


Reachability in Dynamical Systems with Rounding

Christel Baier 

Technische Universität Dresden, Germany

Florian Funke 

Technische Universität Dresden, Germany

Simon Jantsch 

Technische Universität Dresden, Germany

Toghrul Karimov

Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany

Engel Lefauchaux

Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany

Joël Ouaknine 

Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany
Department of Computer Science, Oxford University, UK

Amaury Pouly 

Université de Paris, CNRS, IRIF, F-75006, Paris, France

David Purser 

Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany

Markus A. Whiteland 

Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany

Abstract

We consider reachability in dynamical systems with discrete linear updates, but with fixed digital precision, i.e., such that values of the system are rounded at each step. Given a matrix $M \in \mathbb{Q}^{d \times d}$, an initial vector $x \in \mathbb{Q}^d$, a granularity $g \in \mathbb{Q}_+$ and a rounding operation $[\cdot]$ projecting a vector of \mathbb{Q}^d onto another vector whose every entry is a multiple of g , we are interested in the behaviour of the orbit $\mathcal{O} = \langle [x], [M[x]], [M[M[x]]], \dots \rangle$, i.e., the trajectory of a linear dynamical system in which the state is rounded after each step. For arbitrary rounding functions with bounded effect, we show that the complexity of deciding point-to-point reachability—whether a given target $y \in \mathbb{Q}^d$ belongs to \mathcal{O} —is **PSPACE**-complete for hyperbolic systems (when no eigenvalue of M has modulus one). We also establish decidability without any restrictions on eigenvalues for several natural classes of rounding functions.

2012 ACM Subject Classification Theory of computation

Keywords and phrases dynamical systems, rounding, reachability

Related Version A full version of the paper is available at <http://arxiv.org/abs/2009.13353>.

Funding This work was funded by DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>), the Cluster of Excellence EXC 2050/1 (CeTI, project ID 390696704, as part of Germany's Excellence Strategy), DFG-projects BA-1679/11-1 and BA-1679/12-1, and the Research Training Group QuantLA (GRK 1763).

Joël Ouaknine: Supported by ERC grant AVS-ISS (648701).

Amaury Pouly: Supported by CODYS project ANR-18-CE40-0007.

1 Introduction

A *discrete-time linear dynamical system* in ambient space \mathbb{Q}^d is specified via a linear transformation together with a starting point. The state of the system is then updated at each step by applying the linear transformation, giving rise to an *orbit* (or infinite trajectory) in \mathbb{Q}^d .

One of the most well-known questions for such systems is the *Skolem Problem*, which asks whether the orbit ever hits a given $(d - 1)$ -dimensional hyperplane.¹ This problem has long eluded decidability, although instances of dimension $d \leq 4$ are known to be solvable (see, e.g., the survey [30]). Another natural problem is *point-to-point reachability*², known to be decidable in polynomial time [22]. In both cases, however, one assumes *arbitrary* precision, which arguably is unrealistic for simulations carried out on digital computers. In this paper, we therefore turn our attention to instances of these problems in which the numerical state of the system is rounded to finite precision at each time step. This leads us to the following definition:

► **Problem (Rounded Point-to-Point Reachability (Rounded P2P)).** *Given a matrix $M \in \mathbb{Q}^{d \times d}$, an initial vector $x \in \mathbb{Q}^d$, a target vector $y \in \mathbb{Q}^d$, a granularity $g \in \mathbb{Q}_+$, and a rounding operation $[\cdot]$ projecting a vector of \mathbb{Q}^d onto another vector whose every entry is a multiple of g , let the orbit \mathcal{O} of this system be the infinite sequence $\langle [x], [M[x]], [M[M[x]]], \dots \rangle$, i.e., $x^{(0)} = [x]$ and $x^{(i+1)} = [Mx^{(i)}]$. The **Rounded Point-to-Point Reachability (Rounded P2P) Problem** asks whether $[y] \in \mathcal{O}$.*

Main contributions. We make the following contributions, summarised in Figure 1:

1. We introduce a family of natural problems, Rounded P2P (parameterised by the rounding function), which to the best of our knowledge has not previously been studied.
2. We show that for hyperbolic systems (i.e., those whose associated linear transformation has no eigenvalue of modulus 1) the Rounded P2P Problem is solvable—and is in fact **PSPACE**-complete—for any ‘reasonable’ (i.e., bounded-effect) rounding function. It is interesting to note, in contrast, that exact P2P reachability is known to be solvable in polynomial time. Our approach to solving the Rounded P2P Problem relies on the observation that, outside a ball of exponential size, the change in magnitude of the system state at each step dwarfs any effect due to rounding. It thus suffices to exhaustively examine the effect of the dynamics inside an exponentially bounded state space.
3. In the general case (without any restriction on the magnitude of eigenvalues), the effect of rounding may forever remain non-negligible, requiring a careful analysis. We have not been able to solve the problem in full generality, but we do provide a complete solution for certain natural classes of rounding functions. More precisely, assume that the linear transformation has been converted to Jordan normal form (now requiring us to work with complex algebraic numbers). We can then solve the Rounded P2P Problem under two natural classes of rounding functions:
 - (a) *Polar rounding functions:* given a complex number of the form $Ae^{i\theta}$, such functions round A and θ independently. In such instances we can handle in **EXSPACE** all

¹ The Skolem Problem is usually formulated in terms of linear recurrence sequences, but is equivalent to the description given here.

² Historically this problem has been known as the *orbit problem*, however there are now multiple ‘orbit problems’ (polytope reachability, hyperplane reachability, (semi-)algebraic set reachability,... etc.) and so we specify point-to-point reachability.

Rounding type	Hyperbolic Systems	No restrictions on eigenvalues	
		Jordan normal form (Note: no hardness)	General
Polar $Ae^{i\theta}$	PSPACE -complete, Section 3	EXPSPACE , Section 4.1	Open but PSPACE -hard
Argand truncation or expansion		EXPSPACE , Section 4.2	
Argand minimal error		Open (difficulties highlighted in Section 5)	
Arbitrary bounded-effect		Open (Open Problem 21)	

■ **Figure 1** Decidability and complexity table for the Rounded P2P Problem.

reasonable rounding functions on A , and what we view as the only natural rounding function on θ .

- (b) *Argand rounding*: given a complex number of the form $a + bi$, the *Argand truncation* will round a and b independently downwards (in magnitude), ensuring that the modulus never increases. Similarly, the *Argand expansion* (which rounds a and b independently upwards) guarantees that the modulus can only increase. Under such rounding functions, we show decidability in **EXPSPACE**.
4. We highlight some limitations of our methods, identifying a simple but technically challenging open problem, which points to some of the key difficulties in solving the Rounded P2P Problem in full generality. More precisely, we consider minimal error rounding for a simple rotation in two-dimensional space, for which Rounded P2P is presently open.

► **Remark.** It is worth noting that the rounded versions of the Skolem Problem (does the rounded orbit ever hit a $(d - 1)$ -dimensional hyperplane?) and the Positivity Problem (does the rounded orbit ever hit a d -dimensional half-space?) remain at least as hard as their exact integer counterparts, since over the integers rounding has no effect; the decidability of these problems therefore remains open. However, the rounded versions of reaching a bounded polytope or a bounded semialgebraic set (problems not known to be decidable in the exact setting [14, 4]) reduce to a finite number of Rounded P2P reachability queries (since a bounded set can contain only finitely many rounded points). These observations together motivate our focus, in the present paper, on the Rounded P2P Problem.

It is interesting to consider rounded reachability problems in the stochastic setting, i.e., Markov chains. One observes that the state space $[\mathcal{S}] = \{[x] \in [0, 1]^d \mid x \text{ sub-stochastic}\}$ is finite, which entails decidability of virtually any reachability problem, including Skolem and Positivity. This is somewhat arresting, since without rounding reachability problems are known to be exactly as hard for stochastic systems as for general systems [3]. In any event, one should note that ensuring that for all $x \in [\mathcal{S}]$, $[Mx] \in [\mathcal{S}]$ requires some care, as arbitrary rounding does not necessarily preserve (sub-)stochasticity.

Related work

With the emerging use of numerical computations during the 80s, doubts were raised concerning the transferability of results about dynamical systems obtained by simulation in finite-state machines. In this direction, the sensitivity that a rounding function may have on the long-term behaviour of a dynamical system is studied in [5]. How rounded orbits can be simulated by actual orbits of the dynamical system is investigated in [20, 29].

The series of papers [6, 7, 8, 9] examines which statistical properties of a discrete dynamical system are preserved under the introduction of a rounding function, a good summary of which can be found in Blank’s book [10, Chapter 5]. As the rounding is refined, some properties of the discretized orbits follow probabilistic laws asymptotically, as shown in [16, 17]. The paper [18] studies how volatile statistical notions are in the presence of finite precision (such as the mean distance of two orbits of discrete dynamical systems).

Another line of research focuses on discretized rotations in \mathbb{Z}^2 and higher-dimensional lattices [24, 1]. A connection from roundoff problems in the 2-dimensional case to expanding maps on the p -adic integers is described in [11, 36]. Building on this, [35] conjectures periodicity of all orbits of these discretized rotations in \mathbb{Z}^2 . It is shown in [2] that there are infinitely many periodic orbits, and [31] attempts to concisely describe points leading to periodic orbits.

In the context of model checking, continuous dynamical systems have been translated into discrete models, mainly timed automata that approximate the behaviour of the original system [26, 13, 32]. On a more general level, one can observe a growing interest in the systematic study of roundoff errors inherent in finite precision computations [19, 33, 21, 25, 27, 15].

2 Rounding functions

Let $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{A}$ be the naturals, integers, rationals, reals, and algebraic numbers respectively.

Rounding real numbers

Let $g \in \mathbb{R}_+$ be a granularity. We define our rounding functions taking values to integers, i.e., $g = 1$. For $g \neq 1$ we consider $[x] = g \cdot \lceil x/g \rceil$. Given a set S , we let $[S] = \{[x] \mid x \in S\}$.

The floor function $\lfloor x \rfloor$ and ceiling functions $\lceil x \rceil$ are well-known rounding functions in mathematics and computer science. We recall two further rounding functions:

- *Minimal error rounding* rounds to the *nearest* value: $[x] = \arg \min_{y \in \mathbb{Z}} |x - y|$. If $|x - y| = 0.5$ an arbitrary but deterministic choice must be made (e.g. to round up).
- *Truncation* (‘towards zero rounding’, to cut off the remaining bits): if $x > 0$ then $\lfloor x \rfloor$ else $\lceil x \rceil$, or *expansion*: if $x > 0$ then $\lceil x \rceil$ else $\lfloor x \rfloor$.

Whenever possible, we prefer to analyse the problems without choosing a specific rounding function, relying only upon the property of *bounded effect*:

► **Definition 1.** *A real rounding function $[\cdot]: \mathbb{R} \rightarrow \mathbb{R}$ has bounded effect if there exists Δ such that $|x - [x]| \leq \Delta$ for all x .*

Rounding complex numbers

Complex numbers have both a real and imaginary part. Thus one can consider rounding each of the components separately, which we call *Argand rounding*. Consider $x = a + bi$ with $a, b \in \mathbb{R}$, then let $[x] = [a] + [b]i$, where $[\cdot]$ can be any real rounding function (leading to *Argand truncation*, *Argand expansion* and *Argand minimal error rounding* functions).

However, complex numbers can also be readily represented using polar coordinates as follows: a number is represented as $x = Ae^{i\theta}$, where A is the modulus and θ is the angle between the 2-d coordinates $(1, 0)$ and (a, b) (when represented as $a + bi$). Then, a *polar rounding* function rounds A and θ independently, i.e. $[x] = [A]e^{i[\theta]}$. The rounding of $[A]$ can be any real rounding function. For the rounding of the angle we always assume minimal error rounding. That is, given granularity $\theta_g = \frac{\pi}{R}$ for some $R \in \mathbb{N}$, then $[\theta]$ is a multiple of θ_g with minimal error and arbitrary but deterministic tie breaking.

