





Model Checking Linear Dynamical Systems under Floating-point Rounding^{*}

Engel Lefauchaux¹ , Joël Ouaknine² , David Purser^{3,4} , and
Mohammadamin Sharif⁵ 

¹ University of Lorraine, CNRS, Inria, LORIA, Nancy, France
`engel.lefauchaux@inria.fr`

² Max Planck Institute for Software Systems, Saarland Informatics Campus,
Saarbrücken, Germany `joel@mpi-sws.org`

³ University of Warsaw, Warsaw, Poland

⁴ University of Liverpool, Liverpool, UK `D.Purser@liverpool.ac.uk`

⁵ Sharif University of Technology, Tehran, Iran `sharifim689@gmail.com`

Abstract. We consider linear dynamical systems under floating-point rounding. In these systems, a matrix is repeatedly applied to a vector, but the numbers are rounded into floating-point representation after each step (i.e., stored as a fixed-precision mantissa and an exponent). The approach more faithfully models realistic implementations of linear loops, compared to the exact arbitrary-precision setting often employed in the study of linear dynamical systems.

Our results are twofold: We show that for non-negative matrices there is a special structure to the sequence of vectors generated by the system: the mantissas are periodic and the exponents grow linearly. We leverage this to show decidability of ω -regular temporal model checking against semi-algebraic predicates. This contrasts with the unrounded setting, where even the non-negative case encompasses the long-standing open Skolem and Positivity problems.

On the other hand, when negative numbers are allowed in the matrix, we show that the reachability problem is undecidable by encoding a two-counter machine. Again, this is in contrast with the unrounded setting where point-to-point reachability is known to be decidable in polynomial time.

Keywords: Model Checking · Floating-point · Dynamical Systems.

1 Introduction

Loops are a fundamental staple of any programming language, and the study of loops plays a pivotal role in many subfields of computer science, including automated verification, abstract interpretation, program analysis, semantics, etc. The focus of the present paper is on the algorithmic analysis of simple (i.e., non-nested) linear (or affine) while loops, such as the following:

^{*} A long version of this paper is available as [19].

```

x = 3, y = 4, z = 2
while x+3y+z > 4:
    x = 3x +2z
    y = 3x + y
    z = y + z

```

We are interested in analysing how the loop evolves. A simple reachability query is to decide whether the loop variables ever satisfy a Boolean combination of polynomial inequalities, for example modelling a loop guard. More generally, one might seek to consider significantly more complex temporal properties, such as those expressible in linear temporal logic or monadic second-order logic: this gives rise to a model-checking problem.

Modelling the evolution of such a loop may require unbounded memory. That is, the number of bits needed to represent the numbers x , y , and z may grow larger and larger. However, most computer systems do not represent rational numbers to arbitrary precision, but rather use *floating-point rounding*, in which a number y is stored using two components: the mantissa $m \in \mathbb{Q}$ and the exponent $\alpha \in \mathbb{Z}$, such that $y = m \cdot 10^\alpha$.⁶

Typically floating-point numbers are specified using either 32 or 64 bits, with some of these reserved for the mantissa and some for the exponent, thus bounding both the mantissa and the exponent. **We do not do this**, and only place a bound on the number of bits representing the mantissa, allowing the exponent to grow unboundedly (in either direction). From a theoretical standpoint, bounding the number of bits of both the mantissa and the exponent would necessarily give rise to a finite-state system, for which essentially any decision problem would become decidable (at least in principle, if not necessarily in practice). Due to the unboundedness of exponents in our setting, we do not have to consider overflows (‘NaN’, ‘infinity’ or ‘-infinity’ which are part of most floating-point specifications).

Formally, we model our programs using linear dynamical systems (LDS), which comprise a starting vector representing the initial state of each variable and a matrix describing the evolution of the program. An LDS generates an infinite sequence of vectors (the *orbit* of the system) by multiplying the matrix with the current vector and then applying floating-point rounding to the result.

Our results

We consider the *model-checking* problem for linear dynamical systems evolving under floating-point rounding. More formally, let $Y_1, \dots, Y_k \subseteq \mathbb{R}^d$ be semi-algebraic targets. Given an orbit $(x^{(t)})_{t \in \mathbb{N}}$, we define the characteristic word $w = w_1, w_2, w_3, \dots$ with respect to Y_1, \dots, Y_k over alphabet $2^{\{1, \dots, k\}}$ such that $i \in w_t$ if and only if $x^{(t)} \in Y_i$. The model-checking problem asks whether w is in an ω -regular language, or equivalently satisfies a temporal specification given in monadic second-order logic (MSO).

⁶ We work in base 10 throughout for simplicity of exposition. All our results carry over *mutatis mutandis* in any integer base, including base 2 as typically used in practice.

Our results show that analysing LDS under floating-point rounding is neither clearly easier nor harder than in the standard setting (without rounding). Our first contribution establishes *undecidability* of point-to-point reachability (and *a fortiori* model checking) under floating-point rounding, a surprising outcome given that point-to-point reachability is solvable in polynomial time without rounding [16]. On the other hand, in the standard setting neither decidability nor undecidability are known for full model checking (although mathematical hardness results exist); see [24,18,17].

Theorem 1. *The floating-point point-to-point reachability problem is undecidable.*

However, for non-negative matrices, we show that the full MSO model-checking problem is decidable in our setting, without restrictions on the dimensions of the predicates or the ambient space. This is in stark contrast to the standard setting, where assuming non-negativity does not simplify the problem. Model checking non-negative LDS without rounding would require (at a minimum) solving the longstanding open Skolem and Positivity problems [2].

Theorem 2. *Let (M, x) be a non-negative linear dynamical system, let Y_1, \dots, Y_k be semialgebraic targets and let ϕ be an MSO formula using predicates over Y_1, \dots, Y_k . It is decidable whether the characteristic word under floating-point rounding satisfies ϕ .*

We place no dimension restriction on the predicates; in particular, showing that the Skolem and Positivity problems are *decidable* on non-negative systems under floating-point rounding. At this time we do not however have complexity upper bounds on our model-checking algorithm, or lower bounds on the model-checking problem.

Related work

There is a line of practical tools for the analysis, verification, and invariant synthesis for floating-point loops [7,20,1,22]. These tools typically work well in practice, but do not necessarily work in all cases. The analysis of concrete implementations of floating-point specifications requires careful analysis of edge cases around $\pm\infty$ and ‘NaN’. In contrast to these tools which focus primarily on practical analysis, our work seeks to understand the theoretical possibilities and limitations of the exact analysis of (possibly long-running) floating-point loops in a generalised setting.

The study of linear dynamical systems explores the sequence of vectors induced by a matrix. Model checking is only known to be decidable for certain classes of semialgebraic predicates—in particular those with low dimension [18] or for prefix-independent properties [4]; see also [17]. The well-known Skolem and Positivity problems being special cases of model checking, they place technical limits on the dimensions that can be handled without first resolving longstanding open cases of these problems. Recent progress suggests that the Skolem

problem may be yet be conquered, at least for diagonalisable matrices [8,21], but Positivity requires solving particularly difficult problems in analytic number theory [24,12]. The non-negative case can be used to model sequences of distributions induced by Markov chains [6], although all hardness limitations apply already in the probabilistic setting [2].

Baier et al. [5] consider LDS under rounding to fixed-decimal precision, showing reachability is PSPACE-complete for hyperbolic systems (when no eigenvalue has modulus one) and decidable for certain other constrained classes of rounding. A notable difference of fixed-decimal precision is that it cannot allow arbitrarily small numbers, unlike the floating-point numbers we consider.

