On Ranking Function Synthesis and Termination for Polynomial Programs

Eike Neumann

Department of Computer Science, Oxford University, UK
eike.neumann@cs.ox.ac.uk

Joël Ouaknine

Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany
Department of Computer Science, Oxford University, UK
joel@mpi-sws.org

James Worrell

Department of Computer Science, Oxford University, UK
jbw@cs.ox.ac.uk

Abstract

We consider the problem of synthesising polynomial ranking functions for single-path loops over the reals with continuous semi-algebraic update function and compact semi-algebraic guard set. We show that a loop of this form has a polynomial ranking function if and only if it terminates. We further show that termination is decidable for such loops in the special case where the update function is affine.

2012 ACM Subject Classification Theory of computation → Semantics and reasoning

Keywords and phrases Semi-algebraic sets, Polynomial ranking functions, Polynomial programs

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2020.15

Funding *Joël Ouaknine*: Supported by ERC grant AVS-ISS (648701) and DFG grant 389792660 as part of TRR 248 (see `https://perspicuous-computing.science`).
James Worrell: Supported by EPSRC Fellowship EP/N008197/1.

1 Introduction

The method of proving program termination through ranking functions is one of the oldest and most fundamental ideas in computer science. More recently, the idea of automatically synthesising ranking functions has emerged as an important topic in automated verification and program analysis. Particular attention has focussed on linear ranking functions. Indeed for simple programs, such as linear constraint loops, there are complete methods for synthesising linear ranking functions: such methods find a ranking function whenever one exists, typically by reduction to linear and integer programming [16]. We refer to survey [2] for a thorough discussion of the extensive literature on this topic. More expressive generalisations of linear ranking functions include lexicographic linear ranking functions [3, 9], multiphase linear ranking functions [4], and piecewise linear ranking functions [18].

The advent of powerful techniques for solving non-linear constraints has led to another direction generalising linear ranking functions, namely polynomial ranking functions. Semidefinite programming was used in [10] to synthesise polynomial ranking functions on polynomial loops, while [8] uses cylindrical algebraic decomposition. More recently, polynomial ranking functions [7] have been used to prove termination of probabilistic programs.

In this paper we consider the problem of synthesising polynomial ranking functions for *semi-algebraic loops*: single-path loops over the reals in which the update computed by the loop body is a continuous semi-algebraic function and the loop guard is a semi-algebraic set.

Since we allow polynomials of arbitrary degree as ranking functions, the search for a ranking function cannot immediately be reduced to a constraint satisfaction problem. Our main result shows that if the guard set is compact then a semi-algebraic loop admits a polynomial ranking function if and only if it terminates. The assumption of compactness is essential here: it is straightforward to give examples of terminating loops with non-compact guard that admit no polynomial ranking function. Nevertheless, the class of programs we consider is highly non-trivial. Indeed, to the best of our knowledge, the termination problem for semi-algebraic loops with compact guard set is open (see the discussion below), hence the equivalent problem of deciding the existence of polynomial ranking functions is likewise open. Our main result illustrates the utility and generality of polynomial ranking functions, and offers a potential approach to resolve the decidability of termination for the class of loops in question.

As a second contribution, we show that in the case of semi-algebraic loops in which the guard is compact and the update map is affine, non-termination is equivalent to the existence of a polynomial invariant for the update that is contained in the guard set, and hence decidable. For comparison, recall that Tiwari [17] has given a method to decide termination over the reals of loops with convex polyhedral guard sets and linear update functions (i.e., the same loop dynamics and an incomparable class of loop guards). The case of termination of linear loops with general polyhedral guards (i.e., potentially neither compact nor convex) remains open to the best of our knowledge.

Related Work. It is shown in [1, Lemma 12] and [14, last paragraph in the proof of Theorem 8] that a linear constraint loop with compact polyhedral guard terminates if and only if it has a linear ranking function. This is similar in form to our main result, but involves an incomparable class of transition relations in which non-determinism is allowed but all constraints must be linear.

A partial decidability result for termination of single-path polynomial loops with compact connected guard sets is given in [13]. Under certain semantic assumptions on the loop, this procedure reduces the problem of deciding termination to that of finding fixed points of polynomial maps. Another partial decidability result is [11], which considers a syntactic subclass of polynomial loop programs called *triangular weakly non-linear* and reduces the termination problem to the decision problem for the existential theory of real-closed fields.

For loops with linear updates and convex linear guard conditions, termination is known to be decidable when program variables respectively range over $\mathbb{R}$, $\mathbb{Q}$, and $\mathbb{Z}$ – see [6, 12, 17]. In [19] decidability is shown for loops whose guard sets can be expressed as conjunctions of polynomial inequalities and whose updates are linear functions which satisfy an additional technical condition called the “non-zero minimum property”.

2 The main theorem

Recall that $K \subseteq \mathbb{R}^n$ is said to be *semi-algebraic* if it is definable by a Boolean combination of inequalities $f(x_1, \ldots, x_n) < 0$ and $f(x_1, \ldots, x_n) \leq 0$ with $f \in \mathbb{Z}[x_1, \ldots, x_n]$. A function $g : \mathbb{R}^m \to \mathbb{R}^n$ is moreover said to be *semi-algebraic* if its graph is a semi-algebraic set. Our main result concerns the existence of ranking functions for single-path loops with compact semi-algebraic guard set and semi-algebraic update map:

▶ **Theorem 1.** *Let $g\colon \mathbb{R}^n \to \mathbb{R}^n$ be a continuous semi-algebraic map. Let $K \subseteq \mathbb{R}^n$ be a compact semi-algebraic set. If for all $x \in K$ there exists an $m \in \mathbb{N}$ such that $g^m(x) \notin K$ then there exists a polynomial ranking function for g on K, i.e., a polynomial $f \in \mathbb{R}[x_1, \ldots, x_n]$ satisfying $f(x) - f(g(x)) \geq 1$ and $f(x) \geq 0$ for all $x \in K$.*

Since the existence of a ranking function clearly implies termination, Theorem 1 shows that such a semi-algebraic loop terminates precisely when it has a polynomial ranking function. The theorem clearly fails if we remove the requirement that K be closed: for loop guard $K = (0, 1)$ and transition function $g(x) = 2x$ there cannot be a polynomial ranking function since the time to escape K under the action of g is not bounded from above. Likewise the theorem fails in case K is closed but not bounded. For example, consider $K = \{(x, y) \in \mathbb{R}^2 : x \geq 0, y \geq 1\}$ and $g(x, y) = \left(x - \frac{1}{y}, y + 1\right)$. The loop is terminating since $\sum_{y=1}^{\infty} \frac{1}{y}$ diverges, but from $\sum_{y=1}^{n} \frac{1}{y} < \ln n + 1$ one sees that the time to termination from $(m, 1)$ is at least e^{m-1}, which precludes the existence of a polynomial ranking function.

The proof of Theorem 1 relies on several lemmas. In the rest of this section we give an informal overview of the proof structure, formulate the key lemmas, and finally show how they imply the main result.

According to the Stone-Weierstraß theorem, a continuous function on K can be uniformly approximated arbitrarily closely by polynomials, and hence it suffices to find a continuous function $f : K \cup g(K) \to \mathbb{R}$ satisfying $f(x) - f(g(x)) \geq 1$ for all $x \in K$. We show the stronger assertion that there exists a continuous function f satisfying $f(x) = f(g(x)) + 1$ for all $x \in K$.

