Linear dynamical systems with continuous weight functions

Rajab Aghamov
Technische Universität Dresden, Germany

Christel Baier
Technische Universität Dresden, Germany

Toghrul Karimov
Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany

Joël Ouaknine
Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany

Jakob Piribauer
Technische Universität Dresden, Germany; Technische Universität München, Germany

Abstract

In discrete-time linear dynamical systems (LDSs), a linear map is repeatedly applied to an initial vector yielding a sequence of vectors called the orbit of the system. A weight function assigning weights to the points in the orbit can be used to model quantitative aspects, such as resource consumption, of a system modelled by an LDS. This paper addresses the problems to compute the mean payoff, the total accumulated weight, and the discounted accumulated weight of the orbit under continuous weight functions and polynomial weight functions as a special case. Besides general LDSs, the special cases of stochastic LDSs and of LDSs with bounded orbits are considered. Furthermore, the problem of deciding whether an energy constraint is satisfied by the weighted orbit, i.e., whether the accumulated weight never drops below a given bound, is analysed.

2012 ACM Subject Classification

Keywords and phrases linear dynamical systems, mean payoff, total reward, discounted reward

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Dynamical systems describing how the state of a system changes over time constitute a prominent modelling paradigm in a wide variety of fields. A *discrete-time linear dynamical system* (LDS) in ambient space $\mathbb{R}^d$ starts at some initial point $q \in \mathbb{R}^d$. The dynamics of the system are given by a linear update function in form of a matrix $M \in \mathbb{R}^{d\times d}$ that is applied to the current state of the system at each time step. This gives rise to the *orbit* $(q, Mq, M^2q, \dots)$. Surprisingly, several seemingly simple decidability questions about the orbit of a given LDS have been open for many decades (for an overview, see [11]). For example, two prominent problems about *linear recurrence sequences*, the Positivity Problem and the Skolem Problem, are subsumed by the following problem: given (M, q) and a target set H, decide whether there exists $n \in \mathbb{N}$ such that $M^n q \in H$.

Investigation of algorithmic problems concerning LDSs is a lively area of research in computer science. In order to verify that a system modelled as an LDS satisfies desirable properties, typical formal verification problems such as model-checking problems asking whether the orbit of an LDSs satisfies certain temporal properties have been studied [3, 12]. One important special case of LDSs are *stochastic LDSs*. For a finite-state Markov chain, the sequence of distributions over the state space naturally forms an LDS: The initial distribution can be written as a vector $\iota_{init} \in [0,1]^d$. Afterwards, the transition probability matrix P can be repeatedly applied to obtain the distribution $P^k\iota_{init}$ over states after k steps. In contrast

Table 1 Overview of the results.

	LDS type	weight function	algorithmic results	
mean payoff	arbitrary	polynomial	computable	(Thm. 7)
	bounded orbit	continuous	integral representation computable	(Thm. 11)
	stochastic, irreducible	continuous	computable with polynomial many evaluations of the weight function.	(Thm. 13)
	stochastic, reducible	continuous	computable with exponentially many evaluations of the weight function.	(Thm. 14)
total/discounted weight	arbitrary	polynomial	computable	(Thm. 15)
satisfaction of energy constraints	arbitrary	polynomial	decidable in dimension 3	(Thm. 21)
	stochastic	linear	Positivity-hard	(Thm. 22)

to the path semantics where a probability measure over infinite paths in a Markov chain is defined, the view of a Markov chain as an LDS is also called the *distribution transformer semantics* of Markov chains. In this way, LDSs also play an important role in the analysis of probabilistic systems.

In this paper, we address quantitative verification questions arising when systems are equipped with a weight function. Such a weight function assigns a weight to each state of the system that can be used to model various quantitative aspects of a system, such as resource or energy consumption, rewards or utilities, or execution time for example. To this end, we consider a weight function $w\colon \mathbb{R}^d \to \mathbb{R}$ assigning a weight to each state in the ambient space and obtain a sequence of weights of the states in the orbit $(w(q), w(Mq), w(M^2q), \dots)$. The goal of this paper is to provide algorithmic answers to the following typical questions arising for weighted systems:

a) What is the *mean payoff*, i.e., the average weight collected per step?
b) What is the total accumulated weight of the orbit and what is the so-called discounted accumulated weight, where weights obtained after k time steps are discounted with a factor λ^k for a given $\lambda \in (0,1)$?
c) Is there an $n \in \mathbb{N}$ such that the sum of weights obtained in the first n steps lies below a given bound? This problem is referred to as *satisfaction of an energy-constraint* because it corresponds to determining whether a system ever runs out of energy when weights model the energy used or gained during a step.

▶ **Example 1.** Assume a scheduler assigns tasks to d different processors $P_1, \dots, P_d$ and that the load of the processors at different time steps can be modeled as an LDS with matrix $M \in \mathbb{Q}^{d\times d}$ and orbit $(M^k q)_{k\in\mathbb{N}}$ for a $q \in \mathbb{Q}^d$. Further, assume for each processor P_i there is an optimal load μ_i under which it works most efficiently. To evaluate the scheduler, we want to know how closely the real loads in the long-run match the ideal loads. As a measure for how well a vector x matches the vector μ of ideal loads, we use the average squared distance

$$\delta_\mu(x) = \frac{1}{d}\sum_{i=1}^{d}(x_i - \mu_i)^2.$$

To see how well the scheduler manages to get close to optimal loads in the long-run after a possible intialization phase, we consider the mean payoff of the orbit with respect to the weight function δ_μ, i.e.,

$$\lim_{\ell\to\infty} \frac{1}{\ell+1}\sum_{k=0}^{\ell}\delta_\mu(M^k q).$$

If, on the other hand, we know that the orbit will tend to the optimal loads for $k \to \infty$, we might instead also want to measure the total deviation $\sum_{k=0}^{\infty} \delta_\mu(M^k q)$. If this value is small, the orbit converges to the optimal loads rather quickly without large deviations initially.

Contribution.

We address the problems mentioned above for weighted LDS with rational entries under continuous weight functions. For a general LDS and an arbitrary continuous weight function, not much can be said. We either have to restrict the class of LDSs or the class of weight functions in order to be able to address computational problems. Our contributions are as follows. An overview of the results can also be found in Table 1.

a) Mean payoff: For rational LDSs equipped with a polynomial weight function, we show that it is decidable whether the mean payoff exists, in which case it is rational and computable. We then show how to decide whether the orbit of a rational LDS is bounded. If the orbit of a rational LDS is bounded, we show how to compute the set of accumulation points of the orbit and prove that the mean payoff of the orbit can be expressed as an integral of the weight function over a computable parametrisation of this set. We next consider stochastic LDSs, which constitute a special case of LDSs with bounded orbits. We show that in case the transition matrix is irreducible, then one can compute polynomially many rational points in polynomial time such that the mean payoff is the arithmetic mean of the weight function evaluated at these points. In the reducible case, on the other hand, exponentially many such rational points have to be computed.

b) Total and discounted accumulated weights: For rational LDSs and polynomial weight functions, we prove that the total as well as the discounted accumulated weight of the orbit is computable and rational if finite.

c) Satisfaction of energy constraints: First we prove that it is decidable whether an energy constraint is satisfied by an orbit under a polynomial weight function for LDS of dimension $d = 3$. On the other hand, we show that the problem is at least as hard as the Positivity problem for linear recurrence sequences already for stochastic LDSs and linear weight functions. The decidability status of the Positivity Problem is open. In fact, a decidability result would amount to a major breakthrough in Diophantine approximation.

Related work.

Verification problems for linear dynamical systems have been extensively studied for decades, starting with the question about the decidability of the Skolem [21, 23] and Positivity [19, 20] problems, which are special cases of the reachability problem for LDSs, at low orders. Decidable cases of the more general Model-Checking Problem for LDSs have been studied in [3, 12]. In addition, decidability results for parametric LDSs [4] as well as various notions of robust verification [2, 8] have been obtained. See [11] for a survey of what is decidable about discrete-time linear dynamical systems. Recently, Kelmendi has shown [15] that the *natural density* (which is a notion of frequency) of visits of an LDS in a semialgebraic set always exists and is computable to arbitrary precision.

When it comes to Markov chains viewed as LDSs under the distribution transformer semantics, it is known that Skolem and Positivity-hardness results for general LDS persist [1]. Vahanwala has recently shown [22] that this is the case even for ergodic Markov chains.

2 Preliminaries

We briefly present our notation and introduce the concepts used in the subsequent sections.

2.1 Linear dynamical systems

A (discrete-time) *linear dynamical system* (LDS) (M, q) of dimension $d > 0$ consists of an update matrix $M \in \mathbb{R}^{d\times d}$ and an initial vector $q \in \mathbb{R}^d$. The *orbit* $\mathcal{O}(M, q)$ of (M, q) is the sequence $(M^k q)_{k\in\mathbb{N}}$. We say that the orbit of (M, q) is bounded if there exists $c \in \mathbb{R}$ such that $|M^k q| < c$ for all $k \in \mathbb{N}$. An LDS is called *stochastic* if the matrix M and the initial vector q have only non-negative entries and the entries of each column of M as well as the entries of q sum up to 1. In this case we refer to the matrix M as stochastic too.[1]

2.2 Algebraic numbers

A number $\alpha \in \mathbb{C}$ is *algebraic* if there exists a polynomial $p \in \mathbb{Q}[x]$ such that $p(\alpha) = 0$. Algebraic numbers form a subfield of $\mathbb{C}$ denoted by $\overline{\mathbb{Q}}$. The minimal polynomial of $\alpha \in \overline{\mathbb{Q}}$ is the (unique) monic polynomial $p \in \mathbb{Q}[X]$ of the smallest degree such that $p(\alpha) = 0$. The *degree* of α, denoted by $\deg(\alpha)$, is the degree of the minimal polynomial of α. For each $\alpha \in \overline{\mathbb{Q}}$ there exists a unique polynomial $P_\alpha = \sum_{i=0}^{d} a_i x^i \in \mathbb{Z}[x]$ with $d = \deg(\alpha)$, called the *defining polynomial* of α, such that $P_\alpha(\alpha) = 0$ and $\gcd(a_0, \ldots, a_d) = 1$. The polynomial P_α and the minimal polynomial of α have identical roots, and are *square-free*, i.e., all of their roots appear with multiplicity one. The *(naive) height* of α, denoted by $H(\alpha)$, is equal to $\max_{0\le i\le d} |a_i|$. We represent an algebraic number α in computer memory by its defining polynomial P_α and sufficiently precise rational approximations of $\mathrm{Re}(\alpha), \mathrm{Im}(\alpha)$ to distinguish α from other roots of P_α. We denote by $||\alpha||$ the bit length of a representation of $\alpha \in \overline{\mathbb{Q}}$. We can perform arithmetic effectively on algebraic numbers represented in this way

