21, 23, 25, 34–36]. However, in contrast to the case of affine invariants, as of yet

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

no method is known to compute the strongest polynomial invariant, i.e., (a basis for) the set of all polynomial relations holding at each location of a given affine program. Existing methods are either heuristic in nature, or only known to be complete relative to restricted classes of invariants or programs. For example, it is shown in [30] (see also [34]) that Karr’s algorithm can be applied to compute the smallest polynomial invariant that is specified by polynomial relations of a fixed degree d . (The case of affine invariants corresponds to $d = 1$.) The paper [11] gives a method that finds all polynomial invariants for a highly restricted class of affine programs (in which all linear mappings have positive rational eigenvalues). The approach of [16, 23] via so-called P-solvable loops does not encompass the whole class of affine programs either (although it does allow to handle certain classes of programs with polynomial assignments) [24].

Main Contribution In this paper we give a method to compute the set of *all* polynomial relations that hold at a given location of an affine program, or in other words the strongest polynomial invariant. The output of the algorithm gives for each program location a finite basis of the ideal of all polynomial relations holding at that location. Given such a basis one can test for any given polynomial relation whether or not that relation holds at the given location.

Our main tool is an algebraic result of independent interest: we give an algorithm that, given a finite set of rational square matrices of the same dimension, computes the Zariski closure of the semigroup that they generate. Our algorithm generalises (and uses as a subroutine) an algorithm of Derksen, Jeandel, and Koiran [13] to compute the Zariski closure of a finitely generated group of invertible matrices. Related to this, Corollary 3.7 and Lemma 3.6a in [15] reduce the question of computing the Zariski closure of a finitely generated group of invertible matrices to that of finding multiplicative relations among diagonal matrices. Note that if one begins with rational matrices, then such relations can be found simply using prime decomposition of the entries.

Our procedure for computing the Zariski closure of a matrix semigroup also generalises a result of Mandel and Simon [26] and, independently, Jacob [17, 18], to the effect that it is decidable whether a finitely generated semigroup of rational matrices is finite. Note that a variety that is given as the zero set of a polynomial ideal $I \subseteq \mathbb{K}[x_1, \dots, x_n]$ is finite just in case the quotient $\mathbb{K}[x_1, \dots, x_n]/I$ is finite-dimensional as a vector space over $\mathbb{K}$ [10, Chapter 5, Sec. 3]). The latter condition can be checked by computing a Gröbner basis for I .

As mentioned above, we make use of the result of [13] that one can compute the Zariski closure of the group generated by a finite set of invertible matrices. That result itself relies on several non-trivial mathematical ingredients, including results of Masser [28] on computing multiplicative relations among given algebraic numbers and Schur’s theorem that every finitely generated periodic subgroup of the general linear group $\mathrm{GL}_n(\mathbb{C})$ is finite.

Given a set of matrices $A \subseteq M_n(\mathbb{C})$ we leverage these group-theoretic results to compute the Zariski closure $\overline{\langle A \rangle}$ of the generated semigroup $\langle A \rangle$. To this end we use multilinear algebra and structural properties of matrix semigroups to identify finitely many subgroups of $\langle A \rangle$ that can be used to generate the entire semigroup. Pursuing this approach requires that we first generalise the result of [13] to show that one can compute the Zariski closure of the group

generated by a constructible (as opposed to finite) set of invertible matrices.

It is worth pointing out that whether a particular configuration is reachable at a certain program location of a given affine program is in general an undecidable problem—this follows quite straightforwardly from the undecidability of the membership problem for finitely generated matrix semigroups, discussed shortly. It is therefore somewhat remarkable that the Zariski closure (i.e., the smallest algebraic superset) of the set of reachable configurations at any particular location nevertheless turns out to be a computable object.

Matrix Semigroups and Automata Decision problems for matrix semigroups have also been studied for many decades, independently of program analysis. One of the most prominent such is the Membership Problem, i.e., whether a given matrix belongs to a finitely generated semigroup of integer matrices. An early and striking result on this topic is due to Markov, who showed undecidability of the Membership Problem in dimension 6 in 1947 [27]. This was improved to dimension 3 by Paterson in 1970 [32], and it is astonishing that decidability in dimension 2 remains open to this day. A breakthrough was achieved in 2017 by Potapov and Semukhin, who showed decidability of membership for semigroups generated by *nonsingular* integer 2×2 matrices [33]. By contrast, the Membership Problem was shown to be polynomial-time decidable in any dimension by Babai *et al.* for *commuting* matrices over algebraic numbers [1]. As aptly noted by Stillwell, “noncommutative semigroups are hard to understand” [39]. Matrix semigroup theory also plays a central role in the analysis of weighted automata (such as probabilistic and quantum automata, see, e.g., [4, 13]).

Abstract Interpretation Polynomial invariants are stronger (i.e., more precise) than affine invariants. Various other types of domains have been considered in the setting of abstract interpretation, e.g., intervals, octagonal sets, and polyhedra (see, e.g., [8, 9, 29] and references in [5]). The precision of such domains in general is incomparable to that of polynomial invariants.

2 Mathematical Background

2.1 Linear Algebra

Matrices. Let $\mathbb{K}$ be a field. We denote by $M_n(\mathbb{K})$ the semigroup of square matrices of dimension n with entries in $\mathbb{K}$. We write $\mathrm{GL}_n(\mathbb{K})$ for the subgroup of $M_n(\mathbb{K})$ comprising all invertible matrices. Given a set of matrices $A \subseteq M_n(\mathbb{K})$, we denote by $\langle A \rangle$ the sub-semigroup of $M_n(\mathbb{K})$ generated by A . The rank of a matrix A is denoted by $\mathrm{rk}(A)$, its kernel by $\ker(A)$, and its image by $\mathrm{im}(A)$.

Exterior Algebra and the Grassmannian. Given a vector space V over the field $\mathbb{K}$, its exterior algebra ΛV is a vector space that embeds V and is equipped with an associative, bilinear, and anti-symmetric map

$$\wedge : \Lambda V \times \Lambda V \rightarrow \Lambda V.$$

We can construct ΛV as a direct sum

$$\Lambda V = \Lambda^0 V \oplus \Lambda^1 V \oplus \Lambda^2 V \cdots,$$

where $\Lambda^r V$ denotes the r^{th} -exterior power of V for $r \in \mathbb{N}$, that is, the subspace of ΛV generated by r -fold wedge products $v_1 \wedge \dots \wedge v_r$ for $v_1, \dots, v_r \in V$. If V is finite dimensional, with basis $e_1, \dots, e_n$, then a basis of $\Lambda^r V$ is given by $e_{i_1} \wedge \dots \wedge e_{i_r}$, $1 \leq i_1 < \dots < i_r \leq n$. Thus $\Lambda^r V$ has dimension $\binom{n}{r}$ (where $\binom{n}{r} = 0$ for $r > n$).

A basic property of the wedge product is that given vectors $u_1, \dots, u_r \in V$, $u_1 \wedge \dots \wedge u_r \neq 0$ if and only if $\{u_1, \dots, u_r\}$ is a linearly independent set. Furthermore given $w_1, \dots, w_r \in V$ we have that $u_1 \wedge \dots \wedge u_r$ and $w_1 \wedge \dots \wedge w_r$ are scalar multiples of each other iff $\text{span}(u_1, \dots, u_r) = \text{span}(w_1, \dots, w_r)$.