We generalise non-specific bounded-effect rounding to the complex numbers.

► **Definition 2.** *A complex rounding function $[\cdot]: \mathbb{C} \rightarrow \mathbb{C}$ has bounded effect **on the modulus** if there exists Δ such that $\|x - [x]\| \leq \Delta$ for all x .*

Argand and polar roundings are both defined by applying bounded-effect real rounding functions to each component, and have bounded effect under Definition 2. However, note the distinction with Definition 1; polar rounding can exhibit arbitrary large effects (in the following sense: given any $\Delta > 0$, one can always find $x \in \mathbb{C}$ such that $|x - [x]| > \Delta$), but nevertheless has only bounded effect *on the modulus*.

► **Definition 3** ($[K]$ -Ball). *Given a complex rounding function $[\cdot]$ and an integer K let a $[K]$ -ball be the set of admissible points of modulus at most K , i.e., $\{[x] \mid x \in \mathbb{C}, |[x]| \leq K\}$.*

Rounding vectors

In general, a rounding function on \mathbb{K} induces a rounding function on vectors \mathbb{K}^d , where $[(x_1, \dots, x_d)] = ([x_1], \dots, [x_d])$, although not all rounding functions on vectors need take this form. We generalise non-specific bounded-effect rounding to vectors.

► **Definition 4.** *A rounding function $[\cdot]: \mathbb{K}^d \rightarrow \mathbb{K}^d$ has bounded effect on the modulus if there exists Δ such that $\|x|_k - |[x]|_k\| \leq \Delta$ for all x and every $k \in \{1, 2, \dots, d\}$.*

Finally, we assume that all of our rounding functions can be computed in polynomial time and are fixed (rather than inputs) in our problems, and thus Δ is also a fixed parameter.

3 Hyperbolic systems

In this section we establish our first main result for hyperbolic systems, which we first define:

► **Definition 5** (Hyperbolic System [23, Section 1.2]). *A linear map represented by the matrix $M \in \mathbb{R}^{d \times d}$ is hyperbolic if all of its eigenvalues have modulus different from one.*

► **Theorem 6.** *The Rounded P2P Problem is **PSPACE**-complete for hyperbolic linear maps represented by rational matrices and real rounding functions with bounded effect.*

We first demonstrate that the problem is in **PSPACE** for matrices in Jordan normal form, to which we will reduce the general case in a second step. As the passage to Jordan normal form inevitably introduces complex numbers, **PSPACE** membership will be shown for Jordan normal form matrices over the algebraic numbers and, accordingly, complex rounding functions with bounded effect on the modulus. To complete the picture we show hardness for hyperbolic systems (in fact, the hardness result applies even for non-hyperbolic systems, that is for matrices whose eigenvalues may include 1).

3.1 Membership in PSPACE

We now prove the membership part of Theorem 6 under the additional assumption that the matrices are in Jordan normal form.

► **Lemma 7.** *The Rounded P2P Problem decidable in **PSPACE** for any complex rounding function with bounded effect on the modulus Δ and hyperbolic matrices $M \in \mathbb{A}^{d \times d}$ in Jordan normal form.*

6 Reachability in Dynamical Systems with Rounding

Proof. We consider a single Jordan block of dimension d with eigenvalue λ . If the matrix M has multiple Jordan blocks, the algorithm can be run in lock step³ for each block. Hence, without loss of generality we let

$$M = \begin{bmatrix} \lambda & 1 & & \\ & \lambda & 1 & \\ & & \ddots & \ddots \\ & & & \lambda \end{bmatrix}.$$

The idea will be to show that for $|\lambda| > 1$, for values large enough growth will outstrip the rounding, and the orbit will grow beyond the target, never to return. If $|\lambda| < 1$ and the orbit gets large enough, it will begin to contract again, so we choose a ball large enough to contain the whole orbit. We do not consider the case $|\lambda| = 1$ here.

Formally, in each dimension $k \in \{1, \dots, d\}$ we compute a radius C_k , defining a $[C_k]$ -ball of radius C_k about 0, containing x_k and y_k such that for all z in the orbit \mathcal{O} if $z_k \notin [C_k]$ -ball then $[Mz]_k \notin [C_k]$ -ball. That is, if the orbit has left the ball, it will never come back. The algorithm proceeds by simulating the orbit from x until one of the following occurs.

- y is found, in which case RETURN YES, or
- a point repeats, in which case RETURN NO, or
- a point $x^{(i)}$ is found such that $|(x^{(i)})_k| \geq C_k$ for some k , in which case RETURN NO.

Since $B = [\{x \in \mathbb{R}^d \mid \text{for all } k \ |x_k| \leq C_k\}]$ is finite, one of the three must occur. Remembering all previous points would require too much space. Therefore we record a counter of the number of steps taken and once this exceeds the maximum number of points then we know some point must have been repeated (possibly many times by this point). Let $C = \max_i C_k$, then the bounding hyper-cube of B has $(2C/g)^d$ points, hence B has fewer points. We show this number has at most exponential size in the description length of the input, and hence can be represented in **PSPACE**.

► **Case 1 (suppose $|\lambda| > 1$).** For the d th component we have $(x^{(i+1)})_d = [\lambda(x^{(i)})_d]$. There is a bounded effect of the rounding Δ , ensuring $|(x^{(i+1)})_d| \geq |\lambda| |(x^{(i)})_d| - \Delta$. So when $|\lambda| |(x^{(i)})_d| - \Delta > |(x^{(i)})_d|$, this component must grow. Let $\ell = \max\{1, \Delta, |y_1|, \dots, |y_d|\}$. We define the radius $C_d := \frac{\Delta}{|\lambda|-1} + \ell$, which satisfies the desired property described above.

Now suppose that the radius C_k is defined so that $C_k \leq \ell \sum_{j=0}^{d-k+1} (\frac{2}{|\lambda|-1})^j$ (holds for $k = d$) and assume that $|(x^{(i)})_j| \leq C_j$ for each $j \in \{k, \dots, d\}$. For the $k-1$ th dimension the update is of the form $(x^{(i+1)})_{k-1} = [\lambda(x^{(i)})_{k-1} + 1(x^{(i)})_k]$. Since $|(x^{(i)})_k| \leq C_k$, we have $|(x^{(i+1)})_{k-1}| \geq |\lambda| |(x^{(i)})_{k-1}| - \Delta - C_k$, and there is growth when $|\lambda| |(x^{(i)})_{k-1}| - \Delta - C_k > |(x^{(i)})_{k-1}|$, i.e., when $|(x^{(i)})_{k-1}| > \frac{\Delta + C_k}{|\lambda|-1}$. So, we may define $C_{k-1} := \frac{\Delta + C_k}{|\lambda|-1} + \ell$, which satisfies the property described above, and moreover, $C_{k-1} \leq \frac{2C_k}{|\lambda|-1} + \ell \leq \ell \sum_{j=0}^{d-(k-1)+1} (\frac{2}{|\lambda|-1})^j$ due to our choice of ℓ . Repeat for all remaining components $k-2, \dots, 1$.

Now $C_k \leq \ell \sum_{j=0}^d (\frac{2}{|\lambda|-1})^j \leq \ell(d+1)(1 + (\frac{2}{|\lambda|-1})^d)$ for each k , and the claim follows.

► **Case 2 (suppose $|\lambda| < 1$).** We require the ball to have the property that if the orbit leaves, it will never come back. However for $|\lambda| < 1$, while initially there may be some growth (due to other components), once large enough $|\lambda|$ will dominate and the modulus will decrease. Therefore, we want to ensure we choose the ball large enough that the orbit will never leave the ball in the first place. The following definitions of the radii C_j can easily be altered to furnish this requirement.

³ By running processes in lock step, here and elsewhere, we mean running all of the processes simultaneously (interleaving instructions for each process) until either $x^{(i)} = y$ or one of the processes concludes non-reachability.

Consider the last component d : we have $|(x^{(i+1)})_d| \leq |\lambda| |(x^{(i)})_d| + \Delta$. Set again $\ell = \max\{1, \Delta, |y_1|, \dots, |y_d|\}$ and define $C_d := \frac{\Delta}{1-|\lambda|} + \ell$; if $|(x^{(i)})_d| \leq C_d$, then $|(x^{(i+1)})_d| \leq C_d$.

Having fixed $C_{k'}$ for $k' \in \{k, \dots, d\}$, consider component $k-1$: We have $(x^{(i+1)})_{k-1} = [\lambda(x^{(i)})_{k-1} + (x^{(i)})_k]$, and so $|(x^{(i+1)})_{k-1}| \leq |\lambda| |(x^{(i)})_{k-1}| + |(x^{(i)})_k| + \Delta$. Let us define $C_{k-1} := \frac{C_k + \Delta}{1-|\lambda|} + \ell$. Now if $|(x^{(i)})_{k-1}| \leq C_{k-1}$ then $|(x^{(i+1)})_{k-1}| \leq C_{k-1}$. Repeat for each remaining component. It can be shown, similar to the previous case, that $C_k \leq \ell(d+1)(1 + (\frac{2}{1-|\lambda|})^d)$ for each k , and this concludes the proof. \blacktriangleleft

Reducing the general form to Jordan normal form

In the previous section we assumed that the matrix is always in Jordan normal form, which is a significant restriction. In this section we will not assume Jordan normal form, which means we cannot make any assumption about the rounding, other than being of bounded effect, to prove Theorem 6. After a change of basis properties such as ‘rounding towards zero’ may not be preserved.

Proof (upper bound of Theorem 6). Let Δ be the fixed, bounded effect on the modulus of $[\cdot]$. Let us consider hyperbolic $M = PJP^{-1} \in \mathbb{Q}^{d \times d}$. We ask whether $x^{(i+1)} = y$ for some i . Observe that $x^{(i+1)} = [Mx^{(i)}] = Mx^{(i)} + e(Mx^{(i)})$ where $e(x) := [x] - x \in [-\Delta, \Delta]^d$ for any x since $[\cdot]$ has bounded effect. Now if we define $z^{(i)} := P^{-1}x^{(i)}$ we have that

$$z^{(i+1)} = P^{-1}x^{(i+1)} = P^{-1}(Mx^{(i)} + e(Mx^{(i)})) = Jz^{(i)} + P^{-1}e(PJz^{(i)}) = (\lceil Jz^{(i)} \rceil)$$

where $(\lceil z \rceil) := z + P^{-1}e(Pz)$ for any z . The question $x^{(i)} \stackrel{?}{=} y$ for some i now becomes equivalent to $z^{(i)} \stackrel{?}{=} P^{-1}y$. But note that the system for $z^{(i)}$ is in Jordan normal form and the rounding function $(\lceil \cdot \rceil)$ has bounded effect on the modulus, with bound $\Delta' \leq \max_{1 \leq k \leq d} \max_{e \in [-\Delta, \Delta]^d} (P^{-1}e)_k$. Since Δ is fixed and P^{-1} is computable in polynomial time [12], then Δ' is of polynomial size. Hence, we have produced in polynomial time an instance of the Rounded P2P problem with a matrix in Jordan normal form. As the proof of Lemma 7 shows that this problem is solvable in **PSPACE** even if Δ is given as input, we can conclude that the **PSPACE** upper bound holds also for the general case. \blacktriangleleft

3.2 PSPACE-hardness

We will prove **PSPACE**-hardness (i.e., the lower bound of Theorem 6) by reduction from quantified boolean formula (QBF), which is **PSPACE**-complete [34]. We do this by first encoding a simple programming language into the rounded P2P Problem. Then, we show that reachability in this language can solve QBF. Whilst a direct reduction is possible, we provide exposition via the language for two reasons; first, we will show that the language is robust to choice of rounding function (Remark 9), and secondly the reduction results in an instance where all eigenvalues have modulus 1, but by a small perturbation, we observe that the problem remains hard when all of the eigenvalues do not have modulus 1 (Remark 10).