A recent line of work focusses on linear dynamical systems with perturbations at every step, with a view to understanding the robustness of reachability problems [13,14,3]. However, unlike rounding, the perturbation is chosen in order to assist hitting the target and the perturbation is arbitrarily small.

For linear while loops the reachability problem can be rephrased as a halting problem, asking whether a guard condition is eventually met from a given initial state. The related termination problem asks whether a guard condition is met from *every* initial state [26,10]. Issues arising from implementations using floating-point representations to solve the termination problem of unrounded (arbitrary precision) loops are considered in [27]. In contrast, we are interested in analysing programs in which the intended behaviour is to round the numbers to fixed-precision floating-point numbers at every step of the loop.

Organisation In Section 2, we formalise the model and problems and discuss some of the properties of floating-point rounding. In Section 3, we present our undecidability result for the general case. Finally, in Section 4 we establish some special periodic structure associated with the orbit and use this structure in Section 5 to show that model checking is decidable for non-negative LDS.

2 Preliminaries

2.1 Linear dynamical systems and rounding functions

Definition 1. A d -dimensional linear dynamical system (LDS) (M, x) comprises a matrix $M \in \mathbb{Q}^{d \times d}$ and an initial vector $x \in \mathbb{Q}^d$.

Given a rounding function $[\cdot] : \mathbb{Q}^d \rightarrow \mathbb{Q}^d$, and an LDS (M, x) the rounded orbit \mathcal{O} is the sequence $(x^{(t)})_{t \in \mathbb{N}}$ such that $x^{(0)} = [x]$ and $x^{(t)} = [Mx^{(t-1)}]$ for all $t \geq 1$.

Given $p \in \mathbb{N}$, we say that a number x is a floating-point number with precision p if $x = m \cdot 10^\alpha$ such that $m \in \mathbb{Q}$ is a decimal number in $\{0\} \cup [0.1, 1)$ with p digits in the fractional part (after the decimal point) and $\alpha \in \mathbb{Z}$. In particular, we associate by convention the number with mantissa $m = 0$ to the exponent $-\infty$. Given a number $x = m \cdot 10^\alpha$ we define $\text{mantissa}(x) = m$ and $\text{exponent}(x) = \alpha$.

We are interested in the floating-point rounding function $[\cdot]$ with precision $p \in \mathbb{N}$. Given a real number $x \in \mathbb{R}$, we define $[x]$, the floating-point rounding of

x , as the closest floating-point number with precision p based on the first $p + 1$ digits of x .

Where there are two possible choices, any deterministic choice that is consistent with the properties listed below is acceptable.⁷ We denote by $\mathbb{FP}_{10}[p]$ the subset of \mathbb{Q} representable in base 10 as a floating-point numbers with p digits. We use the following useful properties of the rounding function:

- it is *log-bounded*, i.e. there exists a constant $c \in \mathbb{R}_+$ such that $\forall x \in \mathbb{R}, \frac{|x|}{c} \leq |[x]| \leq c|x|$.
- it is *mantissa-based*, i.e. if $x = 10^\alpha x'$, then $[x] = 10^\alpha [x']$.
- it is $(p + 1)$ -*finite*, i.e. the output of the rounding is not dependent on the i -th digit of the mantissa, for each integer $i > p + 1$. In other words, if x and x' agree on the first $p + 1$ digits then $[x] = [x']$.
- it is *sign preserving*, i.e. $\text{sign}(x) = \text{sign}([x])$. The fact that $[x] = 0$ if and only if $x = 0$ also follows from the log-bounded property.

The floating-point rounding is defined above on a single real. It is extended straightforwardly to a vector x by applying it to each of its components $(x)_i$ where i ranges from 1 to the dimension of the vector. As such, the term $[Mx]$ is obtained by first computing exactly the the vector Mx and then by rounding each component $(Mx)_i$. An alternative approach could be to maintain each sub-computation in p -bits of precision, *but this is not the approach we take*. Such an orbit can be simulated in our setting by increasing the dimension so that operations can be staggered in a way that at most one operation (scalar product or variable addition) is used in each assignment.

2.2 Model checking

We consider the model-checking problem of an LDS over semialgebraic sets.

Definition 2. *A semialgebraic set $Y \subseteq \mathbb{R}^d$ is defined by a finite Boolean combination of polynomial inequalities.*

Let (M, x) be an LDS with rounded orbit \mathcal{O} and $\mathcal{Y} = \{Y_1, \dots, Y_k\}$ be a collection of semialgebraic sets. The characteristic word of \mathcal{O} is $w = w_1 w_2 w_3 \dots \in (2^{\{1, \dots, k\}})^\omega$, such that $j \in w_t$ if and only if $x^{(t)} \in Y_j$.

The model-checking problem asks whether the characteristic word is contained within a given ω -regular language, usually specified in a temporal logic such as monadic second order logic (MSO), or often its LTL fragment. Without loss of generality we assume that the property is given as a Büchi automaton [11].

Problem 1 (Floating-point Model-checking Problem). Given an LDS (M, x) with rounded orbit \mathcal{O} , a collection of semialgebraic sets $\mathcal{Y} = \{Y_1, \dots, Y_k\}$ and an ω -regular specification ϕ , the model-checking problem consists in deciding whether the characteristic word w of \mathcal{O} satisfies the specification ϕ .

⁷ For example, always rounding up, always rounding down, round to even, rounding towards zero, rounding away from zero are acceptable, providing the choice is fixed.

We will also consider the point-to-point reachability problem, which is a subcase of the model-checking problem (Problem 1):

Problem 2 (Floating-point Point-to-point Reachability Problem). Given a d -dimensional LDS (M, x) , and a target vector $y \in \mathbb{Q}^d$, the point-to-point reachability problem consists in deciding whether y belongs to the rounded orbit \mathcal{O} .

Given a target $Y \subseteq \mathbb{R}^d$, we associate the set of hitting times $\mathcal{Z}(Y) = \{t \mid x^{(t)} \in Y\}$. Under this formulation, the reachability problem is reformulated as whether $\mathcal{Z}(Y)$ is empty. However, for model checking we will develop a more comprehensive understanding of the hitting times of each target Y_1, \dots, Y_k .

2.3 Structure of M

Formally, M is a d -dimensional matrix indexed by the elements $\{1, \dots, d\}$. However, we interpret M as an automaton over states $Q = \{q_1, \dots, q_d\}$ and reference the entries of M by pairs of states. That is, we refer to M_{q_1, q_2} rather than $M_{1,2}$.

We denote by G_M the weighted directed graph whose adjacency matrix is M . That is, a graph with vertices Q and with an edge from q_j to q_i weighted by M_{q_i, q_j} if $M_{q_i, q_j} \neq 0$.⁸

Let $S_1, \dots, S_s \subseteq Q$ be the strongly connected components (SCCs) of G_M . Our analysis will consider each strongly connected component separately, thus it will often be useful to consider the entries of $x \in \mathbb{F}_{10}[p]^Q$ corresponding only to one strongly connected component. Without loss of generality, by re-ordering the states where necessary, we assume that the states in Q are ordered so that states within the same SCC appear next to one another, and the strongly connected components are topologically sorted, *i.e.* there is no edge from S_i to S_j where $i > j$. We split a vector x into s smaller vectors, denoted x_{S_1}, \dots, x_{S_s} , each representing the entries of x corresponding to the SCC. Letting $x_{S_j} = (z_{1,j}, \dots, z_{d_j,j})^T$ and $|S_j| = d_j$, we thus have x is partitioned as

$$x = (z_{1,1} \cdots z_{d_1,1}, \dots, z_{1,s} \cdots z_{d_s,s})^T.$$

Moreover, for each pair of SCCs S_i, S_j , we denote by M_{S_i, S_j} the submatrix of M restricted to the rows related to S_i and columns related to S_j , which is a matrix with d_i rows and d_j columns. If $S_i = S_j$, we simply write M_{S_i} . In other words, M_{S_i, S_j} is the matrix that shows the dependency between S_i and S_j , and we have

$$M = \begin{pmatrix} M_{S_1} & M_{S_1, S_2} & \cdots & M_{S_1, S_s} \\ M_{S_2, S_1} & M_{S_2} & \cdots & M_{S_2, S_s} \\ \vdots & \vdots & \ddots & \vdots \\ M_{S_s, S_1} & M_{S_s, S_2} & \cdots & M_{S_s} \end{pmatrix}$$

We say S_i *feeds* S_j , and S_j is *fed by* S_i if there is some edge in G_M from some state in S_i to some state in S_j .

