

Linear Programs with Polynomial Coefficients and Applications to 1D Cellular Automata

Guy Bresler *, Chenghao Guo*, Yury Polyanskiy*

Abstract

Given a matrix A and vector b with polynomial entries in d real variables $\delta = (\delta_1, \dots, \delta_d)$ we consider the following notion of feasibility: the pair (A, b) is *locally feasible* if there exists an open neighborhood U of 0 such that for every $\delta \in U$ there exists x satisfying $A(\delta)x \geq b(\delta)$ entry-wise. For $d = 1$ we construct a polynomial time algorithm for deciding local feasibility. For $d \geq 2$ we show local feasibility is NP-hard.

As an application (which was the primary motivation for this work) we give a computer-assisted proof of ergodicity of the following elementary 1D cellular automaton: given the current state $\eta_t \in \{0, 1\}^{\mathbb{Z}}$ the next state $\eta_{t+1}(n)$ at each vertex $n \in \mathbb{Z}$ is obtained by $\eta_{t+1}(n) = \text{NAND}(\text{BSC}_\delta(\eta_t(n-1)), \text{BSC}_\delta(\eta_t(n)))$. Here the binary symmetric channel BSC_δ takes a bit as input and flips it with probability δ (and leaves it unchanged with probability $1 - \delta$). It is shown that there exists $\delta_0 > 0$ such that for all $0 < \delta < \delta_0$ the distribution of η_t converges to a unique stationary measure irrespective of the initial condition η_0 .

We also consider the problem of broadcasting information on the 2D-grid of noisy binary-symmetric channels BSC_δ , where each node may apply an arbitrary processing function to its input bits. We prove that there exists $\delta'_0 > 0$ such that for all noise levels $0 < \delta < \delta'_0$ it is impossible to broadcast information for any processing function, as conjectured in [MMP21].

Contents

| | | |
|----------|-------------------------------------------|----------|
| 1 | Introduction | 1 |
| 1.1 | Results | 2 |
| 1.2 | Toy Example of 1-PLP | 3 |
| 1.3 | Sketch of Main Ideas | 3 |
| 1.4 | Outline | 5 |
| 1.5 | Related Literature | 6 |
| 1.6 | Future Directions | 7 |
| 2 | Additional Results | 7 |
| 2.1 | PLP Infeasibility | 7 |
| 2.2 | Tractable Variations of Local Feasibility | 8 |
| 2.2.1 | The Case of Equality | 8 |
| 2.2.2 | The Case of Strict Inequality | 8 |
| 3 | Applications | 9 |
| 3.1 | Probabilistic Cellular automata | 9 |
| 3.2 | Broadcasting of Information on a 2D-Grid | 12 |

*Massachusetts Institute of Technology, email: {guy, chenghao, yp}@mit.edu

| | | |
|----------|------------------------------------------------------------------------|-----------|
| 4 | Preliminaries | 13 |
| 4.1 | Polynomial and Rational Functions | 13 |
| 4.2 | Pseudo Inverse | 14 |
| 4.3 | Useful Lemmas in Computational Algebra | 14 |
| 4.4 | Notation | 14 |
| 5 | A Polynomial Time Algorithm for Local Feasibility of 1-PLP | 15 |
| 5.1 | Rational Function Solution | 15 |
| 5.2 | Polynomial Solution and Reformulation of Feasibility Problem | 18 |
| 5.3 | Subspace Elimination | 19 |
| 5.3.1 | Definition of Subspace Elimination | 19 |
| 5.3.2 | Algorithm for Subspace Elimination | 19 |
| 5.4 | Improvement in Running Time | 21 |
| 5.5 | Outputting a Feasible Solution and its Range | 22 |
| 6 | Hardness of Polynomial Linear Programming | 23 |
| 6.1 | Local Feasibility to Everywhere Feasibility | 23 |
| 6.2 | Hardness of Everywhere Feasibility of PLP | 24 |
| 6.3 | Hardness of Everywhere Infeasibility of PLP | 27 |
| 6.4 | Everywhere Feasibility is in co-NP | 28 |
| 7 | Application to Probabilistic Cellular Automata | 29 |
| 7.1 | Coupling and Auxiliary Chain | 30 |
| 7.2 | Potential Function Method | 30 |
| 7.3 | Connection with PLP | 31 |
| 7.4 | Attempt on Soldier's Rule | 33 |
| 8 | Application in Broadcasting of Information on 2D Grid | 34 |
| 8.1 | Auxiliary Chain and Potential Function | 35 |
| 8.2 | Connection with PLP | 35 |
| 9 | Efficient Solution with Only Equality Constraints | 37 |
| 9.1 | Preliminaries in Algebraic Geometry | 37 |
| 9.2 | Algorithm | 38 |
| A | Details of Lemma 25 and Lemma 27 | 40 |
| B | Proof of of Lemma 23 | 41 |
| C | Potential Functions for Ergodicity of PCA | 42 |
| | References | 44 |

1 Introduction

Linear Programming (LP) is one of the central paradigms of optimization, with broad significance in both theory and applications. In this paper we introduce and study a version of linear programming, *polynomial linear programming (PLP)*, where the coefficients are given by polynomials.

Our initial motivation in studying PLPs was in trying to generalize the exciting method of proving ergodicity of cellular automata introduced in [HMM19]. That method relies on finding a certain potential (or a Lyapunov function) that decreases on average. The problem of finding such potential was shown in [MMP21] to be a PLP.

However, PLPs also arise naturally. As an example, consider the max-flow problem, which (as is known classically) can be formulated as an LP. A typical question is whether a given graph with fixed edge capacities can support at least a flow r between a chosen source-sink pair. If the edge capacities and desired flow depend on some external factors δ , e.g. the amount of rainfall in the case of traffic modeling, then the problem may be approximated by a PLP. The question would then be whether the flow $r(\delta)$ is possible even if the amount of rainfall varies slightly. This robustness question is what we call local-feasibility of a PLP and is the subject of this work. This notion of robustness is distinct from the one considered in the field of robust optimization, since (as elaborated upon in Section 1.5) robust optimization aims for a single solution r (not depending on δ) and the uncertainty sets are not algebraic manifolds as we consider.

We now formally define PLPs and the associated computational problems.

Definition 1.1 (*d*-PLP). A *d*-dimensional polynomial linear program (*d*-PLP) is specified by an $m \times n$ matrix $A(\delta)$ and an $m \times 1$ vector $b(\delta)$ of degree- D polynomials in $\delta \in \mathbb{R}^d$. The coefficients of all polynomial functions are rational numbers represented by integer numerators and denominators. The *size* of a PLP is defined to be the number of bits needed to describe A and b (see Definition 4.2 for full detail).

We are interested in two types of PLP feasibility: local and everywhere feasibility.

Definition 1.2 (Local feasibility). The PLP $(A(\delta), b(\delta))$ is said to be *locally feasible* (*locally infeasible*) if there exists $\epsilon > 0$ such that for every $\delta \in B_\epsilon(0)$ there exists (does not exist) an $x \in \mathbb{R}^n$ such that $A(\delta)x \geq b(\delta)$.

Definition 1.3 (Everywhere feasibility). The PLP $(A(\delta), b(\delta))$ is said to be *everywhere feasible* (*everywhere infeasible*) if for every $\delta \in \mathbb{R}^d$ there exists (does not exist) an $x \in \mathbb{R}^n$ such that $A(\delta)x \geq b(\delta)$.

Intuitively, the local and everywhere feasibility of PLP describes whether a linear program is “stable” subject to changes in the parameters. Imagine a system having a property we care about that is determined by the feasibility of an LP $Ax \geq b$, but the system is not known exactly and we are only sure that coefficients of the LP lie on a polynomial manifold $(A(\delta), b(\delta))$, parameterized by δ . The everywhere feasibility then tells us whether the property is satisfied up to arbitrary perturbation on the manifold, while the local feasibility tells us whether the property is robust locally around a fixed point.

Remark 1. For local feasibility, by a simple translation, taking 0 as the center of the local neighborhood ball is computationally equivalent to any other constant point in \mathbb{R}^d .