The first step associates with the loop update function g an injective "covering map" $\widetilde{g}$. Specifically, we construct a space T and continuous surjective map $p : T \to K \cup g(K)$ such that there is $S \subseteq T$ and an injective semi-algebraic map $\widetilde{g} : S \to T$ making the following diagram commute:

$$\begin{tikzcd} S \arrow[r, "\widetilde{g}"] \arrow[d, "p\restriction_S"] & T \arrow[d, "p"] \\ K \arrow[r, "g"] & K \cup g(K) \end{tikzcd}$$

We then construct a continuous function $\widetilde{f} : T \to \mathbb{R}$ that satisfies the functional equation $\widetilde{f}(x) = \widetilde{f}(\widetilde{g}(x)) + 1$ for all $x \in S$ and is constant on each fibre $p^{-1}(x)$, $x \in K \cup g(K)$. To construct $\widetilde{f}$ we find a partition of T into finitely many semi-algebraic pieces $S_1, \ldots, S_m$ such that (i) the boundary of a piece is contained in a union of pieces of lower dimension, (ii) S is a union of pieces, (iii) each piece $S_i \subseteq S$ is mapped by $\widetilde{g}$ onto another piece S_j. We order the pieces of a given dimension by the transitive closure of the relation $S_i = \widetilde{g}(S_j)$. This is a well-founded partial ordering thanks to the termination assumption – that for all $x \in K$ there exists m such that $g^m(x) \notin K$. Ignoring for now the technical requirement that $\widetilde{f}$ be constant on the fibres of p, we can essentially construct $\widetilde{f}$ as follows: On the zero-dimensional pieces S_i which are minimal with respect to this ordering let $\widetilde{f}(S_i) = 0$. Extend $\widetilde{f}$ to the remaining zero-dimensional pieces using the functional equation $\widetilde{f}(x) = \widetilde{f}(\widetilde{g}(x)) + 1$. Now $\widetilde{f}$ can be defined on the minimal one-dimensional pieces by interpolating the boundary values. It can be extended to all one-dimensional pieces using the functional equation. Continuing with the higher-dimensional pieces in the same manner, we eventually obtain a continuous solution to the functional equation $\widetilde{f}(x) = \widetilde{f}(\widetilde{g}(x)) + 1$ on all of T. We will also be able to ensure that it is constant on the fibres of the projection p, so that we obtain a continuous solution to the equation $f(x) = f(g(x)) + 1$ on $K \cup g(K)$ by letting $f(x) = \widetilde{f}(\widetilde{x})$ where $\widetilde{x}$ is any point in T satisfying $p(\widetilde{x}) = x$.

As a preliminary result we observe that if the iteration of a continuous function escapes a compact set then the number of iterations required to escape is bounded independently of the starting point.

▶ **Proposition 2.** *Let $g\colon \mathbb{R}^n \to \mathbb{R}^n$ be a continuous map. Let $K \subseteq \mathbb{R}^n$ be a compact set. Assume that for every $x \in K$ there exists a positive integer m such that $g^m(x) \notin K$. Then there exists a positive integer M such that for all $x \in K$ there exists $i \leq M$ such that $g^i(x) \notin K$.*

Proof. For every positive integer $m \in \mathbb{N}$ let $B_m = \{x \in \mathbb{R}^n \mid g^m(x) \notin K\}$. Since K is closed and g is continuous, each of the sets B_m is open. The assumption says that every $x \in K$ is contained in B_m for some m. In other words, the family $(B_m)_{m\in\mathbb{N}}$ is an open cover for K. Since K is compact this cover has a finite subcover $B_{m_1}, \ldots, B_{m_s}$. Now take $M = \max\{m_1, \ldots, m_s\}$. ◀

Let us now state the required lemmas more formally. The first lemma concerns the existence of a suitable covering space, which allows us to replace g with an injective map.

▶ **Lemma 3.** *Let $g\colon \mathbb{R}^n \to \mathbb{R}^n$ be a continuous semi-algebraic map. Let $K \subseteq \mathbb{R}^n$ be a compact semi-algebraic set. If for all $x \in K$ there exists $m \in \mathbb{N}$ such that $g^m(x) \notin K$ then there exist a compact semi-algebraic set $T \subseteq \mathbb{R}^N$, a continuous semi-algebraic surjection $p\colon T \to K \cup g(K)$, and an injective continuous semi-algebraic map $\widetilde{g}\colon p^{-1}(K) \to T$ such that the following diagram commutes:*

$$\begin{tikzcd} p^{-1}(K) \arrow[r, "\widetilde{g}"] \arrow[d, "p"] & T \arrow[d, "p"] \\ K \arrow[r, "g"] & K \cup g(K) \end{tikzcd}$$

Proof. By Proposition 2 there exists a number $M \geq 0$ such that for all $x \in K$ there exists $i \leq M$ such that $g^i(x) \notin K$. Let

$$H = \left\{(i, x, g(x), g^2(x), \ldots, g^i(x), 0, \ldots, 0) \in \mathbb{R}^{n(M+2)+1} \mid i \in \mathbb{N}, i \leq M+1, x \in K\right\}.$$

Then H is clearly a compact semi-algebraic set.

Let

$$p\colon H \to \mathbb{R}^n,\ p(i, x, g(x), \ldots, g^i(x), 0, \ldots, 0) = g^i(x).$$

Let $T = p^{-1}(K \cup g(K)) \subseteq H$ and $S = p^{-1}(K)$. Then S and T are compact semi-algebraic subsets of H and p maps S onto K and T onto $K \cup g(K)$ (as we allow $i = 0$ in the definition of H).

Let

$$\widetilde{g}\colon S \to T,\ \widetilde{g}(i, x, g(x), g^2(x), \ldots, g^i(x), 0, \ldots, 0) = (i+1, x, g(x), g^2(x), \ldots, g^{i+1}(x), 0, \ldots, 0).$$

Then $\widetilde{g}$ is a continuous, injective semi-algebraic map. It satisfies the equation $p \circ \widetilde{g}(x) = g \circ p(x)$ for all $x \in S$. ◀

We will construct a ranking function for $\widetilde{g}$ that is constant on fibres of the projection. This will be achieved through the construction and refinement of certain semi-algebraic cellular decompositions of the codomain T of $\widetilde{g}$. For the definition of semi-algebraic cellular decomposition compare, e.g., [5, Definition 9.1.11].

▶ **Definition 4.** *A subset $C \subseteq \mathbb{R}^n$ is a* cell of dimension $d \in \mathbb{N}$ *if there is a semi-algebraic homeomorphism $\Phi\colon D \to C$ where D is the open disk in $\mathbb{R}^d$. Given a closed semi-algebraic set $S \subseteq \mathbb{R}^n$, a* semi-algebraic cellular decomposition *of S is a finite partition of S into cells such that for each cell C its closure in $\mathbb{R}^n$ is the union of C and a collection of other cells of strictly smaller dimension than C.*

Based on the above, we introduce the notion of an *invariant semi-algebraic stratification*:

▶ **Definition 5.** *Let $g\colon S \to T$ be a semi-algebraic map between semi-algebraic sets S and T, with $S \subseteq T$ and T closed. A g-*invariant semi-algebraic stratification *of T consists of a partition of T into finitely many semi-algebraic sets $S_1, \dots, S_m$, called* strata*, together with a semi-algebraic cellular decomposition of T, such that*

1. *Each stratum is a finite union of cells, homeomorpic to an open disk.*
2. *The set S is a finite union of strata.*
3. *For all strata $S_i \subseteq S$, the set $g(S_i)$ is a stratum.*

The *boundary* of a stratum S is defined to be $\mathrm{cl}(S) \setminus S$. Note that this is different from the topological boundary of S regarded as a subset of $\mathbb{R}^n$. Since every stratum S is a finite union of cells, and the closure of every cell is also a union of cells, the boundary of S is a finite union of cells.