2.3 Linear recurrence sequences

A sequence $(u_n)_{n\in\mathbb{N}}$ is a *linear recurrence sequence* over a ring $R \subseteq \mathbb{C}$ if there exists a positive integer d and a *recurrence relation* $(a_0, \ldots, a_{d-1}) \in R^d$ such that $u_{n+d} = \sum_{i=0}^{d-1} a_i u_{n+i}$ for all $n \in \mathbb{N}$. The *order* of $(u_n)_{n\in\mathbb{N}}$ is the smallest positive integer d such that $(u_n)_{n\in\mathbb{N}}$ satisfies a recurrence relation in R^d. We will mostly work with sequences over $\mathbb{Q}$. Examples of rational LRS include the Fibonacci sequence, $u_n = p(n)$ for $p \in \mathbb{Q}[x]$, and $u_n = \cos(n\theta)$ where $\theta \in \{\arg(\lambda) : \lambda \in \mathbb{Q}(i)\}$. We refer the reader to the books by Everest et al. [9] and Kauers & Paule [14] for a detailed discussion of linear recurrence sequences.

Let $(u_n)_{n\in\mathbb{N}}$ be a non-zero LRS given by the (minimal) recurrence relation $u_{n+d} = \sum_{i=0}^{d-1} a_i u_{n+i}$. Writing $A = \begin{bmatrix} a_1 & \cdots & a_{d-1} \end{bmatrix}$ and $q = \begin{bmatrix} u_0 & \cdots & u_{d-1} \end{bmatrix}^\top$, the matrix

$$C := \begin{bmatrix} \mathbf{0} & I_{d-1} \\ a_0 & A \end{bmatrix} = \begin{bmatrix} 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \\ a_0 & a_1 & \cdots & a_{d-1} \end{bmatrix} \in R^{d\times d}$$

is called the *companion matrix* of $(u_n)_{n\in\mathbb{N}}$. We have that $C^n q = \begin{bmatrix} u_n & \cdots & u_{n-d+1} \end{bmatrix}^\top$ and $u_n = e_1 C^n s$ for all $n \in \mathbb{N}$, where e_i denotes the ith standard basis vector. Note that as $a_0 \neq 0$, the matrix C is invertible and does not have zero as an eigenvalue.

The *characteristic polynomial* of $(u_n)_{n\in\mathbb{N}}$ is given by $p(x) = x^d - \sum_{i=0}^{d-1} a_i x^i$. Note that p is identical to the characteristic polynomial $\det(xI - C)$ of the companion matrix C. The

[1] Note that – in order to keep the notation in line with the notation for general LDSs – we deviate from the standard convention that rows of stochastic matrices sum up to 1 and that stochastic matrices are applied to distributions by multiplication from the right.

eigenvalues (also called the *roots*) of $(u_n)_{n\in\mathbb{N}}$ are the d (possibly non-distinct) roots $\lambda_1, \dots, \lambda_d$ of the characteristic polynomial p. An LRS is

- *simple* (or *diagonalisable*) if its characteristic polynomial does not have a repeated root,
- *non-degenerate* if (i) all real eigenvalues are non-negative, and (ii) for every pair of distinct eigenvalues λ_1, λ_2, the ratio λ_1/λ_2 is not a root of unity.

For each LRS $(u_n)_{n\in\mathbb{N}}$ there exists effectively computable L such that the sequences $(u_n^{(k)})_{n\in\mathbb{N}}$ for $0 \le k < L$ defined by $u_n^{(k)} = nL + k$ are all non-degenerate [9, Section 1.1.9]. Finally, if $(u_n)_{n\in\mathbb{N}}, (v_n)_{n\in\mathbb{N}}$ are LRS over a field R, and $\circ \in \{+, -, \cdot\}$, then $w_n = u_n \circ v_n$ also defines an LRS [14, Theorem 4.2] over R. Moreover, if $(u_n)_{n\in\mathbb{N}}$ and $(v_n)_{n\in\mathbb{N}}$ are both simple, then so is $(w_n)_{n\in\mathbb{N}}$.

The exponential polynomial representation of an LRS

Every LRS $(u_n)_{n\in\mathbb{N}}$ of order $d > 0$ over $\overline{\mathbb{Q}}$ can be written in the form [9, Chapter 1]

$$u_n = \sum_{j=1}^{m} p_j(n)\lambda_j^n \tag{1}$$

where $m \ge 1$ if $(u_n)_{n\in\mathbb{N}}$ is not identically zero, $\lambda_1, \dots, \lambda_m$ are the distinct non-zero eigenvalues of $(u_n)_{n\in\mathbb{N}}$, and each p_i is a non-zero polynomial with algebraic coefficients. With these conditions, we say that the right-hand side is in the *exponential polynomial form*. The following two lemmas about exponential polynomial solutions of LRS are folklore. For completeness, we give the proofs in Appendix A.

▶ **Lemma 2.** *Let $u_n = \sum_{i=1}^{m} p_i(n)\lambda_i^n$, where all $\lambda_i \in \overline{\mathbb{Q}}$ and $p_i \in \overline{\mathbb{Q}}[x]$ are non-zero, and $\lambda_i \ne \lambda_j$ for $i \ne j$. The sequence $(u_n)_{n\in\mathbb{N}}$ is not identically zero. Specifically, there exists $0 \le n < d$, where $d = \sum_{i=1}^{m}(\deg(p_i) + 1)$, such that $u_n \ne 0$.*

▶ **Lemma 3.** *Let $(u_n)_{n\in\mathbb{N}}$ be as in the statement of Lemma 2. If $u_n \in \mathbb{R}$ for all $n \in \mathbb{N}$, then for every $1 \le i \le m$ there exists j such that $p_j(n) = \overline{p_i}(n)$ and $\lambda_j = \overline{\lambda_i}$.*

Throughout this work we will consider sequences of the form $u_n = p(M^n q)$ where p is a polynomial with rational coefficients. Since $p(M^n q) = p(e_1 M^n q, \dots, e_d M^n q)$, each $u_N^{(k)} = e_k M^n q$ is an LRS over $\mathbb{Q}$ (this can be seen, e.g., by applying the Cayley-Hamilton theorem), and LRS over $\mathbb{Q}$ are closed under addition and multiplication, the sequence $(p(M^n q))_{n\in\mathbb{N}}$ is itself an LRS over $\mathbb{Q}$.

Decision problems about LRS

Sign patterns of LRS have been studied for a long time. Two prominent open problems in this area are the *Skolem Problem* and the *Positivity Problem*. The Skolem Problem is to find an algorithm that, given an LRS u_n, decides if the set $Z = \{n : u_n = 0\}$ is non-empty. The most well-known result in this direction is the celebrated Skolem-Mahler-Lech theorem, which (non-constructively) shows that Z is semilinear. In particular, it shows that a non-degenerate $(u_n)_{n\in\mathbb{N}}$ can have only finitely many zeros. The Positivity Problem, on the other hand, asks to find an algorithm that determines if $u_n \ge 0$ for all n.

2.4 Markov Chains.

A finite-state *discrete-time Markov chain* (DTMC) M is a tuple (S, P, ι_{init}), where S is a finite set of states, $P : S \times S \to [0, 1]$ is the transition probability function where we require $\sum_{s'\in S} P_{ss'} = 1$ for all $s \in S$ and $i_{init} : S \to [0, 1]$ is the initial distribution, such that $\sum_{s\in S} i_{init}(s) = 1$. For algorithmic problems, all transition probabilities are assumed to be rational. A finite path ρ in M is a finite sequence $s_0 s_1 \dots s_n$ of states such that $P(s_i, s_{i+1}) > 0$ for all $0 \le i \le n-1$. We say that a state s is reachable from t if there is a finite path from s to t. If all states are reachable from all other states, we say that M is *irreducible*; otherwise, we say it is *reducible*. A set $B \subseteq S$ of states is called a bottom strongly connected component (BSCC) if it is strongly connected, i.e., all states in B are reachable form all other states in B and if there are no outgoing transitions, i.e., $P(s,t) > 0$ and $s \in B$ implies $t \in B$.

W.l.o.g., we identify S with $\{1, \dots, d\}$ for $d = |S|$. Then, overloading notation, we consider $P \in \mathbb{R}^{d \times d}$ as a matrix with $P_{ij} = P(j,i)$ for $i, j \le d$.[2] Likewise, we consider ι_{init} to be a (column[3]) vector in $\mathbb{R}^d$ with $(\iota_{init})_i = \iota_{init}(i)$ for $i \le d$. Then, the sequence of distributions over states after k steps is given by $P^k \iota_{init}$, which forms a stochastic LDS. We also write $P^{(k)}_{ij}$ for $(P^k)_{ij}$, which is the probability to move from state j to i in exactly k steps. Further, we say that the matrix P is irreducible if the underlying Markov chain is irreducible. The period d_i of a state i is given by: $d_i = \mathbf{gcd}\{m \ge 1 : P^{(m)}_{ii} > 0\}$. If $d_i = 1$, then we call the state i aperiodic. A Markov chain (and its matrix) are aperiodic if and only if all its states are aperiodic. The period of a Markov chain M as well as of its transition probability matrix P is the least common multiple of the periods of the states of M.

A vector $\pi \in \mathbb{R}^d$ is called a stationary distribution of the Markov chain if: a) π is a distribution, i.e., $\pi_j \ge 0$ for all j with $1 \le j \le d$, and $\sum_{j=1}^{d} \pi_j = 1$; b) π is stationary, i.e., $\pi = P\pi$, which is to say that $\pi_i = \sum_{i \in S} P_{ij} \pi_j$ for all $j \in S$. For aperiodic Markov chains, it is known that the sequence of distributions over states $(P^k \iota_{init})_{k \in \mathbb{N}}$ converges to a stationary distribution π, which can be computed in polynomial time (see [16, 5]).