Remark 2. When $d = 1$, everywhere feasibility on \mathbb{R} is equivalent to feasibility on any interval (c, d) . To transform a problem on \mathbb{R} to the interval (c, d) , we can substitute $c + \frac{\delta^2(d-c)}{\delta^2+1}$ for δ and multiply by the denominator on both sides. Similarly, to transform a problem on the interval (c, d) to \mathbb{R} , we can replace δ with $1/(c - \delta) + 1/(d - \delta)$ and multiply by the denominator on both sides.

Remark 3. For ordinary linear programs, optimization and feasibility are equivalent algorithmic tasks. Optimizing an LP can be solved by checking LP feasibility via binary search. The same reasoning can be applied to PLP. Given an oracle for everywhere feasibility, we can use binary search to solve the following optimization problem:

$$\begin{aligned} & \max_{\delta, x(\delta)} c^\top(\delta)x(\delta) \\ & \text{subject to } \forall \delta, A(\delta)x(\delta) \geq b(\delta). \end{aligned}$$

1.1 Results

We close the tractability of both types of PLP feasibility in all dimensions, and apply the results to various problems.

Theorem 1 (Local feasibility of 1-PLP is tractable). *The local feasibility of 1-PLP can be solved in polynomial time. Furthermore, if the PLP is locally feasible, a solution $x(\delta)$ can be given as a rational number at 0 and two rational functions of δ on the positive and negative neighborhoods of 0, respectively. The interval (δ_1, δ_2) on which $x(\delta)$ is a feasible solution is given by δ_1 and δ_2 each being the smallest positive root of an explicit polynomial. The running time is $O(n^3mD^3 \cdot \text{LP}(O(n^2D), O(nmD), L))$, where $\text{LP}(a, b, c)$ stands for the running time of ordinary linear programming with n variables, m constraints and L -bit numbers.*

Theorem 2 (Hardness). *The everywhere feasibility of d -PLP is NP-hard for any d and is co-NP-complete when $d = 1$. The local feasibility of d -PLP is NP-hard for any $d \geq 2$.*

As a consequence of Theorem 1, the max-flow problem discussed in the beginning, where capacity and demand are polynomials in δ which is 1-dimensional and varies locally around 0, can be solved in polynomial time by viewing it as a special case of PLP. However, it is not clear whether the max-flow problem is tractable when δ has higher dimension: The hardness construction of Theorem 2 is an LP that cannot be easily reduced to a flow problem. We leave this question open.

Aside from direct application to problems such as max-flow, the algorithm in Theorem 1 will be used to give computer-assisted proofs solving open problems in probabilistic cellular automata (PCA) and broadcasting on the 2-dimensional grid under the most interesting regime of vanishing (but non-zero) noise. See Section 3 for a summary of these results.

Theorem 3 (Informal). *There exists $\delta_0 > 0$ such that an elementary PCA with NAND function and either vertex binary-symmetric-channel (BSC) noise δ or edge BSC noise δ is ergodic for all $\delta \in (0, \delta_0)$.*

Theorem 4 (Informal). *There exists $\delta_0 > 0$ such that when information is broadcast on the space-homogeneous 2D grid of relays interconnected by BSC(δ), it is impossible to recover any information about the origin from the boundary values for all $\delta \in (0, \delta_0)$.*

The BSC_p channel has binary input and output. For a given noise level $p \in [0, 1/2]$, it outputs the same bit as input with probability $1 - 2p$ and outputs a random bit with probability $2p$.

Remark 4. In Theorems 3 and 4 we are interested in only the feasibility on the positive neighborhood of 0 rather than $B_\epsilon(0)$, which is different than the statement of Theorem 1. However, as discussed in the beginning of Section 5, our algorithm considers the positive side and negative side separately and combines the results. So the algorithm can also solve feasibility on one side.

Remark 5. We note that impossibility of broadcasting via noisy NANDs (or any other symmetric function) was conjectured in [MMP21] for all $\delta \in (0, 1/2]$. Our result above resolves this conjecture only partially. Indeed, intuitively showing impossibility of broadcasting with a small noise (arbitrarily close to 0) as we do is the most challenging setting. However, there are no known monotonicity (in the noise magnitude) results for these problems, so it remains possible in principle that an intermediate amount noise could somehow help information propagate.

In the high noise regime, Evans-Schulman [ES99] show that for $\delta > 2 - \frac{\sqrt{2}}{4} \approx 0.146$ broadcasting is impossible with any arrangement of two-input gates. To improve this estimate for the graph corresponding to a 2D grid we can use a) the connection between information and percolation from [PW17, Theorem 5] and b) a rigorous lower bound $p_{lb} = 0.6231$ for the oriented bond-percolation threshold from [GWS80]. These arguments show that broadcasting on a 2D grid of arbitrary gates and BSC_δ wires is impossible for $\delta > 1 - \frac{\sqrt{p_{lb}}}{2} \approx 0.105$. It remains open to close the gap between our results for small $\delta < \delta_0$ and these ones for large $\delta > \delta_1$. Our method cannot do so, because it is based on feasibility of general PLP: Showing impossibility of broadcasting (or ergodicity of the PCA) in the interval (δ_0, δ_1) for some given fixed δ_1 requires (by Remark 2) solving everywhere feasibility of a PLP, which by Theorem 2 is NP-hard in general.

1.2 Toy Example of 1-PLP

To give intuition for how the feasibility of PLP can vary with δ , we start with a toy example of a 1-PLP. This 1-PLP is feasible when $\delta \notin (0, \epsilon)$ and infeasible when $\delta \in (0, \epsilon)$, where ϵ is an arbitrary small constant. This example demonstrates that local feasibility cannot be directly solved by simply taking a representative δ close enough to 0, as ϵ can be too small for any particular choice of δ . (While here the dependence on ϵ is simple, in general it is not clear how to determine what constitutes a small enough value of δ based on the constraints.)

We take $n = 2$, writing x and y as the variables. There are 5 constraints c_1 to c_5 , and only c_4 and c_5 are varying with δ . The parameter ϵ is left free for now, to be thought of as a small constant. The 1-PLP is given by the constraints

$$\begin{aligned} c_1(x, y) &= y - x \geq 0 \\ c_2(x, y) &= y + x \geq 0 \\ c_3(x, y) &= y - 1 + \epsilon \leq 0 \\ c_4(x, y) &= y - \delta^2 + \epsilon\delta \leq 0 \\ c_5(x, y) &= y + 1 - \delta \geq 0. \end{aligned} \tag{1}$$

Now consider varying δ from $-\infty$ to $+\infty$. We depict the evolution of feasibility in Figure 1.

1.3 Sketch of Main Ideas

Here we present some of the ideas underlying the algorithm in Theorem 1.

Reduction to Polynomial Solutions. Suppose a given PLP is locally feasible. It is not obvious that the solution x can be expressed explicitly: $x(\delta)$ can potentially be any function defined on $B_\epsilon(0)$, and the space of such functions has infinite dimension. In Figure 1, we observe that $x(\delta)$ can be taken to be a piece-wise rational function of δ . This turns out to be true in general, and by carefully using the structure of the problem we will prove that there must exist a rational function $x(\delta)$ locally solving the PLP if it is locally feasible. To this end, we demonstrate that each corner (or point meeting a subset of constraints with equality) of the feasible polyhedron is a rational




Figure 1: Feasible region of the PLP (1) as δ changes. When $\delta \leq 0$, it is feasible, and when $\delta = 0$, there is a unique solution $x = y = 0$. When $\delta \in (0, \epsilon)$, there is no solution because the combination of $c_1 \geq 0, c_2 \geq 0$, and $c_4 \leq 0$. When $\delta \in (\epsilon, 1]$, the solution set is the region between $c_1 = 0, c_2 = 0$, and $c_4 = 0$, and taking any two of these constraints as equality constraints yields a feasible solution. When $\delta > 1$, the solution set is the region between $c_1 = 0, c_2 = 0, c_3 = 0$ and $c_5 = 0$.

function in δ on a small enough open interval, and then argue that feasibility of the point cannot change on the interval using a continuity argument.

We then show that this same statement holds for $x(\delta)$ having polynomial entries, with only a polynomial increase in the magnitude of the coefficients and the degrees. As a result, we can restrict our search space to $x(\delta) = (x_1(\delta), \dots, x_n(\delta))$ being a vector of polynomials, written as

$$x_i(\delta) = \sum_j h_{ij} \delta^j,$$

and view the PLP as a set of constraints on the coefficients $(h_{ij})_{ij}$ of the polynomials. Crucially, the dimension of the search space has been changed from infinite to finite.