The next two lemmas constitute the core of the proof of Theorem 1. They will be proved in Sections 3 and 4 respectively.

▶ **Lemma 6.** *Let $g\colon S \to T$ be an injective continuous semi-algebraic map between compact semi-algebraic sets S and T with $S \subseteq T$. Assume that for all $x \in S$ there exists an integer $i \geq 0$ such that $g^i(x) \notin S$. Then there exists a g-invariant semi-algebraic stratification of T.*

▶ **Lemma 7.** *Let $g\colon \mathbb{R}^n \to \mathbb{R}^n$ be a continuous semi-algebraic map. Let $K \subseteq \mathbb{R}^n$ be a compact semi-algebraic set such that for all $x \in K$ there exists $m \in \mathbb{N}$ such that $g^m(x) \notin K$. Let $T \subseteq \mathbb{R}^N$ be a compact semi-algebraic set such that there exists a continuous semi-algebraic surjection $p\colon T \to K \cup g(K)$ and an injective continuous semi-algebraic map $\widetilde{g}\colon p^{-1}(K) \to T$ such that the following diagram commutes:*

$$\begin{tikzcd} p^{-1}(K) \arrow[r, "\widetilde{g}"] \arrow[d, "p"] & T \arrow[d, "p"] \\ K \arrow[r, "g"] & K \cup g(K) \end{tikzcd}$$

Assume that there exists a $\widetilde{g}$-invariant stratification of T. Then there exists a continuous function $\widetilde{f}\colon T \to \mathbb{R}$ satisfying the functional equation $\widetilde{f}(x) = \widetilde{f}(\widetilde{g}(x)) + 1$ for all $x \in p^{-1}(K)$ which is constant on the fibres of p.

Assuming the lemmas above we can now prove Theorem 1:

Proof of Theorem 1. We will prove that there exists a continuous function $f\colon K \cup g(K) \to \mathbb{R}$ which satisfies the functional equation $f(x) = f(g(x)) + 1$ for all $x \in K$. This suffices to prove the claim, for by the Stone-Weierstraß theorem there exists a polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ satisfying $|p(x) - f(x)| < \frac{1}{4}$ for all $x \in K$. Let $c = \min\{p(x) \mid x \in K\}$. Then the polynomial $h(x) = 2(p(x) - c)$ satisfies $h(x) - h(g(x)) \geq 1$ and $h(x) \geq 0$ for all $x \in K$.

By Lemma 3 there exist compact semi-algebraic sets $p^{-1}(K) = S \subseteq T \subseteq \mathbb{R}^N$, a continuous semi-algebraic surjection $p\colon T \to K \cup g(K)$, which restricts to a surjection $p\colon S \to K$, and an injective continuous semi-algebraic map $\widetilde{g}\colon S \to T$ such that $p \circ \widetilde{g} = g \circ p$. By Lemma 6 there exists a $\widetilde{g}$-invariant stratification of T. By Lemma 7 there exists a continuous function $\widetilde{f}\colon T \to \mathbb{R}$ which satisfies $\widetilde{f}(x) = \widetilde{f}(\widetilde{g}(x)) + 1$ for all $x \in S$ and is constant on fibres of p.

Define on $K \cup g(K)$ the function $f(x) = \widetilde{f}(\widetilde{x})$ where $\widetilde{x}$ is any point in the fibre of x under p. This function is well-defined since $\widetilde{f}$ is constant on fibres. We have for all $x \in K$:

$$f(x) = \widetilde{f}(\widetilde{x}) = \widetilde{f}(\widetilde{g}(\widetilde{x})) + 1 = f(g(x)) + 1.$$

The last equality uses the fact that $p \circ \widetilde{g} = g \circ p$, so that if $p(\widetilde{x}) = x$ then $p(\widetilde{g}(\widetilde{x})) = g(p(\widetilde{x})) = g(x)$.

Finally, f is continuous. Observe that we have $f^{-1}(A) = p(\widetilde{f}^{-1}(A))$ for all sets $A \subseteq \mathbb{R}$. If A is closed, then $\widetilde{f}^{-1}(A)$ is compact, so that $p(\widetilde{f}^{-1}(A))$ is closed. Hence the preimage of any closed set under f is closed, so that f is continuous. ◀

3 Proof of Lemma 6

In this section we prove Lemma 6. Throughout this section S and T are compact semi-algebraic subsets of $\mathbb{R}^n$ with $S \subseteq T$, and $g\colon S \to T$ is a continuous injective semi-algebraic map such that for all $x \in S$ there exists an integer $i \geq 0$ such that $g^i(x) \notin S$.

We show that we can find a g-invariant stratification of T, in the sense of Definition 5. The construction relies on the following standard result in real-algebraic geometry on the existence of cellular decompositions (see, e.g., [5, Proposition 9.1.12]):

▶ **Theorem 8.** *Let $S \subseteq \mathbb{R}^n$ be a compact semi-algebraic set. Let $S_1, \ldots, S_m$ be semi-algebraic subsets of S. Then S admits a semi-algebraic cellular decomposition such that each S_i is a finite union of cells.*

We have a semi-algebraic function

$$\mathrm{rank}\colon T \to \mathbb{N},\ \mathrm{rank}(x) = \min\left\{i \in \mathbb{N} \mid g^i(x) \notin S\right\}.$$

We can decompose T into finitely many semi-algebraic subsets

$$E_i = \left\{x \in T \mid g^i(x) \notin S \text{ and } g^j(x) \in S \text{ for all } j < i\right\}.$$

On these sets the function rank is constant. By Theorem 8 the set T admits a semi-algebraic cellular decomposition such that each E_i is a finite union of cells. Then rank is constant on every cell of this cellular decomposition, so that we can assign to every cell e the number $\mathrm{rank}(e)$.

Let every cell of this decomposition be a stratum. We will keep refining this stratification until g maps strata onto strata.

A stratum A is called an *injury* if it is not mapped by g onto another stratum, *i.e.*, there does not exist another stratum A' such that $g(A) = A'$. If A is an injury which is mapped into another stratum, *i.e.*, there exists another stratum A' such that $g(A) \subseteq A'$, then A is called en injury of the second kind. Otherwise it is called an injury of the first kind.

The *signature* of an injury A is the tuple $(\mathrm{kind}\, A, \dim A, \mathrm{rank}\, A)$. The *injury signature* of a stratification of T is the multiset of all signatures of all injuries. Thus, if n different injuries have the same signature, then the injury signature of the stratification contains n copies of this signature.