The *Grassmannian* $\text{Gr}(r, n)$ is the set of r -dimensional subspaces of $\mathbb{C}^n$. By the above-stated properties of the wedge product there is an injective function

$$\iota : \text{Gr}(r, n) \rightarrow \mathbb{P}(\Lambda^r(\mathbb{C}^n))$$

$$\text{span}(v_1, \dots, v_r) \mapsto [v_1 \wedge \dots \wedge v_r]$$

mapping $\text{Gr}(r, n)$ into projective space over $\Lambda^r(\mathbb{C}^n)$. Given subspaces $W_1, W_2 \subseteq V$ we moreover have $W_1 \cap W_2 = 0$ iff $\iota(W_1) \wedge \iota(W_2) \neq 0$.

2.2 Algebraic Geometry

In this section we summarise some basic notions of algebraic geometry that will be used in the rest of the paper.

Let $\mathbb{K}$ be a field. An *affine variety* $X \subseteq \mathbb{K}^n$ is the set of common zeros of a finite collection of polynomials, i.e., a set of the form

$$X = \{x \in \mathbb{K}^n : p_1(x) = p_2(x) = \dots = p_\ell(x) = 0\},$$

where $p_1, \dots, p_\ell \in \mathbb{K}[x_1, \dots, x_n]$. Given a polynomial ideal $I \subseteq \mathbb{K}[x_1, \dots, x_n]$, by Hilbert's basis theorem the set

$$V(I) = \{x \in \mathbb{K}^n : \forall p \in I, p(x) = 0\}$$

is a variety, called the *variety of I* . The two main varieties of interest to us are $X = M_n(\mathbb{K})$, which we identify with affine space $\mathbb{K}^{n^2}$ in the natural way, and $X = \text{GL}_n(\mathbb{K})$, which we identify with the variety

$$\{(A, y) \in \mathbb{K}^{n^2+1} : \det(A) \cdot y = 1\}.$$

Given an affine variety $X \subseteq \mathbb{K}^n$, the *Zariski topology* on X has as closed sets the subvarieties of X , i.e., those sets $A \subseteq X$ that are themselves affine varieties in $\mathbb{K}^n$. For example, $\{a \in M_n(\mathbb{K}) : \text{rk}(a) < r\}$ is a Zariski closed subset of $M_n(\mathbb{K})$, since for $a \in M_n(\mathbb{K})$ we have $\text{rk}(a) < r$ iff all $r \times r$ minors of a vanish. Given an arbitrary set $S \subseteq X$, we write $\bar{S}$ for its closure in the Zariski topology on X .

It is straightforward that if $X \subseteq \mathbb{C}^n$ is a complex variety then $X \cap \mathbb{R}^n$ is a real variety. It follows that the Zariski topology on $M_n(\mathbb{R})$ coincides with the subspace topology induced on $M_n(\mathbb{R})$ by the Zariski topology $M_n(\mathbb{C})$. In particular, we can compute the Zariski closure of a set of matrices $A \subseteq M_n(\mathbb{R})$ by first computing the Zariski closure of A in the complex variety $M_n(\mathbb{C})$ and then intersecting with $M_n(\mathbb{R})$.

A set $S \subseteq X$ is *irreducible* if for all closed subsets $A_1, A_2 \subseteq X$ such that $S \subseteq A_1 \cup A_2$ we have either $S \subseteq A_1$ or $S \subseteq A_2$. It is well known that the Zariski topology on a variety is Noetherian. In particular, any closed subset A of X can be written as a finite union of *irreducible components*, where an irreducible component of A is a maximal irreducible closed subset of A .

The class of *constructible* subsets of a variety X is obtained by taking all Boolean combinations of Zariski closed subsets. Suppose that the underlying field $\mathbb{K}$ is algebraically closed. Since the first-order theory of algebraically closed fields admits quantifier elimination, the constructible subsets of X are exactly the subsets of X that are first-order definable over $\mathbb{K}$.

Suppose that $X \subseteq \mathbb{K}^m$ and $Y \subseteq \mathbb{K}^n$ are affine varieties. A function $\varphi : X \rightarrow Y$ is called a *regular map* if it arises as the restriction of a polynomial map $\mathbb{K}^m \rightarrow \mathbb{K}^n$. Chevalley's Theorem states that

if $\mathbb{K}$ algebraically closed and $\varphi : X \rightarrow Y$ is a regular map then the image $\varphi(A)$ of a constructible set $A \subseteq X$ under φ is a constructible subset of Y . This result also follows from the fact that the theory of algebraically closed fields admits quantifier elimination.

A regular map of interest to us is matrix multiplication $M_n(\mathbb{K}) \times M_n(\mathbb{K}) \rightarrow M_n(\mathbb{K})$. In particular, we have that for constructible sets of matrices $A, B \subseteq M_n(\mathbb{K})$ the set of products

$$A \cdot B := \{ab : a \in A, b \in B\}$$

is again constructible. Notice also that matrix inversion is a regular map $\text{GL}_n(\mathbb{K}) \rightarrow \text{GL}_n(\mathbb{K})$. Thus if $A \subseteq \text{GL}_n(\mathbb{K})$ is a constructible set then so is $A^{-1} := \{a^{-1} : a \in A\}$. Finally, the projection $(A, y) \mapsto A$ yields an injective regular map $\text{GL}_n(\mathbb{K}) \rightarrow M_n(\mathbb{K})$. Via this map we can identify $\text{GL}_n(K)$ with a constructible subset of $M_n(\mathbb{K})$.

On several occasions we will use the facts that regular maps are continuous with respect to the Zariski topology and that the image of an irreducible set under a regular map is again irreducible. In particular, we have:

Lemma 1. *If $X, Y \subseteq \text{GL}_n(\mathbb{C})$ are irreducible closed sets then $\overline{X \cdot Y}$ is also irreducible.*

2.3 Algorithmic Manipulation of Constructible Sets

In this subsection we briefly recall some algorithmic constructions on constructible subsets of a variety. Here we assume that the underlying field is $\mathbb{C}$ and that all ideals are generated by polynomials with algebraic coefficients.

Representing Constructible Sets. Consider a variety $X \subseteq \mathbb{C}^n$ and let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be the ideal of polynomials that vanish on X . We represent Zariski closed subsets of X as zero sets of ideals in the coordinate ring $\mathbb{C}[X] = \mathbb{C}[x_1, \dots, x_n]/I$ of X . The coordinate ring of $M_n(\mathbb{C})$ is just $\mathbb{C}[x_{1,1}, \dots, x_{n,n}]$ while the coordinate ring of $\text{GL}_n(\mathbb{C})$ is

$$\mathbb{C}[x_{1,1}, \dots, x_{n,n}, y] / (\det(x_{i,j})y - 1).$$

Unions and intersections of Zariski closed subsets of X respectively correspond to products and sums of the corresponding ideals in $\mathbb{C}[X]$. The latter can be computed using Gröbner basis techniques [3]. We furthermore represent constructible subsets of X as Boolean expressions over Zariski closed subsets.

Irreducible Components. Let $A \subseteq X$ denote a Zariski closed set that is given as the variety of an $I \subseteq \mathbb{C}[X]$. If $I = P_1 \cap \dots \cap P_m$ is an irredundant decomposition of I into primary ideals, then $A = V(P_1) \cup \dots \cup V(P_m)$ is a decomposition of A into irreducible components. One can compute the primary decomposition of an ideal using Gröbner basis techniques [3].

Zariski Closure. At several points below, we will need to compute the Zariski closure of a constructible subset of a variety. Now an arbitrary constructible subset of a variety X can be written as a union of differences of closed subsets of X . Thus it suffices to be able to compute the closure of $A \setminus B$ for closed sets $A, B \subseteq X$. Furthermore, by first computing a decomposition of A as a union of irreducible closed sets, we may also assume that A is irreducible. But $A \subseteq A \setminus B \cup (A \cap B)$; thus by irreducibility of A we have $A \setminus B = \emptyset$ if $A \subseteq B$ and otherwise $A \setminus B = A$. An exponential-time algorithm for computing the Zariski closure of a constructible set, essentially following this recipe, is given in [22, Theorem 1].