Solving Feasibility Problem over Coefficients of Polynomials. Given the reduction to polynomial solutions described just above, we now investigate the form of the constraints on the coefficients $(h_{ij})_{ij}$ of the polynomial entries of $x(\delta)$. It turns out to be much easier to understand the constraints in a variant of local feasibility of the PLP where we restrict attention to the positive axis $\delta > 0$. The latter problem is equivalent to each of the m constraint polynomials $f_i(\delta) \stackrel{\text{def}}{=} a_i(\delta)x(\delta) - b_i(\delta)$ being either 0 or having its first non-zero derivative (in δ) be positive at zero. Let us call this condition the “derivative condition”.

We now express the derivative condition geometrically. Observe that $f_i(\delta)$ is a linear function of the $(h_{ij})_j$, hence so too is any order of derivative of $f_i(\delta)$. Each derivative being zero or positive is therefore a linear constraint on $(h_{ij})_j$ plus a halfspace constraint. The derivative condition on a particular $f_i(\delta)$ thus corresponds geometrically to a set S_i given by a union of half affine subspaces (halfspace intersected with an affine subspace), so the derivative condition is described by the intersection of the S_i .

Local feasibility for δ boils down to whether the intersection of the S_i sets is nonempty, in which case there exists a choice of coefficients $(h_{ij})_{ij}$ satisfying the derivative constraints. This problem will in turn be reduced to the *subspace elimination problem*, which asks whether a union of half affine spaces consists of the entire ambient space \mathbb{R}^N .

Subspace elimination is introduced in Section 5, and we show how to exploit the underlying geometry of the problem in order to solve it with ordinary linear programming as a subroutine. At

a high level, the algorithm iteratively processes and removes subspaces that have the potential to change the dimension of the current feasible set, which is updated after each step. Crucially, we can efficiently determine whether a subspace is relevant based on the dimension of its intersection with the current feasible set.

The Obstruction in Dimension $d \geq 2$. A similar approach does not work for local feasibility when δ has higher dimension than 1 and indeed we show that the problem is NP-hard. We now give an example illustrating the root of the issue. Suppose $\delta = (\delta_1, \delta_2)$ is 2-dimensional. It is possible that a PLP is locally feasible but there is no single rational function solution $x(\delta)$ which satisfies the program for all $\delta \in B_\epsilon(0)$. Consider the following simple example consisting of a single constraint:

$$\delta_1 \delta_2 (\delta_1 + \delta_2) x = \delta_1^2 \delta_2^2.$$

For any choice of δ_1 and δ_2 , the linear equation has a solution. But $x = 1/\delta_1 + 1/\delta_2$ is only feasible for $\delta_1 \neq 0$ and $\delta_2 \neq 0$. This example can be generalized so that even locally around the origin one must consider a partition into many regions with each rational function only working on a subset.

Computational Hardness in Dimension $d \geq 2$. To prove the hardness result, Theorem 2, we first show that the everywhere feasibility of 1-PLP can be reduced to the local feasibility of a 2-PLP by a mapping from \mathbb{R} to rays on \mathbb{R}^2 . So it suffices to prove the hardness for everywhere feasibility of 1-PLP,

$$\forall \delta \in \mathbb{R}, \exists x \in \mathbb{R}^n, A(\delta)x \geq b(\delta).$$

To make the format more aligned with a typical optimization problem that asks whether a solution exists, we can take the negation and LP dual, which gives us

$$\exists \delta, y \text{ satisfying } y \geq 0 \text{ and } A^\top(\delta)y = 0 \text{ with } b^\top(\delta)y > 0.$$

We will show through transformation of variables that this problem can be used to approximate an arbitrary *integer linear programming* instance

$$\exists x \in \{0, 1\}^n, Ax \geq b$$

with arbitrary precision. The precision of the approximation is enough to guarantee equivalence between any instance of the maximum independent set problem and everywhere feasibility of a polynomial-size instance of 1-PLP. Hardness of everywhere feasibility of 1-PLP then follows from NP-hardness of maximum independent set.

1.4 Outline

We review related literature in Section 1.5, including cylindrical algebraic decomposition and robust optimization. In Section 2, we give several additional results that are closely related to local feasibility and everywhere feasibility. Results for applications to PCA and broadcasting on 2D grid are in Section 3. Section 4 contains preliminaries and notation.

The proof and algorithm for Theorem 1 is in Section 5 and the proof of Theorem 2 is in Section 6. Detailed proofs and discussion for applications to PCA and broadcasting on 2D grid are given in Section 7 and Section 8.

1.5 Related Literature

Cylindrical Algebraic Decomposition To give intuition for why it is non-trivial to solve local feasibility of a 1-PLP, we now place the problem in the broader setting of a general polynomial system. An algebraic decomposition is the following process. Given a set of polynomials with n variables, decompose \mathbb{R}^n into regions where all polynomials have constant sign.

The local feasibility and everywhere feasibility of PLP is a special case of the above problem. To see this, we can choose the set of polynomials to be $\{a_i(\delta)x - b_i(\delta)\}$ where δ and x are both considered as variables. Then the feasible region of a PLP is the projection of $\{(\delta, x) : a_i(\delta)x - b_i(\delta) \geq 0, \forall i\}$ onto the space of δ .

The general algorithm for computing the decomposition is cylindrical algebraic decomposition (CAD), introduced by Arnon and Collins [ACM84] for quantifier elimination. The algorithm is powerful and solves any polynomial system, which includes local feasibility and everywhere feasibility of a general d -PLP. However, it is known to have doubly exponential running time: If d is the degree and m is the number of input polynomials, then the worst case running time of a general CAD is $dm^{2^{\Omega(n)}}$ [DH88, BD07, ED16]. On the contrary, our algorithm in Theorem 1 has a running time polynomial in n, d and m .

Robust Optimization and Robust Linear Programming The formulation of our problem is closely related to robust optimization, a field which aims to study optimization when parameters are uncertain [BBC11]. In robust optimization, the parameters of the optimization problem are assumed to reside in an uncertainty set, and the result should hold for any possible parameter. This formulation has broad applications in portfolio optimization, learning, and control [BBC11].

Robust linear programming is a basic problem in robust optimization. It is written as

$$\begin{aligned} & \text{minimize } c^\top x \\ & \text{subject to } Ax \geq b, \forall (A, b) \in \mathcal{U}. \end{aligned}$$

The tractability of the problem depends on the choice of \mathcal{U} , and the typical choices are polyhedron and ellipsoid. If \mathcal{U} is a polyhedron, the robust counterpart can be reduced to a normal LP problem [BTN99]. If \mathcal{U} is an ellipsoid, then the robust linear program is equivalent to a second-order cone program, which can again be solved efficiently [BTN99]. Other choices of uncertainty sets have also been studied [BPS04, BS04].

To compare with PLP, we can consider the feasibility of a robust linear program (the optimization of a robust linear program can be solved by binary search and an oracle for feasibility). The feasibility problem is then to answer

$$\begin{aligned} & \text{whether } \exists x \\ & \text{satisfying } Ax \geq b, \forall (A, b) \in \mathcal{U}. \end{aligned}$$

PLP can be written in a similar way, as

$$\begin{aligned} & \forall (A, b) \in \mathcal{A}, \text{ whether } \exists x \\ & \text{satisfying } Ax \geq b, \end{aligned}$$

where \mathcal{A} is a d -dimensional algebraic curve parameterized by $\delta \in \mathbb{R}^d$, or a small enough ball on the curve.

There are two main differences between PLP and robust linear programming. First, in PLP the feasible solution x itself is not required to be robust, and can vary with δ . Second, the uncertainty set is assumed to be a low-dimensional set having structure encoded by $A(\delta), b(\delta)$. In previous works, the uncertainty set is typically a geometrically much simpler set.

1.6 Future Directions

In this section we present some open directions. First, although the algorithm for Theorem 1 has polynomial running time, we believe the running time is far from optimal. So one possible direction is to improve the running time for local feasibility of 1-PLP.

Second, it is discussed in the beginning that the local feasibility of PLP could be used to solve the max-flow problem where edge capacity and demand are parameterized by $\delta \in \mathbb{R}$ if we care about local behavior of δ around 0. But it is not obvious whether it is a tractable problem with δ having higher dimensions.