The language will consist of m instructions, operating over d variables. Each instruction is a boolean map $f_i : [0, 1]^d \rightarrow [0, 1]^d$, where each dimension i is updated using a logical formula of the d inputs. Each of the m instructions is conducted in turn and updating the d variables is *simultaneous* in each step. Thus, references to variable in a function are the evaluation in the previous step. Once the m instructions are complete, the system returns to the first instruction and repeats $(x^{(i)} = (f_m \circ f_{m-1} \circ \dots \circ f_2 \circ f_1)(x^{(i-1)}))$, see also Algorithm 1).

An instruction is encoded into the rounded dynamical system using a map $f_i : \mathbb{N}^d \rightarrow \mathbb{N}^d$ for $0 \leq i \leq m-1$, where instructions are of the form $(f_i(x))_j = \lfloor (p_j \cdot x) \rfloor$ where p_j in

■ **Algorithm 1** System behaviour of the language

Input: $x \in [0, 1]^d$ initial vector, $y \in [0, 1]^d$ target vector

while $x \neq y$ **do**

$x \leftarrow f_1(x)$	}	$x_1 \leftarrow x_2 \vee (x_5 \wedge x_3)$
$x \leftarrow f_2(x)$ <i>e.g.</i>		$x_2 \leftarrow \text{if } (x_1 \vee x_3) \text{ then } x_6 \text{ else } x_2$
\vdots		$x_3 \leftarrow \text{true}$
\vdots		\vdots
$x \leftarrow f_m(x)$		$x_d \leftarrow x_4$

end

\mathbb{Q}^d . We demonstrate how to encode the required logical operations in a rounded dynamical system: and ($x_i \leftarrow x_j \wedge x_k = \lfloor \frac{1+x_j+x_k}{3} \rfloor$), or ($x_i \leftarrow x_j \vee x_k = \lfloor \frac{1+x_j+x_k}{2} \rfloor$), negation ($x_i \leftarrow \neg x_j = \lfloor 1 - x_j \rfloor$), resetting a variable to **false** ($x_i \leftarrow \lfloor 0 \rfloor$), copying a variable without change ($x_i \leftarrow \lfloor x_i \rfloor$) or moving/duplicating a variable ($x_i \leftarrow \lfloor x_j \rfloor$). To enable this, we will assume there is always access to the constant 1 (or **true**) by an implicit dimension, fixed to 1.

In multiple steps any logical formula can be evaluated. This can be done with auxiliary variables to store partial computations, where the instructions will in fact be multi-step instructions making use of a finite collection of auxiliary variables which will not be referenced explicitly. Meanwhile any unused variables can be copied without change. In particular the syntax $x_1 \leftarrow \text{if } (x_2) \text{ then } x_3 \text{ else } x_4$ can be encoded, by equivalence with the logical formula $x_1 \leftarrow ((x_2 \implies x_3) \wedge (\neg x_2 \implies x_4))$.

We ask, given some initial configuration $x^{(0)}$, and a target y : does there exist i such that $x^{(i)} = y$. If there was just one step function, the system dynamics would be a direct instance of the rounded orbit semantics. When there are m functions, we remark the sequence of functions can be encoded by taking m copies of each variable, and each function f_i , can transfer the function from one copy to the next, zeroing the previous set of variables. That is, let

$$M = \begin{bmatrix} 0 & & & & f_m \\ f_1 & 0 & & & \\ & f_2 & 0 & & \\ & & & \ddots & \\ & & & & f_{m-1} & 0 \end{bmatrix}.$$

Then the initial configuration becomes $(x^{(0)}, 0, \dots, 0)$, and the target becomes $(y, 0, \dots, 0)$.

An abstraction of the language is depicted in Algorithm 1. It remains to show that QBF can be encoded in the language.

► **Lemma 8.** *Reachability in this language can solve QBF.*

Proof. Formally we write a program in our language to decide the truth of a formula of the form $\forall x_1 \exists x_2 \forall x_3 \dots \exists x_n \psi(x_1, \dots, x_n)$, where ψ is a quantifier free boolean formula. For convenience we assume it starts with \forall , ends with \exists and alternates. Formulae not in this form can be padded if necessary with variables which do not occur in the formula ψ .

The program will have the following variables: $x_1, \dots, x_n, \hat{\psi}, s_1^0, \dots, s_n^0, s_1^1, \dots, s_n^1$ and c_1, \dots, c_n . The bits x_1, \dots, x_n represent the current allocation to the corresponding bit variables of ψ , and $\hat{\psi}$ will store the current evaluation of $\psi(x_1, \dots, x_n)$. To cycle through all allocations to x_1, \dots, x_n , the variables will be treated as a binary number and incremented by one many times, for this purpose the bits c_1, \dots, c_n represent the carry bits when incrementing x_1, \dots, x_n .

The intuition of s_i^z is the following: for fixed x_1, \dots, x_{i-1} it stores the evaluation of $Qx_{i+1} Q'x_{i+2} \dots \exists x_n \psi(x_1, \dots, x_{i-1}, z, x_{i+1}, \dots, x_n)$ where $Q, Q' \in \{\exists, \forall\}$ as required by the formula. Therefore the overall formula is true if and only if $s_1^0 \wedge s_1^1$ is eventually true.

We define $3+n$ instructions, and each run through $f_1 \rightarrow f_{3+n}$ will cover exactly one allocation to x_1, \dots, x_n , with the next run through covering the next allocation that one gets by incrementing the rightmost bit. Once x_{i+1} has been in both the 1 state and the 0 state for all values below, we have enough information to set s_{i+1}^x . This is set when the carry-bit c_{i+1} is one, which indicates that x_{i+1} has visited both 0 and 1 and is being returned back to 0 (thus setting $x_{i+1} = \dots = x_n$ back to 0).

We let the initial configuration be $(0, 0 \dots, 0)$. Note that this is hiding the implicit dimension that is always 1. Each of the following step functions should be interpreted as copying any variable that is not explicitly set.

Step 1.

Evaluate ψ

$$f_1(\cdot) = \left\{ \hat{\psi} \leftarrow \psi(x_1, \dots, x_n) \right.$$

Step 2.

Update either s_n^0 or s_n^1

$$f_2(\cdot) = \begin{cases} s_n^0 \leftarrow \text{if } (x_n = 0) \text{ then } \hat{\psi} \\ \quad \text{else } s_n^0 \\ s_n^1 \leftarrow \text{if } (x_n = 1) \text{ then } \hat{\psi} \\ \quad \text{else } s_n^1 \end{cases}$$

Step 3.

Start incrementing x_n

$$f_3(\cdot) = \begin{cases} x_n \leftarrow \neg x_n \\ c_n \leftarrow x_n \end{cases}$$

Step $3+n-i$, for $i = n-1$ to 1.

If there is a carry, update s_i^z and continue incrementing

i even (*x_i* universally quantified):

$$f_{3+n-i}(\cdot) = \begin{cases} x_i \leftarrow \text{if } (c_{i+1}) \text{ then } \neg x_i \text{ else } x_i \\ c_i \leftarrow c_{i+1} \wedge x_i \\ c_{i+1} \leftarrow 0 \\ s_i^0 \leftarrow \text{if } (c_{i+1} \wedge \neg x_i) \text{ then } s_{i+1}^0 \wedge s_{i+1}^1 \\ \quad \text{else } s_i^0 \\ s_i^1 \leftarrow \text{if } (c_{i+1} \wedge x_i) \text{ then } s_{i+1}^0 \wedge s_{i+1}^1 \\ \quad \text{else } s_i^1 \end{cases}$$

i odd (*x_i* existentially quantified):

$$f_{3+n-i}(\cdot) = \begin{cases} x_i \leftarrow \text{if } (c_{i+1}) \text{ then } \neg x_i \text{ else } x_i \\ c_i \leftarrow c_{i+1} \wedge x_i \\ c_{i+1} \leftarrow 0 \\ s_i^0 \leftarrow \text{if } (c_{i+1} \wedge \neg x_i) \text{ then } s_{i+1}^0 \vee s_{i+1}^1 \\ \quad \text{else } s_i^0 \\ s_i^1 \leftarrow \text{if } (c_{i+1} \wedge x_i) \text{ then } s_{i+1}^0 \vee s_{i+1}^1 \\ \quad \text{else } s_i^1 \end{cases}$$

Step $3+n$.

Set every variable to 1 if QBF satisfied. After this step, the program returns to f_1 .

$$f_{3+n}(\cdot) = \left\{ v \leftarrow \text{if } (s_1^0 \wedge s_1^1) \text{ then } 1 \text{ else } v \text{ (for all variables } v) \right.$$

The $(3+n)$ th step ensures that configuration $(1, \dots, 1)$ will be reached if and only if the given QBF formula is satisfied. \blacktriangleleft

► **Remark 9 (Choice of rounding function).** The presentation here relies on specific choices of rounding function, but we observe that the language can easily exchange several different natural rounding functions, so the reduction is robust. The rounding is only useful in the **and** and **or** instructions. The floor function can be replaced by essentially any other rounding. For example $x_j \vee x_k = \left\lceil \frac{x_j + x_k}{2} \right\rceil$ and $x_j \wedge x_k = \left\lfloor \frac{-1 + x_j + x_k}{2} \right\rfloor$. Similarly, when $\lceil \cdot \rceil$ is minimal error rounding then $x_j \vee x_k = \left\lceil \frac{1 + x_j + x_k}{3} \right\rceil$ and $x_j \wedge x_k = \left\lfloor \frac{x_j + x_k}{3} \right\rfloor$ (the break point is not used). Thus, the problem will also be hard for any of these roundings.

► **Remark 10** (Perturbation: ensuring the eigenvalues are not modulus 1). Observe that under the perturbation that multiplies each operation by 1.1 (before taking floor) we obtain the same resulting operation. For example $x_i \leftarrow x_j \vee x_k = \left\lfloor \frac{1+x_j+x_k}{2} \right\rfloor$ is equivalent to $x_i \leftarrow x_j \vee x_k = \left\lfloor \left(\frac{1+x_j+x_k}{2} \right) * 1.1 \right\rfloor$. Hence, if the resulting matrix M has eigenvalues 1, taking $1.1M$ (or similar value to 1.1) will result in a matrix that does not with the same orbit; which shows that hardness is retained for matrices in which no eigenvalue has modulus 1.

► **Remark 11** (Dimension). The hardness result needs reachability instances of unbounded dimension. For a QBF formula with n variables and ℓ logical operations, the resulting instance of rounded P2P has dimension $(3n + 1 + \ell)(4n + 15 + \ell)$.

4 Special cases on non-hyperbolic systems

In this section we consider certain cases when the eigenvalues can be of modulus one. In particular we work in the Jordan normal form and show that the problem can be solved for certain types of rounding. We fall short of arbitrary deterministic rounding, which would be required to show the problem in full generality through the Jordan normal form approach.

First, we show decidability for polar-rounding, along with an example with numbers requiring exponential space by the time the system becomes periodic—seeming to imply any ‘wait and see’ approach would require **EXPSpace**. We also show decidability for certain types of Argand rounding, in particular truncation and expansion, but minimal-error rounding remains open (which we discuss further in Section 5).

4.1 Polar rounding with updates in Jordan normal form

We restrict ourselves to a Jordan block M of dimension d , with eigenvalue λ of modulus 1. Since the polar rounding function has bounded effect *on the modulus*, the remaining blocks, which need not be of modulus 1 can be solved (Lemma 7) by running this algorithm in lock step with the algorithm for those blocks. All together, this gives us:

► **Theorem 12.** *The Rounded P2P Problem is decidable in **EXPSpace** for the polar rounding function with $\theta_g = \frac{\pi}{R}$, $R \geq 2$ and matrices $M \in \mathbb{A}^{d \times d}$ in Jordan normal form.*

To prove Theorem 12 we show that each dimension $d, d-1, \dots, 1$ will eventually be periodic on a fixed modulus, or permanently diverge beyond y_k (the target value in dimension k).