3 Mean payoff

In this section, we address the computation of the *mean payoff* of an orbit. The mean payoff is the average weight collected per step in the long-run. For an LDS given by $M \in \mathbb{Q}^{d \times d}$ and $q \in \mathbb{Q}^d$ and a weight function $w \colon \mathbb{R}^d \to \mathbb{R}$, we define the mean payoff of the orbit as

$$MP_w(M,q) \coloneqq \lim_{k \to \infty} \frac{1}{k+1} \sum_{i=0}^{k} w(M^i q).$$

In the sequel, we address the problem of computing the mean payoff of the orbit of an LDS with respect to continuous weight functions. For general LDSs, there is not much we can say without knowing more about the form of the weight function. Hence, we have to restrict either the class of weight functions or the class of LDSs. In Section 3.1 we address the problem for polynomial weight functions. In Sections 3.2 and 3.3 we consider continuous weight functions on two classes of systems: LDSs with bounded orbit and stochastic LDSs.

3.1 Polynomial weight-functions

In order to compute the mean payoff of the orbit of an LDS (M, q) with respect to a polynomial weight function p, we first recall that the sequence $(p(M^n q))_{n \in \mathbb{N}}$ is an LRS. The following lemma states that the sequence of partial sums of the weights is also an LRS.

---

[2] This is the transpose of the transition matrix usually defined so that we are in line with our notation for LDSs.

[3] Also here, usually, this is defined as a row vector.

▶ **Lemma 4.** *Let (M, q) be an LDS with $M \in \mathbb{Q}^{d\times d}$ and $q \in \mathbb{Q}^d$, and let $p \in \mathbb{Q}[X_1, \ldots, X_d]$ be a polynomial weight function with rational coefficients. The sequence*

$$u_n = \sum_{i=0}^{n} p(M^i q)$$

is a rational LRS.

Proof. As discussed in subsection 2.3, $w_n = p(M^i q)$ is a rational LRS. Suppose $(w_n)_{n\in\mathbb{N}}$ satisfies a recurrence relation $w_{n+k} = a_0 w_n + \ldots + a_{k-1} w_{n+k01}$, where $a_0, \ldots, a_{k-1} \in \mathbb{Q}$. Then $u_{n+k+1} = u_{n+k} + a_{k-1}(w_{n+k} - w_{n+k-1}) + \ldots + a_0(w_{n+1} - w_n)$. Hence $(u_n)_{n\in\mathbb{N}}$ itself is an LRS of order at most $k + 1$. ◀

Computing $MP_w(M, q)$ hence boils down to determining whether the limit $\lim_{n\to\infty} u_n/n$ exists for an LRS $(u_n)_{n\in\mathbb{N}}$ and computing the limit in case it exists.

▶ **Theorem 5.** *Let $(u_n)_{n\in\mathbb{N}}$ be an LRS over $\mathbb{Q}$. It is decidable whether $\lim_{n\to\infty} u_n/n$ exists, in which case the limit is rational and effectively computable.*

The proof can be found in the appendix. Its main ideas are as follows. By a fundamental result, $|u_n|$ for an LRS $(u_n)_{n\in\mathbb{N}}$ essentially grows at the rate ρ^n, where $\rho > 0$ is the largest magnitude of an eigenvalue. If $\rho = 1$, then the sequence $(u_n)_{n\in\mathbb{N}}$ exhibits a recurring behaviour, which is also well-understood. Hence $\lim_{n\to\infty} u_n$ exists only in rather specific situations. The sequence $(u_n/n)_{n\in\mathbb{N}}$, in this context, almost behaves like an LRS. Hence similar arguments are applicable.

An immediate corollary that will be useful again in Section 4 is the following.

▶ **Corollary 6.** *For a rational LRS $(u_n)_{n\in\mathbb{N}}$, it is decidable whether $\lim_{n\to\infty} u_n$ exists, in which case the limit is rational and effectively computable.*

Proof. Observe that $\lim_{n\to\infty} u_n = \lim_{n\to\infty} v_n/n$, where $v_n = nu_n$ is a rational LRS. ◀

Furthermore, Theorem 5 puts us into the position to prove the first main result on the computation of the mean payoff:

▶ **Theorem 7.** *Let (M, q) be an LDS with $M \in \mathbb{Q}^{d\times d}$ and $q \in \mathbb{Q}^n$ and let $p \in \mathbb{Q}[X_1, \ldots, X_d]$ be a polynomial weight function with rational coefficients. Then, it is decidable whether the mean payoff*

$$MP_p(M, q) = \lim_{k\to\infty} \frac{1}{k+1} \sum_{i=0}^{k} p(M^i q)$$

exists and, in which case it is rational and computable.

Proof. Immediate by Theorem 5 and Lemma 4. ◀

3.2 Bounded LDSs

If the orbit of an LDS is bounded, we can get our hands on the mean payoff with respect to a continuous weight function. We exploit that the orbit of an LDS approaches a limiting shape – which is the set of accumulation points of the orbit – closer and closer in this case. This allows us to express the mean payoff in terms of an integral of the weight function over this limiting shape. This integral computes the "average" value of the weight function on the limiting shape. Of course, we have to carefully ensure that we also know how "frequently" the orbit approaches different parts of the limiting shape. Let us illustrate this idea first:

▶ **Example 8.** Let $w\colon \mathbb{R}^3 \to \mathbb{R}$ be a continuous weight function and consider the LDS

$$M = \begin{bmatrix} 3/5 & 4/5 & 0 \\ -4/5 & 3/5 & 0 \\ 0 & 0 & 1/2 \end{bmatrix} \quad \text{and} \quad q = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}.$$

Looking only at the first two coordinates a rotation is repeatedly applied in this LDS. In the complex plane, this rotation is given by multiplication with $3/5 - 4/5i$. As $3/5 - 4/5i$ is not a root of unity, the orbit never reaches a point with $(1, 0)$ in the first two coordinates again. In fact, the first two components of the orbit are dense in the unit circle. Furthermore, these components visit each interval of the same length on the circle with the same frequency. The third component is halved at every step and converges to 0. As the weight function is continuous, we can hence treat the third coordinate as equal to 0 when determining the mean payoff. So, the set of accumulation points of the orbit is $L = \{v \in \mathbb{R}^3 \mid v_3 = 0, |v| = 1\}$, which we can parametrise via $T\colon [0, 1) \to \mathbb{R}^3$ with $T\colon \alpha \mapsto \begin{bmatrix}\cos(2\pi\alpha) & \sin(2\pi\alpha) & 0\end{bmatrix}^\top$. As this parametrisation moves through the circle with constant speed reflecting the fact that the orbit is "equally distributed" over the circle in the first two components, we can now express the mean payoff of the orbit with respect to the weight function w as

$$MP_w(M, q) = \int_0^1 w\left(\begin{bmatrix}\cos(2\pi\alpha) & \sin(2\pi\alpha) & 0\end{bmatrix}^\top\right)\, \mathrm{d}\alpha.$$

In the sequel, we work out all the necessary steps to check whether the orbit of an LDS is bounded and to obtain such an expression for the mean payoff as an integral for arbitrary rational LDSs with bounded orbit.

Jordan normal form and boundedness of the orbit

Throughout this section, fix a matrix $M \in \mathbb{Q}^{d\times d}$, an initial vector $q \in \mathbb{Q}^d$, and a continuous weight function $w\colon \mathbb{R}^d \to \mathbb{R}$. We first transform the matrix M into Jordan normal form by computing matrices J and B as well as the inverse B^{-1} with algebraic entries such that

$$M = B \cdot J \cdot B^{-1}$$

where J is in Jordan form with the eigenvalues of M on the diagonal and B is an invertible matrix with generalized eigenvectors of M as columns in polynomial time [7]. Since multiplication with B is a linear bijection, $(M^k \cdot q)_{k\in\mathbb{N}}$ is bounded if and only if the sequence $(J^k \cdot (B^{-1}q))_{k\in\mathbb{N}}$ is bounded. To check whether this is the case, we first simplify the sequence.

We use the notation $J_{\alpha,\ell}$ to denote a Jordan block of size ℓ with α on the diagonal. Observe that multiplying a Jordan block to a vector $q = [q_1, \ldots, q_k, 0, \ldots, 0,]^\top$ in which the last $\ell - k$ components are 0 results in a vector where this is still the case:

$$J_{\alpha,\ell} \cdot q = \begin{bmatrix} \alpha & 1 & 0 & \ldots & 0 \\ 0 & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \alpha & 1 & 0 \\ \vdots & & \ddots & \alpha & 1 \\ 0 & \ldots & \ldots & 0 & \alpha \end{bmatrix} \cdot \begin{bmatrix} q_1 \\ \cdots \\ q_k \\ 0 \\ \cdots \\ 0 \end{bmatrix} = \begin{bmatrix} J_{\alpha,k} \cdot \begin{bmatrix} q_1 \\ \cdots \\ q_k \end{bmatrix} \\ 0 \\ \cdots \\ 0 \end{bmatrix}$$

Looking at the initial vector $B^{-1}q$, this allows us to simplify the LDS by determining the coordinates at which the orbit $(J^k B^{-1} q)_{k\in\mathbb{N}}$ always stays 0. Suppose the Jordan blocks in J end at coordinates $i_1, \ldots, i_m$, respectively, with $1 \le i_1 < i_2 < i_m = d$. Now, let

$$I = \{i \in \{1, \ldots, d\} \mid \text{for some index } h, \text{ all } j \text{ with } i \le j \le i_h \text{ satisfy } (B^{-1}q)_j = 0\}.$$

So, I contains only dimensions j such that $(J^k(B^{-1}q))_j = 0$ for all k. We now set all columns and rows of J with an index in I to 0. This does not affect the orbit $(BJ^kB^{-1}q)_{k\in\mathbb{N}}$. After this simplification, the following condition, which we can assume w.l.o.g., is satisfied.