▶ **Lemma 9.** *If the stratification of T contains an injury of the first kind, then the decomposition can be refined such that the new injury signature is the old injury signature with one signature of the form $(1, m, k)$ removed and finitely many signatures added of the form $(2, l, k)$ with $l \leq m$. Additionally, the subdivision may result in any number of injuries of the second kind of rank $k + 1$ and dimension $\leq m$ becoming injuries of the first kind.*

Proof. Let A be an injury of the first kind with dimension m and rank k.

The sets of the form $g^{-1}(C) \cap A$ where C is a stratum are semi-algebraic subsets of $\mathrm{cl}\,(A)$. By Theorem 8 there exists a semi-algebraic cellular decomposition of $\mathrm{cl}\,(A)$ such that each set of the form $g^{-1}(C) \cap A$ and every cell of the old cellular decomposition of $\mathrm{cl}\,(A)$ is a finite union of cells. Replace the cellular decomposition on $\mathrm{cl}\,(A)$ with this new decomposition.

This defines a new cellular decomposition of T that refines the old one. Remove A from the stratification, and add all cells of the refined decomposition which are contained in A as new strata.

The newly added strata get mapped into strata by construction, but they may still be injuries of the second kind. Since they are contained in A they have the same rank as A and their dimension is at most $\dim A$. Strata that were mapped into but not onto A may no longer be mapped into strata in the new stratification, so that they become injuries of the first kind. Their rank is equal to $\operatorname{rank} A + 1$. Since g is injective their dimension is at most $\dim A$.

It may be the case that there is a stratum C that was mapped onto A. Since g is injective there is at most one such stratum. Its closure $\operatorname{cl}(C)$ is mapped by g onto $\operatorname{cl}(A)$. Its dimension is the same as the dimension of A and its rank is $\operatorname{rank} A + 1$. This stratum would become an injury of the first kind in the new stratification. To remove this injury we proceed as follows: Apply Theorem 8 to obtain a new cellular decomposition of $\operatorname{cl}(C)$ such that each old cell which is contained in $\operatorname{cl}(C)$ and each set of the form $g^{-1}(c)$, where c is a cell of the old cellular decomposition of $\operatorname{cl}(A)$, is a finite union of cells. Remove C from the stratification. For each of the strata C_i that have replaced A in the new stratification add the set $g^{-1}(C_i)$ to the stratification. Then every new stratum contained in the old stratum C gets mapped onto some C_i by construction. The new strata are homeomorphic to disks, as they are homeomorphic images of disks under g^{-1}. Note that the refinement of the cellular decomposition may change the cellular decomposition of strata that are not contained in C, but this does not introduce new injuries, as the underlying set of these strata does not change.

If there is a stratum C' that was mapped onto C, repeat this procedure with C' playing the role of C and C playing the role of A. Continue in this manner until all injuries of the first kind that arise in this way are removed. This happens after finitely many steps as the rank of C increases by one with each repetition. ◀

▶ **Lemma 10.** *If the stratification of T contains an injury of the second kind, then the decomposition can be refined such that the new injury signature is the old injury signature with at least one signature of the form $(2, m, k)$ removed and finitely many signatures added of the form $(i, l, k-1)$.*

Proof. Let A be an injury of the second kind with dimension m and rank k. By assumption it is mapped into a stratum C.

Let $A_1, \ldots, A_s$ be all other injuries of the second kind that get mapped into the same stratum C. By Theorem 8 there exists a semi-algebraic cellular decomposition of $\operatorname{cl}(C)$ such that each of the sets $g(A)$ and $g(A_i)$, and each of the cells of the old cellular decomposition of $\operatorname{cl}(C)$ is a finite union of cells. Replace the cellular decomposition on $\operatorname{cl}(C)$ with this new decomposition. This defines a new cellular decomposition of T that refines the old one. Then C is the disjoint union of $g(A)$, the $g(A_i)$'s, and a finite union of cells $c_1, \ldots, c_k$. Remove C from the stratification and add $g(A)$, the $g(A_i)$'s, and $c_1, \ldots, c_k$ as new strata.

Then A and the A_i's are no longer injuries, as they are mapped onto strata by construction. The c_i are potentially new injuries with rank $k-1$. As A is mapped into C, different strata are disjoint, and g is injective, no stratum was mapped onto C before. For all other strata, neither their underlying set nor the underlying set of their image has changed, so that no further injuries are added. ◀

Now all injuries can be removed by repeatedly applying the two previous lemmas in the correct order:

▶ **Lemma 11.** *Consider the following algorithm:*

1. *Find a semi-algebraic cellular decomposition of T such that each cell has constant rank. Let each cell of the decomposition be a stratum.*
2. *Sort the injuries of the stratification lexicographically, comparing first by dimension, then by negative rank, and then by kind.*
3. *Pick a minimal element with respect to this ordering. If it is of the first kind, remove it using the first lemma. If it is of the second kind, remove it using the second lemma.*
4. *If there are no injuries left, then output the stratification. Otherwise go to (2).*

Then this algorithm terminates in finite time. It returns a g-invariant stratification of T.

Proof. Let N be the highest rank among all injuries. Consider the $N \times n$ matrix $A = ((a_{i,j}, b_{i,j}))$ whose entry at index (i, j) is the pair $(a_{i,j}, b_{i,j}) \in \mathbb{N} \times \mathbb{N}$ where $a_{i,j}$ is the number of injuries of the first kind of rank $N-i$ and dimension j, and $b_{i,j}$ is the number of injuries of the second kind of rank $N-i$ and dimension j. Thus, the j^{th} column of the matrix records all injuries of dimension j and the i^{th} row of the matrix records all injuries of rank $N-i$. In each step the algorithm will process an injury which is recorded in the first non-zero entry in the first non-zero column of the matrix A. Define the rank of a non-zero column to be the highest rank of all injuries that are recorded in the column.

We will now show by induction on the index of the first non-zero column that the algorithm will make the first non-zero column into a zero column in finitely many steps without introducing injuries whose dimension is smaller than or equal to the column index or whose rank is greater than or equal to the column rank. In particular it will make all columns into a zero column in finitely many steps, proving termination.

If the first non-zero column is the first column of the matrix, then the algorithm will first remove all injuries of the first kind and of highest rank. This potentially introduces new injuries of the second kind of the same dimension and the same rank, but no further injuries. It will then remove all injuries of the second kind and of highest rank. This potentially introduces new injuries of any kind and any dimension, but of a rank that is strictly lower than the column rank. Once the first non-zero entry of the column has been made into zero, the algorithm will proceed with the next entry of the column and continue in the same manner until the column is made into a zero column. All injuries introduced by the end of this process have rank strictly less than the initial column rank. Their dimension is higher than the column index since the column was assumed to be the first column in the matrix.