One natural generalization from feasibility to optimization that is not discussed in this paper is the following problem when $\delta \in B_\epsilon(0)$:

$$\begin{aligned} & \max_{x(\delta)} c^\top(\delta)x(\delta) \\ & \text{subject to } A(\delta)x(\delta) \geq b(\delta). \end{aligned}$$

Here we want to get an optimal $x(\delta)$ as a function of δ . We know that the problem is NP-hard for dimension higher than 2 as it is harder than deciding feasibility. But for $d = 1$, can the answer be expressed efficiently? If so, is there a polynomial-time algorithm that solves it?

A conjecture in the field of probabilistic cellular automata (PCA) is the ergodicity of soldier's rule under arbitrarily small noise, discussed in Section 7.4. Our approach fails to decide the ergodicity for soldier's rule because the algorithm takes too much time on the PLP instance produced by soldier's rule. It remains open whether we can improve the running time for this specific instance of PLP and prove that the soldier's rule is ergodic under BSC noise.

PLP is a generalization of linear programming obtained by substituting coefficients with polynomials of δ . Other types of optimization problems may also admit such generalizations and it is interesting to ask whether the generalizations can be solved efficiently.

2 Additional Results

In this section, we discuss some additional results that are closely related to local feasibility and everywhere feasibility of PLP.

2.1 PLP Infeasibility

Unlike ordinary linear programming, where a program is feasible or infeasible, a PLP can be locally (everywhere) feasible, locally (everywhere) infeasible, or neither. Therefore, it also makes sense to ask whether the problems of local and everywhere *infeasibility* are tractable. It turns out that feasibility and infeasibility have similar computational complexities.

Theorem 5. *The local infeasibility of 1-PLP can be decided in polynomial time.*

This is a consequence of the proof of Theorem 1. In Section 5, we will show that $B_\epsilon(0)$ can be divided into 3 parts, the negative side, the origin, and the positive side. In each part, the PLP is either always feasible or always infeasible. The algorithm solves the three parts separately, which answers whether a PLP is locally feasible, locally infeasible, or neither.

We also have a corresponding hardness result, proved in Section 6.3.

Theorem 6. *The everywhere infeasibility of d -PLP is NP-hard for any d . The local infeasibility of d -PLP is NP-hard for any $d \geq 2$.*

2.2 Tractable Variations of Local Feasibility

Although the local feasibility of general d -PLP is NP-hard for $d \geq 2$, there are important special cases or variations where the problem becomes tractable.

2.2.1 The Case of Equality

Consider restricting every constraint of the PLP to be equality, yielding

$$A(\delta)x = b(\delta).$$

Although it is NP-hard to decide the local feasibility of 2-PLP, it is tractable in the equality case.

Theorem 7. *Let d be constant. Given as input an $m \times n$ matrix A and an m -dimensional vector b of polynomials in $\delta \in \mathbb{R}^d$, there is a polynomial time algorithm answering the question: is there an $\epsilon > 0$ such that for all $\delta \in B_\epsilon(0)$, there exists an x satisfying $A(\delta)x = b(\delta)$?*

The algorithm is given in Algorithm 3 in Section 9. It partitions the space of δ into a polynomial number of regions so that the pseudoinverse of A can be expressed as a rational function on each region.

The hardness result, Theorem 2, immediately tells us that a similar partition cannot be done for a general PLP. In fact, in Section 5, we will see that to solve the general local feasibility of PLP with a similar approach, all subsets of rows of (A, b) must be considered and this leads to an exponential number of possible regions.

2.2.2 The Case of Strict Inequality

Another natural way to modify the local feasibility problem is to require that all constraints be strict inequalities, yielding

$$A(\delta)x > b(\delta).$$

This modification again results in local feasibility being a tractable problem.

Theorem 8. *Let d be constant. Given as input an $m \times n$ matrix A and an m -dimensional vector b of polynomials in $\delta \in \mathbb{R}^d$, there is a polynomial time algorithm answering the question: is there an $\epsilon > 0$ such that for all $\delta \in B_\epsilon(0)$, there exists an x satisfying $A(\delta)x > b(\delta)$?*

The theorem is a corollary of the following observation.

Lemma 1. *$A(\delta)x > b(\delta)$ is locally feasible if and only if $A(0)x > b(0)$.*

Proof. One direction is immediate, since local feasibility implies feasibility for all δ in a neighborhood of 0 (and this includes 0). Now we prove that if $A(0)x > b(0)$, then $A(\delta)x > b(\delta)$ is locally feasible. Fix x such that $A(0)x > b(0)$. Let $f_i(\delta) = A_i(\delta)x - b_i(\delta)$, and note that this is a polynomial in δ and hence continuous. Since $f_i(0) > 0$, there exists $\epsilon > 0$ such that $f_i(\delta) > 0$ for all $\delta \in B_\epsilon(0)$. This holds for all i , which proves the lemma. \square

Now, checking whether $A(0)x > b(0)$ is feasible is a linear program with strict inequality constraints. This is a special case of the subspace elimination problem discussed in Section 5.3.

3 Applications

Aside from being a natural way of modeling uncertainty in linear programming, PLP can serve as a tool for solving theoretical problems. We will discuss two such problems. First, the ergodicity of elementary PCA with NAND function and BSC noise. Second, whether broadcasting of information on a 2D grid with arbitrary function and BSC noise allows recovery at infinity. We will use our algorithm for the local feasibility of 1-PLP to perform computer assisted proofs answering both questions.

3.1 Probabilistic Cellular automata

A probabilistic cellular automaton (PCA) is specified by a random rule for updating cells on an infinite grid. The update rule is local and homogeneous in time and space. At each time-step each site updates its state based on the states of nearby sites using the local rule. The deterministic version of PCA, cellular automata, was first introduced by John von Neumann and S. Ulam as a biologically motivated model for computing in the 1940s [NB66]. A PCA resembles the human brain in the sense that the computation is massively parallel and local. PCA has been viewed as an important model for fault-tolerant computation [Too95, Gác86], and it also has broad applications in modeling physical and biological processes [CD98, EEK93].

One of the main questions about a PCA is whether it is ergodic [MM14, LN18, TVS⁺90]. A PCA is ergodic if there exists a unique stationary distribution so that starting with any initial configuration, the PCA converges to this distribution. In this section we use the methodology of polynomial linear programming developed in this paper to prove a new ergodicity result for a certain class of PCA.

One-dimensional PCA. In this paper we consider one-dimensional discrete-time PCA. The PCA has alphabet A , so at any time $t \in \mathbb{N}$, the configuration of the PCA is denoted by $\eta_t \in A^{\mathbb{Z}}$ with $\eta_t(n)$ the value on site $n \in \mathbb{Z}$ at time t . The distribution of η_t is a Markov chain as t changes. Given η_t , the next configuration η_{t+1} is obtained by updating each site $n \in \mathbb{Z}$ independently. The update rule is local, random, and homogeneous. That is,

$$\eta_{t+1}(n) = f(\eta_t(n-a), \eta_t(n-a+1), \dots, \eta_t(n+a)),$$

for some a and where f is a random function with a fixed distribution independent of t and n .

A special case is *elementary PCA* (see Figure 2a). In this case the alphabet is $\{0, 1\}$ and at each time step, the state at each location is set to a random Boolean function of its left neighbor and itself, i.e.,

$$\eta_{t+1}(n) = f(\eta_t(n-1), \eta_t(n)).$$

Another important example is *soldier's rule*. In soldier's rule, the alphabet is $\{-1, 1\}$, indicating the direction a soldier is facing. At each time step, each soldier sets its new direction based the majority vote of the direction of itself and the direction of the first neighbor and the third neighbor in its direction (see Figure 2b). Thus,

$$\eta_{t+1}(n) = \text{Maj}(\eta_t(n), \eta_t(n + \eta_t(n)), \eta_t(n + 3\eta_t(n))).$$

The probabilistic version of soldier's rule adds noise to the outcome, so each soldier has probability p of being -1 and probability q of being 1 at each time step regardless of the configuration in the previous time step.