Let $\langle a, b \rangle$ be the smallest angle between vectors a and b – this is a value in $[0, \pi]$ and, in particular, it is always positive. It is used as a measure of alignment: the more a and b are aligned the smaller $\langle a, b \rangle$ is. We will assume that the system will round up if $[x] - x = 0.5$. The remaining case can be adapted by suitably adjusting the relevant inequalities. We say that a dimension $k \in \{1, \dots, d\}$ is *just rotating* after position N , if for all $i \geq N$: $(x^{(i+1)})_k = \lfloor \lambda(x^{(i)})_k \rfloor$. Note that dimension d is just rotating after 0, by definition. Our goal is to show that every dimension k will eventually be just rotating (for which we would require it to have modulus $|y_k|$) or reach a point that lets us conclude it has permanently diverged past y_k . So we assume, henceforth, that dimension k is just rotating.

We let $\phi(i) = \langle \lambda(x^{(i)})_{k-1}, (x^{(i)})_k \rangle$. As $(x^{(i+1)})_{k-1} = \lfloor \lambda(x^{(i)})_{k-1} + (x^{(i)})_k \rfloor$, small values of $\phi(i)$ (between 0 and $\pi/2$) lead to an increase in modulus of $(x^{(i+1)})_{k-1}$, whereas large values (between $\pi/2$ and π) lead to a decrease when $|(x^{(i)})_{k-1}|$ is sufficiently large relative to $|(x^{(i)})_k|$. Our analysis relies on the fact that $\phi(i)$ can never increase:

► **Lemma 13.** *Suppose that dimension k is just rotating after step N . Then, for all $i \geq N+1$: $\phi(i) \geq \phi(i+1)$.*

If dimension $k-1$ repeats its relative angle to k and its modulus in some step, we can conclude that $k-1$ is just rotating:

► **Lemma 14.** *Suppose that dimension k is just rotating after step N , that $\phi(N) = \phi(N+1)$ and $|(x^{(N)})_{k-1}| = |(x^{(N+1)})_{k-1}|$. Then, dimension $k-1$ is just rotating after N .*

If the precondition of Lemma 14 holds, we move to the next dimension $k-2$. Otherwise, we want to give a bound such that whenever $|(x^{(i)})_{k-1}|$ exceeds it, we can conclude that it never decreases back to $|y_{k-1}|$. We first introduce the angle $\gamma(i) = \langle \lambda(x^{(i)})_{k-1} + (x^{(i)})_k, \lambda(x^{(i)})_{k-1} \rangle$. The angle $\gamma(i)$ decreases with increasing $|(x^{(i)})_{k-1}|$, as dimension k is just rotating and hence does not change in modulus. We observe that $\gamma(i) \leq \phi(i)$ for all i . The following shows that an increase in modulus caused by crossing an ‘axis’ (i.e. if $\gamma(i) > \pi/2$) can only happen once, as in the next step, the angle will have decreased.

► **Lemma 15.** *Let $a = \lambda(x^{(i)})_{k-1}$ and $b = (x^{(i)})_k$. Suppose that $\theta_g \leq \pi/2$, $\gamma(i) > \pi/2$ and $|a+b| > |a|$. Then $\langle \lambda[a+b], [\lambda b] \rangle \leq \pi/2$, entailing $\gamma(i+1) \leq \phi(i+1) \leq \pi/2$.*

Furthermore, a decrease cannot be followed by an increase, unless the angle changes:

► **Lemma 16.** *Suppose dimension k is just rotating after N . It is not possible for $i-1 \geq N$, to have $\phi(i-1) = \phi(i) = \phi(i+1) > \pi/2$ and $|(x^{(i+1)})_{k-1}| > |(x^{(i)})_{k-1}| < |(x^{(i-1)})_{k-1}|$.*

Finally, we place a limit on the number of consecutive increases until we can decide that dimension $k-1$ will not decrease below the current modulus in the future:


► **Lemma 17.** *Let $a = \lambda(x^{(N)})_{k-1}$ and $b = (x^{(N)})_k$ for some $N > 0$. Suppose that k is just rotating after N , $|a+b| > |a| + 0.5$ and $|a| \geq \frac{1}{\sqrt{2}}|b|$. Then, for all $i > N$: $|(x^{(i)})_{k-1}| > |a|$.*

With Lemmata 14-17 we are in a position to prove Theorem 12 (the proofs of the preceding lemmata, and the **EXPSpace** analysis can be found in the appendix).

Proof of Theorem 12. As described above, we consider each dimension separately, starting with $k = d$, and assume by induction that the previous dimension is just rotating. We describe an algorithm that tracks the value of ϕ and operates according to Figure 2. Each realisable value of ϕ relates to a copy of Figure 2 (we only draw one example of ϕ satisfying $\phi > \pi/2$ and $\phi \leq \pi/2$ respectively). For $\phi > \pi/2$ two states are used, one which encodes that the previous transition was decrementing the modulus (ϕD), the other which indicates the previous was not decrementing (including first arrival) (ϕI).

The algorithm moves on each update step according to the arrow, which denotes whether the update is modulus increasing $M \uparrow$, decreasing $M \downarrow$ or stationary MS . Similarly ϕ may decrease $\phi \downarrow$ or stay stationary ϕS , but never increase (Lemma 13). Whenever ϕ decreases we make progress through the DAG to a lower value of ϕ . All combinations $\{M \uparrow, M \downarrow, MS\} \times \{\phi \downarrow, \phi S\}$ are accounted for at each state.

Progress is made whenever we move through the DAG towards a stopping criterion. For self-loops a bound is provided (in blue) on the maximum time spent in this state. Since for each dimension we will ultimately end up in just rotating, or be able to stop early, the problem is decidable. ◀



■ **Figure 2** State diagram for ϕ whilst considering dimension $k-1$, assuming k is just rotating.


► **Example 18** (System requiring **EXPSpace** to be periodic). If $\phi \leq \pi/2$, the considered dimension will either diverge at some point, or become periodic. This depends, essentially, on whether $|(x^{(i)})_k| \cos(\phi) < 0.5$, in which case the rounding will not lead to an increase when $|(x^{(i)})_{k-1}|$ is sufficiently large relative to $|(x^{(i)})_k|$. We give an example where $|(x^{(i)})_{k-1}|$ grows to $|(x^{(i)})_k|^2$, and requires numbers of doubly exponential size (and exponential space) in d before becoming periodic. We assume that $\theta_g = \pi/2$ (so there are four possible angles) and integer modulus granularity. Let M be a single Jordan block of dimension d with eigenvalue $\lambda = e^{i\pi/2}$. The angle $\phi(i)$ remains constant, but the modulus grows while $|(x^{(i)})_k| < |(x^{(i)})_{k-1}| \leq |(x^{(i)})_k|^2$. We start at the point $x^{(0)} = ((3 + d, 0), \dots, (6, 0), (5, 0), (4, 0))$, using the representation that $Ae^{i\theta}$ is written (A, θ) . This system is periodic, with maximal component $x^{(N)} = ((4^{2^{d-1}}, 0), \dots, (4^{2 \cdot 2 \cdot 2}, 0), (4^{2 \cdot 2}, 0), (4^2, 0), (4, 0))$. Note that 4^{2^4} is larger than a 32-bit number. This idea is illustrated in Figure 3, where y represents $(x^{(i)})_k$ and is just rotating, and x represents $(x^{(i)})_{k-1}$, which grows to $|y|^2$.

Despite Example 18, which shows that waiting until becoming periodic may need exponential space, we conjecture the Rounded P2P can be solved in **PSPACE**. This is because if $(x^{(i)})_1$ exceeds a value representable in polynomial space we expect it will never return to the target y_1 (a value representable in polynomial space). However, we are unable to show at the moment that it never gets very large and subsequently returns to a small value.

4.2 Argand truncation or expansion in Jordan normal form

We now consider Argand truncation based rounding showing decidability in **EXPSpace**. The rounding function is of the form $[a + bi] = [a] + [b]i$ where, for $x \in \mathbb{R}$, $[x] = \lfloor x \rfloor$ if $x \geq 0$ and $[x] = \lceil x \rceil$ if $x < 0$, which has a non-increasing effect on the modulus.

► **Theorem 19.** *The Rounded P2P Problem is decidable in **EXPSpace** for deterministic Argand rounding function with a non-increasing effect on the modulus and matrices $M \in \mathbb{A}^{d \times d}$ in Jordan normal form.*



■ **Figure 3** Example where the system may become large before being periodic (see Example 18).

As a key ingredient of Theorem 19 we will make use of the following theorem:

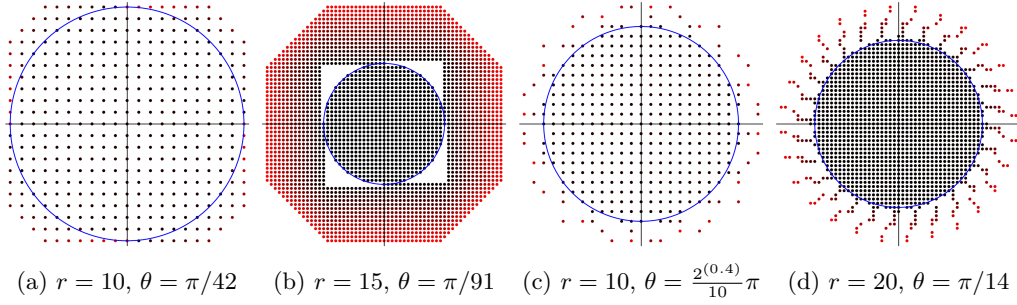
► **Theorem 20** ([28, Corollary 3.12, p.41]). *Both x is a rational multiple of π and $\sin(x)$ is rational only at $\sin(x) = 0, \frac{1}{2},$ or 1 . Both x is a rational multiple of π and $\cos(x)$ are rational only at $\cos(x) = 0, \frac{1}{2},$ or 1 . Both x is a rational multiple of π and $\tan(x)$ are rational only at $\tan(x) = 0,$ or ± 1 .*

Proof sketch of Theorem 19. Without loss of generality we consider only a single Jordan block with $|\lambda| = 1$, as the remaining blocks can be handled in lock step (using the algorithm of Lemma 7 if the eigenvalue is not of modulus one). Consider the d th component. At each step, whenever rounding takes place, then there is some decrease in the modulus. Thus, either the coordinate hits zero (and stays forever), or it stabilises and becomes periodic (with no rounding ever occurring again). The d th coordinate can be simulated until this happens. At this point, if its modulus is not $|y_d|$, y will not be reached in the future and we return NO.

If dimension x_d reaches zero, then this dimension from some point on becomes irrelevant and the instance can be reduced to an instance of dimension $d - 1$. Note that this case must occur if $\arg(\lambda)$ is not a root of unity as an irrational point is found infinitely often.

In the case where x_d does not reach zero, then it is periodic at some modulus. This implies it never rounds again, and so surely hits integer points at every step. We show that this can only occur if $\arg(\lambda)$ is a multiple of $\pi/2$. Assume that $\arg(\lambda)$ is not a multiple of $\pi/2$: the rotation of a point with integer coordinate to integer coordinate leads to the conclusion of either rational tangent or rational sine and cosine. By Theorem 20 a rational tangent alongside a rational angle ($\arg(\lambda)$ is a root of unity) implies that the angle must be a multiple of $\pi/4$. It is not $\pi/4$, as there is no Pythagorean triangle with angle $\pi/4$. By Theorem 20 rational sine and cosine and rational angle concludes the angle must be a multiple of $\pi/2$. Finally, we show that when $\arg(\lambda)$ is a multiple of $\pi/2$ the system surely diverges at dimension $d - 1$, and hence we can put a bound on how far we need to simulate. ◀

► **Remark (Argand expansion in Jordan Normal Form).** Instead of considering the rounding function to always decrease the modulus, we consider the rounding function to always increase the modulus. Then, by the same rationality argument either $\arg(\lambda)$ is a multiple of $\pi/2$ (so no rounding occurs and standard methods can be applied), or $\arg(\lambda)$ is not a multiple of $\pi/2$



■ **Figure 4** Rotational examples. We start with all points in the circle of radius r , and consider the effect of rotating every point by θ , followed by minimal error rounding. This can be seen as viewing the combined orbits, starting at several points. Redder points are added in later generations.

and rounding is applied infinitely often. We observe that rounding infinitely often results in divergence. Suppose instead the modulus converges, in supremum, to C . However the $[C]$ -ball is finite, thus rounding infinitely often must eventually exhaust the set, contradicting supremacy. Since divergence occurs in the d th component the system can be iterated until either $x^{(i)} = y$ or $(x^{(i)})_d$ exceeds y_d . (Unless $(x^{(0)})_d = y_d = 0$, in which case the d th component can be deleted.)