Images under Regular Maps. One can use an algorithm for quantifier elimination for the theory of algebraically closed fields

in order to compute the image of a constructible set under a regular map. An explicit algorithm for this task, using Gröbner bases, is given in [37, Sec. 4]. A special case of this operation is that of computing $\phi(A)$ for a closed set $A \subseteq X$. This reduces to computing the pullback of an ideal along a regular map [3].

Real Zariski Closure. Given a complex variety $V \subseteq \mathbb{C}^n$, the intersection $V \cap \mathbb{R}^n$, which is a real variety, can be computed effectively. Indeed if V is represented by the ideal I then $V \cap \mathbb{R}^n$ is represented by the ideal generated by $\{p_{\mathbb{R}}, p_{i\mathbb{R}} : p \in I\}$ where $p_{\mathbb{R}}$ and $p_{i\mathbb{R}}$ respectively denote the real and imaginary parts of the polynomial p .

3 Algebraic Invariants for Affine Programs

An affine function $f : \mathbb{Q}^n \rightarrow \mathbb{Q}^n$ is a function of the form $f(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$, where $A \in M_n(\mathbb{Q})$ and $\mathbf{b} \in \mathbb{Q}^n$. We write $\text{Aff}_n(\mathbb{Q})$ for the set of affine functions on $\mathbb{Q}^n$.

An *affine program* of dimension n is a tuple $\mathcal{A} = (Q, E, q_{\text{init}})$, where Q is a finite set of *program locations*, $E \subseteq Q \times \text{Aff}_n(\mathbb{Q}) \times Q$ is a finite set of *edges*, and $q_{\text{init}} \in Q$ is the *initial location*.

The *collecting semantics* of an affine program $\mathcal{A}$ assigns to each location q the set $S_q \subseteq \mathbb{Q}^n$ of all those vectors that occur at location q in some execution of the program. The family $\{S_q : q \in Q\}$ can be characterised as the least solution of the following system of inclusions (see [30]):

$$\begin{aligned} S_{q_{\text{init}}} &\supseteq \{\mathbf{0}\} \\ S_q &\supseteq f(S_p) \quad \text{for all } (p, f, q) \in E. \end{aligned}$$

Given $P \in \mathbb{R}[x_1, \dots, x_n]$ we say that the polynomial relation $P = 0$ holds at a program location q if P vanishes on S_q . We are interested in the problem of computing for each location $q \in Q$ a finite set of polynomials that generate the ideal $I_q := \mathbf{I}(S_q) \subseteq \mathbb{R}[x_1, \dots, x_n]$ of all polynomial relations holding at location q . The variety corresponding to ideal I_q is $V_q := \mathbf{V}(I_q) = \overline{S_q} \subseteq \mathbb{R}^n$, i.e., V_q is the Zariski closure of S_q regarded as a subset of real affine space. Thus the problem of computing the family of ideals I_q is equivalent to the problem of computing the (family of ideals representing the) Zariski closure of the collecting semantics.

The indexed collection of varieties $\{V_q : q \in Q\}$ defines an *invariant* in that for every edge $(p, f, q) \in E$ we have $f(V_p) \subseteq V_q$. This follows from the facts that $f(S_p) \subseteq S_q$ and that f is Zariski continuous. By construction we have that $\{V_q : q \in V\}$ is the smallest algebraic invariant of the program $\mathcal{A}$ such that $\mathbf{0} \in V_{q_{\text{init}}}$.

In the remainder of this section we reduce the problem of computing the Zariski closure of the collecting semantics of an affine program to that of computing the Zariski closure of a related semigroup of matrices. The idea of this reduction is first to replace each affine assignment by a corresponding linear assignment by adding an extra dimension to the program. One then simulates a general affine program by a program with a single location.

Consider an affine program $\mathcal{A} = (Q, E, q_{\text{init}})$, where the set of locations is $Q = \{q_1, \dots, q_m\}$ and $q_{\text{init}} = q_1$. For each edge $e = (q_j, f, q_i)$ we define a square matrix $M^{(e)} \in M_{m(n+1)}(\mathbb{Q})$ comprising an $m \times m$ array of blocks, with each block a matrix in $M_{n+1}(\mathbb{Q})$. If the affine map f is given by $f(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$ then the (i, j) -th block of $M^{(e)}$ is

$$\begin{pmatrix} A & \mathbf{b} \\ 0 & 1 \end{pmatrix},$$

while all other blocks are zero. Notice that for $\mathbf{x} \in \mathbb{Q}^n$ we have

$$\begin{pmatrix} A & \mathbf{b} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ 1 \end{pmatrix} = \begin{pmatrix} A\mathbf{x} + \mathbf{b} \\ 1 \end{pmatrix} = \begin{pmatrix} f(\mathbf{x}) \\ 1 \end{pmatrix}. \quad (1)$$

Given $i \in \{1, \dots, m\}$, define the projection $\Pi_i : \mathbb{C}^{m(n+1)} \rightarrow \mathbb{C}^{n+1}$ by $\Pi_i(\mathbf{x}_1, \dots, \mathbf{x}_m) = \mathbf{x}_i$ and define the injection $\text{in}_i : \mathbb{C}^n \rightarrow \mathbb{C}^{m(n+1)}$ by

$$\text{in}_i(\mathbf{x}) = (\mathbf{0}, \dots, (\mathbf{x}, 1), \dots, \mathbf{0}) \in \mathbb{C}^{m(n+1)},$$

where $(\mathbf{x}, 1)$ occurs in the i -th block. We denote $\text{in}_1(\mathbf{0})$ by $\mathbf{v}_{\text{init}}$.

Theorem 2. *Let $\mathcal{M}$ be the semigroup generated by the set of matrices $\{M^{(e)} : e \in E\}$. Then for $i = 1, \dots, m$ we have $V_{q_i} = \overline{U_i}$, where*

$$U_i = \left\{ \mathbf{x} \in \mathbb{R}^n : \text{in}_i(\mathbf{x}) \in \{M\mathbf{v}_{\text{init}} : M \in \overline{\mathcal{M}}\} \right\}.$$

Proof. For an edge $e = (q_i, f, q_j)$ of the affine program $\mathcal{A}$ we have

$$M^{(e)} \text{in}_i(\mathbf{x}) = \text{in}_j(f(\mathbf{x}))$$

and

$$M^{(e)} \text{in}_k(\mathbf{x}) = \mathbf{0}$$

for $k \neq i$. Now consider a sequence of edges

$$e_1 = (q_{i_1}, f_1, q_{j_1}), (q_{i_2}, f_2, q_{j_2}), \dots, e_\ell = (q_{i_\ell}, f_\ell, q_{j_\ell}).$$

If this sequence is a legitimate execution of $\mathcal{A}$, i.e., $i_1 = 1$ and $j_k = i_{k+1}$ for $k = 1, \dots, \ell - 1$, then we have

$$M^{(e_\ell)} \dots M^{(e_1)} \mathbf{v}_{\text{init}} = \text{in}_{j_\ell}(f_\ell(\dots f_1(\mathbf{0}) \dots)).$$