⁸ Note that the orientation of the edge may appear switched from the reader's expectation. This is due to the convention that M is pre-multiplied with x at every step.

3 Undecidability of point-to-point reachability

In this section, we give a sketch of the proof of the undecidability of Problem 2 (and thus of Problem 1) in the general case. The full proof can be found in the long version of this paper [19].

Theorem 1. *The floating-point point-to-point reachability problem is undecidable.*

This result is obtained by reduction from the termination of a two-counter Minsky machine. We recall the definition of this model:

Definition 3. *A two-counter Minsky machine is defined by a finite set of states ℓ_1, \dots, ℓ_m , a distinguished starting state (w.l.o.g. ℓ_1), a distinguished halting state (w.l.o.g. ℓ_m), two natural integer counters, here denoted as x and y , and a mapping deterministically associating to each state transition a particular action.*

Each transition takes one of the following forms: for $z \in \{x, y\}$,

increment $\text{inc}_z(\ell_j)$: *add 1 to counter z , move to state ℓ_j .*

decrement $\text{dec}_z(\ell_j)$: *remove 1 from counter z if $z > 0$, move to state ℓ_j .*

zero test $\text{zero?}_z(\ell_j, \ell_k)$: *if $z = 0$ move to state ℓ_j else move to state ℓ_k .*

The configuration of a two-counter Minsky machine consists of the current state and the values of x and y .

Without loss of generality (by first using a zero test), one can assume a decrementation operation is never used in a configuration where the counter to be decreased has value 0, hence removing the need to check whether $z > 0$.

The halting problem asks whether, starting in configuration $(\ell_1, 0, 0)$, that is, in the distinguished starting state with both counters set to 0, whether the state ℓ_m is reached. The problem is undecidable [23].

We build an LDS with mantissa length $p = 1$ and base 10 that simulates a run of a given Minsky machine. The reduction happens to maintain the invariant that each mantissa always has the value 0 or 1 after rounding (although, as we operate in base 10, there are 10 possible values the mantissa could have taken). For ease of readability, we describe this LDS using variables to represent the dimensions and linear functions to represent the transition matrix. For each state of the Minsky machine, we use two variables corresponding to the two counters. Throughout the simulation, if the Minsky machine is in state j , the counter values are stored in the exponents of the variables associated with state j , and all other variables are zero.

The crux of our reduction lies in the handling of the zero test. More precisely, suppose we need to branch depending on whether x is equal to 0, then we need to define linear transitions that transfer the values of the two counters from one pair of variables to the appropriate new pair of variables. This is done using filter functions: the function $\text{filter}_+(u, v)$ (resp. $\text{filter}_-(u, v)$) is equal to v if $v \geq u$ (resp. $v < u$) and to 0 otherwise. We end this sketch with the construction of these functions and proof that they operate as advertised.

Lemma 1. *Given u, v of the form 10^c with $c \in \mathbb{N}$, one can compute the value $w = \text{filter}_+(u, v)$ in three linear operations with floating-point rounding.*

Proof. We compute $w = \text{filter}_+(u, v)$ in three successive operations using two temporary variables, $temp$ and $temp2$, initially set at 0 (recall, rounding is applied after each step):

$$\begin{aligned} temp &\leftarrow u + v \\ temp2 &\leftarrow temp - u \\ w &\leftarrow 1.1 * temp2 \end{aligned}$$

Let $c_1, c_2 \in \mathbb{N}$ such that $u = 10^{c_1}$ and $v = 10^{c_2}$. Recall that the notation $[\cdot]$ is the floating-point rounding function.

First observe that if $c_1 = c_2$:

$$\begin{aligned} temp &\leftarrow [10^{c_1} + 10^{c_2}] = 2 \cdot 10^{c_1} \\ temp2 &\leftarrow [2 \cdot 10^{c_1} - 10^{c_1}] = 10^{c_1} (= v) \\ w &\leftarrow [1.1 \cdot 10^{c_1}] = 10^{c_1} = v \quad \text{as required.} \end{aligned}$$

Secondly, assume that $u > v$, and thus $c_1 > c_2$:

$$\begin{aligned} temp &\leftarrow [10^{c_1} + 10^{c_2}] = 10^{c_1} = u \\ temp2 &\leftarrow [10^{c_1} - 10^{c_1}] = 0 \\ w &\leftarrow [1.1 \cdot 0] = 0 \quad \text{as required.} \end{aligned}$$

We split the case that $v > u$, thus $c_2 > c_1$, into two cases. Suppose $c_2 > c_1 + 1$:

$$\begin{aligned} temp &\leftarrow [10^{c_1} + 10^{c_2}] = 10^{c_2} = v \\ temp2 &\leftarrow [10^{c_2} - 10^{c_1}] = [0.\underbrace{99\dots99}_{c_2-c_1 \geq 2} \cdot 10^{c_2}] = 1 \cdot 10^{c_2} = v \\ w &\leftarrow [1.1 \cdot 10^{c_2}] = 10^{c_2} = v \quad \text{as required.} \end{aligned}$$

Finally, $c_2 = c_1 + 1$:

$$\begin{aligned} temp &\leftarrow [10^{c_1} + 10^{c_2}] = 10^{c_2} = v \\ temp2 &\leftarrow [10^{c_2} - 10^{c_1}] = [0.9 \cdot 10^{c_2}] = 9 \cdot 10^{c_2-1} \\ w &\leftarrow [1.1 \cdot 9 \cdot 10^{c_2-1}] = [9.9 \cdot 10^{c_2-1}] = 10 \cdot 10^{c_2-1} = 10^{c_2} = v \\ &\hspace{15em} \text{as required.} \quad \square \end{aligned}$$

Corollary 1. *Given u, v of the form 10^c with $c \in \mathbb{N}$, one can compute the value $w = \text{filter}_-(u, v)$ in four linear operations with floating-point rounding.*

Proof. Observe that $\text{filter}_-(u, v) = v - \text{filter}_+(u, v)$, which can be encoded in four steps by first computing $\text{filter}_+(u, v)$ in three steps. \square

4 Pseudo-periodic orbits of non-negative LDS

We shift our focus to proving that model checking is decidable for systems with non-negative matrices. We first establish the behaviour of the system in this section and then complete the proof of Theorem 2 in Section 5. Our main result is that the rounded orbit of an LDS is periodic in the following sense, which we call *pseudo-periodic*.