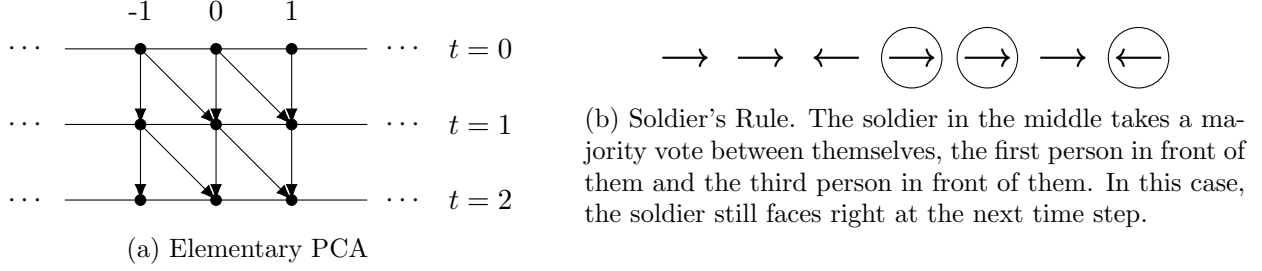


Figure 2: Probabilistic Cellular Automata

Ergodicity. Now we formally define ergodicity. Let F denote the transition of the Markov chain based on the update rule for η_{t+1} in terms of η_t .

Definition 3.1. A PCA F is said to be *ergodic* if there is a unique attractive invariant measure μ , i.e., there exists μ such that

1. $\mu F = \mu$
2. for any probability measure ν over configurations, νF^n converges weakly to μ as $n \rightarrow \infty$.

Asking whether a PCA is ergodic is similar to the question of reconstruction for broadcasting on trees (see, e.g., [EKPS00]). Recall that in broadcasting on trees, the root node starts with a random state which is then propagated outward on the tree according to a given Markov transition on each edge. The reconstruction question asks whether information about the root propagates infinitely far. If we view the layers of a tree as time steps, then the reconstruction problem equivalently asks whether all initial configurations converge to the same distribution as time goes to infinity.

Despite their similarity, unfortunately, techniques from broadcasting on trees do not directly apply to PCA. Given an observation at some level of the tree, the posterior at the root can be computed recursively via belief propagation, and in general the marginal distribution at any node can be expressed as a function of its children's marginals (see, e.g., [BRZ95]). In contrast, in PCA each node is affected by multiple inputs. Moreover, given a fixed noise model, the ergodicity of a PCA still depends on the specific form of the update rule, while broadcasting on trees depends only on the tree structure.

A central open problem. It has been an important open problem to decide the ergodicity of all symmetric elementary PCAs [TVS⁺90, Chapter 7]. It is conjectured that all non-deterministic PCAs are ergodic [DPSV69]. There are three parameters for a symmetric elementary PCA: the probability to output 1 when the input is 00, 01 or 11. The conjecture was proven for most of the region of possible parameters [TVS⁺90, Chapter 7]. Specifically, regions around any deterministic functions except NAND function leads to ergodicity. Therefore, one important region of parameters where the question remained open was the corner (i.e., small noise) around the NAND function.

In [HMM19], a potential function method was proposed for analyzing the ergodicity of PCA. The update rule they considered was a noisy NAND function with vertex noise: with probability $1 - 2p$ the output is the NAND of the input, and with probability $2p$ the output is a random bit. In other words,

$$\eta_{t+1}(n) = \text{BSC}_p(\text{NAND}(\eta_t(n-1), \eta_t(n))), \quad (2)$$

where BSC_p stands for the binary symmetric channel with crossover probability p .

The function in Equation (2) can be thought of as vertex noise, as the noise is added after the NAND function. An alternate version is that noise occurs as the information is passing from time t to time $t + 1$. In other words,

$$\eta_{t+1}(n) = \text{NAND}(\text{BSC}_p(\eta_t(n-1)), \text{BSC}_p(\eta_t(n))). \quad (3)$$

There is no obvious way to relate which noise (edge or vertex) leads to faster or more likely mixing. However, it is known that mixing for PCA with vertex noise of level p is implied by mixing of the PCA with edge noise of level $p/2$ (but a slightly different processing function), see Lemma 3.1 of [DO77]. Thus, intuitively, showing ergodicity for arbitrarily small edge-noise should be harder.

Potential function method. In [HMM19], the potential function essentially counts a weighted sum of the probability that certain patterns appear. The idea is to use a coupling argument to control the distributional convergence as follows. Starting two chains with different initial conditions, one uses an auxiliary Markov chain to keep track of the locations where the two chains have coupled or have not coupled. If one can show that the density of uncoupled locations tends to zero then this implies convergence, and the potential is used to count the number of uncoupled locations by keeping track of certain patterns. The crucial property of a good potential function is that it is non-increasing and decreases in value over time when the chosen pattern appears.

While we know some sufficient conditions on a potential function that leads to ergodicity, the design of a good potential function can be challenging. Although the potential function method is not restricted to the specific choice of random function, the design of a potential function in [HMM19] was tailored specifically for NAND function with vertex noise. Therefore, a natural question is whether there exists a standard procedure to search for potential functions, regardless of function in the update rule of the PCA.


In this paper our approach is to encode the condition into a PLP which can then be solved on a computer to yield a good potential function. The idea of reducing the condition on such potential functions to a linear program is first discussed in [MMP21]. The key observation is that transition of the PCA is linear for a fixed noise level p , thus the evolution of the potential function is also linear. If the coefficients of the potential function are viewed as variables, these conditions turn out to be linear inequalities with parameter p , and the dependence on p is polynomial. Therefore, the feasibility of a PLP corresponds to the existence of such potential function. The PLP is then fed to our algorithm for the local feasibility, which certifies the existence of a potential that meets all conditions and outputs a potential that in principle can be checked by hand.

Our results. We prove the ergodicity of PCA that uses NAND function with both vertex noise and edge noise for small enough error probability.


Theorem 9. *For the PCAs described in (2) and (3), there is an $\epsilon > 0$ such that the PCA is ergodic for any p on $(0, \epsilon)$.*

The proof of the theorem, based on analysis of the potential function method and its connection to feasibility of PLP, is in Section 7 and the (\mathbb{Q} -rational) potentials are given on Appendix C.

It is worth noting that our approach of finding potential functions is not limited to PCA with two inputs, but is also a possible method to try on any 1-dimensional PCAs. In Section 7.4, we will also discuss an attempt to prove ergodicity for soldier’s rule and the difficulties encountered.



(a) Broadcasting of Information on 2D Grid



(b) We can consider just one of the quadrants without loss of generality.

Figure 3: Broadcasting on 2D grid

3.2 Broadcasting of Information on a 2D-Grid

Broadcasting of information on a grid is a model inspired by broadcasting on trees. In the setting of broadcasting on trees, there is an information source that generates one bit of information. The information spreads through a noisy tree with bounded degree where edges are binary symmetric channels with crossover probability δ . The central question is whether the information propagates out to infinity, i.e., whether the mutual information between the root and the frontier of broadcasting converges to 0. If the information is not completely lost, then we say that *reconstruction* is possible. It is shown in a series of results [BRZ95, EKPS00] that reconstruction is possible for small enough noise if and only if $(1 - 2\delta)^2 \text{br}(T) > 1$, where $\text{br}(T)$ is the branching factor of the tree.

Broadcasting on the grid. Broadcasting of information was generalized from trees to general directed acyclic graphs (DAGs) in [MMP19], with the case of the grid being an important specialization. However, because each node is affected by multiple inputs, like PCA, the recursive approach in [BRZ95] also does not apply to broadcasting on 2D grid.

Let us describe the problem formally. Figure 3a illustrates how information is broadcasted on 2D grid. Since there is no edge from one quadrant to another, we can restrict attention to the non-negative quadrant. Suppose there is one bit of information at the origin and we want to broadcast the bit to the entire infinite 2-dimensional grid. Each node on the grid is indexed by a pair (t, i) ($t \geq 0, 0 \leq i \leq t$). At any time $t \geq 1$, nodes with index (t, i) , where $0 < i < t$, receives the bit from $(t-1, i-1)$ and $(t-1, i)$ and computes its own bit with a fixed Boolean function f . Nodes with index $(t, 0)$ and (t, t) receive the bit from $(t-1, 0)$ and $(t-1, t-1)$, respectively, and keep what it receives as its own bit. When a bit is passed from one node to another, it passes through a BSC channel with parameter p .