5 Discussion of open problems

In this section we consider the following open problem, which already exhibits a technical difficulty for a relatively simple instance.

► **Open Problem 21.** *Under which deterministic bounded-effect rounding functions does the Rounded P2P Problem become decidable (even when restricted to Jordan normal form)?*

In particular we emphasize that even decidability of the Rounded P2P Problem in the case of a 2D *rotation* matrix remains open. This should be compared to the papers [24, 11, 36, 35, 2, 31], which consider linear maps on \mathbb{R}^2 that are close to rotations, and the floor rounding $\lfloor \cdot \rfloor$ is used to induce discretized maps on \mathbb{Z}^2 . The conjecture made in [35] that all orbits of these maps are eventually periodic (and thus finite) is, to the best of our knowledge, still open in general. This lack of understanding of the dynamics of rotations even on a 2-dimensional lattice is striking and hints at an intrinsic level of difficulty in dealing with eigenvalues of modulus 1.

We ran experiments on the behaviour of rounded orbits induced by rotations in the plane. Four prototypical results are depicted in Figure 4. We note that in every one of our examples the orbits eventually become periodic. Moreover, all experiments fall into the four categories of Figure 4, i.e., where the resulting set consists of (a) a square with cut-off corners, (b) this same square, but with a central square cut out, and (c) all points within the circle with some seemingly randomly added points outside (in the case of an irrational multiple of π), (d) the initial circle with added ‘tentacles’ occurring in intervals corresponding to the rotational angle (in the case of a rational multiple of π). We have been unable to construct a rotation with an infinite rounded orbit.

One could hope that other kinds of rounding functions simplify the analysis of the orbits. We have shown truncation based rounding, for example, either helps converge towards zero, or diverge towards infinity, and this can be exploited (particularly at the bottom dimension of a Jordan block). However, roundings which may either round up or down greatly complicate the

analysis. Nevertheless, we conjecture that all rounded orbits obtained by rotation eventually become periodic.

Random rounding functions Orbit problems for rounding functions which behave probabilistically are another line of open problems and are a natural candidate for future work.

References

- 1 Shigeki Akiyama, Tibor Borbély, Horst Brunotte, Attila Pethő, and Jörg Thuswaldner. Generalized radix representations and dynamical systems. I. *Acta Mathematica Hungarica*, 108:207 – 238, 08 2005. doi:10.1007/s10474-005-0221-z.
- 2 Shigeki Akiyama and Attila Pethő. Discretized rotation has infinitely many periodic orbits. *Nonlinearity*, 26(3):871–880, 2013. doi:10.1088/0951-7715/26/3/871.
- 3 S Akshay, Timos Antonopoulos, Joël Ouaknine, and James Worrell. Reachability problems for Markov chains. *Information Processing Letters*, 115(2):155–158, 2015.
- 4 Shaull Almagor, Joël Ouaknine, and James Worrell. The semialgebraic orbit problem. In Rolf Niedermeier and Christophe Paul, editors, *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, volume 126 of *LIPICs*, pages 6:1–6:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- 5 C. Beck and G. Roepstorff. Effects of phase space discretization on the long-time behavior of dynamical systems. *Physica D: Nonlinear Phenomena*, 25(1):173 – 180, 1987. doi:10.1016/0167-2789(87)90100-X.
- 6 Michael Blank. Ergodic properties of discretizations of dynamic systems. *Dokl. Akad. Nauk SSSR*, 278(4):779 – 782, 1984.
- 7 Michael Blank. Ergodic properties of a method of numerical simulation of chaotic dynamical systems. *Mathematical Notes of the Academy of Sciences of the USSR*, 45:267–273, 1989. doi:10.1007/BF01158885.
- 8 Michael Blank. Small perturbations of chaotic dynamical systems. *Russian Mathematical Surveys*, 44(6):1–33, dec 1989. doi:10.1070/rm1989v044n06abeh002302.
- 9 Michael Blank. Pathologies generated by round-off in dynamical systems. *Physica D: Nonlinear Phenomena*, 78(1):93 – 114, 1994. doi:10.1016/0167-2789(94)00103-0.
- 10 Michael Blank. *Discreteness and Continuity in Problems of Chaotic Dynamics*. Translations of mathematical monographs. American Mathematical Society, 1997.
- 11 D. Bosio and F. Vivaldi. Round-off errors and p -adic numbers. *Nonlinearity*, 13(1):309–322, 1999. doi:10.1088/0951-7715/13/1/315.
- 12 Jin-yi Cai. Computing Jordan normal forms exactly for commuting matrices in polynomial time. *Int. J. Found. Comput. Sci.*, 5(3/4):293–302, 1994. doi:10.1142/S0129054194000165.
- 13 Rebekah Carter and Eva M. Navarro-López. Dynamically-driven timed automaton abstractions for proving liveness of continuous systems. In Marcin Jurdziński and Dejan Ničković, editors, *Formal Modeling and Analysis of Timed Systems*, pages 59–74, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. doi:10.1007/978-3-642-33365-1_6.
- 14 Ventsislav Chonev, Joël Ouaknine, and James Worrell. The polyhedron-hitting problem. In Piotr Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 940–956. SIAM, 2015.
- 15 Eva Darulova, Anastasiia Izycheva, Fariha Nasir, Fabian Ritter, Heiko Becker, and Robert Bastian. Daisy - framework for analysis and optimization of numerical programs (tool paper). In Dirk Beyer and Marieke Huisman, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 270–287, Cham, 2018. Springer International Publishing. doi:10.1007/978-3-319-89960-2_15.

- 16 Phil Diamond and Igor Vladimirov. Asymptotic independence and uniform distribution of quantization errors for spatially discretized dynamical systems. *International Journal of Bifurcation and Chaos*, 8:1479–1490, 1998. doi:10.1142/S0218127498001133.
- 17 Phil Diamond and Igor Vladimirov. Set-valued Markov chains and negative semitrajectories of discretized dynamical systems. *Journal of Nonlinear Science*, 12:113–141, 2002. doi:10.1007/s00332-001-0450-4.
- 18 S.P. Dias, L. Longa, and E. Curado. Influence of the finite precision on the simulations of discrete dynamical systems. *Communications in Nonlinear Science and Numerical Simulation*, 16(3):1574 – 1579, 2011. doi:10.1016/j.cnsns.2010.07.003.
- 19 Eric Goubault and Sylvie Putot. Static analysis of finite precision computations. In Ranjit Jhala and David Schmidt, editors, *Verification, Model Checking, and Abstract Interpretation*, pages 232–247, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. doi:10.1007/978-3-642-18275-4_17.
- 20 Stephen Hammel, James Yorke, and Celso Grebogi. Numerical orbits of chaotic processes represent true orbits. *Bull. Amer. Math. Soc.*, 19:465–, 04 1988. doi:10.1090/S0273-0979-1988-15701-1.
- 21 Anastasiia Izycheva and Eva Darulova. On sound relative error bounds for floating-point arithmetic. In *Proceedings of the 17th Conference on Formal Methods in Computer-Aided Design, FMCAD '17*, page 15–22, Austin, Texas, 2017. FMCAD Inc. doi:10.23919/FMCAD.2017.8102236.
- 22 Ravindran Kannan and Richard J. Lipton. Polynomial-time algorithm for the orbit problem. *J. ACM*, 33(4):808–821, 1986. doi:10.1145/6490.6496.
- 23 Anatole Katok and Boris Hasselblatt. *Introduction to the Modern Theory of Dynamical Systems*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1995. doi:10.1017/CB09780511809187.
- 24 John Lowenstein, Spyros Hatjispyros, and Franco Vivaldi. Quasi-periodicity, global stability and scaling in a model of Hamiltonian round-off. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 7(1):49–66, 1997. doi:10.1063/1.166240.
- 25 Victor Magron, George Constantinides, and Alastair Donaldson. Certified roundoff error bounds using semidefinite programming. *ACM Trans. Math. Softw.*, 43(4), 2017. doi:10.1145/3015465.
- 26 Oded Maler and Grégory Batt. Approximating continuous systems by timed automata. In Jasmin Fisher, editor, *Formal Methods in Systems Biology*, pages 77–89, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. doi:10.1007/978-3-540-68413-8_6.
- 27 Mariano Moscato, Laura Titolo, Aaron Dutle, and César A. Muñoz. Automatic estimation of verified floating-point round-off errors via static analysis. In Stefano Tonetta, Erwin Schoitsch, and Friedemann Bitsch, editors, *Computer Safety, Reliability, and Security*, pages 213–229, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-66266-4_14.
- 28 Ivan Niven. *Irrational Numbers*. Number 11 in The Carus Mathematical Monographs. The Mathematical Association of America, 1956. doi:10.5948/9781614440116.
- 29 Helena E. Nusse and James A. Yorke. Is every approximate trajectory of some process near an exact trajectory of a nearby process? *Comm. Math. Phys.*, 114(3):363–379, 1988. doi:10.1007/BF01242136.
- 30 Joël Ouaknine and James Worrell. On linear recurrence sequences and loop termination. *ACM SIGLOG News*, 2(2):4–13, 2015.
- 31 Attila Pethő, Jörg M. Thuswaldner, and Mario Weitzer. The finiteness property for shift radix systems with general parameters. *Integers*, 19:A50, 2019. URL: <http://math.colgate.edu/%7Eintegers/t50/t50.Abstract.html>.
- 32 Stefano Schivo and Rom Langerak. Discretization of continuous dynamical systems using UPPAAL. In Joost-Pieter Katoen, Rom Langerak, and Arend Rensink, editors, *ModelEd, TestEd, TrustEd - Essays Dedicated to Ed Brinksma on the Occasion of His 60th Birthday*,

- volume 10500 of *Lecture Notes in Computer Science*, pages 297–315. Springer, 2017. doi:10.1007/978-3-319-68270-9_15.
- 33 Alexey Solovyev, Charles Jacobsen, Zvonimir Rakamarić, and Ganesh Gopalakrishnan. Rigorous estimation of floating-point round-off errors with symbolic Taylor expansions. In Nikolaj Bjørner and Frank de Boer, editors, *FM 2015: Formal Methods*, pages 532–550, Cham, 2015. Springer International Publishing. doi:10.1007/978-3-319-19249-9_33.
 - 34 Larry J. Stockmeyer and Albert R. Meyer. Word problems requiring exponential time: Preliminary report. In Alfred V. Aho, Allan Borodin, Robert L. Constable, Robert W. Floyd, Michael A. Harrison, Richard M. Karp, and H. Raymond Strong, editors, *Proceedings of the 5th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1973, Austin, Texas, USA*, pages 1–9. ACM, 1973. doi:10.1145/800125.804029.
 - 35 Franco Vivaldi. The arithmetic of discretized rotations. *AIP Conference Proceedings*, 826, 03 2006. doi:10.1063/1.2193120.
 - 36 Franco Vivaldi and Igor Vladimirov. Pseudo-randomness of round-off errors in discretized linear maps on the plane. *International Journal of Bifurcation and Chaos*, 13(11):3373–3393, 2003. doi:10.1142/S02181274030008557.