Definition 4. *A sequence $(x^{(t)})_{t \in \mathbb{N}}$ of d -dimensional vectors of floating-point numbers is called *pseudo-periodic* if and only if there exists a starting point $N \in$*

\mathbb{N} , period $T \in \mathbb{N}$ and growth rates $\alpha_1, \dots, \alpha_d \in \mathbb{Z}$ such that

$$\forall t \geq N, \forall j \in \{1, \dots, d\}, (x^{(t+T)})_j = 10^{\alpha_j} (x^{(t)})_j.$$

We say the sequence is effectively pseudo-periodic if the defining constants $N, T, \alpha_1, \dots, \alpha_d$ can be computed.

Theorem 3. *Let (M, x) be a d -dimensional LDS where M is non-negative and let $(x^{(t)})_{t \in \mathbb{N}}$ be its rounded orbit.*

The rounded orbit $(x^{(t)})_{t \in \mathbb{N}}$ is effectively pseudo-periodic.

In order to establish this result, we will find some partitions of the graph associated to M such that each part is effectively pseudo-periodic with the same increasing rate α for every state in the partition.

4.1 Preprocessing periodicity

The core of our approach is to show that, within each SCC of the graph associated to M , the values associated with states are of similar magnitude. This is however only true if the SCC is aperiodic. When a state is in a periodic SCC its value could change drastically depending on which phase the system is in. For example, consider a simple alternation between two states, in which the value is very large in one state and very small in the other; the states will alternate between big and small values.

We “hide” these periodic behaviours by blowing up the system so that each SCC of the new system describes only one of the periodic subsequence and we will subsequently show that the value of each state in an SCC is either zero or of a similar magnitude.

We apply the following construction to our system. Let P be the period, defined as the least common multiple of the length of every simple cycle in the graph. Let Q be the indices of M (*i.e.* the states of the generated automaton). We define new states $Q' = Q \times \{0, \dots, P - 1\}$ by annotating each state in Q with the phase. To avoid cluttering notation we will regularly refer to states in Q' in the form $(q, i + \ell)$ for $\ell \in \mathbb{Z}$, on the understanding that the phase, $i + \ell$, is normalised into $\{0, \dots, P - 1\}$ by taking the residue modulo P if necessary. We define a new matrix M' over the states Q' such that $M'_{(q, i+1), (q', i)} = M_{q, q'}$ for $i \in \{0, \dots, P - 1\}$, and zero otherwise. We initialise a new starting vector $x_{(q, 0)}^{(0)} = x_q^{(0)}$ and $x_{(q, i)}^{(0)} = 0$ for $i \in \{1, \dots, P - 1\}$.

Intuitively, at each time step t the vector generated by the original system is equal to the vector of the new system restricted to the states indexed by $i \equiv t \pmod{P}$ and every state with another index is equal to 0.

Let $S \subseteq Q$ be a strongly connected component. In Q' there exists strongly connected components $S'_1, \dots, S'_k \subseteq Q'$ with $k \leq |S|$ such that $\bigcup_{i=1}^k S'_i = S \times \{0, \dots, P - 1\}$. Each set S'_j is periodic, with period P .

Henceforth in the rest of this section we work on the system (M', x') implicitly over states Q' which, by overloading of notation, we rename (M, x) over Q to avoid cluttering notation.

Note that this transformation also requires to marginally complicate the targets. Indeed, consider a set $Y \subseteq \mathbb{R}^Q$. We define the sets Y/i for $i < P$ such that $Y/i = \{y \in \mathbb{R}^Q \mid \exists y' \in Y : y_{(q,i)} = y'_q \text{ for } q \in Q \text{ and } y'_{(q,j)} = 0 \text{ for } j \neq i\}$. The hitting times of Y , $\mathcal{Z}(Y)$, in the original LDS can then be obtained in the new LDS as the disjoint union: $\bigcup_{i \in \{0, \dots, P-1\}} \mathcal{Z}(Y/i)$. It suffices to characterise the hitting times for each Y/i .

4.2 Pseudo-periodicity within top SCCs

Let us first consider top SCCs, these are SCCs with no incoming edges from states of other SCC, and therefore the value of each variable at each step depends only on the value of states in the same SCC.

Lemma 2. *Let S_j be a strongly connected component of (M, x) . Let $S_{j,i} = \{(q, i) \in S_j\}$ be the states associated with S_j from the i -th phase.*

There exists $C \leq Pd^2$, such that, for every i, j , $(M^C)_{S_{j,i}}$ is positive.

Proof. The matrix $(M^P)_{S_{j,i}}$ is non-negative, irreducible (*i.e.*, its graph is strongly connected) and of period 1. As such, $(M^P)_{S_{j,i}}$ is primitive [9] which means that a power C' of this matrix is positive. The theorem follows with $C = PC'$. Moreover, C' is at most $d^2 - 2d + 2$ [25]. \square

Our goal is to show that within an SCC, each of the non-zero entries are of a similar magnitude due to the presence of a relatively short path (C) between any two states in the SCC. To do this we introduce the notion of closeness and observe some useful properties.

Definition 5. *We say two numbers $x, x' \in \mathbb{FP}_{10}[p]$ are δ -close, denoted by $x \approx_\delta x'$ if $|\text{exponent}(x) - \text{exponent}(x')| < \delta$. In particular, for every $\delta > 0$, zero is assumed to be δ -close only to itself.*

We extend the notion to vectors $y, y' \in \mathbb{FP}_{10}[p]^S$, indexed by $S \subseteq Q$, such that $y \approx_\delta y'$ if all entries of the same phase are δ -close to one another across both y and y' , that is, for each phase $i \in \{0, \dots, P-1\}$ and all $(q, i), (q', i) \in S$: $y_{(q,i)} \approx_\delta y'_{(q',i)}$, $y_{(q,i)} \approx_\delta y_{(q',i)}$ and $y'_{(q,i)} \approx_\delta y'_{(q',i)}$.

Proposition 1. *Let $x, x' \in \mathbb{FP}_{10}[p]$ be non-zero floating-point numbers.*

- (1) *If $x \approx_\delta x'$ then $10^{-\delta-1} \leq x/x' \leq 10^{\delta+1}$.*
- (2) *If $10^{-\delta} \leq x/x' \leq 10^\delta$ then $x \approx_{\delta+2} x'$.*
- (3) *If $x \approx_\delta x'$ and $x' \approx_\eta x''$ then $x \approx_{\delta+\eta+4} x''$.*

Lemma 3. *Let S_j be a top strongly connected component of (M, x) , and let C be as given by Lemma 2.*

There exists $\beta \in \mathbb{N}$ such that for all $(q, i), (q', i) \in S_j$ and every $t \geq C$ then

- if $t \not\equiv i \pmod{P}$, then $x_{(q,i)}^{(t)} = 0$,*
- otherwise, $x_{(q,i)}^{(t)} \approx_\beta x_{(q',i)}^{(t)}$.*

Proof. Let $t \in \mathbb{N}$. If $t \not\equiv i \pmod{P}$ then $x_{(q,i)}^{(t)} = 0$ for all $(q, i) \in S_{j,i}$ by construction.

Otherwise, let $m \geq \max_{q,q' \in Q: M_{q,q'} \neq 0} \max(M_{q,q'}, (M_{q,q'})^{-1})$ be a constant larger than all values occurring in M and so that $\frac{1}{m}$ is smaller than all non-zero values appearing in M . Let c be the constant from the log bounded property of the rounding function $[\cdot]$ and d be the dimension of M .