If the sequence is not a legitimate execution of $\mathcal{A}$ then we have

$$M^{(e_\ell)} \dots M^{(e_1)} \mathbf{v}_{\text{init}} = \mathbf{0}.$$

From the above it follows that for all $i \in \{1, \dots, m\}$,

$$\begin{aligned} S_{q_i} &= \left\{ \mathbf{x} \in \mathbb{R}^n : \text{in}_i(\mathbf{x}) \in \{M\mathbf{v}_{\text{init}} : M \in \mathcal{M}\} \right\} \\ &= \left\{ \mathbf{x} \in \mathbb{R}^n : (\mathbf{x}, 1) \in \Pi_i(\{M\mathbf{v}_{\text{init}} : M \in \mathcal{M}\}) \right\} \end{aligned}$$

Since the projection Π_i and the map $M \mapsto M\mathbf{v}_{\text{init}}$ are both Zariski continuous, we have

$$\begin{aligned} S_{q_i} &\subseteq \left\{ \mathbf{x} \in \mathbb{R}^n : (\mathbf{x}, 1) \in \Pi_i(\{M\mathbf{v}_{\text{init}} : M \in \overline{\mathcal{M}}\}) \right\} \\ &\subseteq \overline{S_{q_i}}, \end{aligned}$$

that is, $S_{q_i} \subseteq U_i \subseteq \overline{S_{q_i}}$ where U_i is as defined in the statement of the theorem. It follows that $V_{q_i} = \overline{S_{q_i}} = \overline{U_i}$. $\square$

Theorem 2 allows us to compute the invariant $\{V_q : q \in Q\}$. Note however that U_i , being a set of real numbers, uses the real Zariski closure of $\mathcal{M}$, whereas Theorem 15 in Sec. 5.4 allows us to compute the complex Zariski closure. This is not a problem as we can let

$$Z_i = \left\{ \mathbf{x} \in \mathbb{C}^n : \text{in}_i(\mathbf{x}) \in \{M\mathbf{v}_{\text{init}} : M \in \overline{\mathcal{M}}^{\mathbb{C}}\} \right\}$$

where $\overline{\mathcal{M}}^{\mathbb{C}}$ is the complex Zariski closure of $\mathcal{M}$, which is computable by Theorem 15. It follows that Z_i is definable and computable, and $V_{q_i} = U_i = Z_i \cap \mathbb{R}^n$. Finally, as noted in Sec. 2.3, the intersection with $\mathbb{R}^n$ is effective.

4 Zariski Closure of a Subgroup of $\text{GL}_n(\mathbb{C})$

In this section we show how to compute the Zariski closure of the subgroup of $\text{GL}_n(\mathbb{C})$ generated by a given constructible subset of $\text{GL}_n(\mathbb{C})$. We show this by reduction to the problem of computing the Zariski closure of a finitely generated subgroup of $\text{GL}_n(\mathbb{C})$. An algorithm for the latter problem was given by Derken, Jeandel, and Koiran [13].

Recall that for $X \subseteq \text{GL}_n(\mathbb{C})$ we use $\langle X \rangle$ to denote the *sub-semigroup* of $\text{GL}_n(\mathbb{C})$ generated by X . But we have:

Lemma 3 ([13]). *A closed subsemigroup of $\text{GL}_n(\mathbb{C})$ is a subgroup.*

In particular, if $X \subseteq \text{GL}_n(\mathbb{C})$ then $\overline{\langle X \rangle}$ is a subgroup of $\text{GL}_n(\mathbb{C})$. Our aim is to generalise the following result.

Theorem 4 ([13]). *Given matrices $a_1, \dots, a_k \in \text{GL}_n(\mathbb{C})$ with algebraic entries, we can compute the closed subgroup $\overline{\langle a_1, \dots, a_k \rangle}$.*

The first generalisation is as follows.

Corollary 5. *Let $a_1, \dots, a_k \in \text{GL}_n(\mathbb{C})$ and let $Y \subseteq \text{GL}_n(\mathbb{C})$ be an irreducible variety containing the identity I_n . Assume that $a_1, \dots, a_k$ have algebraic entries and that Y is presented as the zero set of a finite collection of polynomials with algebraic coefficients. Then $\overline{\langle a_1, \dots, a_k, Y \rangle}$ is computable from Y and the a_i .*

Proof. Let $G = \overline{\langle a_1, \dots, a_k \rangle}$ and let H be the smallest Zariski closed subgroup of $\text{GL}_n(\mathbb{C})$ that contains Y and is closed under conjugation by $a_1, \dots, a_k$ (i.e., such that $a_i H a_i^{-1} \subseteq H$ for $i = 1, \dots, k$). We claim that $\overline{\langle a_1, \dots, a_k, Y \rangle} = \overline{G \cdot H}$.

To prove the claim, note that since H is closed under conjugation by $a_1, \dots, a_k$ then H is also closed under conjugation by any $g \in \langle a_1, \dots, a_k \rangle$. Moreover, since the map $g \mapsto ghg^{-1}$ is Zariski continuous for each fixed $h \in H$, we have that H is closed under conjugation by any $g \in G = \overline{\langle a_1, \dots, a_k \rangle}$. It follows that $G \cdot H$ is a sub-semigroup of $\text{GL}_n(\mathbb{C})$ and so $\overline{G \cdot H}$ is a group by Lemma 3. But

$$\{a_1, \dots, a_k\} \cup Y \subseteq G \cdot H \subseteq \overline{\langle a_1, \dots, a_k, Y \rangle}$$

and hence $\overline{G \cdot H} = \overline{\langle a_1, \dots, a_k, Y \rangle}$.

It remains to show that we can compute $\overline{G \cdot H}$. Now we can compute G by Theorem 4. To compute H we use the following algorithm:

Procedure FinPlusIrredClosure($a_1, \dots, a_k, Y$)

input : Irreducible variety $Y \subseteq \text{GL}_n(\mathbb{C})$ containing I_n

input : $a_1, \dots, a_k \in \text{GL}_n(\mathbb{C})$

1 $H := Y$

2 $S = \{a_1, \dots, a_k, I_n\}$

3 **repeat**

4 $H_{old} := H$

5 **for** $y \in S$ **do**

6 $H := \overline{H \cdot yHy^{-1}}$

7 **until** $H_{old} = H$

output : H

We show that Algorithm FinPlusIrredClosure computes the smallest subgroup H of $\text{GL}_n(\mathbb{C})$ that is Zariski closed, contains Y , and is closed under conjugation by $a_1, \dots, a_k$. To this end, notice that since Y contains the identity the successive values taken

by H in the algorithm form an increasing chain of sub-varieties of $\text{GL}_n(\mathbb{C})$. Moreover by Lemma 1 this chain is in fact an increasing chain of *irreducible* sub-varieties. But such a chain has bounded length since $\text{GL}_n(\mathbb{C})$ has finite dimension and hence the algorithm must terminate.

We know that $Y \subseteq H$ on termination. Moreover, from the loop termination condition, it clear that on termination H must be closed under conjugation by $a_1, \dots, a_k$, and be a Zariski closed sub-semigroup of $\text{GL}_n(\mathbb{C})$ (and hence a sub-group of $\text{GL}_n(\mathbb{C})$ by Lemma 3). Finally, by construction, H is the smallest such subgroup of $\text{GL}_n(\mathbb{C})$. This concludes the proof. $\square$

We can now prove the main result of this section.