Now assume that we have shown the result for all column indexes strictly smaller than $j > 0$. Assume that the j^{th} column is the first non-zero column, and let $(a_{i,j}, b_{i,j})$ be its first non-zero entry. If $a_{i,j} > 0$ then an application of one step of the algorithm decreases $a_{i,j}$ by one, increases $b_{i,j}$ by an arbitrary amount, and potentially introduces new injuries of the second kind in the i^{th} row and to the left of (i, j). The algorithm will now proceed to remove non-zero columns with smaller index than j. These columns have rank at most $N - i$, so that by induction hypothesis any injuries added in the process of removing these columns have rank strictly smaller than $N - i$, so that in particular the entry at index (i, j) does not change throughout. Therefore, after finitely many steps, the first non-zero entry in the first non-zero column will be $(a_{i,j} - 1, b_{i,j} + c)$ with $c \geq 0$. This shows that the value of the first entry at index (i, j) is strictly decreasing and will hence eventually be equal to zero. Note that at no point injuries of rank greater than or equal to $N - i$ were introduced. If $a_{i,j} = 0$ and $b_{i,j} > 0$, then an application of one step of the algorithm decreases $b_{i,j}$ by at least one, and potentially introduces new injuries of lower rank and any dimension. The algorithm will proceed to remove non-zero columns of lower index, leaving the entries with index (i, j) unaffected by induction hypothesis. Thus, after finitely many steps, the first non-zero entry in the first non-zero column will be $(0, b'_{i,j})$ with $b'_{i,j} < b_{i,j}$. It follows that the entry is made into zero after finitely many steps. Again, no injuries of rank greater than or equal to $N - i$ were introduced. The algorithm will proceed to process the rest of the entries of the column in the same manner until the column and all columns to the left of it are made into zero, not introducing any injuries of rank $N - i$ or greater. ◀

4 Proof of Lemma 7

In this section we prove Lemma 7. Throughout this section, let $g\colon \mathbb{R}^n \to \mathbb{R}^n$ be a continuous semi-algebraic map, let $K \subseteq \mathbb{R}^n$ be a compact semi-algebraic set such that for all $x \in K$ there exists $m \in \mathbb{N}$ such that $g^m(x) \notin K$, let $T \subseteq \mathbb{R}^N$ be a compact semi-algebraic set, let $p\colon T \to K \cup g(K)$ be a continuous semi-algebraic surjection, and let $\widetilde{g}\colon p^{-1}(K) \to T$ be an injective continuous semi-algebraic map such that the following diagram commutes:

$$\begin{array}{ccc} p^{-1}(K) & \xrightarrow{\widetilde{g}} & T \\ \big\downarrow{\scriptstyle p} & & \big\downarrow{\scriptstyle p} \\ K & \xrightarrow{g} & K \cup g(K) \end{array}$$

Write $S = p^{-1}(K)$. Fix a $\widetilde{g}$-invariant stratification of T.

Define a well-founded partial ordering on the strata of T as follows: $S_1 \leq S_2$ if and only if $S_1 = \widetilde{g}^m(S_2)$ for some $m \geq 0$. Arrange the strata in a linear chain $S_1, S_2, \ldots, S_m$ such that $i < j$ implies $\dim S_i \leq \dim S_j$ and $S_j \not\leq S_i$.

Let $F_1 = \{x \in T \mid \exists x' \in S_1 . p(x') = p(x)\}$. For all $k = 2, \ldots, m$, define the set $F_k = \{x \in T \mid \exists x' \in S_k \cup F_{k-1} . p(x') = p(x)\}$. Note that we have $F_k \supseteq F_i$ and $F_k \supseteq S_i$ for all $i \leq k$.

Before we begin with the construction of $\widetilde{f}$, we record some basic properties of the sets S_k and F_k:

▶ **Lemma 12.** *The sets F_k and $F_k \cup S_{k+1}$ are closed for all k.*

Proof. The set S_1 is a zero-dimensional disk and hence a singleton. The set $F_1 = p^{-1}(p(S_1))$ is closed, as S_1 is compact and p is continuous.

Assume now that F_k is closed for a given k. The boundary of S_{k+1} is contained in the union of all strata of dimension $< \dim S_{k+1}$, which by definition are contained in F_k. Hence $F_k \cup S_{k+1}$ is closed. It then follows, using compactness of $F_k \cup S_{k+1}$ and continuity of p, that the set $F_{k+1} = p^{-1}(p(F_k \cup S_{k+1}))$ is closed as well. ◀

▶ **Lemma 13.** *For all $k = 1, \ldots, m$, the set F_k contains all fibres of p that intersect it.*

Proof. Assume that there exists $x' \in p^{-1}(x) \cap F_k$. Then there exists $x'' \in S_k \cup F_{k-1}$ with $p(x'') = x$. Let $x''' \in p^{-1}(x)$. Then $p(x''') = x = p(x'')$ with $x'' \in S_k \cup F_{k-1}$, so that $x''' \in F_k$. ◀

▶ **Lemma 14.** *For all $k = 1, \ldots, m$, if $x \in F_k \cap S$ then $\widetilde{g}(x) \in F_k$.*

Proof. We prove this by induction on k. We have $S_1 \cap S = \emptyset$, for if this were not the case then we could apply $\widetilde{g}$ to S_1 to obtain a stratum $S_j = \widetilde{g}(S_1)$ with $j > 1$. This would imply $S_j \leq S_1$ which directly contradicts the definition of the linear ordering on strata. It follows that $p(S_1) \cap K = \emptyset$ and hence $F_1 \cap S = p^{-1}(p(S_1)) \cap p^{-1}(K) = \emptyset$. This establishes the claim for $k = 1$.

Assume that the claim is true for $k-1$. Let $x \in F_k \cap S$. By definition of F_k there exists $x' \in S_k \cup F_{k-1}$ with $p(x') = p(x)$. Since x is assumed to be contained in $S = p^{-1}(K)$ we have $p(x') \in K$ and thus $x' \in S$, so that $\widetilde{g}$ can be applied to x'.

If $x' \in F_{k-1}$ then $\widetilde{g}(x') \in F_{k-1}$. Since $p \circ \widetilde{g}(x') = p \circ \widetilde{g}(x)$ and F_{k-1} contains all fibres that intersect it, it follows that $\widetilde{g}(x) \in F_{k-1}$. As $F_{k-1} \subseteq F_k$ the claim follows.

Now assume that $x' \in S_k$. Since all strata are disjoint and S is a union of strata, the stratum S_k is either contained in S or disjoint from S. We have seen above that $x' \in S$, so that $S_k \subseteq S$. It follows that $\widetilde{g}$ can be applied to S_k. Since the stratification is $\widetilde{g}$-invariant, the set $\widetilde{g}(S_k)$ is again a stratum, say, $\widetilde{g}(S_k) = S_j$ for some j. By definition of the ordering on strata we have $S_j \leq S_k$ which implies $j \leq k$, and hence $S_j \subseteq F_k$ by definition of F_k. Thus, $\widetilde{g}(x') \in F_k$. Since $p \circ \widetilde{g}(x') = p \circ \widetilde{g}(x)$ and F_k contains all fibres that intersect it, it follows that $\widetilde{g}(x) \in F_k$. ◀

We prove by induction on k that there exists a continuous function $\widetilde{f} \colon F_k \to \mathbb{R}$ which satisfies the functional equation $\widetilde{f}(x) = \widetilde{f}(\widetilde{g}(x)) + 1$ for all $x \in F_k \cap S$ and is constant on all fibres of p which intersect F_k. Note that by Lemma 14, if $x \in F_k \cap S$ then $\widetilde{g}(x) \in F_k$, so that $\widetilde{f}(\widetilde{g}(x))$ is defined, and it makes sense to ask that $\widetilde{f}$ satisfy the functional equation in x.

The stratum S_1 is necessarily a 0-cell and minimal with respect to the ordering on strata. In particular we have $S_1 \cap S = \emptyset$ and hence $F_1 \cap S = \emptyset$ (see also the proof of Lemma 14 above). Put $\widetilde{f}(x) = 0$ on F_1. Then clearly $\widetilde{f}$ has all the required properties, as the functional equation is only required to hold for points in S.