Observe that for all $t \in \mathbb{N}$ with $t = i \pmod{P}$ we have

$$\begin{aligned} x_{(q,i)}^{(t)} &= \left[\sum_{(q',i-1)} M_{(q,i),(q',i-1)} x_{(q',i-1)}^{(t-1)} \right] \\ &\geq \frac{1}{c} \sum_{(q',i-1)} M_{(q,i),(q',i-1)} x_{(q',i-1)}^{(t-1)} \quad (\text{by log bounded}) \\ &\geq \frac{1}{cm} \max_{(q',i-1) \text{ s.t. } M_{(q,i),(q',i-1)} > 0} x_{(q',i-1)}^{(t-1)} \quad (\text{by defn of } m) \end{aligned}$$

In particular

$$x_{(q,i)}^{(t)} \geq \frac{1}{cm} x_{(q',i-1)}^{(t-1)} \text{ for all } (q', i-1) \text{ s.t. } M_{(q,i),(q',i-1)} > 0$$

Using induction we obtain:

$$x_{(q,i+k)}^{(t+k)} \geq \frac{1}{(cm)^{k-1}} x_{(q',i+1)}^{(t+1)} \geq \frac{1}{(cm)^k} x_{(q'',i)}^{(t)}$$

for all $(q', i+1), (q'', i)$ such that $M_{(q,i+k),(q',i+1)}^{k-1} > 0$ and $M_{(q',i+1),(q'',i)} > 0$.

In particular, we have $x_{(q,i)}^{(t+C)} \geq \frac{1}{(cm)^C} x_{(q',i)}^{(t)}$ for all q' (since $M_{(q,i),(q',i)}^C > 0$ for all q' by the previous lemma).

On the other hand we have

$$x_{(q,i+1)}^{(t+1)} = \left[\sum_{q': M_{(q,i+1),(q',i)} > 0} M_{(q,i+1),(q',i)} x_{(q',i)}^{(t)} \right] \leq mcd \max_{(q',i) \in S_j} x_{(q',i)}^{(t)}.$$

By induction we get that $x_{(q,i)}^{(t+C)} \leq (mcd)^C \max_{(q',i) \in S_j} x_{(q',i)}^{(t)}$. Hence, for all $q, q' \in S_j$ we have

$$\frac{1}{(mc)^C} \max_{(q'',i) \in S_j} x_{(q'',i)}^{(t)} \leq x_{(q',i)}^{(t+C)} \quad \text{and} \quad x_{(q,i)}^{(t+C)} \leq (mcd)^C \max_{(q'',i) \in S_j} x_{(q'',i)}^{(t)}.$$

Hence $\frac{x_{(q,i)}^{(t+C)}}{x_{(q',i)}^{(t+C)}} \leq d^C (mc)^{2C}$.

Setting $\gamma = \lceil \log_{10} d^C (mc)^{2C} \rceil$, we thus have that $10^{-\gamma} x_{(q',i)}^{(t+C)} \leq x_{(q,i)}^{(t+C)} \leq 10^\gamma x_{(q',i)}^{(t+C)}$ for all $(q, i), (q', i) \in S_{j,i}$ and $t \in \mathbb{N}$. Then $x_{(q',i)}^{(t)}$ and $x_{(q,i)}^{(t)}$ are $\beta = \gamma + 2$ close by Proposition 1. \square

Lemma 4. *Let S_j be a top strongly connected component of (M, x) . Then the sequence $(x_{S_j}^{(t)})_{t \in \mathbb{N}}$ is effectively pseudo-periodic.*

Proof. Let β and C be as in Lemma 3. Denote q_1, \dots, q_m the states of S_j . We define the sequence $(y^{(t)})_{t \geq C}$ such that for all $t \geq C$ and $q \in S_j$ denoting $(p^{(t)})_q = \text{mantissa}([x_q^{(t)}])$ and $(\alpha^{(t)})_q = \text{exponent}([x_q^{(t)}])$ we have that $y^{(t)} = (p_{q_1}, 0, p_{q_2}, \alpha_{q_2} - \alpha_{q_1}, \dots, p_{q_m}, \alpha_{q_m} - \alpha_{q_1})$. Note that this sequence can only take finitely many values as the mantissas have a precision of p decimals and by Lemma 3, for all $k \leq m$, $\alpha_{q_k} - \alpha_{q_1} \in \{-\beta, \dots, \beta\}$. As a consequence, the sequence $(y^{(t)})_{t \geq C}$ takes the same value multiple times. Let k_1 and k_2 be the two distinct minimal integers such that $y^{(k_1)} = y^{(k_2)}$. Setting $\alpha = \alpha_{q_1}^{(k_2)} - \alpha_{q_1}^{(k_1)}$ We have that $x^{(k_1)} = x^{(k_2)} \cdot 10^\alpha$. Since $[\cdot]$ is mantissa-based, one can show by induction that for all $t \geq 0$, $x^{(k_1+t)} = x^{(k_2+t)} \cdot 10^\alpha$. Therefore the sequence $(x_{S_j}^{(t)})_{t \in \mathbb{N}}$ is effectively pseudo-periodic with period $T = k_2 - k_1$ and starting point $N = C + k_1$.

Moreover, as the maximum number of different values taken by $(y^{(t)})_{t \geq C}$ is known, we can deduce that both k_1 and $k_2 - k_1$ are smaller than $10^{pm}(2\beta+1)^m + 1$. \square

Note that the increasing rate is the same for every state of the strongly connected component.

4.3 Pseudo-periodicity within lower SCCs

We consider a strongly connected component S_{me} , which is fed by at least one strongly connected components F_1, \dots, F_ℓ , $\ell \geq 1$. We let $S_F = F_1 \cup \dots \cup F_\ell$ and assume every F_i is pseudo-periodic.

In this section we show

Theorem 4. *S_{me} is effectively pseudo-periodic and the growth rate of S_{me} is the same for all $q \in S_{me}$.*

We first observe that the difference between values in S_{me} is bounded. This is achieved with a proof similar to the one of Lemma 2 and Lemma 3 (though having to combine considerations of S_{me} and S_F).

Lemma 5. *There exists $\eta, N' \in \mathbb{N}$, such that for all $(q, i), (q', i) \in S_{me}$, all $t \geq N'$ and all $i \in \{0, \dots, P-1\}$ then*

- if $t \not\equiv i \pmod{P}$, then $x_{(q,i)}^{(t)} = 0$,
- otherwise, $x_{(q,i)}^{(t)} \approx_\eta x_{(q',i)}^{(t)}$.

Definition 6. *We say that $x_q^{(t)}$ is influenced by S_F if*

$$x_q^{(t)} = \left[\sum_{q' \in S_F} M_{q,q'} x_{q'}^{(t-1)} + \sum_{q' \in S_{me}} M_{q,q'} x_{q'}^{(t-1)} \right] \neq \left[\sum_{q' \in S_{me}} M_{q,q'} x_{q'}^{(t-1)} \right]$$

and in particular $x_q^{(t)}$ is influenced by $u \in S_F$ if:

$$\left[\sum_{q' \in S_F \cup S_{me}} M_{q,q'} x_{q'}^{(t-1)} \right] \neq \left[\sum_{q' \in S_F \cup S_{me} \setminus \{u\}} M_{q,q'} x_{q'}^{(t-1)} \right].$$

We can restrict S_F to the F_i in S_F with the maximum growth rate. Indeed, from some point on, any F_i with non-maximal growth rate is much smaller than the maximal ones, and as by the proof of Lemma 5 the values within S_{me} are close to (or greater than) the maximum value within S_F , this F_i would not influence with any $x_q^{(t)}$ with $q \in S_{me}$. Let N_1 be the point from which we can assume, that the elements of S_F are much larger than any other feeding SCCs and are thus the only ones potentially influencing of S_{me} .