Theorem 6. *Given a constructible subset A of $\text{GL}_n(\mathbb{C})$, we can compute $\overline{\langle A \rangle}$.*

Proof. Let $X_1, \dots, X_k$ be the irreducible components of $\overline{A}$, which are computable from A . For each i , compute a point $a_i \in X_i$ with algebraic entries (using, e.g., the procedure of [2, Chapter 12.6]). Form $Y_i = a_i^{-1} X_i$ which is an irreducible variety containing the identity and let $Y = \overline{Y_1 \cdot Y_2 \cdots Y_k}$ which by Lemma 1 is also an irreducible variety containing the identity. We then have that $\overline{\langle A \rangle} = \overline{\langle a_1, \dots, a_k, Y \rangle}$. Indeed, it is clear that $\langle A \rangle = \langle a_1, \dots, a_k, Y_1 \cdot Y_2 \cdots Y_k \rangle$ and thus

$$\begin{aligned} \overline{\langle A \rangle} &= \overline{\langle a_1, \dots, a_k, Y_1 \cdot Y_2 \cdots Y_k \rangle} \\ &= \overline{\langle a_1, \dots, a_k, \overline{Y_1 \cdot Y_2 \cdots Y_k} \rangle}. \end{aligned}$$

We can compute the closure of $\langle a_1, \dots, a_k, Y \rangle$ thanks to Corollary 5. $\square$

5 Zariski Closure of a Finitely Generated Matrix Semigroup

In this section we give a procedure to compute the Zariski closure of a finitely generated matrix semigroup. We proceed by induction on the rank of the generators. To this end, it is useful to generalise from finite sets of generators to constructible sets of generators. In particular, we will use Theorem 6 on the computability of the Zariski closure of the group generated by a constructible set of matrices.

We first introduce a graph structure on the set of generators that allows us to reason about all products of generators that have a given rank.

5.1 A Generating Graph

Given integers n and r , let $A \subseteq M_n(\mathbb{C})$ be a set of matrices of rank r . We define a labelled directed graph $\mathcal{K}(A)$ as follows:

- There is a vertex (U, V) for each pair of subspaces $U, V \subseteq \mathbb{C}^n$ such that $\dim(V) = r$, $\dim(U) = n - r$, and $U \cap V = \{0\}$.
- There is a labelled edge $(U, V) \xrightarrow{a} (U', V')$ for each pair of vertices (U, V) and (U', V') , and each matrix $a \in A$ such that $\ker(a) = U$ and $\text{im}(a) = V'$.

We note in passing that $\mathcal{K}(A)$ can be seen as an edge-induced subgraph of the *Karoubi envelope* [38] of the semigroup $M_n(\mathbb{C})$.

A *path* in $\mathcal{K}(A)$ is a non-empty sequence of consecutive edges

$$(U_0, V_0) \xrightarrow{a_1} (U_1, V_1) \xrightarrow{a_2} (U_2, V_2) \xrightarrow{a_3} \dots \xrightarrow{a_m} (U_m, V_m).$$

The length of such a path is m and its *label* is the product $a := a_m \cdots a_1$. Matrix a has rank r since $\ker(a_{i+1}) \cap \text{im}(a_i) = 0$ for $i = 1, \dots, m-1$. It is moreover clear that $\{a \in \langle A \rangle : \text{rk}(a) = r\}$ is precisely the set of labels over all paths in $\mathcal{K}(A)$. We will denote that there is a path from (U, V) to (U', V') with label a by writing $(U, V) \xrightarrow{a} (U', V')$.

The following sequence of propositions concerns the structure of the SCCs in $\mathcal{K}(A)$. The respective proofs make repeated use of the fact that for each vertex (U, V) of $\mathcal{K}(A)$ we have $\iota(U) \wedge \iota(V) \neq 0$ and that $\dim \Lambda^r(\mathbb{C}^n) = \binom{n}{r}$ (cf. Sec. 2). We say that an SCC of $\mathcal{K}(A)$ is *non-trivial* if it contains a vertex (U, V) such that there is a path from (U, V) back to itself. Figure 1 summarises the structural results on $\mathcal{K}(A)$.

Proposition 7. $\mathcal{K}(A)$ has at most $\binom{n}{r}$ non-trivial SCCs.

Proof. Let $(U_1, V_1), \dots, (U_m, V_m)$ be an arbitrary finite set of vertices drawn from distinct non-trivial SCCs of $\mathcal{K}(A)$. To prove the proposition it suffices to show that $m \leq \binom{n}{r}$.

Assume that the vertices $(U_1, V_1), \dots, (U_m, V_m)$ are given according to a topological ordering of SCCs—so that there is no path from (U_j, V_j) back to (U_i, V_i) for $i < j$. By assumption, for $i = 1, \dots, m$ there exists a path $(U_i, V_i) \xrightarrow{a_i} (U_i, V_i)$.

We claim that for all $1 \leq i < j \leq m$ we have $\iota(U_i) \wedge \iota(V_j) = 0$ (equivalently, $U_i \cap V_j \neq 0$)—for otherwise there would be path

$$(U_j, V_j) \xrightarrow{a_j} (U_i, V_j) \xrightarrow{a_i} (U_i, V_i),$$

contrary to the topological ordering. On the other hand we have that $\iota(U_j) \wedge \iota(V_j) \neq 0$ (equivalently, $U_j \cap V_j = 0$) for all $j \in \{1, \dots, m\}$ by definition of $\mathcal{K}(A)$. It follows that

$$\iota(U_j) \notin \text{span}\{\iota(U_i) : i = 1, \dots, j-1\}$$

for all $j \in \{1, \dots, m\}$. Indeed, by the claim, any element U in this span satisfies $\iota(U) \wedge \iota(V_j) = 0$ by bilinearity of the wedge product. We conclude that

$$\dim \text{span}\{\iota(U_i) \in \Lambda^r(\mathbb{C}^n) : i = 1, \dots, j\} = j$$

for all $1 \leq j \leq m$ and hence $m \leq \dim \Lambda^r(\mathbb{C}^n) = \binom{n}{r}$, as we wished to prove. $\square$

Proposition 8. If there is a path from (U, V) and (U', V') in $\mathcal{K}(A)$, then there is a path from (U, V) to (U', V') of length at most $\binom{n}{r} + 1$.

Proof. Let

$$(U_0, V_0) \xrightarrow{a_1} (U_1, V_1) \xrightarrow{a_2} \dots \xrightarrow{a_m} (U_m, V_m) \quad (2)$$

be a shortest path from $(U_0, V_0) = (U, V)$ to $(U_m, V_m) = (U', V')$. By construction we have that $U_i \cap V_i = 0$ for $i = 0, \dots, m$. Furthermore we have $U_j \cap V_i \neq 0$ for all $0 < i < j < m$, for otherwise we would have a shortcut

$$(U_{i-1}, V_{i-1}) \xrightarrow{a_i} (U_j, V_i) \xrightarrow{a_{j+1}} (U_{j+1}, V_{j+1}),$$

contradicting the minimality of (2). But then $\iota(V_j) \notin \text{span}\{\iota(V_i) : 1 \leq i < j\}$ for $j = 1, \dots, m-1$: indeed any element V in this span satisfies $\iota(U_j) \wedge \iota(V) = 0$ by bilinearity of the wedge product, but we know that $\iota(U_j) \wedge \iota(V_j) \neq 0$. We conclude that

$$\dim \text{span}\{\iota(V_i) \in \Lambda^r(\mathbb{C}^n) : i \in \{1, \dots, j\}\} = j$$

for all $j = 1, \dots, m-1$. It follows that $m-1 \leq \binom{n}{r}$. $\square$

Proposition 9. Given any path $(U_0, V_0) \xrightarrow{a_1} (U_1, V_1) \xrightarrow{a_2} \dots \xrightarrow{a_m} (U_m, V_m)$ in $\mathcal{K}(A)$, where $m = 2\binom{n}{r}$, some vertex (U_i, V_i) lies in a non-trivial SCC.