▶ **Assumption 1.** The LDS given by $M \in \mathbb{Q}^{d\times d}$ and $q \in \mathbb{Q}^d$ has the following property: For the Jordan normal form $M = B \cdot J \cdot B^{-1}$ of M and $v \stackrel{\text{def}}{=} B^{-1}q$, we have that $v_i \neq 0$ for any coordinate $1 \leq i \leq d$ at which a non-zero Jordan block of J ends.

▶ **Proposition 9.** *Under Assumption 1, the orbit $(J^kq)_{k\in\mathbb{N}}$ is bounded if and only if all eigenvalues on the diagonal of J have modulus at most 1 and the Jordan blocks in J with an eigenvalue α with $|\alpha| = 1$ have size 1.*

We delegate the proof to the appendix. Proposition 9 allows us to decide whether the orbit of the LDS given by M and v is bounded. From now on, we assume that it is bounded. We now further simplify the LDS by removing all eigenvalues with modulus less than 1: For a Jordan block $J_{\alpha,\ell}$ with $|\alpha| < 1$, we know $J_{\alpha,\ell} \to 0$ for $k \to \infty$. As we apply the function B viewed as a linear map and the *continuous* function w to the points in the orbit and as the mean payoff does not depend on a prefix of the orbit, we can set all such Jordan blocks to 0 without affecting the mean payoff. So, w.l.o.g. we can work under the following assumption after this simplification because the Jordan blocks with eigenvalues with modulus 1 have size 1 in the light of Proposition 9:

▶ **Assumption 2.** The matrix M of the rational LDS (M, q) is diagonalisable and all non-zero eigenvalues have modulus 1. So, there is a computable algebraic matrix B with computable inverse B^{-1} and a computable algebraic diagonal matrix D whose entries all have modulus 1 or 0 with $M = B \cdot D \cdot B^{-1}$.

Multiplicative relations between the eigenvalues

Before we can parametrise the set of accumulation points of the orbit, we have to detect *multiplicative relations* between the elements on the diagonal of D. Before defining (the group of) multiplicative relations, let us illustrate this concept in an example:

▶ **Example 10.** Consider the matrix $D = \begin{bmatrix} \lambda & 0 \\ 0 & \bar{\lambda} \end{bmatrix}$ for an algebraic number λ with $|\lambda| = 1$ that is not a root of unity. Then, $\lambda \cdot \bar{\lambda} = 1$ is a multiplicative relation between λ and $\bar{\lambda}$. Further, $(\lambda^k)_{k\in\mathbb{N}}$ is dense in the torus $\mathbb{T} \coloneqq \{x \in \mathbb{C} \mid |x| = 1\}$. Now, the sequence $\left(\lambda^k, \bar{\lambda}^k\right)_{k\in\mathbb{N}}$ is dense in $L \coloneqq \{(x, y) \in \mathbb{T}^2 \mid x \cdot y = 1\}$, but not in $\mathbb{T}^2$. So, for an initial vector v, the set of accumulation points of $(D^kv)_{k\in\mathbb{N}}$ is $L \cdot v$ and not $\mathbb{T}^2 \cdot v$.

We follow an approach also taken in [15] to detect multiplicative relations between the algebraic numbers $\lambda_1, \ldots, \lambda_d \in \overline{\mathbb{Q}}$. We work under Assumption 2 and we first reorder the coordinates such that the entries on the diagonal of D are $\lambda_1, \ldots, \lambda_\ell, \lambda_{\ell+1}, \ldots, \lambda_d$ where λ_i is not 0 or 1 for $i \leq \ell$ and the entries λ_j with $j > \ell$ are all equal to 0 or 1. The group

$$G \coloneqq G(\lambda_1, \ldots, \lambda_\ell) = \{(m_1, \ldots, m_\ell) \in \mathbb{Z}^\ell \ \lambda_1^{m_1} \cdots \lambda_\ell^{m_\ell} = 1\}$$

is called the *group of multiplicative relations* between $\lambda_1, \ldots, \lambda_\ell$. If this group is consists only of the neutral element, we say that $\lambda_1, \ldots, \lambda_\ell$ are *multiplicatively independent*.

Note that G is a free abelian group, and has a basis of at most ℓ elements from $\mathbb{Z}^\ell$. By a deep result of Masser [17], G has a basis B such that for each $v \in B$, $||v||_\infty < p(||\lambda_1|| + \ldots + ||\lambda_\ell||)^\ell$, where p is an absolute polynomial. Hence a basis of G can be computed in polynomial space (given $\lambda_1, \ldots, \lambda_\ell$) by simply enumerating all possible bases satisfying Masser's bound. As described in detail in [15], each element $(b_1, \dots, b_\ell) \in B$ of the basis allows us to express one of the eigenvalues in terms of the others: Suppose $b_j \neq 0$. Then, the equation $\lambda_1^{b_1} \cdots \lambda_\ell^{b_\ell} = 1$, allows us to conclude

$$\lambda_j^{b_j} = \prod_{i \neq j} \lambda_i^{-b_i} \qquad \text{and hence} \qquad \lambda_j = \rho_j \prod_{i \neq j} \lambda_i^{-b_i/b_j}$$

where ρ_j is a b_jth root of unity. Applying this procedure consecutively to all elements of the basis B, we can divide and reorder the eigenvalues $\lambda_1, \dots, \lambda_\ell$ as $\lambda_1, \dots, \lambda_m, \lambda_{m+1}, \dots, \lambda_\ell$ such that $\lambda_1, \dots, \lambda_m$ are multiplicatively independent and such that each λ_j with $m+1 \leq j \leq \ell$ is not 1 and can be written as

$$\lambda_j = \rho_j \cdot \prod_{i=1}^{m} \lambda_i^{q_{j,i}}$$

where ρ_j is a root of unity and $q_{j,i} \in \mathbb{Q}$ for $1 \leq i \leq m$.

Subsequences without periodicity

The fact that expression for the eigenvalues λ_j with $m+1 \leq j \leq \ell$ contains the b_jth root of unity ρ_j introduces a periodic behavior to the sequence $(\lambda_j^k)_{k \in \mathbb{N}}$. In order to eliminate this periodic behavior, we divide the orbit into subsequences as follows: We let P be the least common multiple of the values b_j for $m+1 \leq j \leq \ell$. As ρ_j is a b_jth root of unity, $\rho_j^P = 1$ for all j with $m+1 \leq j \leq \ell$. We now split the sequence $(D^k)_{k \in \mathbb{N}}$ into the P subsequences of the form $(D^{Pk+r})_{k \in \mathbb{N}}$ for $r \in \{0, \dots, P-1\}$. The diagonal entries of D^{kP} are

$$\lambda_1^{Pk}, \dots, \lambda_m^{Pk}, \prod_{i=1}^{m} (\lambda_i^{Pk})^{q_{m+1,i}}, \dots, \prod_{i=1}^{m} (\lambda_i^{Pk})^{q_{\ell,i}}, \lambda_{\ell+1}, \dots, \lambda_d.$$

Recall here that $\lambda_{\ell+1}, \dots, \lambda_d$ are all 0 or 1.

We can now express any point in the orbit $BD^{Pk+r}B^{-1}q$ in terms of $\lambda_1^k, \dots, \lambda_m^k$ and D^r. To this end, we define the map

$$\begin{aligned} T_r \colon \mathbb{T}^m &\to \mathbb{R}^d \\ (\mu_1, \dots, \mu_m) &\mapsto BD^r \operatorname{diag}\left(\mu_1^P, \dots, \mu_m^P, \prod_{i=1}^{m} (\mu_i^P)^{q_{m+1,i}}, \dots, \prod_{i=1}^{m} (\mu_i^P)^{q_{m+1,i}}, \lambda_{\ell+1}, \dots, \lambda_d\right) B^{-1}q \end{aligned}$$

where $\operatorname{diag}(x_1, \dots, x_d)$ denotes a diagonal matrix with entries $x_1, \dots, x_d$ on the diagonal. The map T is chosen such that

$$T_r(\lambda_1^k, \dots, \lambda_m^k) = BD^{Pk+r}B^{-1}q.$$

This is also the reason why T_r maps into $\mathbb{R}^d$.

Parametrising the set of accumulation points

For a real x, we define $x \mod 1 := x - \lfloor x \rfloor$. For $1 \leq j \leq m$, we define the number $\alpha_j \in [0,1)$ as the unique number with $\lambda_j = e^{2\pi i \alpha_j}$. Let $S \colon [0,1)^m \to \mathbb{T}^m$ (recall that $\mathbb{T} := \{x \in \mathbb{C} \mid |x| = 1\}$) be the map

$$(\beta_1, \dots, \beta_m) \mapsto (e^{2\pi i \beta_1}, \dots, e^{2\pi i \beta_m}).$$

So, we get $(\lambda_1^k, \dots, \lambda_m^k) = S(k\alpha_1 \mod 1, \dots, k\alpha_m \mod 1)$ and hence

$$BD^{Pk+r}B^{-1}q = T_r(S(k\alpha_1 \mod 1, \dots, k\alpha_m \mod 1)).$$

Following the exposition in [15], we can now apply an equidistribution theorem by Weyl [24]. First, observe that the fact that $\lambda_1, \dots, \lambda_m$ are multiplicatively independent means that the values $1, \alpha_1, \dots, \alpha_m$ are linearly independent over $\mathbb{Q}$: If there were a non-zero vector $c_0, c_1, \dots, c_m$ with $c_0 + \sum_{j=1}^{m} c_j \alpha_j = 0$, this vector would witnesses a multiplicative relation between $\lambda_1, \dots, \lambda_m$. In [24], it is now shown that for any measurable set $U \subseteq [0,1)^m$, we have

$$\lim_{n\to\infty} \frac{|\{0 \le k \le n \mid (k\alpha_1 \mod 1, \dots, k\alpha_m \mod 1) \in U\}|}{n+1} = \mathcal{L}(U) \qquad (*)$$

where $\mathcal{L}$ is the Lebesgue measure. For more details, we also refer to the exposition of this argument in [15].

This means that the sequence of arguments $((k\alpha_1 \mod 1, \dots, k\alpha_m \mod 1))_{k\in\mathbb{N}}$ is dense and "equally distributed" in the cube $[0,1)^m$, and hence the sequence $((\lambda_1^k, \dots, \lambda_m^k))_{k\in\mathbb{N}}$ is dense and "equally distributed" in the m-dimensional torus $\mathbb{T}^m$ where "equally distributed" means that every subset of the same size is hit equally often in the sense of Equation $(*)$.