Since each F_i is assumed to be pseudo-periodic, we have that S_F pseudo-periodic. Let T be the period of S_F , N_2 be the starting point and α be the growth rate of every state of S_F (meaning the exponent of every state changes by α every T starting from the N -th step.) Let $N = \max\{N_1, N_2\}$, that is, the point from which we can assume S_F is both pseudo-periodic and dominating non-maximal SCCs feeding S_{me} .

As a direct consequence of having the same growth rate, the non-zero terms within S_F are close:

Proposition 2. *If a sequence of non-zero floating-point vectors $(v^{(t)})_{t \in \mathbb{N}}$ is pseudo-periodic with the same growth rate within a set Q , then there exists δ such that for all $q, q' \in Q$ and all $t \geq N$, $v_q^{(t)} \approx_\delta v_{q'}^{(t)}$.*

Moreover, either S_F does not influence S_{me} , or they are close.

Lemma 6. *There exists $\beta, N \in \mathbb{N}$ such that:*

For $t \geq N$ and $(q, i) \in S_{me}$, if $x_{(q,i)}^{(t)}$ is influenced by $(q', i-1) \in S_F$, then $x_{(r,i)}^{(t)} \approx_\beta x_{(r',i)}^{(t)}$ for all $(r, i), (r', i) \in S_{me} \cup S_F$.

We will show Theorem 4 through the following observation:

Observation 1. Observe that S_F either influences S_{me} infinitely many times or finitely many times. We have two cases:

- If S_F influences S_{me} infinitely often, then they are infinitely often β -close by Lemma 6. Then we will observe through a simultaneous version of Lemma 4 that S_{me} is pseudo-periodic.
- If S_F influences S_{me} only finitely often, then clearly from some point on S_{me} behaves like a top SCC, and thus is pseudo-periodic directly by Lemma 4.

It will then remain to show that we can detect which of the two cases applies, and place a bound on the time to detect this, which will effectively reveal the constants of the pseudo-periodic behaviour.

We now present a version of Lemma 4 to observe that if S_F and S_{me} are infinitely often β -close then S_{me} is pseudo-periodic:

Lemma 7. *Suppose $x_{S_F}^{(t)} \approx_\beta x_{S_{me}}^{(t)}$ for infinitely many t . Then there exists $t_1 < t_2$, such that $x_{S_F}^{(t_1)} \approx_\beta x_{S_{me}}^{(t_1)}$ and $x_{S_F}^{(t_2)} \approx_\beta x_{S_{me}}^{(t_2)}$, $x_{S_F}^{(t_2)} = 10^\gamma x_{S_F}^{(t_1)}$ and $x_{S_{me}}^{(t_2)} = 10^\gamma x_{S_{me}}^{(t_1)}$. In particular, the sequence $(x_{S_{me}}^{(t)})_{t \in \mathbb{N}}$ is pseudo-periodic with period $(t_2 - t_1)$, starting from t_1 with growth rate of γ in every state.*

Proof. At a time t such that $x_{S_F}^{(t)} \approx_\beta x_{S_{me}}^{(t)}$, we denote the vectors $x_{S_F}^{(t)} \in \mathbb{FP}_{10}[p]^{|S_F|}$ and $x_{S_{me}}^{(t)} \in \mathbb{FP}_{10}[p]^{|S_{me}|}$ respectively

$$(m_1^{(t)} 10^{\gamma_1^{(t)}}, m_2^{(t)} 10^{\gamma^{(t)} + \alpha_2^{(t)}}, \dots, m_{|S_F|}^{(t)} 10^{\gamma^{(t)} + \alpha_{|S_F|}^{(t)}}) \text{ and} \\ (n_1^{(t)} 10^{\gamma^{(t)} + \zeta_1^{(t)}}, \dots, n_{|S_{me}|}^{(t)} 10^{\gamma^{(t)} + \zeta_{|S_{me}|}^{(t)}}),$$

where m_i, n_i are taken from the finite set of mantissa values expressible in p bits, $\gamma^{(t)} \in \mathbb{Z}$ and $\alpha_i, \zeta_i \in \mathbb{Z} \cap [-\beta, \beta]$ denote the offset from $\gamma^{(t)}$.

Let F bound the number of possible values $m_i, n_i, \alpha_i, \zeta_i$ can take on, where $F \leq 10^{p(|S_F| + |S_{me}|)} \cdot (2\beta + 1)^{|S_F| + |S_{me}| - 1}$. By the pigeonhole principle, after at most $F + 1$ times in which $x_{S_F}^{(t)} \approx_\beta x_{S_{me}}^{(t)}$ there must exist two times $t_1 < t_2$ where the values of $m_i, n_i, \alpha_i, \zeta_i$'s are equal (although the value of γ could be different), thus $x_{S_F \cup S_{me}}^{(t_2)} = \frac{10^{\gamma^{(t_2)}}}{10^{\gamma^{(t_1)}}} x_{S_F \cup S_{me}}^{(t_1)}$.

Since the rounding function is mantissa-based, the system evolution from $x^{(t_1)}$ is equivalent to the systems evolution from $x^{(t_2)} = 10^\gamma x^{(t_1)}$, where γ is the growth rate, $\gamma^{(t_2)} - \gamma^{(t_1)}$. \square

We can in fact decide whether $x_{S_F}^{(t)} \approx_\beta x_{S_{me}}^{(t)}$ for the last time:

Lemma 8. *Let β, N be defined as in Lemma 6. If $t \geq N$ then it is decidable whether there exists $t' > t$ such that $x_{S_F}^{(t')} \approx_\beta x_{S_{me}}^{(t')}$.*

Proof Sketch (Full proof available in [19]). If we considered S_{me} in isolation, without the effect of S_F , we know it would be pseudo-periodic. We can simulate one period of S_{me} with and without the effect of S_F and determine if S_F influences S_{me} within one period. If it does then they must be close at this point. If S_F does not influence S_{me} we know that S_{me} will behave pseudo-periodically at least until S_F is close to S_{me} again; having established a growth rate for S_{me} , we can compare the growth rates of S_F and S_{me} to see if S_{me} will ever be close to S_F again in the future. \square

Finally to conclude the proof of Theorem 4, we refine Observation 1 to show that the period is bounded and thus the growth rates are computable:

- either S_F is β -close to S_{me} infinitely often, in particular if they become close $F + 1$ times then by Lemma 7 it is pseudo-periodic.
- or the system is pseudo-periodic because it behaves like a top-SCC, in which Lemma 4 gives effective computation of the constants.

Which of these occurs is determined by at most $F + 1$ applications of Lemma 8.

5 Decidability of model checking

In this section we use the results obtained in the previous section to show that model checking is decidable. We use pseudo-periodicity to show that the characteristic word is eventually periodic, a case for which model checking is decidable.

Theorem 2. *Let (M, x) be a non-negative linear dynamical system, let Y_1, \dots, Y_k be semialgebraic targets and let ϕ be an MSO formula using predicates over Y_1, \dots, Y_k . It is decidable whether the characteristic word under floating-point rounding satisfies ϕ .*

Consider a semialgebraic target Y , which can be expressed as a Boolean combination of polynomial inequalities over variables representing the dimensions. That is $Y = \{(x_1, \dots, x_d) \mid \bigwedge_i \bigvee_j P_{ij}(x_1, \dots, x_n) \triangleright_{ij} 0\}$, where $\triangleright_{ij} \in \{\geq, >, =\}$.

Given a linear dynamical system (M, x) defining the rounded orbit $(x^{(n)})_{n=1}^\infty$, recall that $\mathcal{Z}(Y) = \{n \mid x^{(n)} \in Y\}$ are the hitting times of Y . We claim that this set is semi-linear (equivalently eventually periodic) for semialgebraic Y .