Proof. The set of $\binom{n}{r} + 1$ vectors $\{\iota(U_0), \iota(U_2), \iota(U_4), \dots, \iota(U_m)\}$ is linearly dependent since $\dim \Lambda^r(\mathbb{C}^n) = \binom{n}{r}$. Thus there must exist $i \in \{0, \dots, m\}$ such that $\iota(U_i) \in \text{span}\{\iota(U_j) : j \leq i-2\}$. Now by definition of $\mathcal{K}(A)$ we have $U_i \cap V_i = 0$ and hence $\iota(U_i) \wedge \iota(V_i) \neq 0$. Thus by bilinearity of the wedge product there must exist $j \leq i-2$ such that $\iota(U_j) \wedge \iota(V_i) \neq 0$, that is, $U_j \cap V_i = 0$. But then we have a path

$$(U_{i-1}, V_{i-1}) \xrightarrow{a_i} (U_j, V_i) \xrightarrow{a_{j+1}} (U_{j+1}, V_{j+1}),$$

showing that (U_{i-1}, V_{i-1}) and (U_{j+1}, V_{j+1}) lie in the same (necessarily non-trivial) SCC. Indeed, recall that $j \leq i-2$ so either $(U_{j+1}, V_{j+1}) \Rightarrow (U_{i-1}, V_{i-1})$ or $(U_{j+1}, V_{j+1}) = (U_{i-1}, V_{i-1})$ in the original path. $\square$

5.2 Adding Pseudo-Inverses

We now focus on individual SCCs within $\mathcal{K}(A)$. Let C be such a non-trivial SCC. For each edge $(U, V) \xrightarrow{a} (U', V')$ in C , define its *pseudo-inverse* to be a directed edge $(U', V') \xrightarrow{a^+} (U, V)$, where $a^+ \in M_n(\mathbb{C})$ is the unique matrix such that $\ker(a^+) = U'$, $\text{im}(a^+) = V$, $a^+ a v = v$ for all $v \in V$, and $a a^+ v = v$ for all $v \in V'$. We write C^+ for the graph obtained from C by adding pseudo-inverses of every edge in C .

The graph C^+ can be seen as the generator of a groupoid in which the above-defined pseudo-inverse matrices are *bona fide* inverses. We do not develop this idea, except to observe that not only edges but also paths in C have pseudo-inverses in C^+ . Specifically, given a path $(U, V) \xrightarrow{a} (U', V')$ in C , one obtains a path $(U', V') \xrightarrow{a^+} (U, V)$ in C^+ by taking the pseudo-inverse of each constituent edge. In the remainder of this section we show that the pseudo-inverses of all paths in C are already present in the Zariski closure $\overline{\langle A \rangle}$.

Proposition 10. Let (U, V) be a vertex of C and let $B \subseteq M_n(\mathbb{C})$ be a constructible set of matrices such that there is a path $(U, V) \xrightarrow{b} (U, V)$ in C for all $b \in B$. Then $\overline{\langle B \rangle}$ is computable from B and for every $b \in \langle B \rangle$ the pseudo-inverse $(U, V) \xrightarrow{b^+} (U, V)$ is such that $b^+ \in \overline{\langle B \rangle}$.

Proof. By construction, all elements of B have kernel U and image V , where $U \oplus V = \mathbb{C}^n$. Thus there is an invertible matrix $y \in \text{GL}_n(\mathbb{C})$ such that for every $b \in B$ there exists $c \in \text{GL}_r(\mathbb{C})$ with

$$y^{-1} b y = \begin{bmatrix} c & 0 \\ 0 & 0 \end{bmatrix}.$$

Let

$$C := \left\{ c \in \text{GL}_r(\mathbb{C}) : \exists b \in B. y^{-1} b y = \begin{bmatrix} c & 0 \\ 0 & 0 \end{bmatrix} \right\},$$

which is constructible. We can compute $\overline{\langle C \rangle}$ (the Zariski closure of $\langle C \rangle$ in the variety $\text{GL}_r(\mathbb{C})$) using Theorem 6. But then

$$\left\{ y \begin{bmatrix} c & 0 \\ 0 & 0 \end{bmatrix} y^{-1} : c \in \overline{\langle C \rangle} \right\}$$

is a constructible subset of $M_n(\mathbb{C})$ whose closure equals $\overline{\langle B \rangle}$. Note that we are using the fact that $\overline{\langle C \rangle}$ is a subvariety of $\text{GL}_n(\mathbb{C})$ thus it is constructible in $M_n(\mathbb{C})$. Finally, if $b = y \begin{bmatrix} c & 0 \\ 0 & 0 \end{bmatrix} y^{-1} \in \langle B \rangle$ then




Figure 1. Graphical representation of $\mathcal{K}(A)$, vertex and edge labels omitted for clarity. Note that the graph can have infinitely many vertices. Propositions 7 and 8 respectively show that there are only finitely many nontrivial SCCs and they have finite diameter. Proposition 9 shows that paths avoiding nontrivial SCCs must be short. All paths in $\mathcal{K}(A)$ are labelled by rank r matrices. Dotted arrows represent products in the semigroup where the rank becomes less than r : those products do not correspond to labels in $\mathcal{K}(A)$ and need to be handled separately.

$b^+ = y \begin{bmatrix} c^{-1} & 0 \\ 0 & 0 \end{bmatrix} y^{-1} \in \overline{\langle B \rangle}$ since $c^{-1} \in \overline{\langle C \rangle}$ (which is a group by Lemma 3). $\square$

Corollary 11. Suppose that $(U, V) \xRightarrow{a} (U', V')$ is a path in C with pseudo-inverse $(U', V') \xRightarrow{a^+} (U, V)$. Then $a^+ \in \overline{\langle A \rangle}$.

Proof. Since C is strongly connected, there is a path $(U', V') \xRightarrow{b} (U, V)$. Consider the path $(U, V) \xRightarrow{ba} (U, V)$ and its pseudo-inverse $(U, V) \xRightarrow{(ba)^+} (U, V)$. By Proposition 9 we have $(ba)^+ \in \overline{\langle A \rangle}$. We moreover have $a^+ = a^+ b^+ b = (ba)^+ b$ and hence $a^+ \in \overline{\langle A \rangle}$, since $\overline{\langle A \rangle}$ is a semigroup. $\square$

5.3 Maximum-Rank Matrices in the Closure

Let C be a non-trivial SCC in $\mathcal{K}(A)$. Write $B \subseteq M_n(\mathbb{C})$ for the set of labels of all paths in C^+ of length at most $\binom{n}{r} + 2$. Moreover fix a vertex (U_*, V_*) in C^+ and write B_* for the set of labels of all paths in C^+ of length at most $2\binom{n}{r} + 3$ that are self-loops on (U_*, V_*) .

Proposition 12. Let $\langle C \rangle$ denote the set of labels of all paths in C . Then

$$\langle C \rangle \subseteq B \overline{\langle B_* \rangle} B \subseteq \overline{\langle A \rangle}$$

Proof. By Corollary 11 we have that $B, B_* \subseteq \overline{\langle A \rangle}$. Thus the right-hand inclusion follows from the fact that $\overline{\langle A \rangle}$ is a semigroup.