Assume that $\widetilde{f} \colon F_k \to \mathbb{R}$ has been defined. We want to extend $\widetilde{f}$ to F_{k+1}. Either S_{k+1} is minimal with respect to the ordering or it is mapped by $\widetilde{g}$ onto some S_i with $i \leq k$.

If S_{k+1} is minimal with respect to the ordering on strata, then S_{k+1} is not contained in S. Extend $\widetilde{f}$ continuously from $F_k \cap \operatorname{cl}(S_{k+1})$ to all of S_{k+1}, ensuring that $p(x) = p(x')$ implies $\widetilde{f}(x) = \widetilde{f}(x')$. This can be achieved as follows: By assumption the map $\widetilde{f}$ is defined and constant on all fibres of p which intersect F_k, so that we obtain a well-defined continuous map $\varphi \colon p(F_k \cap \operatorname{cl}(S_{k+1})) \to \mathbb{R}$ by letting $\varphi(p(x)) = \widetilde{f}(x)$. By the Tietze extension theorem the map φ extends continuously to $p(\operatorname{cl}(S_{k+1}))$. Then letting $\widetilde{f}(x) = \varphi(p(x))$ on S_{k+1} yields the desired function. This extension is continuous by definition. It still satisfies the functional equation since no points in S were added to its domain.

We hence have a continuous map $\widetilde{f} \colon F_k \cup S_{k+1} \to \mathbb{R}$ which factors through a continuous map $f \colon p(F_k \cup S_{k+1}) \to \mathbb{R}$. This yields a continuous extension to F_{k+1}, by letting $\widetilde{f}(x) = f \circ p(x)$ for all $x \in F_{k+1}$.

The extension still satisfies the functional equation: Let $x \in F_{k+1}$. Then there exists $x' \in S_{k+1} \cup F_k$ with $p(x') = p(x)$. If $x' \in S_{k+1}$ then $x' \notin S$ by minimality of S_{k+1}, so there is no constraint on $\widetilde{f}(x')$. If $x' \in F_k$ then the functional equation $\widetilde{f}(x') = \widetilde{f}(\widetilde{g}(x')) + 1$ is satisfied by induction hypothesis. This yields the functional equation for x as $\widetilde{f}$ is constant on fibres and $\widetilde{g}$ sends fibres to fibres.

This concludes the case where S_{k+1} is minimal with respect to the ordering on strata. Let us now assume that $\widetilde{g}(S_{k+1}) = S_i$ for some i.

On S_{k+1}, put $\widetilde{f}(x) = \widetilde{f}(\widetilde{g}(x)) + 1$. Then clearly $\widetilde{f}$ is continuous and satisfies the functional equation on S_{k+1}. To show that it is continuous on $S_{k+1} \cup F_k$ it suffices to show that it is continuous on the closure of S_{k+1}. Thus, let $(x_n)_n$ be a sequence in S_{k+1} which converges to $x \in F_k$. The set S is closed, so that $x \in S$. As $\widetilde{g}(x_n) \in F_k$ and F_k is closed, we have $\widetilde{g}(x) \in F_k$. By definition we have $\widetilde{f}(x_n) = \widetilde{f}(\widetilde{g}(x_n)) + 1$. Since $\widetilde{f}$ is continuous on F_k by induction hypothesis we have $\widetilde{f}(\widetilde{g}(x_n)) \to \widetilde{f}(\widetilde{g}(x))$. Hence, $\widetilde{f}(x_n) \to \widetilde{f}(\widetilde{g}(x)) + 1 = \widetilde{f}(x)$ where the equality holds because the functional equation is satisfied on F_k by induction hypothesis.

On F_{k+1}, let $\widetilde{f}(x) = \widetilde{f}(x')$ where $x' \in S_{k+1}$ and $p(x') = p(x)$. As above, the extension to F_{k+1} is well-defined and continuous.

It remains to show that the extension satisfies the functional equation on $F_{k+1} \cap S$. Let $x \in F_{k+1} \cap S$. Then $p(x) = p(x')$ where $x' \in S_{k+1} \cup F_k$. If $x' \in F_k$ then the functional equation $\widetilde{f}(x) = \widetilde{f}(\widetilde{g}(x)) + 1$ is satisfied by the same argument as above. If $x' \in S_{k+1}$ then the functional equation $\widetilde{f}(x') = \widetilde{f}(\widetilde{g}(x')) + 1$ is satisfied by construction. By construction, $\widetilde{f}(x') = \widetilde{f}(x)$ and $\widetilde{f}(\widetilde{g}(x')) = \widetilde{f}(\widetilde{g}(x))$ since $\widetilde{g}(x') \in F_k$, the set F_k contains all fibres that intersect it, and $\widetilde{f}$ is constant on fibres.

5 Decidability of termination in the case of linear updates

Theorem 1 raises the question whether the existence of a ranking function or, equivalently, termination is decidable for compact semi-algebraic K and continuous semi-algebraic update $g\colon \mathbb{R}^n \to \mathbb{R}^n$. We can answer this affirmatively in the case where g is linear:

▶ **Theorem 15.** *There exists an algorithm which receives as input a compact semi-algebraic set $K \subseteq \mathbb{R}^n$, encoded as a finite boolean combination of rational polynomial inequalities, and a linear map $A\colon \mathbb{R}^n \to \mathbb{R}^n$, encoded as a rational matrix, and decides whether for all $x \in K$ there exists $i \in \mathbb{N}$ such that $A^i x \notin K$.*

Proof. If every point of K escapes K under A in finitely many steps, then by Proposition 2 there exists a constant $M \in \mathbb{N}$ such that every point of K escapes K in at most M steps. For a fixed $m \in \mathbb{N}$, the statement that every point of K escapes K in at most m steps is a sentence in the first-order theory of the reals and hence decidable. It follows that we can semi-decide if every point escapes K under A.

We will show in Theorem 18 below that the existence of a point in K which does not escape K under A implies the existence of a semi-algebraic invariant for A in K, *i.e.*, a semi-algebraic set $S \subseteq K$ with $A(S) \subseteq S$. The statement that there exists a semi-algebraic invariant which can be expressed as a boolean combination of m inequalities involving polynomials of degree at most d is a sentence in the first-order theory of the reals and hence decidable. It follows that the existence of a non-escaping point is semi-decidable as well. ◀

More generally, termination is decidable for affine maps, *i.e.*, maps of the form $g(x) = Ax + b$ where A is an $n \times n$-matrix and b is an n-dimensional vector:

▶ **Corollary 16.** *There exists an algorithm which receives as input a compact semi-algebraic set $K \subseteq \mathbb{R}^n$, encoded as a finite boolean combination of rational polynomial inequalities, and an affine map $A\colon \mathbb{R}^n \to \mathbb{R}^n$, encoded by a rational matrix, and decides whether for all $x \in K$ there exists $i \in \mathbb{N}$ such that $A^i x \notin K$.*

Proof. Let $A(x) = B(x) + c$. Apply Theorem 15 to the compact semi-algebraic set $K \times \{1\} \subseteq \mathbb{R}^{n+1}$ and the linear map $\widetilde{A}(x, z) = (B(x) + cz, z)$. ◀

Note that the decision procedure in the proof of Theorem 15 combines two unbounded searches, one of which is guaranteed to terminate. We hence do not obtain any non-trivial upper bound on the computational complexity of deciding termination.