Use η_t^+ and η_t^- to denote the Markov chain containing bits (t, i) for $0 \leq i \leq t$ that starts with bit 1 and bit 0 respectively at the origin. We are interested in whether the information of the initial bit is lost at infinity, i.e.,

$$\lim_{t \rightarrow \infty} d_{\text{TV}}(P_{\eta_t^+}, P_{\eta_t^-}) = 0.$$

If we view the coordinate t as time, the problem is similar to elementary PCA because each node at level t receives information from two adjacent nodes from $t-1$. The only difference is that

broadcasting on the grid has a bounded-length configuration, and the behavior is different at the boundary.

The potential method in [HMM19] can be extended to this problem, as discussed in [MMP21]. Instead of counting the probability of certain pattern appearing, the potential now counts the expected number of appearance of a certain pattern at time t . [MMP21] also showed the close relation between LP and finding the potential function. The condition on the potential is linear in the transition probabilities. Transition probabilities are polynomial in the noise parameter so the condition can be written as the feasibility of a PLP. We will utilize our algorithm for local feasibility to find a potential function and prove the following theorem.

Theorem 10. *For any function $f : \{0, 1\}^2 \rightarrow \{0, 1\}$, there exists $\epsilon > 0$ such that for any $p \in (0, \epsilon)$,*

$$\lim_{t \rightarrow \infty} d_{\text{TV}}(P_{\eta_t^+}, P_{\eta_t^-}) = 0.$$

The theorem is proved in Section 8 and the computer-discovered potentials (with entries in \mathbb{Q}) are given in (28) and (29).

One interesting open question is whether recovery is possible for higher dimensional grids. For grids with dimension 3 or more, each node has more outputs, so recovery might be easier analogously to the situation for broadcasting on trees of higher degree. Indeed, recovery is conjectured to be possible in this setting using the Majority function, in Section II.D of [MMP21].

4 Preliminaries

4.1 Polynomial and Rational Functions

We now provide several preliminary definitions and results that will be useful in the proofs.

Definition 4.1. A degree- D *polynomial* in $\delta \in \mathbb{R}^d$ is a function with the following form

$$p(\delta) = \sum_{\sum_{j=1}^d i_j \leq D, i_j \in \mathbb{Z}^+ \cup \{0\}} C_{i_1, i_2, \dots, i_d} \prod_{j=1}^d \delta_j^{i_j}.$$

Here $(C_{i_1, i_2, \dots, i_d})$ are the coefficients of the polynomial.

A degree- D *rational function* is a function with the following form

$$h(\delta) = \frac{p(\delta)}{q(\delta)},$$

where p and q are degree- D polynomials, defined on $\mathbb{R}^d \setminus \{\delta | q(\delta) = 0\}$.

Throughout the paper, all polynomials and rational functions are assumed to have rational coefficients.

Definition 4.2. The *size* of a polynomial or rational function is the number of bits used to describe it. If a polynomial or rational function is a function of n variables, its degree is bounded by D , and the number of bits of any coefficient is bounded by L , then its size is $O(nDL)$.

4.2 Pseudo Inverse

For a matrix $A \in \mathbb{R}^{n \times m}$, its pseudo inverse is denoted by A^+ .

Definition 4.3. $A^+ \in \mathbb{R}^{m \times n}$ is a pseudo inverse of A if and only if

1. $AA^+A = A$
2. $A^+AA^+ = A^+$
3. AA^+ and A^+A are symmetric

The pseudo inverse of a matrix tells us whether a linear system is feasible.

Lemma 2. *A linear system $Ax = b$ has a solution if and only if $AA^+b = b$, in which case A^+b is a solution.*

Proof. If $AA^+b = b$ holds then A^+b is a solution by definition. Conversely, assume $Ax = b$ has solution x_0 . By definition of pseudo inverse, $b = Ax_0 = AA^+Ax_0 = AA^+b$. \square

The pseudo inverse can be expressed as a limit, which will be useful in our proof (see, for example [Alb72]).

Lemma 3. *For any $A \in \mathbb{R}^{n \times m}$, its pseudo inverse A^+ always exists and is unique. Also,*

$$A^+ = \lim_{y \rightarrow 0} (A^\top A + yI)^{-1} A^\top.$$

4.3 Useful Lemmas in Computational Algebra

We will use the following standard result in computational algebra.

Lemma 4 ([BM04]). *If x_1 and x_2 are two distinct roots of a polynomial p with size s , then $|x_1 - x_2| \geq 2^{-\text{poly}(s)}$.*

Lemma 5 ([CR88]). *Let δ_0 be a real root of an irreducible polynomial p (irreducible in \mathbb{Q}). δ_0 is uniquely determined by p and signs of all derivatives of p at δ_0 . The sign of $q(\delta_0)$ for another polynomial q can be checked in time polynomial in size of p and q .*

4.4 Notation

For an $m \times n$ matrix A we denote by A_{-i} the $(m-1) \times n$ matrix obtained by removing row i of A . For a subset $J \subseteq [m]$, the $|J| \times n$ matrix A_J consists of the rows of A indexed by J .

The *kernel* of a matrix $C \in \mathbb{R}^{m \times n}$, denoted by $\ker(C)$, consists of those $w \in \mathbb{R}^n$ such that $Cw = 0$. The *row space* of matrix C is denoted by $\text{row}(C)$, which is the space spanned by all row vectors of C . The subspaces $\text{row}(C)$ and $\ker(C)$ are orthogonal complements to each other by definition.

5 A Polynomial Time Algorithm for Local Feasibility of 1-PLP

In this section we prove Theorems 1 and 5 by giving a polynomial-time algorithm for local feasibility of 1-PLP. We begin in Sections 5.1 and 5.2 by showing that although the solution x can in principle be any function of the parameter δ , it suffices to focus on x being polynomials. Then in Section 5.3 we introduce a problem called subspace elimination, reduce determining local feasibility of 1-PLP to feasibility of subspace elimination, and give an efficient algorithm for the latter problem. Section 5.4 improves the runtime of the algorithm from general subspace elimination by using the structure specific to local feasibility of PLP. Section 5.5 shows how to find a locally feasible solution for the PLP as a rational function.

We now briefly recall the problem formulation and introduce a simplification. A is an $m \times n$ polynomial matrix and b is an $m \times 1$ polynomial vector. Suppose the degree of polynomials in A and b are bounded by D with a 1-dimensional parameter δ . We want to decide whether there exists $x(\delta)$ that satisfies the program

$$A(\delta) \cdot x(\delta) - b(\delta) \geq 0, \quad (4)$$

for all δ in an arbitrarily small neighborhood $(-\epsilon, \epsilon)$ of 0, and output $x(\delta)$ when the answer is yes.

It turns out to be simpler, and sufficient, to derive an algorithm that decides the PLP is feasible on the positive side of the origin, i.e., whether there exists an ϵ such that the PLP is feasible in $(0, \epsilon)$. The same algorithm can be applied to the negative side. The two results can then be combined with a normal LP at the origin to decide if the PLP is feasible or infeasible in $B_\epsilon(0)$ for a small enough ϵ . We will prove the following lemma and Theorems 1 and 5 follow as corollaries.

Lemma 6. *The problem of whether a 1-PLP is feasible on $(0, \epsilon)$ for a small enough ϵ can be solved in polynomial time. Furthermore, if the PLP is locally feasible, a solution $x(\delta)$ can be given as a rational function of δ on the positive neighborhoods of 0. The interval $(0, \epsilon)$ on which $x(\delta)$ is a feasible solution is given by ϵ being the smallest positive root of an explicit polynomial.*

5.1 Rational Function Solution

The following theorem states that feasibility implies existence of a rational function solution, and gives a constructive way of finding such a function.

Theorem 11. *If there exists $\epsilon > 0$ such that for any $\delta \in (0, \epsilon)$ there exists a $x_1(\delta)$ such that (4) is satisfied, then there exists a rational function $x(\delta)$ that satisfies (4) for $\delta \in (0, \epsilon')$ for a positive ϵ' . Further, we can take $x(\delta)$ equal to $A_J^+ b_J$ for a subset of constraints $J \subseteq [m]$, where A_J^+ denotes the pseudoinverse of the matrix A_J consisting of the rows of A indexed by J .*

The numerators and least common multiple (LCM) of denominators for all entries in $x(\delta)$ has degree at most $2nD$.