To establish the left-hand inclusion, consider a path

$$(U_0, V_0) \xrightarrow{a_1} (U_1, V_1) \xrightarrow{a_2} (U_2, V_2) \xrightarrow{a_3} \dots \xrightarrow{a_n} (U_n, V_n)$$

within C . Proposition 8 ensures that for each vertex (U_i, V_i) there is a path $(U_*, V_*) \xRightarrow{f_i} (U_i, V_i)$ in C of length at most $\binom{n}{r} + 1$. Such a path has a pseudo-inverse $(U_i, V_i) \xRightarrow{f_i^+} (U_*, V_*)$ in C^+ . Now by the definition of a pseudo-inverse we have $a_i f_{i-1}^+ f_i^+ = a_i$ for all $i \in \{1, \dots, n\}$. Thus

$$\begin{aligned} a_n \dots a_2 a_1 &= a_n f_{n-1}^+ f_{n-1}^+ a_{n-1} f_{n-2}^+ f_{n-2}^+ \dots f_2^+ f_2^+ a_2 f_1^+ f_1^+ a_1 \\ &= a_n f_{n-1}^+ (f_{n-1}^+ a_{n-1} f_{n-2}^+) \dots (f_2^+ a_2 f_1^+) f_1^+ a_1. \end{aligned}$$

The result follows from the observation that $a_n f_{n-1}^+$ and $f_1^+ a_1$ are both elements of B and that $f_i^+ a_i f_{i-1}^+ \in B_*$ for $i = 2, \dots, n-1$. $\square$

Recall from Proposition 8 that the graph $\mathcal{K}(A)$ has at most $\binom{n}{r}$ non-trivial SCCs. Let $C_1, \dots, C_\ell$ be a list of the non-trivial SCCs in $\mathcal{K}(A)$ and write

$$P := A \cup \langle C_1 \rangle \cup \dots \cup \langle C_\ell \rangle. \quad (3)$$

Lemma 13. Given $a \in \langle A \rangle$ with $\text{rk}(a) = r$, we have $a \in P \cup P^2 \cup \dots \cup P^\kappa$, where $\kappa = 2\binom{n}{r}^2$.

Proof. Suppose that a is the label of a path

$$(U_0, V_0) \xrightarrow{a_1} (U_1, V_1) \xrightarrow{a_2} (U_2, V_2) \xrightarrow{a_3} \dots \xrightarrow{a_m} (U_m, V_m) \quad (4)$$

in $\mathcal{K}(A)$. The vertices along this path can be partitioned into maximal blocks of contiguous vertices all lying in the same SCC of $\mathcal{K}(A)$. By Proposition 9 there are at most $\binom{n}{r}$ such blocks corresponding to non-trivial SCCs. The remaining blocks, corresponding to trivial SCCs, are singletons. By Proposition 9 there can be at most $2\binom{n}{r}$ consecutive blocks corresponding to trivial SCCs. Thus there are at most $\kappa = 2\binom{n}{r}^2$ blocks in total.

Now we can factor the path into single edges that connect vertices in different blocks and sub-paths all of whose vertices lie in the same block. There are at most κ such factors (the same as the number of blocks) and the label of each factor lies in the set P defined in (3). This completes the proof. $\square$

Let $R_r = \{x \in M_n(\mathbb{C}) : \text{rk}(x) = r\}$ which is a definable set, and $R_{<r} = \{x \in M_n(\mathbb{C}) : \text{rk}(x) < r\}$ which is closed.

Proposition 14. Let $A \subseteq M_n(\mathbb{C})$ be a constructible set of matrices, all of rank r . Then we can compute $\overline{\langle A \rangle} \cap R_r$ from A .

Proof. By Proposition 12, for $i = 1, \dots, \ell$ we can compute a constructible set $E_i \subseteq M_n(\mathbb{C})$ such that $\langle C_i \rangle \subseteq E_i \subseteq \overline{\langle A \rangle}$. Writing $E := A \cup E_1 \cup \dots \cup E_\ell$, we have $P \subseteq E \subseteq \overline{\langle A \rangle}$.

By Lemma 13 we have $\langle A \rangle \cap R_r \subseteq X$, where $X := E \cup E^2 \cup \dots \cup E^{2\binom{n}{r}^2}$. Now

$$\langle A \rangle \cap R_r \subseteq X \subseteq \overline{\langle A \rangle}$$

$$\begin{aligned} \overline{\langle A \rangle \cap R_r} &\subseteq \overline{X} \subseteq \overline{\langle A \rangle} \\ \overline{\langle A \rangle \cap R_r} \cap R_r &\subseteq \overline{X} \cap R_r \subseteq \overline{\langle A \rangle} \cap R_r. \end{aligned}$$

We claim that

$$\overline{\langle A \rangle \cap R_r} \cap R_r = \overline{\langle A \rangle} \cap R_r \quad (5)$$

which shows that

$$\overline{\langle A \rangle} \cap R_r = \overline{X} \cap R_r$$

is constructible and computable. It remains to see why (5) holds. Since all matrices in A have rank r , all matrices in $\langle A \rangle$ have rank r or less, thus

$$\begin{aligned} \langle A \rangle &= (\langle A \rangle \cap R_r) \cup (\langle A \rangle \cap R_{<r}) \\ \overline{\langle A \rangle} &= \overline{\langle A \rangle \cap R_r} \cup \overline{\langle A \rangle \cap R_{<r}} \\ \overline{\langle A \rangle} \cap R_r &= \underbrace{(\overline{\langle A \rangle \cap R_r} \cap R_r) \cup (\overline{\langle A \rangle \cap R_{<r}} \cap R_r)}_{=\emptyset}. \end{aligned}$$

Indeed, $\langle A \rangle \cap R_{<r} \subseteq R_{<r}$ thus $\overline{\langle A \rangle \cap R_{<r}} \subseteq R_{<r}$ because $R_{<r}$ is closed, and $R_{<r} \cap R_r = \emptyset$.

□

5.4 Computing the Closure

We now present the main result of the paper.

Theorem 15. *Given a constructible set of matrices $A \subseteq M_n(\mathbb{C})$, one can compute $\overline{\langle A \rangle}$ —the Zariski closure of the semigroup generated by A .*

Proof. The proof is by induction on the maximum rank r of the matrices in A . The base case $r = 0$ is trivial. For the induction step, write $A_r := \{a \in A : \text{rk}(a) = r\}$ for the subset of matrices in A of maximum rank and $B := \{a \in \langle A_r \rangle : \text{rk}(a) = r\}$. Now B is computable by Proposition 14.

We claim that $\overline{\langle A \rangle} = \overline{B} \cup \overline{\langle C \rangle}$, where

$$C = \{a \in A \cup BA \cup AB \cup BAB : \text{rk}(a) < r\}.$$

The theorem follows from the claim since $\overline{\langle C \rangle}$ is computable by the induction hypothesis.

It remains to prove the claim. For the right-to-left inclusion notice that since $A, B \subseteq \overline{\langle A \rangle}$ and $\overline{\langle A \rangle}$ is a Zariski-closed semigroup, then $\overline{\langle A \rangle}$ contains both $\overline{B}$ and $\overline{\langle C \rangle}$.