Mean payoff as integral

Now, we are in the position to prove the main result of this subsection: The mean payoff of a bounded orbit wrt a continuous weight function can be expressed as an integral.

▶ **Theorem 11.** *Let $M \in \mathbb{Q}^{d\times d}$ be a matrix and $q \in \mathbb{Q}^d$ an initial vector satisfying Assumption 2. Let $w\colon \mathbb{R}^d \to \mathbb{R}$ be a continuous weight function. Let $P \in \mathbb{N}$ and $T_r\colon \mathbb{T}^m \to \mathbb{R}^d$ for $r < P$, and $S\colon [0,1)^m \to \mathbb{T}^m$ be as above. Then, for each r with $0 \le r < P$, the mean payoff of the sub-orbit $(M^{kP+r}q)_{k\in\mathbb{N}}$ wrt w exists and can be expressed as*

$$MP_w(M^P, M^r q) = \lim_{k\to\infty} \frac{1}{k+1} \sum_{i=0}^{k} w(M^{kP+r}q) = \int_{[0,1)^m} w \circ T_r \circ S \, \mathrm{d}\mathcal{L}$$

where $\mathcal{L}$ is the Lebesgue measure on $[0,1)^m$. The mean payoff of the original orbit is then the arithmetic mean

$$MP_w(M,q) = \frac{\sum_{r=0}^{P-1} MP_w(M^P, M^r q)}{P}.$$

Proof. Let $\alpha_1, \dots, \alpha_m \in [0,1)$ be such that $\lambda_j = e^{2\pi i \alpha_j}$ as above. For $r < P$, we have constructed S and T_r such that

$$M^{kP+r}q = T_r(S(k\alpha_1 \mod 1, \dots, k\alpha_m \mod 1))$$

for all k. As w is continuous, it can be written as sum of Lebesgue measurable step functions $w = \sum_{j=0}^{\infty} f_j \cdot \mathbb{1}_{A_j}$ where, for all j, the coefficient f_j is in $\mathbb{R}$, the set $A_j \subseteq \mathbb{R}^d$ is measurable, and $\mathbb{1}_{A_j}$ is 1 on points in A_j and 0 otherwise. For $\mathbb{1}_{A_j}$, we now observe

$$\begin{aligned}\lim_{k\to\infty} \frac{1}{k+1} \sum_{i=0}^{k} \mathbb{1}_{A_j}(M^{kP+r}q) &= \lim_{k\to\infty} \frac{1}{k+1} \sum_{i=0}^{k} \mathbb{1}_{A_j}(T_r(S(k\alpha_1 \mod 1, \dots, k\alpha_m \mod 1))) \\ &= \lim_{k\to\infty} \frac{|\{i \le k \mid T_r(S(i\alpha_1 \mod 1, \dots, i\alpha_m \mod 1)) \in A_j\}|}{k+1} = \mathcal{L}((T_r \circ S)^{-1}(A_j))\end{aligned}$$

where the last equality follows from equation $(*)$ that is stated above and shown in [24]. But, we also have

$$\int_{[0,1)^m} \mathbb{1}_{A_j} \circ T_r \circ S \, \mathrm{d}\mathcal{L} = \mathcal{L}((T_r \circ S)^{-1}(A_j)).$$

Putting this together, we obtain

$$MP_w(M^P, M^r q) = \lim_{k\to\infty} \frac{1}{k+1} \sum_{i=0}^{k} w(M^{kP+r}q) = \sum_{j=0}^{\infty} f_j \cdot \lim_{k\to\infty} \frac{1}{k+1} \sum_{i=0}^{k} \mathbb{1}_{A_j}(M^{kP+r}q)$$

$$= \sum_{j=0}^{\infty} f_j \cdot \int_{[0,1)^m} \mathbb{1}_{A_j} \circ T_r \circ S \,\mathrm{d}\mathcal{L} = \int_{[0,1)^m} w \circ T_r \circ S \,\mathrm{d}\mathcal{L}.$$

This finishes the proof of the first claim. The claim that the mean payoff $MP_w(M, q)$ can now be expressed as the arithmetic mean is obvious. ◀

3.3 Stochastic LDSs

Stochastic LDSs are a special case of LDSs with a bounded orbit. In this section, we will show that in the case of stochastic LDSs, we can compute the mean payoff of the orbit under a continuous weight function by evaluating the weight function on finitely many points. In the aperiodic case, the orbit even converges to a single point so that it suffices to evaluate the weight function once:

▶ **Lemma 12.** *Let $P \in \mathbb{Q}^{d\times d}$ be a stochastic, aperiodic matrix and $\iota_{init} \in \mathbb{Q}^d$ an initial distribution. Furthermore, let $w\colon \mathbb{R}^d \to \mathbb{R}$ be a continuous weight function. Then,*

$$MP_w(P, \iota_{init}) = w(\pi)$$

where π is a stationary distribution of P computable in polynomial time.

Proof. As described in Section 2.4, we know that the orbit $(P^k \iota_{init})_{k\in\mathbb{N}}$ converges to a stationary distribution π in this case, which can be computed in polynomial time [16, 5]. So, $\lim_{k\to\infty} P^k \iota_{init}$ exists and, as w is continuous, we know $\lim_{k\to\infty} w(P^k \iota_{init}) = w(\pi)$. It is straightforward to observe that

$$MP_w(P, \iota_{init}) \stackrel{\text{def}}{=} \lim_{k\to\infty} \frac{1}{k+1} \sum_{i=0}^{k} w(P^i \iota_{init}) = w(\lim_{k\to\infty} P^k \iota_{init}) = w(\pi).$$ ◀

Hence the computation of the mean payoff boils down to evaluating the function w once on a rational point computable in polynomial time in this case. We next address the periodic case by splitting up the orbit into subsequences.

For an irreducible and periodic Markov chain with period L, we have that P^L is aperiodic and $L \le d$ by [18, Theorem 1.8.4]. Together with Lemma 12, this allows us to compute $MP_w(P^L, P^r \iota_{init})$, which is the mean payoff of $(P^{Lk+r} \iota_{init})_{k\in\mathbb{N}}$. We conclude

$$MP_w(P, \iota_{init}) = \frac{1}{L} \sum_{r=0}^{L-1} MP_w(P, P^r \iota_{init}).$$

So, for irreducible stochastic LDSs, we can divide $(P^{Lk+r} \iota_{init})_{k\in\mathbb{N}}$ into L equally spaced subsequences and compute the mean payoff $MP_w(P, \iota_{init})$ as the arithemetic mean of the mean payoffs of these subsequences.

▶ **Theorem 13.** *Let $P \in \mathbb{Q}^{d\times d}$ be a stochastic, irreducible matrix and $\iota_{init} \in \mathbb{Q}^d$ an initial distribution. Let $w\colon \mathbb{R}^d \to \mathbb{R}$ be a continuous weight function. Then, we can compute points $\pi_0, \ldots, \pi_{L-1} \in \mathbb{Q}^d$ in polynomial time for some $L \le d$ such that $MP_w(P, \iota_{init}) = \frac{1}{L} \sum_{i=0}^{L-1} w(\pi_i)$.*

If the weight function w can be evaluated in polynomial time on rational inputs, Theorem 13 implies that the mean payoff $MP_w(P, \iota_{init})$ can be computed in polynomial time.

When a Markov chain is reducible, the states can be renamed in a way such that, the matrix representation of the Markov chain will contain distinct blocks corresponding to the bottom strongly connected components (BSCCs) on the diagonal along with additional columns at the right representing states that do not belong to any BSCC:

$$\begin{bmatrix} \square & 0...0 & 0...0 & * & * \\ 0...0 & \square & 0...0 & * & * \\ 0...0 & 0...0 & \square & * & * \\ 0...0 & 0...0 & 0....0 & * & * \end{bmatrix}$$

Each block representing a BSCC constitutes an irreducible Markov chain. Assume we have k blocks with periods $L_1, L_2, ..., L_k$ correspondingly. Let l be the least common multiple of the periods. Now we will have l subsequences of the orbit each of which will converge. The convergence of the rows in the bottom is a result of the fact that Markov chain will enter a BSCC with probability 1. So, in general, we have l subsequences of the orbit, all of which converge. We observe that $l \leq d^d$, from which the following result follows:

▶ **Theorem 14.** *Let $P \in \mathbb{Q}^{d\times d}$ be a stochastic matrix and $\iota_{init} \in \mathbb{Q}^d$ an initial distribution. Let $w \colon \mathbb{R}^n \to \mathbb{R}$ be a continuous weight function. Then, we can compute points $\pi_0, \dots, \pi_{l-1} \in \mathbb{Q}^d$ in exponential time for some $l \leq d^d$ such that $MP_w(P, \iota_{init}) = \frac{1}{L}\sum_{i=0}^{L-1} w(\pi_i)$.*

4 Total (discounted) reward and satisfaction of energy constraints

In this section, again let $M \in \mathbb{Q}^d$ be a matrix, $q \in \mathbb{Q}^d$ be an initial vector, and $w \colon \mathbb{R}^d \to \mathbb{R}$ be a polynomial weight function with rational coefficients. We define the *total reward* as

$$\mathrm{tr}(M, q, w) := \sum_{k=0}^{\infty} w(M^k q).$$

Likewise, for a rational discount factor $\delta \in (0, 1)$ we define the *total discounted reward* as

$$\mathrm{dr}(M, q, w, \delta) := \sum_{k=0}^{\infty} \delta^k w(M^k q).$$

Both of these quantities, when they exist, can be determined effectively.