Note that A_J^+ is not always a rational function, but we will show in Lemma 8 that this is indeed the case if we choose a small enough ϵ .

We begin with the observation that when a linear program is feasible, a solution can always be found by changing a subset of the inequality constraints to equalities and dropping the rest. The solution can then be expressed as a rational function of the coefficients via the pseudoinverse.

Lemma 7. *For any feasible LP $L : Ax \geq b$, there exists a subset of constraints J such that $\{x : A_J x = b_J\}$ is a subset of the solution set $\{x : Ax \geq b\}$.*

Proof. The proof is obtained by inductively applying the following claim. Let S be a set defined by linear inequalities $Ax \geq b$ and equations $Cx = d$. We claim that either $S = \{x : Cx = d\}$, i.e. none of the inequalities are active, or there exists an inequality $a_i x \geq b_i$ such that

$$S' = S \cap \{x : a_i x = b_i\} = \{x : A_{-i}x \geq b_{-i}, Cx = d, a_i x = b_i\} \text{ is nonempty.}$$

In words, for any feasible set of constraints we can find an inequality and change it to an equality while preserving feasibility.

Now we prove the claim. We first suppose that for all $w \in \ker(C) \setminus \{0\}$, $\max_{x \in S} w^\top x = +\infty$ and show that in this case $\{x : Cx = d\} = S$. Note that $\max_{x \in S} -a_i x \leq -b_i < +\infty$, so for any i , $a_i \notin \ker(C)$. Because the kernel and row space are orthogonal complements, $a_i = a_i^0 + a_i^1$ where $a_i^0 \in \ker(C)$ and $a_i^1 \in \text{row}(C)$. Suppose $a_i^1 = \sum_j \alpha_j C_j$. Taking any x satisfying $Cx = d$, $a_i^1 x = \sum_j \alpha_j C_j x = \sum_j \alpha_j d_j$. Hence $\max_{x \in S} -a_i^0 x = \max_{x \in S} (-a_i + a_i^1)x \leq -b_i + \sum_j \alpha_j d_j < +\infty$ and we conclude that $a_i^0 = 0$. Now, since S is non-empty, we may take a $y \in S$ to conclude that $b_i \leq a_i y = a_i^1 y = \sum_j \alpha_j d_j$ for all i . So for any x in $\{x : Cx = d\}$ and any i , we have $a_i x = a_i^1 x = \sum_j \alpha_j d_j \geq b_i$, which implies that $\{x : Cx = d\} = S$.

We henceforth assume that there exists $w_0 \in \ker(C)$ such that $\max_{x \in S} w_0^\top x < +\infty$. Let $x^* \in S$ achieve this maximum. Note that there must exist an i such that $a_i x^* = b_i$, otherwise we can increase $w_0^\top x$ by choosing $x = x^* + \epsilon w_0$ for a small enough $\epsilon > 0$. So $x^* \in \{x : A_{-i}x \geq b_{-i}, Cx = d, a_i x = b_i\} =: S'$, which means that S' is nonempty. \square

Corollary 1. *Under the condition of Theorem 11, there exists $\epsilon > 0$ that for any $\delta \in (0, \epsilon)$ there exists a subset of constraints J (possibly depending on δ) such that $A_J^+ b_J$ is a solution.*

Proof. The proof follows by combining Lemmas 7 and 2. \square

Informally, Corollary 1 tells us that for fixed δ , x can indeed be chosen as a rational function. The conclusion can be extended to a small enough open interval, as shown in the following lemma.

Lemma 8. *For a small enough ϵ , it holds for every $J \subset [m]$ that $A_J^+(\delta) b_J(\delta)$ is a rational function of δ on $(0, \epsilon)$. The numerators and LCM of denominators for all entries in $A_J^+(\delta) b_J(\delta)$ has degree at most $2nD$.*

Proof. By Lemma 3, $A_J^+ = A_J^+(\delta)$ has the following limit form:

$$A_J^+ = \lim_{y \rightarrow 0} (A_J^\top A_J + yI)^{-1} A_J^\top.$$

Now, A_J is a $|J| \times n$ matrix, so $A_J^\top A_J + yI$ is an $n \times n$ matrix. Each entry of $A_J^\top A_J + yI$ is a polynomial of δ and y with degree at most $2D$. From the analytic formula for the inverse of a matrix, it follows that entries of $(A_J^\top A_J + yI)^{-1}$ are rational functions with denominator of degree at most $2nD$ and numerator of degree at most $2(n-1)D$. Thus, $(A_J^\top A_J + yI)^{-1} A_J^\top$ is an $n \times |J|$ matrix and the entries are rational functions of δ and y with denominator and numerator of degree at most $2nD$.

It remains to take the limit $y \rightarrow 0$. Taking entry i, j of $(A_J^\top A_J + yI)^{-1} A_J^\top$, write it as

$$f(\delta, y) = \frac{(-1)^{i+j} M_{j,i}}{\det(A_J^\top A_J + yI)}$$

where $M_{j,i}$ stands for the minor of matrix $A_J^\top A_J + yI$ in the j th row and i th column. Suppose

$$(-1)^{i+j} M_{j,i} = y^{a_{ij}} (p_0^{(ij)}(\delta) + p_1^{(ij)}(\delta)y + \dots), \quad \text{and}$$

$$\det(A_J^\top A_J + yI) = y^a(q_0(\delta) + q_1(\delta)y + \dots).$$

We have $a_{ij} \geq a$ since the limit always exists (by virtue of the pseudoinverse existing). Therefore,

$$\lim_{y \rightarrow 0} f(\delta, y) = \frac{\mathbb{I}\{a = a_{ij}\} p_0^{(ij)}(\delta)}{q_0(\delta)} \quad (5)$$

for δ such that $q_0(\delta) \neq 0$. $q_0(\delta)$ is not the zero function, so there exists an interval $(0, \epsilon)$ such that for $\delta \in (0, \epsilon)$, $q_0(\delta) \neq 0$. We get that on $(0, \epsilon)$, $A_J^+(\delta)b_J(\delta)$ is a vector of rational functions with degree at most $2nD$, given by

$$\frac{\mathbb{I}\{a = a_{ij}\} p_0^{(ij)}(\delta)}{q_0(\delta)}.$$

q_0 is a common multiple for all denominators in the vector, so the LCM for denominators of all entries has degree bounded by $2nD$. \square

Lemma 9. *For a small enough ϵ , there exists a fixed set $J \subset [m]$ such that that $x(\delta) = A_J^+(\delta)b_J(\delta)$ is a solution of (4) for all $\delta \in (0, \epsilon)$.*

Proof. Define a set of functions

$$g_J(\delta) \triangleq \begin{cases} 1, & \text{if } A_J^+(\delta)b_J(\delta) \text{ is a solution} \\ 0, & \text{otherwise.} \end{cases}$$

Now let $S(\delta) = \{J \subset [m] : g_J(\delta) = 1\}$. From Corollary 1, $S(\delta)$ is always non-empty for $\delta \in (0, \epsilon)$. From Lemma 10 below, $S(\delta)$ can change only finite number of times on \mathbb{R}^+ . So we can pick ϵ small enough such that $S(\delta)$ is non-empty and remains constant on $(0, \epsilon)$. Pick one J from $S(\delta)$, and $A_J^+(\delta)b_J(\delta)$ is the desired solution. \square

Proof of Theorem 11. By Lemma 9, we can choose a small enough ϵ_1 so that there is a set $J \subset [m]$ for which $x(\delta) = A_J^+(\delta)b_J(\delta)$ is a solution for all $\delta \in (0, \epsilon_1)$. By Lemma 8, there exists an ϵ_2 such that $A_J^+(\delta)b_J(\delta)$ is a rational function for $\delta \in (0, \epsilon_2)$. Now let ϵ' be the smaller of ϵ_1, ϵ_2 . \square

The following lemma was used in the proof of Lemma 9 above.

Lemma 10. *For any $J \subset [m]$, $g_J(\delta)$ has a finite number of value changes as δ varies over \mathbb{R} .*

Proof. Note that $g_J(\delta) = 1$ is equivalent to satisfying for each $i \in [m]$ the inequality

$$a_i A_J^+ b_J - b_i \geq 0.$$

By Equation 5, $a_i A_J^+ b_J - b_i$ is a rational function of δ except at a finite number of points, so its sign can only change a finite number of times. When g_J changes value, the sign of $a_i A_J^+ b_J - b_i$ changes for at least one i , which means g_J has a finite number of value changes. \square

One of the consequences of this lemma is that a PLP must be either feasible or infeasible on the positive neighborhood of 0.