For the left-to-right inclusion it suffices to show that $\langle A \rangle \subseteq \overline{B} \cup \overline{\langle C \rangle}$. To this end, consider a non-empty product $a := a_1 a_2 \cdots a_m$, where $a_1, \dots, a_m \in A$. Suppose first that $\text{rk}(a) = r$. Then $a_1, \dots, a_m \in A_r$ and hence $a \in B$. Suppose now that $\text{rk}(a) < r$. We show that $a \in \overline{\langle C \rangle}$ by induction on m . Let $a_1 \cdots a_\ell$ be a prefix of minimum length that has rank less than r . Clearly such a prefix lies in $A \cup BA$. Moreover the corresponding suffix $a_{\ell+1} \cdots a_m$ is either empty, has rank r (and hence is in B), or has rank $< r$ and hence is in $\langle C \rangle$ by induction. In all cases we have that $a \in \overline{\langle C \rangle}$. □

Corollary 16. *Given a constructible set of matrices $A \subseteq M_n(\mathbb{R})$, one can compute $\overline{\langle A \rangle}^{M_n(\mathbb{R})}$ —the real Zariski closure of the semigroup generated by A .*

Proof. For any set $X \subseteq \mathbb{R}^n$, we have $\overline{X}^{M_n(\mathbb{R})} = \overline{X}^{M_n(\mathbb{C})} \cap M_n(\mathbb{R})$ (see Secs. 2.2 and 2.3). □

References

- [1] 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.
- [2] S. Basu, R. Pollack, and M.-F. Roy. 2006. *Algorithms in Real Algebraic Geometry* (2nd ed.). Springer.
- [3] T. Becker and V. Weispfenning. 1993. *Gröbner bases: a computational approach to commutative algebra*. Graduate Texts in Mathematics, Vol. 141. Springer-Verlag, New York.
- [4] V. D. Blondel, E. Jeandel, P. Koiran, and N. Portier. 2005. Decidable and Undecidable Problems about Quantum Automata. *SIAM J. Comput.* 34, 6 (2005), 1464–1473.
- [5] Aaron R. Bradley and Zohar Manna. 2007. *The calculus of computation - decision procedures with applications to verification*. Springer.
- [6] David Cachera, Thomas P. Jensen, Arnaud Jobin, and Florent Kirchner. 2014. Inference of polynomial invariants for imperative programs: A farewell to Gröbner bases. *Sci. Comput. Program.* 93 (2014), 89–109.
- [7] Michael Colón. 2007. Polynomial approximations of the relational semantics of imperative programs. *Sci. Comput. Program.* 64, 1 (2007), 76–96.
- [8] P. Cousot and R. Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages, Los Angeles, California, USA, January 1977*. 238–252.
- [9] P. Cousot and N. Halbwachs. 1978. Automatic Discovery of Linear Restraints Among Variables of a Program. In *Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages, Tucson, Arizona, USA, January 1978*. 84–96.
- [10] David A. Cox, John B. Little, and Donal O’Shea. 1997. *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra* (2nd ed.). Springer-Verlag.
- [11] Steven de Oliveira, Saddek Bensalem, and Virgile Prevosto. 2016. Polynomial Invariants by Linear Algebra. In *Automated Technology for Verification and Analysis - 14th International Symposium, ATVA 2016, Chiba, Japan, October 17-20, 2016, Proceedings*. Vol. 9938. 479–494.
- [12] S. de Oliveira, S. Bensalem, and V. Prevosto. 2016. Polynomial Invariants by Linear Algebra. In *Automated Technology for Verification and Analysis - 14th International Symposium, ATVA 2016, Chiba, Japan, October 17-20, 2016, Proceedings*. 479–494.
- [13] H. Dersken, E. Jeandel, and P. Koiran. 2005. Quantum automata and algebraic groups. *J. Symb. Comput.* 39, 3-4 (2005), 357–371.
- [14] Sumit Gulwani and George C. Necula. 2003. Discovering affine equalities using random interpretation. In *Conference Record of POPL 2003: The 30th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, New Orleans, Louisiana, USA, January 15-17, 2003*. ACM, 74–84.
- [15] Ehud Hrushovski. 2002. Computing the Galois Group of a Linear Differential Equation. 58 (2002).
- [16] A. Humenberger, M. Jaroschek, and L. Kovács. 2018. Invariant Generation for Multi-Path Loops with Polynomial Assignments. In *Verification, Model Checking, and Abstract Interpretation - 19th International Conference, VMCAI 2018, Los Angeles, CA, USA, January 7-9, 2018, Proceedings (Lecture Notes in Computer Science)*, Vol. 10747. Springer, 226–246.
- [17] Gérard Jacob. 1977. Un Algorithme Calculant le Cardinal, Fini ou Infini, des Demi-Groupes de Matrices. *Theor. Comput. Sci.* 5, 2 (1977), 183–204.
- [18] Gérard Jacob. 1978. La finitude des représentations linéaires des semi-groupes est décidable. *Journal of Algebra* 52, 2 (1978), 437–459.
- [19] Deepak Kapur. 2013. Elimination Techniques for Program Analysis. In *Programming Logics - Essays in Memory of Harald Ganzinger (Lecture Notes in Computer Science)*, Vol. 7797. 194–215.
- [20] Michael Karr. 1976. Affine Relationships Among Variables of a Program. *Acta Inf.* 6 (1976), 133–151.
- [21] Z. Kincaid, J. Cyphert, J. Breck, and T. W. Reps. 2018. Non-linear reasoning for invariant synthesis. *PACMPL* 2, POPL (2018), 54:1–54:33.
- [22] Pascal Koiran. 2000. The Complexity of Local Dimensions for Constructible Sets. *J. Complexity* 16, 1 (2000), 311–323.
- [23] Laura Kovács. 2008. Reasoning Algebraically About P-Solvable Loops. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings (Lecture Notes in Computer Science)*, Vol. 4963. Springer, 249–264.
- [24] L. Kovacs. 2018. personal communication. (2018).
- [25] Laura Ildikó Kovács and Tudor Jebelean. 2005. An Algorithm for Automated Generation of Invariants for Loops with Conditionals. In *Seventh International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC 2005), 25-29 September 2005, Timisoara, Romania*. IEEE Computer Society, 245–249.
- [26] Arnaldo Mandel and Imre Simon. 1977. On Finite Semigroups of Matrices. *Theor. Comput. Sci.* 5, 2 (1977), 101–111.
- [27] A. Markov. 1947. On certain insoluble problems concerning matrices. *Doklady Akad. Nauk SSSR* 57, 6 (1947), 539–542.
- [28] D. W. Masser. 1988. Linear Relations on Algebraic Groups. In *New Advances in Transcendence Theory*. Cambridge University Press.

- [29] A. Miné. 2001. The Octagon Abstract Domain. In *Proceedings of the Eighth Working Conference on Reverse Engineering, WCRE'01, Stuttgart, Germany, October 2-5, 2001*.
- [30] Markus Müller-Olm and Helmut Seidl. 2004. A Note on Karr's Algorithm. In *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings (Lecture Notes in Computer Science)*, Vol. 3142. Springer, 1016–1028.
- [31] Markus Müller-Olm and Helmut Seidl. 2004. Precise interprocedural analysis through linear algebra. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2004, Venice, Italy, January 14-16, 2004*. ACM, 330–341.
- [32] M. Paterson. 1970. Unsolvability in 3×3 matrices. *Studies in Appl. Math.* 49, 1 (1970), 105–107.
- [33] I. Potapov and P. Semukhin. 2017. Decidability of the Membership Problem for 2×2 integer matrices. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*. 170–186.
- [34] Enric Rodríguez-Carbonell and Deepak Kapur. 2007. Automatic generation of polynomial invariants of bounded degree using abstract interpretation. *Sci. Comput. Program.* 64, 1 (2007), 54–75.
- [35] Enric Rodríguez-Carbonell and Deepak Kapur. 2007. Generating all polynomial invariants in simple loops. *J. Symb. Comput.* 42, 4 (2007), 443–476.
- [36] Sriram Sankaranarayanan, Henny Sipma, and Zohar 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*. ACM, 318–329.
- [37] Peter Schauenburg. 2007. A Gröbner-based Treatment of Elimination Theory for Affine Varieties. *Journal of Symbolic Computation* 9 (2007), 859–870.
- [38] Benjamin Steinberg. 2016. *Representation Theory of Finite Monoids*. Springer.
- [39] J. Stillwell. 2016. *Elements of Mathematics: From Euclid to Gödel*. Princeton University Press.