A Additional material for Section 3.2, PSPACE-hardness

A.1 Perturbation

We expand on Remark 10, observing that multiplying by 1.1 (or similar) maintains the logical equivalence required for all of our update functions. Assuming x_i, x_j, x_k are each in $\{0, 1\}$ we have,

- $1 \leftarrow \lfloor 1 * 1.1 \rfloor$
- $x_i \leftarrow x_j \vee x_k = \left\lfloor \frac{1+x_j+x_k}{2} \right\rfloor = \left\lfloor \left(\frac{1+x_j+x_k}{2} \right) * 1.1 \right\rfloor$
- $x_i \leftarrow x_j \wedge x_k = \left\lfloor \frac{1+x_j+x_k}{3} \right\rfloor \left\lfloor \left(\frac{1+x_j+x_k}{3} \right) * 1.1 \right\rfloor$
- $x_i \leftarrow \neg x_j = \lfloor 1 - x_j \rfloor = \lfloor (1 - x_j) * 1.1 \rfloor$.
- $x_i \leftarrow \lfloor 0 \rfloor = \lfloor (0) * 1.1 \rfloor$
- $x_i \leftarrow \lfloor x_i \rfloor = \lfloor (x_i) * 1.1 \rfloor$
- $x_i \leftarrow \lfloor x_j \rfloor = \lfloor (x_j) * 1.1 \rfloor$.

A.2 Dimension of Rounded P2P instance in proof of PSPACE-hardness

It is clear that the required reduction is polynomial, we precisely characterise the dimension of the resulting system here.

► **Proposition 22.** *The resulting instance of rounded P2P has dimension $(3n + 1 + \ell)(4n + 15 + \ell)$, if ψ has ℓ logical operations.*

Proof of Proposition 22. The functions f_1, \dots, f_{3+n} hide the inner workings of the reduction to the Rounded P2P Problem, by contracting steps and auxiliary variables and illustrating the effect using logic, rather than the floor of a linear combination.

Following the routine steps to translate the logical commands into the Rounded P2P Problem, we see that:

- f_1 depends on the formula ψ to evaluate. If ψ is a formula with ℓ logical operators we have:
 - ℓ steps, resolving each logical operator according to topological ordering
 - ℓ auxiliary variables to store partial computations.

- f_2 takes two steps and four extra variables.
- f_{3+n-i} takes 3 steps for each i and 8 extra variables. The 8 variables can be shared for all functions.
- f_{3+n} takes 3 steps and $2 \times t$ extra variables, where t is the total number of main variables. However it can be simplified to 2 steps, and 1 extra variable (by noticing it is equivalent to $v \leftarrow (s_1^0 \wedge s_1^1) \vee v$).

Thus the total number of steps is $\ell + 2 + 3(n-1) + 2 = 3n + 1 + \ell$ steps. The total number of variables is $\ell + t + 4 + 8 + 1$ plus 1 to store **true**, so total of $t + 14 + \ell$. Note that $t = 4n + 1$, total $4n + 15 + \ell$. Thus when exploded as per Section 3.2, there are $(3n + 1 + \ell)(4n + 15 + \ell)$ dimensions in the Rounded P2P Problem. ◀

B Additional material for Section 4.1, Polar rounding in Jordan normal form

► **Lemma 23.** *Assume $b = [b]$, then $\langle [a], b \rangle \geq \langle [a + b], b \rangle$*

Proof of Lemma 23. Assume without loss of generality (by rotation) that $\arg(b) = 0$. Thus $b = (x, 0)$ for $x \geq 0$. A point at $a = (u, v)$ is translated to $(u + x, v)$, thus the angle between the x -axis is smaller. Hence $\langle a, b \rangle \geq \langle a + b, b \rangle$.

Now assume first that $\arg(a) \in [0, \pi]$. Then $\arg(a) = \langle a, b \rangle$ and $\arg(a) \geq \langle a + b, b \rangle = \arg(a + b) \geq 0$. From the fact that $b = [b]$ it follows that $\arg(b) = 0$ is a viable angle in our rounding. As we assume minimal error rounding on the angle, it follows that $\arg([a]) \geq \arg([a + b])$ and hence $\langle [a], b \rangle \geq \langle [a + b], b \rangle$.

If $\arg(a) \in [-\pi, 0]$, we have $-\arg(a) = \langle a, b \rangle$ and $\arg(a) \leq -\langle a + b, b \rangle = \arg(a + b) \leq 0$. As before, we can conclude $\langle [a], b \rangle \geq \langle [a + b], b \rangle$. ◀

► **Lemma 13.** *Suppose that dimension k is just rotating after step N . Then, for all $i \geq N + 1$: $\phi(i) \geq \phi(i + 1)$.*

Proof of Lemma 13. We show that for all $i \geq N + 1$:

$$\left\langle \lambda(x^{(i)})_{k-1}, (x^{(i)})_k \right\rangle \geq \left\langle \lambda(x^{(i+1)})_{k-1}, (x^{(i+1)})_k \right\rangle$$

We first make the following calculating:

$$\begin{aligned} \left\langle [\lambda(x^{(i)})_{k-1}], (x^{(i)})_k \right\rangle &\geq \left\langle [\lambda(x^{(i)})_{k-1} + (x^{(i)})_k], (x^{(i)})_k \right\rangle && \text{(Lemma 23)} \\ &= \left\langle (x^{(i+1)})_{k-1}, (x^{(i)})_k \right\rangle && \text{(by definition)} \\ &= \left\langle [\lambda(x^{(i+1)})_{k-1}], [\lambda(x^{(i)})_k] \right\rangle && (1) \\ &= \left\langle [\lambda(x^{(i+1)})_{k-1}], (x^{(i+1)})_k \right\rangle && (k \text{ is just rotating}) \end{aligned}$$

Equation (1): Since $(x^{(i+1)})_{k-1}$ and $(x^{(i)})_k$ are both at admissible angles, their rotations are at the same point between two admissible angles. Thus the rotation-effect of the rounding will be the same for both values.

Observe that $[\lambda(x^{(i+1)})_{k-1}]$ and $(x^{(i+1)})_k$ both lie on admissible angles. Consider $\theta = \langle [\lambda[a]], \lambda[a] \rangle$, observe this angle (and the direction of the angle) is the same, no matter the value of a . This angle corresponds with $\langle [\lambda(x^{(i+1)})_{k-1}], \lambda(x^{(i+1)})_{k-1} \rangle$ and $\langle [\lambda(x^{(i)})_{k-1}], \lambda(x^{(i)})_{k-1} \rangle$. Further $\theta \leq \theta_g/2$ the maximum effect of the rounding.

► **Case 1** (Suppose $\langle [\lambda(x^{(i)})_{k-1}], (x^{(i)})_k \rangle < \pi$). The calculation above also shows that $\arg([\lambda(x^{(i)})_{k-1}])$ relative to $\arg((x^{(i)})_k)$ is positive if and only if $\arg([\lambda(x^{(i+1)})_{k-1}])$ relative to $\arg((x^{(i+1)})_{k-1})$ is positive. That is, not only does the angle-distance decrease, but the relative position of the two points stays the same.

Because of this, and the fact that both $(x^{(i)})_{k-1}$ and $(x^{(i+1)})_{k-1}$ lie on admissible angles, the angle effect of rounding $\lambda(x^{(i)})_{k-1}$ relative to $(x^{(i)})_k$ is the same as the angle effect of rounding $\lambda(x^{(i+1)})_{k-1}$ relative to $(x^{(i+1)})_k$. Hence we can conclude:

$$\langle \lambda(x^{(i)})_{k-1}, (x^{(i)})_k \rangle \geq \langle \lambda(x^{(i+1)})_{k-1}, (x^{(i+1)})_k \rangle$$

► **Case 2** (Suppose $\langle [\lambda(x^{(i)})_{k-1}], (x^{(i)})_k \rangle = \pi$ and $\langle [\lambda(x^{(i+1)})_{k-1}], (x^{(i+1)})_k \rangle < \pi$). Given $R\theta_g = \pi = \langle [\lambda(x^{(i)})_{k-1}], (x^{(i)})_k \rangle$ and $\langle [\lambda(x^{(i)})_{k-1}], \lambda(x^{(i)})_{k-1} \rangle \leq \theta_g/2$ we have

$$\phi(i) = \langle [\lambda(x^{(i)})_{k-1}], (x^{(i)})_k \rangle - \langle [\lambda(x^{(i)})_{k-1}], \lambda(x^{(i)})_{k-1} \rangle \geq (R-1)\theta_g + \theta_g/2$$

Given $\langle [\lambda(x^{(i+1)})_{k-1}], (x^{(i+1)})_k \rangle \leq (R-1)\theta_g$ we have

$$\phi(i+1) \leq \langle [\lambda(x^{(i+1)})_{k-1}], (x^{(i+1)})_k \rangle + \langle [\lambda(x^{(i+1)})_{k-1}], \lambda(x^{(i+1)})_{k-1} \rangle \leq (R-1)\theta_g + \theta_g/2$$

Hence $\phi(i) \geq (R-1)\theta_g + \theta_g/2 \geq \phi(i+1)$.

► **Case 3** (Suppose $\langle [\lambda(x^{(i)})_{k-1}], (x^{(i)})_k \rangle = \langle [\lambda(x^{(i+1)})_{k-1}], (x^{(i+1)})_k \rangle = \pi$). The rotation by $\langle [\lambda(x^{(i)})_{k-1}], \lambda(x^{(i)})_{k-1} \rangle$ in either direction results in the angle (when renormalised into $[0, \pi]$) of

$$\phi(i) = \langle [\lambda(x^{(i)})_{k-1}], (x^{(i)})_k \rangle - \langle [\lambda(x^{(i)})_{k-1}], \lambda(x^{(i)})_{k-1} \rangle$$

Similarly, since the effect is the same at $i+1$ we have

$$\phi(i+1) = \langle [\lambda(x^{(i+1)})_{k-1}], (x^{(i+1)})_k \rangle - \langle [\lambda(x^{(i+1)})_{k-1}], \lambda(x^{(i+1)})_{k-1} \rangle$$

Since $\langle [\lambda(x^{(i)})_{k-1}], (x^{(i)})_k \rangle = \langle [\lambda(x^{(i+1)})_{k-1}], (x^{(i+1)})_k \rangle$ and $\langle [\lambda(x^{(i)})_{k-1}], \lambda(x^{(i)})_{k-1} \rangle = \langle [\lambda(x^{(i+1)})_{k-1}], \lambda(x^{(i+1)})_{k-1} \rangle$ we have $\phi(i) = \phi(i+1)$. ◀

► **Lemma 14.** *Suppose that dimension k is just rotating after step N , that $\phi(N) = \phi(N+1)$ and $|(x^{(N)})_{k-1}| = |(x^{(N+1)})_{k-1}|$. Then, dimension $k-1$ is just rotating after N .*


Proof of Lemma 14. We show that for all $i \geq N$: $(x^{(i+1)})_{k-1} = [\lambda(x^{(i)})_{k-1}]$. First let $i = N$. We have

$$\begin{aligned} \langle \lambda(x^{(N+1)})_{k-1}, (x^{(N+1)})_k \rangle &= \langle \lambda(x^{(N)})_{k-1}, (x^{(N)})_k \rangle && \text{(by assumption)} \\ &= \langle \lambda[\lambda(x^{(N)})_{k-1}], [\lambda(x^{(N)})_k] \rangle && (1) \\ &= \langle \lambda[\lambda(x^{(N)})_{k-1}], (x^{(N+1)})_k \rangle && (k \text{ is just rotating}) \end{aligned}$$

Equation (1) holds because $(x^{(N)})_{k-1}$ and $(x^{(N)})_k$ are both at admissible points, and hence rounding after rotating by λ has the same effect on both sides.

It follows that $\arg(\lambda(x^{(N+1)})_{k-1}) = \arg(\lambda[\lambda(x^{(N)})_{k-1}])$. As $|(x^{(N+1)})_{k-1}| = |(x^{(N)})_{k-1}|$ by assumption, we can conclude that $(x^{(N+1)})_{k-1} = [\lambda(x^{(N)})_{k-1}]$.