Definition 7. *A 1-dimensional linear-set, defined by a base $b \in \mathbb{N}$ and period $p \in \mathbb{N}$, is the set $\{x \mid \exists k \in \mathbb{N} : x = b + k \cdot p\}$. A semi-linear set is the finite union of a finite set $F \subseteq \mathbb{N}$ and linear sets. It can be assumed that each linear-set has the same period. Hence a 1-dimensional semi-linear set X is defined by a finite set $F \subseteq \mathbb{N}$ and integers $m, p, b_1, \dots, b_m \in \mathbb{N}$ such that $x \in X$ if and only if $x \in F$ or $x = b + k \cdot p$ for some $k \in \mathbb{N}$ and $b \in \{b_1, \dots, b_m\}$.*

Theorem 5. *Let Y be a semialgebraic target, $\mathcal{Z}(Y)$ is a semi-linear set.*

Theorem 5 essentially completes the proof of Theorem 2. It is almost immediate that the characteristic word is eventually periodic (see the long version [19] for a formal proof) and thus the model-checking problem can be decided by checking $A \cap \bar{B} = \emptyset$, where A is an automaton representing the characteristic word and B encodes the language of ϕ .

It is standard that semi-linear sets are closed under intersection, union, and complementation (see [15] for a nice introduction to semi-linear sets). Thus in order to express the hitting times of $\mathcal{Z}(Y)$ it is sufficient to express the hitting times of $\{(x_1, \dots, x_d) \mid P(x_1, \dots, x_n) \geq 0\}$ for a finitely many polynomials P . Conjunction is found by taking the intersection of the hitting times, and disjunction by taking union. The hitting times of $P(x_1, \dots, x_n) > 0$ can be rewritten as the complement of the hitting times of $-P(x_1, \dots, x_n) \geq 0$. The hitting times of $P(x_1, \dots, x_n) = 0$ is the conjunction (intersection) of $P(x_1, \dots, x_n) \geq 0$ and $-P(x_1, \dots, x_n) \geq 0$. Thus Theorem 5 is a consequence of the following lemma.

Lemma 9. *Assume $x^{(t)} = (z_1^{(t)}, \dots, z_d^{(t)})_{i=1}^\infty$, is a pseudo-periodic sequence with start point N , period T and growth rates $\alpha_1, \dots, \alpha_n$ and $P \in \mathbb{Q}[x_1, \dots, x_d]$ a rational polynomial in d variables.⁹ Then, $\{i \in \mathbb{N} \mid P(z_1^{(t)}, \dots, z_d^{(t)}) \geq 0\}$ is a semi-linear set.*

⁹ Some variables may be redundant, that is, if the polynomial does not depend on all dimensions of $x^{(t)}$ then some of the variables may not appear in P .

Proof. First, we show that pseudo-periodicity is closed under product. Suppose $x_i^{(N+Tn)} = m_i 10^{\beta_i + \alpha_i \cdot n}$ and $x_j^{(N+Tn)} = m_j 10^{\beta_j + \alpha_j \cdot n}$. Observe that $x_i^{(N+Tn)} \cdot x_j^{(N+Tn)} = m_i \cdot 10^{\beta_i + \alpha_i n} m_j \cdot 10^{\beta_j + \alpha_j n} = m_i m_j \cdot 10^{\beta_i + \beta_j + n(\alpha_i + \alpha_j)}$. We conclude that the vector $(x_i \cdot x_j)^{(t)}$ is pseudo-periodic with growth rate $\alpha_i + \alpha_j$. Observe that the mantissa precision increase by at most 2.

Secondly, we show that if two pseudo-periodic sequences have the same growth rate, then their sum is also pseudo-periodic with the same growth rate. Suppose $x_i^{(N+Tn)} = m_i 10^{\beta_i + \alpha \cdot n}$, and $x_j^{(N+Tn)} = m_j 10^{\beta_j + \alpha \cdot n}$. Observe that $(x_i + x_j)^{(N+Tn)} = m_i 10^{\beta_i + \alpha \cdot n} + m_j 10^{\beta_j + \alpha \cdot n} = (m_i + m_j \cdot 10^{\beta_j - \beta_i}) 10^{\beta_i + \alpha \cdot n}$. Observe that the mantissa precision increased by at most $10^{|\beta_j - \beta_i|}$.

Let $P(x_1, \dots, x_n) = \sum_{i=1}^N c_i Z_i$, where Z_i is a product of x_1, \dots, x_n . Consider each monomial Z_i occurring in P , since product preserves pseudo-periodicity, we conclude that Z_i is pseudo-periodic. $P^{(t)}$ is thus a linear combination of these pseudo-periodic vectors. Note our prior observation does not immediately imply that $P^{(t)}$ is pseudo-periodic as we required taking the sum of elements with the same growth rate. However, from some point on, we are only interested in those with the maximal growth rate.

Without loss of generality, let Z_1, \dots, Z_r have the maximum-growth rate, and Z_{r+1}, \dots, Z_N have strictly smaller growth rate. For every $L \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $t > N$, $\text{exponent}(Z_1^{(t)}) - \text{exponent}(Z_{r+1}^{(t)}) > L$.

Hence there exists $N \in \mathbb{N}$ such that for all $t > N$ if $\sum_{i=1}^r c_i Z_i > 0$ if and only if $\sum_{i=1}^N c_i Z_i = \sum_{i=1}^r c_i Z_i + \sum_{i=r+1}^N c_i Z_i > 0$ because $\left| \sum_{i=r+1}^N c_i Z_i \right| < \left| \sum_{i=1}^r c_i Z_i \right|$ from some point on. Hence $\text{sign}(\sum_{i=1}^N c_i Z_i^{(t)}) = \text{sign}(\sum_{i=1}^r c_i Z_i^{(t)})$.

Thus we restrict our attention to $\sum_{i=1}^r c_i Z_i^{(t)}$. Since each of the Z_i for $i \in \{1, \dots, r\}$ have the same growth rate, we know that $\sum_{i=1}^r c_i Z_i^{(t)}$ is pseudo-periodic. Since $\text{sign}(\sum_{i=1}^r c_i Z_i^{(t)})$ does not depend on the exponent, only the periodic mantissa, we have that the sign is periodic. The hitting times for $t \leq N$ can be determined exhaustively and included in the finite set of the semi-linear set. \square

Acknowledgements Partially funded by DFG grant 389792660 as part of TRR 248 – CPEC, see `perspicuous-computing.science`. Joël Ouaknine is also affiliated with Keble College, Oxford as `emmy.network` Fellow. David Purser was partially supported by the ERC grant INFSYS, agreement no. 950398.