To complete the proof of Theorem 15 we require the following version of Kronecker's theorem, as stated in [15, Corollary 3.1]:

▶ **Lemma 17.** *Let $\lambda_1, \ldots, \lambda_f$ be complex numbers of modulus one. Consider the free abelian group*

$$L = \left\{ (v_1, \ldots, v_f) \in \mathbb{Z}^f \mid \lambda_1^{v_1} \cdot \cdots \cdot \lambda_f^{v_f} = 1 \right\}.$$

Let $\{\ell_1, \dots, \ell_p\}$ be a basis of this group. Then L is a free abelian group with a finite basis $\{\ell_1, \dots, \ell_p\}$. Let

$$T = \left\{(z_1, \dots, z_f) \in \mathbb{C}^f \mid |z_1| = \dots = |z_f| = 1, (z_1 \cdot \dots \cdot z_f)^{\ell_i} = 1 \text{ for all } i = 1, \dots, p\right\}.$$

Then the sequence $(\lambda_1^s, \dots, \lambda_f^s)_{s\in\mathbb{N}}$ is dense in T.

▶ **Theorem 18.** *Let $K \subseteq \mathbb{R}^n$ be a compact semi-algebraic set and let $A \colon \mathbb{R}^n \to \mathbb{R}^n$ be a linear map. Assume that there exists $x \in K$ such that $A^i x \in K$ for all $i \in \mathbb{N}$. Then there exists a closed semi-algebraic set $S \subseteq K$ with $A(S) \subseteq S$.*

Proof. By choosing an appropriate basis of $\mathbb{R}^n$ we may assume that A is given as a matrix in real Jordan normal form, *i.e.*, A can be written as

$$\begin{pmatrix} A_1 & & & \\ & A_2 & & \\ & & \ddots & \\ & & & A_l \end{pmatrix}$$

where the A_k's are Jordan blocks of the form

$$A_k = \begin{pmatrix} \Lambda_k & I & & & \\ & \Lambda_k & I & & \\ & & \ddots & \ddots & \\ & & & \Lambda_k & I \\ & & & & \Lambda_k \end{pmatrix}$$

where, by slight abuse of notation, Λ_k is either a real eigenvalue or a 2×2-matrix corresponding to a pair of complex conjugate eigenvalues, and I is either the number 1 or the 2×2-identity matrix. The s^{th} iterate of A_k is given by:

$$A_k^s = \begin{pmatrix} \Lambda_k^s & \binom{s}{1}\Lambda_k^{s-1} & \binom{s}{2}\Lambda_k^{s-2} & \cdots & \binom{s}{d_k-1}\Lambda_k^{s-d_k+1} \\ & \Lambda_k^s & \binom{s}{1}\Lambda_k^{s-1} & \cdots & \binom{s}{d_k-2}\Lambda_k^{s-d_k+2} \\ & & \ddots & \vdots & \vdots \\ & & & \Lambda_k^s & \binom{s}{1}\Lambda_k^{s-1} \\ & & & & \Lambda_k^s \end{pmatrix}.$$

Here, d_k denotes the size of the Jordan block, *i.e.*, the number of Λ_k's.

Let $x \in K$ be a point whose orbit under A does not escape K. Fix a Jordan block A_k as above. Let $x^k = (x_0^k, \dots, x_{d_k-1}^k)$ be the corresponding component of x. Again by slight abuse of notation, x_i^k denotes a number if Λ_k is a number and a 2-dimensional vector if Λ_k is a 2×2-matrix.

If Λ_k is real, let $|\Lambda_k|$ be the absolute value of Λ_k. If Λ_k is a 2×2-matrix, then we have $\Lambda_k = \begin{pmatrix} a & -b \\ b & a \end{pmatrix}$. In this case, let $|\Lambda_k| = \sqrt{a^2 + b^2}$.

If $|\Lambda_k| > 1$ then we claim that $x^k = (0, \dots, 0)$. Indeed, if $x_i^k \neq 0$, then the i^{th} component of $A_k^s(x^k)$ is of the form $\Lambda_k^s x_i^k + O(|\Lambda_k^{s-1}|)$. Hence, the euclidean distance of $A^s x$ to 0 is unbounded as $s \to \infty$, contradicting the assumption that the sequence $A^s x$ stays in the compact set K.

If $|\Lambda_k| < 1$ then each of the expressions $\binom{s}{i}\Lambda_k^{s-i}$ converges to zero as $s \to \infty$, so that $A_k^s x^k \to 0$ as $s \to \infty$.

If $|\Lambda_k| = 1$ then we claim that $x^k = (x_0^k, 0, \ldots, 0)$. Indeed, the first component of $A^s x^k$ is equal to

$$\Lambda_k^s x_0 + \sum_{i=1}^{d_k-1} \binom{s}{i} \Lambda_k^{s-i} x_i^k.$$

For large s the term $\binom{s}{i}\Lambda_k^{s-i} x_i^k$ with the largest i where $x_i^k \neq 0$ dominates the whole expression. If $i \neq 0$ then the expression is unbounded as $s \to \infty$. The claim follows.

Permute the basis vectors such that A is given in the new basis as

$$A = \left(\begin{array}{cccc|c} \Lambda_1 & & & & \\ & \Lambda_2 & & & \\ & & \ddots & & \\ & & & \Lambda_f & \\ \hline & & & & \text{A'} \end{array}\right)$$

Where $\Lambda_1, \ldots, \Lambda_f$ correspond to all the real and complex eigenvalues of A of modulus one, so that $(A^s x)_i \to 0$ for all $i \geq f + 1$.

Let

$$\bar{A} = \begin{pmatrix} \Lambda_1 & & & & & & \\ & \Lambda_2 & & & & & \\ & & \ddots & & & & \\ & & & \Lambda_f & & & \\ & & & & 0 & & \\ & & & & & \ddots & \\ & & & & & & 0 \end{pmatrix}$$

Then we have $A\bar{A}^s x = \bar{A}^{s+1} x$, so that the orbit of x under $\bar{A}$ is invariant under A. Since A is continuous, the closure of the orbit is still invariant under A. We will now show that the closure of the orbit is a semi-algebraic set.

We can associate with each Λ_i with $i \leq f$ a complex algebraic number λ_i of modulus one. Let

$$L = \left\{ (v_1, \ldots, v_f) \in \mathbb{Z}^f \mid \lambda_1^{v_1} \cdot \cdots \cdot \lambda_f^{v_f} = 1 \right\}.$$

Then L is a free abelian group with a finite basis $\{\ell_1, \ldots, \ell_p\}$. Let

$$T = \left\{ (z_1, \ldots, z_f) \in \mathbb{C}^f \mid |z_1| = \cdots = |z_f| = 1, (z_1 \cdot \cdots \cdot z_f)^{\ell_i} = 1 \text{ for all } i = 1, \ldots, p \right\}$$

Then T is a complex algebraic set and by Lemma 17 the sequence $(\lambda_1^s, \ldots, \lambda_f^s)$ is dense in T. This yields a real algebraic subset T' of $\mathbb{R}^{n \times n}$ such that $\bar{A}^s$ is dense in T'. The matrix evaluation map $\mathbb{R}^{n \times n} \times \mathbb{R}^n \to \mathbb{R}^n$ is a polynomial map, so that the image of $T' \times \{x\}$ under this map is a semi-algebraic subset of $\mathbb{R}^n$. But this is the same as the closure of the orbit of x under $\bar{A}$.