▶ **Theorem 15.** *It is decidable whether the series $\sum_{k=0}^{\infty} w(M^k q)$ and $\sum_{k=0}^{\infty} \delta^k w(M^k q)$ converge, in which case their value is rational and can be computed.*

Proof. Let $u_n = \sum_{k=0}^{n} w(M^k q)$. As discussed in subsection 3.1, $(u_n)_{n\in\mathbb{N}}$ is a rational LRS, and we can apply Corollary 6. Similarly, let $v_n = \sum_{k=0}^{\infty} \delta^k w(M^k q)$. As $(\delta^n)_{n\in\mathbb{N}}$ is itself a (rational) LRS and such LRS are closed under pointwise multiplication, v_n is also a rational LRS. We again apply Corollary 6. ◀

We next discuss *energy constraints*. We say that a series of real weights $(w_i)_{i\in\mathbb{N}}$ satisfies the energy constraint with budget B if

$$\sum_{i=0}^{k} w_i \geq -B$$

for all $k \in \mathbb{N}$. We will prove that for LDS (M, q) of dimension at most 3, satisfaction of energy constraints is decidable. The proof is based on the fact that three-dimensional systems are tractable thanks to Baker's theorem [13]. For higher-dimensional systems, no such tractability result is known. We will show that deciding satisfaction of energy constraints is, in general, at least as hard as the Positivity Problem, already with linear weight functions.

4.1 Baker's theorem and its applications

A *linear form in logarithms* is an expression of the form $\Lambda = b_1 \operatorname{Log} \alpha_1 + \ldots + b_m \operatorname{Log} \alpha_m$ where $b_i \in \mathbb{Z}$ and $\alpha_i \in \overline{\mathbb{Q}}$ for all $1 \leq i \leq m$. Here Log denotes the principal branch of the complex logarithm. The celebrated theorem of Baker places a lower bound on $|\Lambda|$ in case $\Lambda \neq 0$. Baker's theorem, as well as its p-adic analogue, play a critical role in the proof of [21] that the Skolem Problem is decidable for LRS of order at most 4, as well as decidability of the Positivity Problem for low-order LRS.

▶ **Theorem 16** (Special case of the main theorem in [25]). *Let $\Lambda = b_1 \operatorname{Log} \alpha_1 + \ldots + b_m \operatorname{Log} \alpha_m$ be as above, $D = [\mathbb{Q}(\alpha_1, \ldots, \alpha_m) : \mathbb{Q}]$, and suppose $A, B \geq e$ are such that $A > H(\alpha_i)$ and $B > |b_i|$ for all $1 \leq i \leq m$. If $\Lambda \neq 0$, then*

$$\log |\Lambda| > -(16mD)^{2(m+2)} (\log A)^m \log B.$$

A direct consequence of Baker's theorem is the following [20, Corollary 8]. Recall that $\mathbb{T}$ denotes $\{z \in \mathbb{C} : |z| = 1\}$.

▶ **Lemma 17.** *Let $\alpha \in \mathbb{T} \cap \overline{\mathbb{Q}}$ and $\beta \in \overline{\mathbb{Q}}$. For all $n \geq 2$, if $\alpha^n \neq \beta$ then $|\alpha^n - \beta| > n^{-C}$ where C is an effective constant that depends on α and β.*

If α is not a root of unity, $\alpha^n = \beta$ holds for at most one n and n can be effectively bounded.

▶ **Lemma 18.** *Let $\alpha, \beta \in \overline{\mathbb{Q}}$ be non-zero, and suppose α is not a root of unity. There exists effectively computable $N \in \mathbb{N}$ such that $\alpha^n \neq \beta$ for all $n \in \mathbb{N}$ with $n > N$.*

Combining the two lemmas above, we obtain the following.

▶ **Theorem 19.** *Let $\alpha \in \mathbb{T}$, $\beta \in \mathbb{Q}$, and suppose α is not a root of unity. There exists effectively computable $N, C \in \mathbb{N}$ such that for $n > N$, $|\alpha - \beta| > n^{-C}$.*

The next lemma summarises the family of LRS to which we can apply Baker's theorem. For reasons of space we delegate the proof to the appendix.

▶ **Lemma 20.** *Let $\gamma \in \mathbb{T}$ be not a root of unity, $r_1, \ldots, r_\ell \in \mathbb{R}$ be non-zero, and*

$$u_n = \sum_{i=1}^{m} c_i \Lambda_i^n$$

be an LRS over $\mathbb{R}$ where the right-hand side is in the exponential-polynomial form, $c_i, \Lambda_i \in \overline{\mathbb{Q}}$ for all i, and each Λ_i is in the multiplicative group generated by $\{\gamma, r_1, \ldots, r_\ell\}$. Suppose $m > 0$, i.e. $(u_n)_{n \in \mathbb{N}}$ is not identically zero.

(a) There exists effectively computable N_1 such that $u_n \neq 0$ for all $n > N_1$.

(b) For $n > N_1$, $|u_n| > L^n n^{-C}$, where $L = \max_i |\Lambda_i|$ and C is an effectively computable constant.

(c) It is decidable whether $u_n \geq 0$ for all n.

4.2 Satisfaction of energy constraints

Before giving our decidability result, we need one final ingredient about partial sums of LRS. Let $w_n = n^k\lambda^n$ and $u_n = \sum_{k=0}^{n} w_k$. If $\lambda = 1$, then $u_n = p(n)$, where p is a polynomial of degree $k+1$. If $\lambda \neq 1$, then $u_n = q(n)\lambda^n$, where $q(n)$ is a polynomial of degree at most k. To see this, observe that $q(n)$ can be chosen as the solution of $\lambda q(n) - q(n-1) = n^k$. It follows that if the LRS $(w_n)_{n\in\mathbb{N}}$ has only real eigenvalues, then so does the sequence given by $u_n = \sum_{k=0}^{n} w_k$. Similarly, if $(w_n)_{n\in\mathbb{N}}$ is diagonalisable and does not have 1 as an eigenvalue, then the same applies to $(u_n)_{n\in\mathbb{N}}$. In fact, the eigenvalues of $(u_n)_{n\in\mathbb{N}}$ form a subset of the eigenvalues of $(w_n)_{n\in\mathbb{N}}$.

▶ **Theorem 21.** *Let $M \in \mathbb{Q}^{3\times 3}$, $q \in \mathbb{Q}^3$, $\delta \leq 1$, and $w : \mathbb{R}^3 \to \mathbb{R}$ be a polynomial weight function with rational coefficients. For $B \in \mathbb{Q}_{\geq 0}$, it is decidable whether the weights $(w(M^n q))_{n\in\mathbb{N}}$ satisfy the energy constraint with budget B.*

Proof. Let $w_n = \delta^n w(M^n q)$ and $u_n = B + \sum_{i=0}^{n} w(M^i q)$. We have to decide if $u_n \geq 0$ for all n. First suppose M has only real eigenvalues. Then w_n and u_n are both LRSs with only real eigenvalues. By taking subsequences if necessary, we can assume $(u_n)_{n\in\mathbb{N}}$ is non-degenerate. We can write $u_n = \sum_{i=1}^{m} p_i(n)\rho_i^n$ where the right-hand side is in the exponential-polynomial form. In particular, for all i, p_i is not the zero polynomial. Since $(u_n)_{n\in\mathbb{N}}$ is non-degenerate, wlog we can assume $\rho_1 > \ldots > \rho_m > 0$. If $p_1(n)$ is negative for sufficiently large n, then the energy constraint is not satisfied. Otherwise, we can compute N such that for all $n > N$, $u_n > 0$. It remains to check whether $u_n \geq 0$ for $0 \leq n \leq N$. Next, suppose M has non-real eigenvalues $\lambda, \overline{\lambda}$, and a real eigenvalue ρ. Write $\gamma = \lambda/|\lambda|$ and $r = |\lambda|$. Then u_n is of the form

$$u_n = cn + \sum_{i=1}^{m} c_i \Lambda_i^n \coloneqq cn + v_n$$

where $\Lambda_1, \ldots, \Lambda_m$ are pairwise distinct and in the multiplicative group generated by r, ρ, δ, γ. Wlog we can assume $c_i \neq 0$ for all i, but c may be zero. If γ is a root of unity of order $k > 0$ (i.e. $\gamma^k = 1$), then we can take subsequences $(u_n^{(0)})_{n\in\mathbb{N}}, \ldots, (u_n^{(k-1)})_{n\in\mathbb{N}}$, where $u_n^{(j)} = u_{nk+j}$ for $n \in \mathbb{N}$ and $0 \leq j < k$, and each $(u_n^{(j)})_{n\in\mathbb{N}}$ has only real eigenvalues. We can then apply the analysis above. Hereafter we assume γ is not a root of unity.

Suppose $c = 0$. Then Lemma 20 (c) applies and we can decide if u_n is positive. Next, suppose $c \neq 0$ and $L \leq 1$. We can compute N_2 such that $|cn| > |v_n|$ for all $n > N_2$. Hence in this case $u_n \geq 0$ for all n iff $c > 0$ and $u_n > 0$ for $0 \leq n \leq N_2$. Finally, suppose $c \neq 0$ and $L \geq 1$. Applying Lemma 20 (b), there exists N_3 such that $|u_n| > |cn|$ for $n > N_3$. Hence $u_n \geq 0$ for all n iff $u_n \geq 0$ for $0 \leq n \leq N_3$ and the sequence $v_n = u_{n+N_3}$ is positive. The latter can be decided by observing that $v_n = \sum_{i=1}^{m}(c_i\Lambda_i^{N_3})\Lambda_i^n$ and applying Lemma 20 (c). ◀

4.3 Positivity-hardness

Recall that the energy satisfaction problem is to decide, given $M \in \mathbb{Q}^{d\times d}$, $q \in \mathbb{Q}^d$, $B \in \mathbb{Q}$, and a polynomial p, whether there exists n such that $\sum_{k=0}^{n} p(M^k q) < B$. This problem is Positivity-hard already for LDS that are Markov chains; see the appendix for the proof.

▶ **Lemma 22.** *The Positivity Problem can be reduced to the energy satisfaction problem above restricted to a Markov chain (M, q) and a linear weight function w.*

5 Conclusion

We have shown how to compute the mean-payoff for arbitrary LDS equipped with a polynomial weight function and how to find an integral expression for the mean payoff in bounded LDS with a continuous weight function. In the special case of stochastic LDSs, which always have a bounded orbit, we could go further and compute finitely many points such that the mean payoff of the orbit is the arithmetic mean of the weight function evaluated at these points. For energy constraints, we showed decidability for three-dimensional systems by utilising the results about LRS based on Baker's theorem.