References

1. Abbasi, R., Schiffel, J., Darulova, E., Ulbrich, M., Ahrendt, W.: Deductive verification of floating-point java programs in key. In: Groote, J.F., Larsen, K.G. (eds.) Tools and Algorithms for the Construction and Analysis of Systems - 27th International Conference, TACAS 2021, Part of ETAPS 2021. Part II. Lecture Notes in Computer Science, vol. 12652, pp. 242–261. Springer (2021). https://doi.org/10.1007/978-3-030-72013-1_13

2. Akshay, S., Antonopoulos, T., Ouaknine, J., Worrell, J.: Reachability problems for Markov chains. *Inf. Process. Lett.* **115**(2), 155–158 (2015). <https://doi.org/10.1016/j.ipl.2014.08.013>
3. Akshay, S., Bazille, H., Genest, B., Vahanwala, M.: On robustness for the Skolem and Positivity problems. In: Berenbrink, P., Monmege, B. (eds.) 39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022. LIPIcs, vol. 219, pp. 5:1–5:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022). <https://doi.org/10.4230/LIPIcs.STACS.2022.5>
4. Almagor, S., Karimov, T., Kelmendi, E., Ouaknine, J., Worrell, J.: Deciding ω -regular properties on linear recurrence sequences. *Proc. ACM Program. Lang.* **5**(POPL), 1–24 (2021). <https://doi.org/10.1145/3434329>
5. Baier, C., Funke, F., Jantsch, S., Karimov, T., Lefauchaux, E., Ouaknine, J., Pouly, A., Purser, D., Whiteland, M.A.: Reachability in dynamical systems with rounding. In: 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2020. LIPIcs, vol. 182, pp. 36:1–36:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). <https://doi.org/10.4230/LIPIcs.FSTTCS.2020.36>
6. Baier, C., Funke, F., Jantsch, S., Karimov, T., Lefauchaux, E., Ouaknine, J., Purser, D., Whiteland, M.A., Worrell, J.: Parameter Synthesis for Parametric Probabilistic Dynamical Systems and Prefix-Independent Specifications. In: Klin, B., Lasota, S., Muscholl, A. (eds.) 33rd International Conference on Concurrency Theory (CONCUR 2022). Leibniz International Proceedings in Informatics (LIPIcs), vol. 243, pp. 10:1–10:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2022). <https://doi.org/10.4230/LIPIcs.CONCUR.2022.10>
7. Becker, H., Panchekha, P., Darulova, E., Tatlock, Z.: Combining tools for optimization and analysis of floating-point computations. In: Havelund, K., Peleska, J., Roscoe, B., de Vink, E.P. (eds.) Formal Methods - 22nd International Symposium, FM 2018, Held as Part of the Federated Logic Conference, FloC 2018. Lecture Notes in Computer Science, vol. 10951, pp. 355–363. Springer (2018). https://doi.org/10.1007/978-3-319-95582-7_21
8. Bilu, Y., Luca, F., Nieuwveld, J., Ouaknine, J., Purser, D., Worrell, J.: Skolem meets Schanuel. In: Szeider, S., Ganian, R., Silva, A. (eds.) 47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022. LIPIcs, vol. 241, pp. 20:1–20:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022). <https://doi.org/10.4230/LIPIcs.MFCS.2022.20>
9. Boyle, M.: Notes on the Perron-Frobenius theory of nonnegative matrices (2005)
10. Braverman, M.: Termination of integer linear programs. In: Ball, T., Jones, R.B. (eds.) Computer Aided Verification, 18th International Conference, CAV 2006 Proceedings. Lecture Notes in Computer Science, vol. 4144, pp. 372–385. Springer (2006). https://doi.org/10.1007/11817963_34
11. Büchi, J.R.: On a decision method in restricted second order arithmetic. In: The collected works of J. Richard Büchi, pp. 425–435. Springer (1990)
12. Chonev, V., Ouaknine, J., Worrell, J.: On the complexity of the orbit problem. *J. ACM* **63**(3), 23:1–23:18 (2016). <https://doi.org/10.1145/2857050>
13. D’Costa, J., Karimov, T., Majumdar, R., Ouaknine, J., Salamati, M., Soudjani, S., Worrell, J.: The pseudo-Skolem problem is decidable. In: Bonchi, F., Puglisi, S.J. (eds.) 46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021. LIPIcs, vol. 202, pp. 34:1–34:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021). <https://doi.org/10.4230/LIPIcs.MFCS.2021.34>

14. D’Costa, J., Karimov, T., Majumdar, R., Ouaknine, J., Salamati, M., Worrell, J.: The pseudo-reachability problem for diagonalisable linear dynamical systems. In: Szeider, S., Ganian, R., Silva, A. (eds.) 47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022. LIPIcs, vol. 241, pp. 40:1–40:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022). <https://doi.org/10.4230/LIPIcs.MFCS.2022.40>
15. Haase, C.: A survival guide to Presburger arithmetic. *ACM SIGLOG News* **5**(3), 67–82 (2018). <https://doi.org/10.1145/3242953.3242964>
16. Kannan, R., Lipton, R.J.: Polynomial-time algorithm for the orbit problem. *J. ACM* **33**(4), 808–821 (1986). <https://doi.org/10.1145/6490.6496>
17. Karimov, T., Kelmendi, E., Ouaknine, J., Worrell, J.: What’s decidable about discrete linear dynamical systems? In: Raskin, J., Chatterjee, K., Doyen, L., Majumdar, R. (eds.) Principles of Systems Design - Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday. *Lecture Notes in Computer Science*, vol. 13660, pp. 21–38. Springer (2022). https://doi.org/10.1007/978-3-031-22337-2_2
18. Karimov, T., Lefauchaux, E., Ouaknine, J., Purser, D., Varonka, A., Whiteland, M.A., Worrell, J.: What’s decidable about linear loops? *Proc. ACM Program. Lang.* **6**(POPL), 1–25 (2022). <https://doi.org/10.1145/3498727>
19. Lefauchaux, E., Ouaknine, J., Purser, D., Sharifi, M.: Model checking linear dynamical systems under floating-point rounding. *CoRR* **abs/2211.04301** (2022). <https://doi.org/10.48550/arXiv.2211.04301>
20. Lohar, D., Jeangoudoux, C., Sobel, J., Darulova, E., Christakis, M.: A two-phase approach for conditional floating-point verification. In: Groote, J.F., Larsen, K.G. (eds.) Tools and Algorithms for the Construction and Analysis of Systems - 27th International Conference, TACAS 2021, Part of ETAPS 2021. Part II. *Lecture Notes in Computer Science*, vol. 12652, pp. 43–63. Springer (2021). https://doi.org/10.1007/978-3-030-72013-1_3
21. Luca, F., Ouaknine, J., Worrell, J.: Algebraic model checking for discrete linear dynamical systems. In: Bogomolov, S., Parker, D. (eds.) Formal Modeling and Analysis of Timed Systems - 20th International Conference, FORMATS 2022. *Lecture Notes in Computer Science*, vol. 13465, pp. 3–15. Springer (2022). https://doi.org/10.1007/978-3-031-15839-1_1
22. Maurica, F., Mesnard, F., Payet, E.: Optimal approximation for efficient termination analysis of floating-point loops. In: 2017 1st International Conference on Next Generation Computing Applications (NextComp). pp. 17–22. IEEE (2017)
23. Minsky, M.L.: *Computation*. Prentice-Hall Englewood Cliffs (1967)
24. Ouaknine, J., Worrell, J.: Positivity problems for low-order linear recurrence sequences. In: Chekuri, C. (ed.) Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014. pp. 366–379. SIAM (2014). <https://doi.org/10.1137/1.9781611973402.27>
25. Schneider, H.: Wielandt’s proof of the exponent inequality for primitive nonnegative matrices. *Linear Algebra and its Applications* **353**(1), 5–10 (2002)
26. Tiwari, A.: Termination of linear programs. In: Alur, R., Peled, D.A. (eds.) Computer Aided Verification, 16th International Conference, CAV 2004 Proceedings. *Lecture Notes in Computer Science*, vol. 3114, pp. 70–82. Springer (2004). https://doi.org/10.1007/978-3-540-27813-9_6
27. Xia, B., Yang, L., Zhan, N., Zhang, Z.: Symbolic decision procedure for termination of linear programs. *Formal Aspects Comput.* **23**(2), 171–190 (2011). <https://doi.org/10.1007/s00165-009-0144-5>