The next step from $N+1$ to $N+2$ is just a rotation of this case, and hence we can conclude by induction. ◀



■ **Figure 5** Situation in the upper left quadrant ($\phi(i) > \pi/2$) after an increasing step. $\alpha \leq \pi/4$, and so $\phi(i+1) \leq \pi/2$ from this point on.

► **Lemma 24.** For $R \geq 3$: $\left\lceil \frac{\pi/4}{\theta_g} \right\rceil \theta_g + \frac{\theta_g}{2} \leq \pi/2$, where $\theta_g = \frac{\pi}{R}$.

Proof of Lemma 24. $\left\lceil \frac{\pi/4}{\theta_g} \right\rceil \theta_g + \frac{\theta_g}{2} \rightarrow \pi/4$ as $\theta_g \rightarrow 0$ (or $R \rightarrow \infty$). Enumeration of the first 100 cases concludes less than $\pi/2$ before being close to $\pi/4$. ◀

► **Lemma 15.** Let $a = \lambda(x^{(i)})_{k-1}$ and $b = (x^{(i)})_k$. Suppose that $\theta_g \leq \pi/2$, $\gamma(i) > \pi/2$ and $|a+b| > |a|$. Then $\langle \lambda[a+b], [\lambda b] \rangle \leq \pi/2$, entailing $\gamma(i+1) \leq \phi(i+1) \leq \pi/2$.

Proof of Lemma 15. Refer to Figure 5, let $\alpha = \langle a+b, b \rangle$, $\beta = \langle a, b \rangle - \pi/2$ and $\theta = \pi/2 - \beta = \pi - \langle a, b \rangle$.

First we claim $\alpha \leq \pi/4$: If $\theta > \pi/4$ and since $\alpha \leq \pi/2 - \theta$, we have $\alpha \leq \pi/4$. Otherwise suppose $\theta \leq \pi/4$ and observe that $\theta' = \alpha$. Because $|a+b| > |a|$ we have $\theta' < \theta$ so $\alpha \leq \pi/4$.

Suppose $\theta_g = \pi/2$, which implies $\theta_g/2 = \pi/4$ and $\alpha \leq \theta_g/2$. Since $\langle a+b, b \rangle < \theta_g/2$ then $\langle [a+b], b \rangle = 0$ since b is already rounded and $[\lambda a+b]$ is within $\theta_g/2$. then $\langle \lambda[a+b], [\lambda b] \rangle \leq \theta_g/2$.

Instead suppose $\theta_g = \pi/R$, $R \geq 3$. Then $\langle a+b, b \rangle \leq \pi/4$ and hence $\langle [a+b], b \rangle \leq \left\lceil \frac{\pi/4}{\theta_g} \right\rceil \theta_g$ and $\langle \lambda[a+b], [\lambda b] \rangle \leq \left\lceil \frac{\pi/4}{\theta_g} \right\rceil \theta_g + \theta_g/2$, then by Lemma 24 $\left\lceil \frac{\pi/4}{\theta_g} \right\rceil \theta_g + \theta_g/2 \leq \pi/2$. ◀

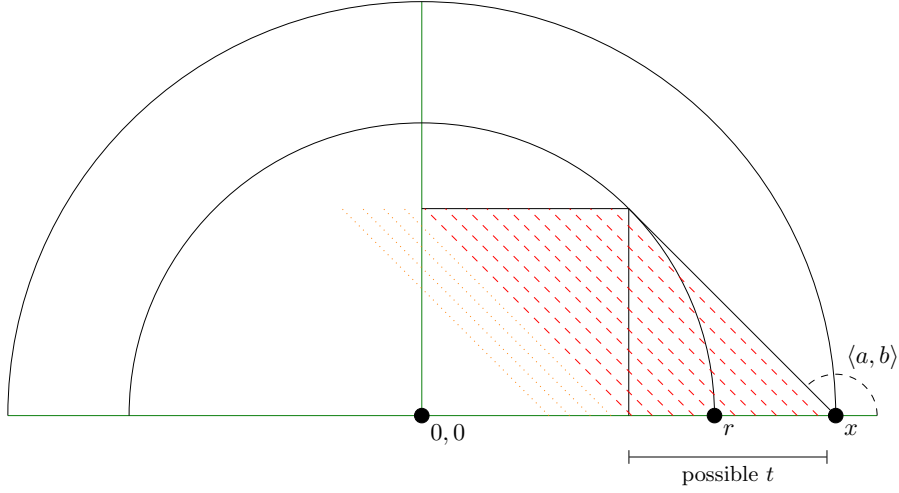
► **Lemma 25.** Let $a_i = \lambda(x^{(i)})_{k-1}$ and $b_i = (x^{(i)})_k$.

Suppose $\langle a_i, b_i \rangle > \pi/2$ and $\langle a_i + b_i, a_i \rangle \leq \pi/2$ and $|a_{i+1}| = |[a_i + b_i]| < |a_i|$ (i.e. $|(x^{(i+1)})_{k-1}| < |(x^{(i)})_{k-1}|$).

If $\langle a_{i+1}, b_{i+1} \rangle = \langle a_i, b_i \rangle$ and $\langle a_{i+1} + b_{i+1}, a_{i+1} \rangle \leq \pi/2$ then $|a_{i+2}| \leq |a_{i+1}|$ (entailing $|(x^{(i+2)})_{k-1}| \leq |(x^{(i+1)})_{k-1}|$).

Proof of Lemma 25. The situation is depicted in Figure 6. By rotational symmetry, assume $a_i = (x, 0)$. By applying the same rotational normalisation to b_i let $b = (-l, m)$. Let r, t be such that: $|a_i + b| = r$ (inner black circle) and $|[a_i + b]| = t < |a_i| (= x)$.

Then by rotational symmetry again (using $\langle a_{i+1}, b_{i+1} \rangle = \langle a_i, b_i \rangle$), assume $a_{i+1} = (t, 0)$. Then observe that $|(t, 0) + b| \leq r$ (red/dashed lines), and since $[r] = t$ and minimal error rounding is used $|(t, 0) + b| \leq t$. We have $a_{i+2} = \lambda[a_{i+1} + b] = \lambda[(t, 0) + b]$, hence $|a_{i+2}| \leq t$. ◀



■ **Figure 6** Situation in the upper right quadrant after a decreasing step, indicating that if the next step stays inside the upper right quadrant, it cannot increase. Orange/dotted can be excluded as then $\langle a + b, a \rangle > \pi/2$.

► **Lemma 16.** *Suppose dimension k is just rotating after N . It is not possible for $i - 1 \geq N$, to have $\phi(i - 1) = \phi(i) = \phi(i + 1) > \pi/2$ and $|(x^{(i+1)})_{k-1}| > |(x^{(i)})_{k-1}| < |(x^{(i-1)})_{k-1}|$.*

Proof of Lemma 16. Recall $\gamma(i) = \langle \lambda(x^{(i)})_{k-1} + (x^{(i)})_k, \lambda(x^{(i)})_{k-1} \rangle$.

Suppose following a modulus decreasing transition (hence $\phi(i) \geq \pi/2$) there is $\phi(i) = \phi(i + 1)$ and $|(x^{(i+1)})_{k-1}| > |(x^{(i)})_{k-1}|$


- the last step was a decrease with $\gamma(i - 1) \leq \pi/2$ and $\gamma(i) \leq \pi/2$ then by Lemma 25 it is not increasing (contradicting $|(x^{(i+1)})_{k-1}| > |(x^{(i)})_{k-1}|$).
- the last step was a decrease with $\gamma(i - 1) > \pi/2$ and this step has $\gamma(i) \leq \pi/2$, this cannot happen without a change of angle (contradicting $\phi(i - 1) = \phi(i)$).
- this step has $\gamma(i) > \pi/2$, then by Lemma 15 an increase ($|(x^{(i+1)})_{k-1}| > |(x^{(i)})_{k-1}|$) would cause the angle to decrease (contradicting $\phi(i) = \phi(i + 1)$). ◀

► **Lemma 26.** *Let $a = \lambda(x^{(N)})_{k-1}$ and $b = (x^{(N)})_k$ for some $N > 0$, and let $\gamma = \langle a + b, a \rangle$. Suppose that k is just rotating after N , $\gamma < \pi/2$ and $|a + b| > |a| + 0.5$.*

Then, for all $i > N$: $|(x^{(i)})_{k-1}| > |a|$.

Proof. Recall that $\phi(i) = \langle \lambda(x^{(i)})_{k-1}, (x^{(i)})_k \rangle$ and let $\phi = \phi(N)$. By Lemma 13 we have: $\phi(i) \leq \phi$ for all $i \geq N$. We show the claim by induction on i . If $i = N + 1$, it holds by assumption. Else, we can assume that $|(x^{(i-1)})_{k-1}| > |a|$. We let $d = (x^{(i-1)})_{k-1}$, $e = (x^{(i-1)})_k$ and we aim to show that $|d + e| \geq |a + b|$.

$$\begin{aligned}
|d + e| &\geq |a + b| \\
\iff |d + e|^2 &\geq |a + b|^2 \\
\iff |d|^2 + |b|^2 - 2|d||b|\cos(\pi - \phi) &\geq |a|^2 + |b|^2 - 2|a||b|\cos(\pi - \phi) & (*) \\
\iff (|a| + C)^2 - 2(|a| + C)|b|\cos(\pi - \phi) &\geq |a|^2 - 2|a||b|\cos(\pi - \phi) & (**) \\
\iff 2|a|C + C^2 - 2C|b|\cos(\pi - \phi) &\geq 0 \\
\iff 2|a|C + C^2 &\geq 2C|b|\cos(\pi - \phi) \\
\iff |a| &\geq |b|\cos(\pi - \phi)
\end{aligned}$$



■ **Figure 7**

To see that $|a| \geq |b| \cos(\pi - \phi)$ holds, it is enough to observe that $|a| = |b| \cos(\pi - \phi) + |a + b| \cos(\gamma)$ (see Figure 7), and $0 < \gamma < \frac{\pi}{2}$. The step (*) is valid as k is just rotating, which implies $|b| = |e|$, together with $\phi(i-1) \leq \phi$, which implies $\cos(\pi - \phi) \geq \cos(\pi - \phi(i-1))$. In (***) we use that $|d| = |(x^{(i-1)})_{k-1}| = |a| + C$, for some positive C .

Now, from $|d + e| \geq |a + b| > |a| + 0.5$ it follows that $||d + e|| = |(x^{(i)})_{k-1}| > |a|$. ◀

► **Lemma 27.** *Let a, b be algebraic numbers, and assume that $|a + b| > |a| \geq \frac{1}{\sqrt{2}} |b|$. Then, $\langle a + b, a \rangle \leq \pi/2$.*

Proof. Let $\gamma = \langle a + b, a \rangle$ and assume, for contradiction, that $\pi/2 < \gamma \leq \pi$. We have:

$$|b|^2 = |a + b|^2 + |a|^2 - 2|a||a + b| \cos(\gamma) > 2|a|^2 - 2|a||a + b| \cos(\gamma) > 2|a|^2$$

The last step follows by $\cos(\gamma) < 0$. But then we have $|b| > \sqrt{2}|a|$, which contradicts the assumptions. ◀

► **Lemma 17.** *Let $a = \lambda(x^{(N)})_{k-1}$ and $b = (x^{(N)})_k$ for some $N > 0$. Suppose that k is just rotating after N , $|a + b| > |a| + 0.5$ and $|a| \geq \frac{1}{\sqrt{2}} |b|$. Then, for all $i > N$: $|(x^{(i)})_{k-1}| > |a|$.*

Proof of Lemma 17. Direct corollary of Lemma 26 and Lemma 27. ◀

► **Proposition 28.** *The problem (of Theorem 12) is in **EXPSPACE**.*

Proof of Proposition 28. This proof gives a more detailed analysis of the algorithm as presented in Theorem 12, showing that it only uses exponential space. As in Theorem 12 we consider each dimension separately, starting with dimension d . For each dimension k we will establish upper bounds T_k on the number of steps that we need while considering dimension k , and U_k on the maximum numeric value in dimension k that we need to consider.