Instead of continuous weight functions, also functions w assigning a weight to each semialgebraic set in a collection of semialgebraic sets $\mathcal{T}_1, \ldots, \mathcal{T}_m$ constitute an interesting class of weight functions. Here, several interesting questions can be asked. E.g., given an LDS $(M, q) \in \mathbb{Q}^{d\times d} \times \mathbb{Q}^d$ and w as above, compare the (discounted) total reward/mean-payoff to a given threshold. Here at time n the reward received is $\sum_{i=1}^{m} \mathbb{1}(M^n q \in T_i) w(T_i)$. This problem appears to be difficult with deep connections to Diophantine approximation.

References

1 S. Akshay, Timos Antonopoulos, Joël Ouaknine, and James Worrell. Reachability problems for markov chains. *Inf. Process. Lett.*, 115(2):155–158, 2015. doi:10.1016/j.ipl.2014.08.013.

2 S. Akshay, Hugo Bazille, Blaise Genest, and Mihir Vahanwala. On robustness for the skolem and positivity problems. In Petra Berenbrink and Benjamin Monmege, editors, *39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15-18, 2022, Marseille, France (Virtual Conference)*, volume 219 of *LIPIcs*, pages 5:1–5:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.STACS.2022.5.

3 S. Almagor, T. Karimov, E. Kelmendi, J. Ouaknine, and J. Worrell. Deciding ω-regular properties on linear recurrence sequences. *Proc. ACM Program. Lang.*, 5(POPL), 2021.

4 Christel Baier, Florian Funke, Simon Jantsch, Toghrul Karimov, Engel Lefaucheux, Florian Luca, Joël Ouaknine, David Purser, Markus A. Whiteland, and James Worrell. The orbit problem for parametric linear dynamical systems. In Serge Haddad and Daniele Varacca, editors, *32nd International Conference on Concurrency Theory, CONCUR 2021, August 24-27, 2021, Virtual Conference*, volume 203 of *LIPIcs*, pages 28:1–28:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.CONCUR.2021.28.

5 Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.

6 Mark Braverman. Termination of integer linear programs. In *Proceedings of the 18th International Conference on Computer Aided Verification*, CAV'06, page 372–385, Berlin, Heidelberg, 2006. Springer-Verlag. doi:10.1007/11817963_34.

7 Jin-Yi Cai. Computing jordan normal forms exactly for commuting matrices in polynomial time. *International Journal of Foundations of Computer Science*, 05(03n04):293–302, 1994. arXiv:https://doi.org/10.1142/S0129054194000165, doi:10.1142/S0129054194000165.

8 Julian D'Costa, Toghrul Karimov, Rupak Majumdar, Joël Ouaknine, Mahmoud Salamati, and James Worrell. The pseudo-reachability problem for diagonalisable linear dynamical systems. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022, August 22-26, 2022, Vienna, Austria*, volume 241 of *LIPIcs*, pages 40:1–40:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.MFCS.2022.40.

9 Graham Everest, Alf van der Poorten, Igor Shparlinski, and Thomas Ward. *Recurrence Sequences*. 2003. URL: http://dx.doi.org/10.1090/surv/104, doi:10.1090/surv/104.

10 Toghrul Karimov, Edon Kelmendi, Joris Nieuwveld, Joël Ouaknine, and James Worrell. The power of positivity. *To appear in the Proceedings of LICS23*, 2023.

11 Toghrul Karimov, Edon Kelmendi, Joël Ouaknine, and James Worrell. What's decidable about discrete linear dynamical systems? In Jean-François Raskin, Krishnendu Chatterjee, Laurent

Doyen, and Rupak Majumdar, editors, *Principles of Systems Design - Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday*, volume 13660 of *Lecture Notes in Computer Science*, pages 21–38. Springer, 2022. doi:10.1007/978-3-031-22337-2_2.

12 Toghrul Karimov, Engel Lefaucheux, Joël Ouaknine, David Purser, Anton Varonka, Markus A. Whiteland, and James Worrell. What's decidable about linear loops? *Proc. ACM Program. Lang.*, 6(POPL), 2022.

13 Toghrul Karimov, Joel Ouaknine, and James Worrel. On ltl model checking for low-dimensional discrete linear dynamical systems. In *45th International Symposium on Mathematical Foundations of Computer Science, MFCS 2020, LIPIcs 170*, 2020.

14 Manuel Kauers and Peter Paule. *The Concrete Tetrahedron*. Springer Vienna, 2011. URL: https://doi.org/10.1007%2F978-3-7091-0445-3, doi:10.1007/978-3-7091-0445-3.

15 Edon Kelmendi. Computing the density of the positivity set for linear recurrence sequences. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12, 2022.

16 Vidyadhar G. Kulkarni. *Modeling and Analysis of Stochastic Systems*. Chapman & Hall, Ltd., 1995.

17 David W Masser. Linear relations on algebraic groups. *New Advances in Transcendence Theory*, pages 248–262, 1988.

18 J. R. Norris. *Markov Chains*. Cambridge University Press, 2009.

19 Joël Ouaknine and James Worrell. Positivity problems for low-order linear recurrence sequences. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, 12 2013. doi:10.1137/1.9781611973402.27.

20 Joël Ouaknine and James Worrell. On the positivity problem for simple linear recurrence sequences,. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming*, pages 318–329, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.

21 R. Tijdeman, M. Mignotte, and T.N. Shorey. The distance between terms of an algebraic recurrence sequence. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1984(349):63–76, 1984. doi:10.1515/crll.1984.349.63.

22 Mihir Vahanwala. Skolem and positivity completeness of ergodic markov chains. *CoRR*, abs/2305.04881, 2023. arXiv:2305.04881, doi:10.48550/arXiv.2305.04881.

23 N. K. Vereshchagin. The problem of appearance of a zero in a linear recurrence sequence (in Russian). *Mat. Zametki*, 38(2), 1985.

24 Hermann Weyl. Über die gleichverteilung von zahlen mod. eins. *Mathematische Annalen*, 77(3):313–352, 1916.

25 G. Wüstholz and A. Baker. Logarithmic forms and group varieties. *Journal für die reine und angewandte Mathematik*, 442:19–62, 1993. URL: http://eudml.org/doc/153550.

A Omitted proofs

Proof of Lemma 2

▶ **Lemma 2.** *Let* $u_n = \sum_{i=1}^m p_i(n)\lambda_i^n$*, where all* $\lambda_i \in \overline{\mathbb{Q}}$ *and* $p_i \in \overline{\mathbb{Q}}[x]$ *are non-zero, and* $\lambda_i \neq \lambda_j$ *for* $i \neq j$*. The sequence* $(u_n)_{n\in\mathbb{N}}$ *is not identically zero. Specifically, there exists* $0 \le n < d$*, where* $d = \sum_{i=1}^m(\deg(p_i)+1)$*, such that* $u_n \neq 0$*.*

Proof. Suppose $\deg(p_k) \ge 1$ for some $1 \le k \le m$. Consider the sequence $v_n = u_{n+1} - \lambda_k u_n$. It will be of the form

$$v_n = \sum_{i\in I} q_i(n)\lambda_i^n$$

where $I \subseteq \{1,\dots,m\}$ with $k \in I$, $\deg(q_k) < \deg(p_k)$, and for all $i \in I$, q_i is not identically zero with $\deg(q_i) \le \deg(p_i)$. Observe that if $(u_n)_{n\in\mathbb{N}}$ is identically zero, then so is $(v_n)_{n\in\mathbb{N}}$.

Moreover, if v_n is non-zero, then either u_n or u_{n+1} is non-zero. Repeating the process of constructing v_n from u_n at most $\sum_{i=1}^m \deg(p_i)$ times, we obtain

$$w_n = \sum_{i=1}^{m} c_i \lambda_i^n$$

that is identically zero if u_n is identically zero, where each c_i is an algebraic number and $c := [c_1 \ \cdots \ c_m]^\top \neq \mathbf{0}$.

It remains to argue that w_n cannot be identically zero. Consider the system of equations

$$\sum_{i=1}^{m} x_i \lambda_i^n = 0 \qquad \text{for } 0 \leq n < m.$$

We can write it as $Mx = \mathbf{0}$, where $x = [x_1 \ \cdots \ x_m]^\top$ and M is a Vandermonde matrix with $\det(M) = \prod_{i \neq j}(\lambda_i - \lambda_j)$. Since $\lambda_1, \ldots, \lambda_m$ are distinct by assumption, M is invertible and $Mx = 0$ if and only if $x = \mathbf{0}$. Since $c \neq \mathbf{0}$, it follows that $w_n \neq 0$ for some $0 \leq n < m$. Hence there exists $n' \leq n + \sum_{i=1}^m \deg(p_i) = n + (d - m) < d$ such that $u_{n'} \neq 0$. ◀

Proof of Lemma 3

▶ **Lemma 3.** *Let $(u_n)_{n\in\mathbb{N}}$ be as in the statement of Lemma 2. If $u_n \in \mathbb{R}$ for all $n \in \mathbb{N}$, then for every $1 \leq i \leq m$ there exists j such that $p_j(n) = \overline{p_i}(n)$ and $\lambda_j = \overline{\lambda_i}$.*

Proof. If $m = 0$, the statement is (vacuously) true. Suppose $m > 1$, and consider

$$v_n = u_n - \overline{u_n} = \sum_{i=1}^{m} p_i(n)\lambda_i^n - \sum_{j=1}^{m} \overline{p_j}(n)\overline{\lambda_j}^n.$$

Since $v_n = 0$ for all n, and p_i, p_j is non-zero for all i, j, there must be $1 \leq i_1, j_1 \leq m$ such that $\lambda_{i_1} = \overline{\lambda_{j_1}}$. Hence

$$v_n = \underbrace{\sum_{i \neq i_1} p_i(n)\lambda_i^n - \sum_{j \neq j_1} \overline{p_j}(n)\overline{\lambda_j}^n}_{w_n} + (p_{i_1}(n) - \overline{p_{j_1}}(n))\lambda^n$$

where $\lambda = \lambda_{i_1} = \lambda_{j_1}$. Since $\lambda_i \neq \lambda_j$ for $i \neq j$, for all $i \neq i_1$ and $j \neq j_1$ we have $\lambda_i, \lambda_j \neq \lambda$. Hence $p_{i_1}(n) - \overline{p_{j_1}}(n) = 0$. We therefore have $\lambda_{j_1} = \overline{\lambda_{i_1}}$ and $p_{j_1}(n) = \overline{p_{i_1}}(n)$. It remains to observe that w_n is also identically zero and repeat the argument above until for every $1 \leq i \leq m$ a value j with the required property has been found. ◀

Proof of Theorem 5

▶ **Theorem 5.** *Let $(u_n)_{n\in\mathbb{N}}$ be an LRS over $\mathbb{Q}$. It is decidable whether $\lim_{n\to\infty} u_n/n$ exists, in which case the limit is rational and effectively computable.*

Proof. Write $u_n = \sum_{i=1}^m p_i(n)\lambda_i^n$, where the right-hand side is in the exponential-polynomial form, and suppose $|\lambda_1| \geq \ldots \geq |\lambda_m| > 0$. If $m = 0$, then the sequence is identically zero. Suppose $m > 0$. By [10, Theorem 2], for every $\epsilon > 0$, $|u_n| > (|\lambda_1| - \epsilon)^n$ for sufficiently large n. Hence if $|\lambda_1| > 1$, then the limit does not exist. Similarly, if $|\lambda_1| < 1$, then the limit is zero. Suppose therefore $|\lambda_1| = 1$. Let k be the largest integer such that $|\lambda_i| = 1$ for all $1 \leq i \leq k$, and define $v_n = \sum_{i=1}^k p_i(n)\lambda_i^n$. It suffices to consider $\lim_{n\to\infty} v_n/n$ as $\lim_{n\to\infty} \sum_{i=k+1}^m p_i(n)\lambda_i^n = 0$.