We have now shown that the closure S of the orbit of x under $\bar{A}$ is a semi-algebraic invariant for A. It remains to show that S is contained in K. The above argument establishes that the iterates $\bar{A}^s x$ are dense in S. Now, A can be written as $A = \bar{A} + B$ with $A^s = \bar{A}^s + B^s$

and $B^s x \to 0$ as $s \to \infty$. Since the iterates $\bar{A}^s x$ are dense in S, for every $y \in S$ there exists a sequence $(s_k)_k$ with $\bar{A}^{s_k} x \to y$ as $k \to \infty$. It follows that $A^{s_k} x = \bar{A}^{s_k} x + B^{s_k} x \to y$ as $k \to \infty$. By assumption $A^{s_k} x \in K$ for all k, and since K is closed it follows that $y \in K$. We conclude that $S \subseteq K$. ◀

6 Conclusion

We have shown, by non-constructive means, that a single-path loop with continuous semi-algebraic update function and compact semi-algebraic guard set terminates over the reals if and only if it has a polynomial ranking function. In the case of an affine update we have shown the existence of a polynomial ranking function to be decidable by proving that any non-terminating loop of this form admits a semi-algebraic invariant.

This naturally suggests the question whether the existence of a polynomial ranking function can be decided for non-linear updates as well. A sufficient condition for decidability which may be of independent interest is whether an analogue of Theorem 18 holds true for certain classes of non-linear maps, say for instance whether any non-terminating polynomial loop with compact guard admits a semi-algebraic invariant.

Further, it would be interesting to study the computational complexity of deciding termination in the case of affine updates. Our decidability proof unfortunately does not yield any non-trivial complexity bounds.

Another direction of future research is to ascertain whether Theorem 1 generalises to set-valued semi-algebraic updates with appropriate continuity properties. Also the boundedness assumption on the guard set deserves to be further scrutinised. While the examples after Theorem 1 show that closedness is necessary for the existence of a continuous ranking function and that boundedness is necessary for the existence of a polynomial ranking function, it seems reasonable to conjecture that a terminating semi-algebraic loop with a closed but not necessarily bounded guard has, say, a piecewise-linear ranking function.

References

1. Amir M. Ben-Amram, Jesús J. Doménech, and Samir Genaim. Multiphase-Linear Ranking Functions and Their Relation to Recurrent Sets. In Bor-Yuh Evan Chang, editor, *Static Analysis*, pages 459–480. Springer International Publishing, 2019.
2. Amir M. Ben-Amram and Samir Genaim. Ranking Functions for Linear-Constraint Loops. *J. ACM*, 61(4):26:1–26:55, 2014.
3. Amir M. Ben-Amram and Samir Genaim. Complexity of Bradley-Manna-Sipma Lexicographic Ranking Functions. In *Computer Aided Verification*, volume 9207 of *Lecture Notes in Computer Science*, pages 304–321. Springer, 2015.
4. Amir M. Ben-Amram and Samir Genaim. On Multiphase-Linear Ranking Functions. In *Computer Aided Verification - 29th International Conference, CAV 2017*, volume 10427 of *Lecture Notes in Computer Science*, pages 601–620. Springer, 2017.
5. Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real Algebraic Geometry*. Springer-Verlag Berlin Heidelberg, 1998.
6. Mark Braverman. Termination of Integer Linear Programs. In *Computer Aided Verification, 18th International Conference, CAV, Proceedings*, volume 4144 of *Lecture Notes in Computer Science*, pages 372–385. Springer, 2006.
7. Krishnendu Chatterjee, Hongfei Fu, and Amir Kafshdar Goharshady. Termination Analysis of Probabilistic Programs Through Positivstellensatz's. In Swarat Chaudhuri and Azadeh Farzan, editors, *Computer Aided Verification - 28th International Conference, CAV 2016, Proceedings, Part I*, volume 9779 of *Lecture Notes in Computer Science*, pages 3–22. Springer, 2016.

8 Yinghua Chen, Bican Xia, Lu Yang, Naijun Zhan, and Chaochen Zhou. Discovering Non-linear Ranking Functions by Solving Semi-algebraic Systems. In *Theoretical Aspects of Computing - ICTAC 2007, 4th International Colloquium, 2007, Proceedings*, volume 4711 of *Lecture Notes in Computer Science*, pages 34–49. Springer, 2007.

9 Byron Cook, Abigail See, and Florian Zuleger. Ramsey vs. Lexicographic Termination Proving. In *Tools and Algorithms for the Construction and Analysis of Systems - 19th International Conference, TACAS 2013*, volume 7795 of *Lecture Notes in Computer Science*, pages 47–61. Springer, 2013.

10 Patrick Cousot. Proving Program Invariance and Termination by Parametric Abstraction, Lagrangian Relaxation and Semidefinite Programming. In *Verification, Model Checking, and Abstract Interpretation, 6th International Conference, VMCAI 2005, Proceedings*, volume 3385 of *Lecture Notes in Computer Science*, pages 1–24. Springer, 2005.

11 Florian Frohn, Marcel Hark, and Jürgen Giesl. On the Decidability of Termination for Polynomial Loops. *CoRR*, abs/1910.11588, 2019. `arXiv:1910.11588`.

12 Mehran Hosseini, Joël Ouaknine, and James Worrell. Termination of Linear Loops over the Integers. In *46th International Colloquium on Automata, Languages, and Programming, ICALP*, volume 132 of *LIPIcs*, pages 118:1–118:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

13 Yi Li. Termination of Single-Path Polynomial Loop Programs. In *Theoretical Aspects of Computing - ICTAC, Proceedings*, volume 9965 of *Lecture Notes in Computer Science*, pages 33–50, 2016.

14 Yi Li, Wenyuan Wu, and Yong Feng. On ranking functions for single-path linear-constraint loops. *Int J Softw Tools Technol Transfer*, 2019. `doi:10.1007/s10009-019-00549-9`.

15 Joël Ouaknine and James Worrell. Positivity Problems for Low-Order Linear Recurrence Sequences. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 366–379. SIAM, 2014.

16 Andreas Podelski and Andrey Rybalchenko. A Complete Method for the Synthesis of Linear Ranking Functions. In *Verification, Model Checking, and Abstract Interpretation, 5th International Conference, VMCAI 2004, Proceedings*, volume 2937 of *Lecture Notes in Computer Science*, pages 239–251. Springer, 2004.

17 Ashish Tiwari. Termination of Linear Programs. In Rajeev Alur and Doron A. Peled, editors, *Computer Aided Verification, 16th International Conference, CAV 2004, Proceedings*, volume 3114 of *Lecture Notes in Computer Science*, pages 70–82. Springer, 2004.

18 Caterina Urban. The Abstract Domain of Segmented Ranking Functions. In *Static Analysis - 20th International Symposium, SAS 2013, Proceedings*, volume 7935 of *Lecture Notes in Computer Science*, pages 43–62. Springer, 2013.

19 Bican Xia and Zhihai Zhang. Termination of linear programs with nonlinear constraints. *Journal of Symbolic Computation*, 45(11):1234–1249, 2010. `doi:10.1016/j.jsc.2010.06.006`.