We assume that we are at step N and have concluded that dimension k just rotates at the right modulus (to conclude this for dimension d we need only one step). We use the same stopping criteria as in Theorem 12 to conclude that we no longer reach y . Observe that in Theorem 12 the value of $|(x^{(i)})_{k-1}|$ never increases beyond the value $K = \max(|y_k|, |y_{k-1}|, |(x^{(N)})_{k-1}|)$; thus if $|(x^{(i)})_{k-1}| > K$ we can conclude that we can

stop. This uses that the previous dimension is rotating on the “right” modulus, that is $|y_k| = |(x^{(i)})_k|$ holds.

The $[K]$ -ball (see Definition 3) has $K \cdot \frac{2\pi}{\theta_g}$ admissible points and hence, by the above argument, after this many steps we will either have: left the ball and concluded that we can stop, become just rotating in the current dimension, or decreased $\phi(i)$. As there are $\frac{2\pi}{\theta_g}$ possible values of ϕ , in total we will spend at most $K \cdot \left(\frac{2\pi}{\theta_g}\right)^2$ steps in dimension $k-1$.

To ease calculations, we want to assume that the time spent in each dimension increases with respect to the previous one, and so we set $T_{k-1} = K \cdot \left(\frac{2\pi}{\theta_g}\right)^2 + T_k$. This lets us assume that N , the step at which we start considering dimension $k-1$, satisfies: $N \leq d \cdot T_k$. We note that $|(x^{(N)})_{k-1}|$ is at most $|(x^{(0)})_{k-1}| + N \cdot U_k$, as U_k is an upper bound on dimension k . Hence we put $U_{k-1} = |(x^{(0)})_{k-1}| + d \cdot T_k \cdot U_k$

Using these upper bounds we now calculate how large the values U_j may get with decreasing j (we start at $j = d$). In order not to distinguish the initial and target values of each dimension, we overestimate by using $y_s = \sum_{j=0}^d |y(j)|$ and $i_s = \sum_{j=0}^d |(x^{(0)})_j|$.

$$\begin{aligned} T_d &= 1, & U_d &= i_s \\ T_{k-1} &= K \cdot \left(\frac{2\pi}{\theta_g}\right)^2 + T_k \leq (y_s + U_{k-1}) \cdot \left(\frac{2\pi}{\theta_g}\right)^2 + T_k \\ U_{k-1} &= i_s + d \cdot T_k \cdot U_k \\ &\leq i_s + d \cdot \left((y_s + U_k) \cdot \left(\frac{2\pi}{\theta_g}\right)^2 + T_{k+1} \right) \cdot U_k \\ &\leq \underbrace{i_s \cdot 3 \cdot d \cdot y_s \cdot \left(\frac{2\pi}{\theta_g}\right)^2}_{F} \cdot U_k^2 \end{aligned}$$

The last step uses that $T_{k+1} \leq U_k$. Given the input, F is fixed and of pseudopolynomial size. We can conclude that:

$$U_{d-j} \leq (F \cdot i_s)^{2^j}$$

As $F \cdot i_s$ is exponential in the input, it follows that U_0 is at most double-exponential in the input. Hence, it requires at most exponentially many bits to express. ◀

C Additional material for Section 4.2, Argand truncation or expansion in Jordan normal form

In this subsection we assume all angles are given in degrees as we will make use of rationality arguments on the angles. This is simply a stylistic choice, since it would be equivalent to consider rational multiples of π .

► **Theorem 19.** *The Rounded P2P Problem is decidable in **EXPSpace** for deterministic Argand rounding function with a non-increasing effect on the modulus and matrices $M \in \mathbb{A}^{d \times d}$ in Jordan normal form.*

Proof of Theorem 19. Without loss of generality we consider only the case where $|\lambda| = 1$, since if $|\lambda| \neq 1$ the algorithm of Theorem 6 can be used on each such block in lock step.

Consider the d th component. At each step, whenever rounding takes place, then there is some decrease in the modulus. Thus, either the coordinate hits zero (and stays forever), or it stabilises and becomes periodic (with no rounding ever occurring again).

The d th coordinate can be simulated until this happens. Then clearly it must match the target y_d occasionally, otherwise the answer is no.

In the following we argue either it stabilises at zero in which case it is trivial. Or it becomes periodic with non-zero modulus and that this occurs if and only if $\arg(\lambda)$ is a multiple of 90 degrees (and otherwise must reduce to zero).

► **Case 1** ($(x^{(i)})_d$ reaches zero). In this case it is stable, and the next coordinate is not effected by this coordinate (from some point on). Then the problem can be reduced to a smaller instance, by deleting the d th coordinate.

► **Case 2** (Periodic at modulus > 0 and $\arg(\lambda)$ is not a root of unity). This case does not occur. If $\arg(\lambda)$ is not a root of unity then $\lambda^i x$ is dense on the circle of radius $|x|$, and must eventually hit a point with non-integer coordinates. Such points must be rounded, decreasing the modulus, contradicting stability, and thus periodicity.

► **Case 3** (Periodic at modulus > 0 and $\arg(\lambda)$ is a root of unity). If $\arg(\lambda)$ is a root of unity it is rational ($\lambda^n = 1$ implies $n \arg(\lambda) = k360$ for some $k \in \mathbb{N}$). We show the only angle that does not tend to zero is a multiple of 90.

The following arguments assume we start at $a + bi$ and move to $c + di$ by a rotation of $\arg(\lambda)$, the proofs will be based on the rationality/irrationality of the angle and \tan, \sin, \cos of the angle. To do this we assume both are in the upper right quadrant, as by rotating both by 90, 180 or 270 to get it there will have the same argument regarding the irrationality.

Suppose we move from $a + 0i$ to $c + di$. Recall the modulus is fixed, so $c^2 + d^2 = a^2$. The angle formed by this is $\arg(\lambda)$ and the tangent is $\frac{d}{c}$. Since c, d are rounded to a rational (but no rounding takes place) then the tangent is rational. By Theorem 20, the only point with $\arg(\lambda)$ rational and $\tan(\arg(\lambda))$ rational are $\arg(\lambda) = 45$ and 90. Note that it cannot be 45, because then $c + di = c + ci$, and we have $\sqrt{(c^2 + c^2)} = a$. There is no integer solution to this equation (no Pythagorean triangle has angle 45 degrees). Hence to move from axis to non-axis the only acceptable angle is a multiple of 90 degrees (which indeed is not non-axis).

However, as a result of the finite period before stabilising and becoming periodic, the orbit could already be at a non-axis point, and move entirely within non-axis points. Suppose we move from $a + bi$ to $c + di$, with angle $\arg(\lambda)$. Note that $a + bi = C \exp(i\theta)$ and $c + di = C \exp(i(\theta + \arg(\lambda)))$ and hence $c + di = (a + bi)(\exp(i \arg(\lambda))) = (a + bi)(\cos(\arg(\lambda)) + i \sin(\arg(\lambda)))$. Then we have

$$c = \operatorname{Re}(c + di) = \operatorname{Re}((a + bi)(\cos(\arg(\lambda)) + i \sin(\arg(\lambda)))) = a \cos(\arg(\lambda)) - b \sin(\arg(\lambda))$$

and

$$d = \operatorname{Im}(c + di) = \operatorname{Im}((a + bi)(\cos(\arg(\lambda)) + i \sin(\arg(\lambda)))) = b \cos(\arg(\lambda)) + a \sin(\arg(\lambda)).$$

Note then that

$$c + \frac{bd}{a} = (a + \frac{b^2}{a}) \cos(\arg(\lambda)) \text{ and } d - \frac{bc}{a} = (a + \frac{b^2}{a}) \sin(\arg(\lambda)),$$

but since a, b, c, d are rational we have $\cos(\arg(\lambda))$ and $\sin(\arg(\lambda))$ rational. Recall, by Theorem 20, the only point $\cos(\arg(\lambda))$ and $\arg(\lambda)$ are rational is $\arg(\lambda)$ multiple of 30 (but not 60) or 90 and the only point $\sin(\arg(\lambda))$ and $\arg(\lambda)$ are rational is $\arg(\lambda)$ multiple of 60 or 90. Thus $\arg(\lambda)$ is a multiple of 90.

In this case there is *no rounding* whatsoever. Indeed in this case the final coordinate, x_d , is periodic after the first rounding step, with period at most 4. If it starts at $a + bi$, it goes through (at most) $-b + ai$, $-a + bi$, $b - ai$ before returning back to $a + bi$. Then we show the penultimate coordinate, x_{d-1} , grows from some point on giving a stopping criterion (either $x^{(i)} = y$ at some point, or $(x^{(i)})_{d-1} > y_{d-1}$ and never comes back). The initial point is invariant-under-rounding (i.e. when $x = [x]$), it is rotated by 90 degrees to another invariant-under-rounding point and then adds a point from the previous component (which is already invariant-under-rounding), resulting in an invariant-under-rounding point. Therefore we can use standard techniques to show that $(x^{(n)})_{d-1} = \lambda(x^{(n-1)})_{d-1} + (x^{(n-1)})_d$ must grow; To see this note that $(x^{(n)})_{d-1} = \lambda^n(x^{(0)})_{d-1} + n\lambda^{n-1}(x^{(0)})_d$, which diverges as $n \rightarrow \infty$. Thus the analysis of components $1 \dots d-2$ is not necessary.

To see that this algorithm needs at most exponential space, we use a similar argument as in Proposition 28. First, we observe that the value in dimension d never increases. Hence, an upper bound for the value in this dimension is $U_d = |(x^{(0)})_d|$. This implies that we never exit the $[U_d]$ -ball, and hence, after at most $T_d = \lceil [U_d] \rceil \leq (2U_d/g)^d$ steps we can conclude whether $(x^{(0)})_d$ becomes periodic at 0, or some other modulus. In the latter case we must be in Case 3, and we can conclude that dimension $d-1$ diverges. Hence it is enough to simulate the system inside the $[y_{d-1}]$ -ball, which is of single exponential size.

In the first case, we proceed to dimension $d-1$, to which the same analysis applies, as now dimension d is at modulus 0 and does not influence the dynamics any more. Dimension $d-1$ may have grown to at most $U_{d-1} = |(x^{(0)})_{d-1}| + T_d \cdot U_d$.

To simplify the calculations, we want to assume that $T_{k-1} > T_k$ and hence set: $T_{k-1} = \lceil [U_{k-1}] \rceil + T_k \leq (2U_{k-1}/g)^d + T_k$. This gives us $\sum_{j=k}^d T_j \leq d \cdot T_k$. We will use $d \cdot T_k$ as an overestimate for the number of steps that were taken before reaching dimension $k-1$. In order not having to distinguish the different initial values, we overestimate by assuming value $i_s = \sum_{j=0}^d |(x^{(0)})_j|$ in every dimension. Overall, this leads us to the following equations for U_k and T_k :

$$U_d = i_s, \quad T_d = [U_d] \leq (2U_d/g)^d \quad (1)$$

$$T_{k-1} = [U_{k-1}] + T_k \leq \left(\frac{2U_{k-1}}{g} \right)^d + T_k \quad (2)$$

$$U_{k-1} = i_s + d \cdot T_k \cdot U_k \quad (3)$$

$$\leq i_s + d \cdot \left(\left(\frac{2U_k}{g} \right)^d + T_{k+1} \right) \cdot U_k \quad (4)$$

$$\leq \underbrace{i_s \cdot d \cdot (2/g)^d}_{F} \cdot 2(U_k)^{d+1} \quad (5)$$

Step 5 uses that $T_{k+1} < U_k$ (see Step 3) and $d > 0$. Given the input, F is fixed and single exponential. It follows that:

$$U_{d-j} \leq (F \cdot i_s)^{(d+1)^j}$$

As $F \cdot i_s$ is single exponential, it follows that U_0 is at most double exponential in the input and hence expressible in single exponential space. \blacktriangleleft