Write $v_n = \sum_{i=1}^{l} n^i \sum_{j=1}^{k_i} c_{i,j}\lambda_{i,j}^n$ where $\sum_{j=1}^{k_i} c_{i,j}\lambda_{i,j}^n$ is in the exponential-polynomial form for all i. If $l = 0$, then $\lim_{n\to\infty} v_n/n = 0$. Hence suppose $l \geq 1$. Let

$$w_n = \sum_{j=1}^{k_l} c_{l,j}\lambda_{l,j}^n.$$

By Lemma 2, $(w_n)_{n\in\mathbb{N}}$ is not identically zero. Applying [6, Lemma 4], if $\lambda_{l,j} \neq 1$ for some j then there exist $a, b \in \mathbb{R}$ such that $a < b$, $w_n < a$ for infinitely many n, and $w_n > b$ for infinitely many n. Hence $\lim_{n\to\infty} v_n/n$ can exist only if $k_l = 1$ and $\lambda_{l,1} = 1$. Under this assumption, $\lim_{n\to\infty} v_n/n$ exists and is equal to $c_{l,1}$ if and only if $l = 1$.

Suppose the limit above exists. To see that it must be rational, observe that by construction there exists $1 \leq i \leq k$ such that $\lambda_i = 1$ and $p_i(n)$ is equal to either $c_{l,1}$ or $nc_{l,1}$. Since $(u_n)_{n\in\mathbb{N}}$ takes rational values, $\sigma(u_n) = u_n$ for all $n \in \mathbb{N}$ and σ an automorphism of $\mathbb{C}$. By the uniqueness of the exponential-polynomial representation, $\sigma(c_{l,1}\lambda_{l,1}^n) = c_{l,1}\lambda_{l,1}^n$ for all n and σ, which implies that $c_{l,1}$ is rational. ◀

Proof of Proposition 9

▶ **Proposition 9.** *Under Assumption 1, the orbit $(J^k q)_{k\in\mathbb{N}}$ is bounded if and only if all eigenvalues on the diagonal of J have modulus at most 1 and the Jordan blocks in J with an eigenvalue α with $|\alpha| = 1$ have size 1.*

Proof. First, assume there is an eigenvalue α on the diagonal of J with $|\alpha| > 1$. Let i be a coordinate at which a Jordan block containing α ends. Then, $q_i \neq 0$ and hence $(J^k q)_i = \alpha^k q_i$ is not bounded. Next, assume there is a Jordan block of size $\ell > 1$ with an eigenvalue α with $|\alpha| = 1$. W.l.o.g. assume that this is the first Jordan block. Then, the first ℓ coordinates of $J^k q$ are given by $J_{\alpha,\ell}^k \cdot \begin{bmatrix} q_1 & \dots & q_\ell \end{bmatrix}^\top$ and we have $q_\ell \neq 0$. We compute

$$(J^k q)_{\ell-1} = \alpha^k \cdot q_{\ell-1} + k \cdot \alpha^{k-1} \cdot q_\ell,$$

which diverges for $k \to \infty$.

For the other direction, observe that $J_{\alpha,\ell}^k$ tends to 0 if $|\alpha| < 1$. Further, the powers of Jordan blocks of size 1 with an eigenvalue α with $|\alpha| = 1$ are bounded as they simply contain α^k, which has modulus 1. ◀

Proof of Lemma 20

▶ **Lemma 20.** *Let $\gamma \in \mathbb{T}$ be not a root of unity, $r_1, \dots, r_\ell \in \mathbb{R}$ be non-zero, and*

$$u_n = \sum_{i=1}^{m} c_i \Lambda_i^n$$

be an LRS over $\mathbb{R}$ where the right-hand side is in the exponential-polynomial form, $c_i, \Lambda_i \in \overline{\mathbb{Q}}$ for all i, and each Λ_i is in the multiplicative group generated by $\{\gamma, r_1, \dots, r_\ell\}$. Suppose $m > 0$, i.e. $(u_n)_{n\in\mathbb{N}}$ is not identically zero.

(a) *There exists effectively computable N_1 such that $u_n \neq 0$ for all $n > N_1$.*

(b) *For $n > N_1$, $|u_n| > L^n n^{-C}$, where $L = \max_i |\Lambda_i|$ and C is an effectively computable constant.*

(c) *It is decidable whether $u_n \geq 0$ for all n.*

Proof. Define $\mathcal{D} = \{i : |\Lambda_i| = L\}$ and $\mathcal{R} = \{i : |\Lambda_i| < L\}$. The terms $c_i\Lambda_i^n$ for $i \in \mathcal{D}$ are called *dominant*. We have

$$u_n = \underbrace{\sum_{i\in\mathcal{D}} c_i\Lambda_i^n}_{v_n} + \underbrace{\sum_{i\in\mathcal{R}} c_i\Lambda_i^n}_{z_n}.$$

We next investigate $|v_n|$ as $n \to \infty$. Recall that each Λ_i is of the form $\gamma^{m_0} r_1^{m_1} \cdots r_\ell^{m_\ell}$, where $m_0, \ldots, m_\ell \in \mathbb{Z}$. In particular, for all i, $\Lambda_i = |\Lambda_i|\gamma^{k_i}$ for some $k_i \in \mathbb{Z}$. Hence we can write

$$v_n = L^n \sum_{i=-K}^{K} b_i \gamma^{in}$$

where each b_i is equal to some c_j. We have

$$v_n = \gamma^{-Kn} L^n \sum_{i=0}^{2K} b_{-K+i}\gamma^{in} = \gamma^{-Kn} L^n \prod_{i=0}^{2K} (\gamma^n - \alpha_i)$$

where $\alpha_0, \ldots, \alpha_{2K} \in \overline{\mathbb{Q}}$ are the zeros of the polynomial $p(z) = \sum_{i=0}^{2K} b_{-K+i} z^i$. Since γ is not a root of unity, we can apply Theorem 19 to each factor $(\gamma^n - \alpha_i)$ to conclude that there exist effectively computable N_1, C such that $|v_n| > L^n n^{-C}$ for al $n > N_1$. Since $|\Lambda_i| < L$ for all $i \in \mathcal{R}$, there exists (effectively computable) N_2 such that $|v_n| > |z_n|$ for $n > N_2$. We have proven (a) and (b).

Since u_n is real-valued, as discussed in for each $1 \le i \le m$ there exists $1 \le j \le m$ such that $c_j\Lambda_j = \overline{c_i\Lambda_i}$. Hence v_n, z_n are both real-valued. By the analysis above $\mathrm{sign}(u_n) = \mathrm{sign}(v_n)$ for $n > N_2$. Hence to check if u_n is positive we have to check if $u_n \ge 0$ for $0 \le n \le N_2$ and $v_n \ge 0$ for $n > N_2$. To do the latter, let $f(z) = z^{-K}p(z)$ and consider $Z \coloneqq f(\mathbb{T}) \subset \mathbb{R}$. Observe that Z is equal to the closure of $\{\gamma^{-Kn}p(\gamma^n) \mid n \in \mathbb{N}\}$ and hence is compact. If Z contains a negative number, then by Kronecker's theorem, v_n is negative for infinitely many n, in which case u_n is not positive. Otherwise, v_n and hence u_n are both positive. This concludes the proof of (c). ◀

Proof of Lemma 22

▶ **Lemma 22.** *The Positivity Problem can be reduced to the energy satisfaction problem above restricted to a Markov chain (M, q) and a linear weight function w.*

Proof. It is known from [1, 22] that the Positivity Problem for arbitrary LRS over $\mathbb{Q}$ can be reduced to the following problem: given a Markov chain (M, q), decide if there exists n such that $e_1 M^n q \ge 1/2$. We reduce the latter to the energy satisfaction problem. Given a Markov chain $(M, q) \in \mathbb{Q}^{d\times d} \times d$, let

$$P = \begin{bmatrix} M & \mathbf{0} \\ \mathbf{0} & M \end{bmatrix}$$

and $t = (1/2q, 1/2Mq) \in \mathbb{Q}^{2d}$. Observe that (P, t) is also a Markov chain. Moreover, $P^n t = (1/2M^n q, 1/2M^{n+1}q)$. We choose the weight function $w(x_1, \ldots, x_{2d}) = 2(x_{d+1} - x_1)$ and $B = 1/2 - e_1 \cdot q$. Then $w(P^n q) = e_1 M^{n+1} q - e_1 M^n q$, and $u_n \coloneqq \sum_{k=0}^{n} w(P^n q) \ge B$ if and only if $e_1 M^{n+1} q \ge 1/2$. Hence there does not exist n such that $e_1 M^n q \ge 1/2$ if and only if $e_1 \cdot q < 1/2$ and there does not exist n such that $u_n < B$. ◀