Corollary 2. *Given a 1-PLP, we can divide \mathbb{R} into a finite number of intervals such that the PLP is always feasible or infeasible when δ is within an interval. Therefore, there exists an $\epsilon > 0$ such that the PLP is either feasible or infeasible on $(0, \epsilon)$.*

Proof. By Lemma 10, for each $J \in [m]$, $g_J(\delta)$ has a finite number of value changes on \mathbb{R} . So we can divide \mathbb{R} into a finite number of intervals such that on each interval $g_J(\delta)$ is constant for all J . Then by Lemma 7, whether the PLP is feasible is also fixed. \square

5.2 Polynomial Solution and Reformulation of Feasibility Problem

In this subsection we show that when the 1-PLP is locally feasible, it has a polynomial solution $x(\delta)$. We then use this to reformulate the feasibility problem in terms of a condition on the coefficients of the polynomials.

Polynomial Solution Starting from the conclusion of Theorem 11, let $x(\delta)$ be the rational solution of (4). Suppose the LCM of all denominators of $x_i(\delta)$ is $q(\delta)$ and $p(\delta) = q(\delta) \cdot x(\delta)$. The degree of q is bounded by $2nD$.

Assume $q(\delta) = \delta^c(1 + \delta q_1(\delta))$ where $0 \leq c \leq 2nD$. We can rewrite (4) in the form

$$A(\delta) \cdot p(\delta) - \delta^{c+1} \cdot q_1(\delta) \cdot b(\delta) - \delta^c \cdot b(\delta) \geq 0, \quad (6)$$

By Theorem 11, it is a vector of polynomials of degree bounded by $2nD$. Now we have a new instance of the problem which is guaranteed to have a polynomial solution (for at least one c). So we can solve the problem in the space of the coefficients of $p(\delta)$ and $q_1(\delta)$. The total number of variables then becomes $O(n^2D)$.

Reformulation of Local Feasibility We will solve the local feasibility of Equation 6 for all possible c in $0 \leq c \leq 2nD$, the original problem would be feasible if and only if one of the instances is feasible. By letting

$$x'(\delta) = (p(\delta), q_1(\delta))^\top$$

and

$$A'(\delta) = \begin{pmatrix} A(\delta) & 0 \\ 0 & -\delta^{c+1}b(\delta) \end{pmatrix},$$

Equation 6 can be simplified to

$$A'(\delta)x'(\delta) - \delta^c \cdot b(\delta) \geq 0 \quad (7)$$

Let

$$x'_i(\delta) = \sum_{j=0}^{2nD} h_{ij} \delta^j,$$

the goal is to decide feasibility in the possible space of $(h_{ij})_{ij}$.

Let $f_i(\delta) = A'_i(\delta) \cdot x'(\delta) - \delta^c b_i(\delta)$, and $[f_i]_j$ be the coefficient of δ^j in f_i . $[f_i]_j$ is a linear function of $(h_{kl})_{kl}$. Then $f_i \geq 0$ around $\delta = 0$ is equivalent to that the first non-zero coefficient of f_i is larger than 0 at 0 or $f_i = 0$. This translates to the OR of some linear constraints over $(h_{kl})_{kl}$. Thus, considering all the constraints, we get that the PLP is locally feasible if and only if

$$\begin{aligned} \text{For every } i \in [m] : \quad & [f_i]_0 > 0 \\ & \text{OR } [f_i]_0 = 0, [f_i]_1 > 0 \\ & \text{OR } [f_i]_0 = 0, [f_i]_1 = 0, [f_i]_2 > 0 \\ & \vdots \\ & \text{OR } [f_i]_0 = 0, [f_i]_1 = 0, \dots, [f_i]_{d-1} = 0, [f_i]_d' \geq 0, \end{aligned} \quad (8)$$

where D' is the total degree of $f_i(\delta)$, $D' \leq 2nD$.

The algorithmic task is reduced to the following: For all c between 0 and $2nD$, check whether there exists $(h_{kl})_{kl}$ satisfying (8).

5.3 Subspace Elimination

In this section, we will reduce the problem in Section 5.2 to a new problem called subspace elimination and solve the problem in polynomial time.

5.3.1 Definition of Subspace Elimination

Definition 5.1. In the *subspace elimination* problem, the input is a collection B_1, \dots, B_m of non-empty open half-spaces of subspaces in \mathbb{R}^n , where each B_i is defined by the following equalities and inequality:

$$\begin{aligned} a_{i,1}x &= b_{i,1} \\ a_{i,2}x &= b_{i,2} \\ &\vdots \\ a_{i,k_i}x &= b_{i,k_i} \\ c_i x &< d_i. \end{aligned}$$

The goal is to decide whether or not

$$\mathbb{R}^n \setminus \left(\bigcup_{i=1}^m B_i \right)$$

is empty.

Note that this definition is a generalization of LP feasibility. It is also a generalization of the problem we have in Section 5.2.

Lemma 11. *The feasibility of (8) is a special case of the subspace elimination problem.*

Proof. To see this, let us take the negation of (8), which becomes

$$\begin{aligned} \exists i, [f_i]_0 &< 0 \\ \text{OR } [f_i]_0 &= 0, [f_i]_1 < 0 \\ \text{OR } [f_i]_0 &= 0, [f_i]_1 = 0, [f_i]_2 < 0 \\ &\vdots \\ \text{OR } [f_i]_0 &= 0, [f_i]_1 = 0, \dots, [f_i]_{d-1} = 0, [f_i]'_d < 0. \end{aligned} \tag{9}$$

One can check that this is indeed the negation of (8) by taking the union of corresponding terms. Then the negation of the AND of (8) for all i becomes the OR of (9) for all i , which coincides with the definition of subspace elimination. \square

Therefore, it is sufficient to find a polynomial-time algorithm for subspace elimination.

5.3.2 Algorithm for Subspace Elimination

Theorem 12. *The subspace elimination problem defined in Definition 5.1 can be solved by Algorithm 1 with running time $O(m^2 K \cdot \text{LP}(n, mK, L))$, where $K = \max_i k_i$ is the maximum number of equalities among all subspaces, $\text{LP}(a, b, c)$ stands for the running time of ordinary linear programming with n variables, m constraints and L -bit numbers*

Algorithm 1 Subspace Elimination

```

1: input:  $\mathcal{B} = \{B_i : i \in [m]\}$ 
2:  $P \leftarrow \mathbb{R}^n, \mathcal{A} \leftarrow \mathcal{B}$ 
3: while  $P \neq \emptyset$  and  $\exists B \in \mathcal{A}, P \subseteq \text{aff}(B)$  do
4:   Find  $B \in \mathcal{A}$  such that  $P \subseteq \text{aff}(B)$ 
5:    $P \leftarrow P \setminus B$ 
6:    $\mathcal{A} \leftarrow \mathcal{A} \setminus \{B\}$ 
7: end while
8: if  $P = \emptyset$  then
9:   return False
10: else
11:   return True
12: end if

```

Let $\text{aff}(A)$ denote the affine hull of a set A . The algorithm for subspace elimination is Algorithm 1.

Let us explain some details of the algorithm. During the process, P is always represented by a linear program, i.e., a set of linear constraints. First, we start with P having no constraints. Whenever we change P in line 5, the chosen B_i in line 3 satisfies $P \subseteq \text{aff}(B_i)$. So $P \setminus B_i$ is simply P with a new constraint $c_i x \geq d_i$. To check whether $P \subset \text{aff}(B_i)$ in line 3, notice that $\text{aff}(B_i)$ is the set of all x satisfying

$$\begin{aligned}
a_{i,1}x &= b_{i,1} \\
a_{i,2}x &= b_{i,2} \\
&\vdots \\
a_{i,k_i}x &= b_{i,k_i}.
\end{aligned}$$

We can check whether $P \subset \text{aff}(B_i)$ by solving whether

$$\max_{x \in P} a_{i,j}x = \min_{x \in P} a_{i,j}x = b_{i,j}$$

for all $1 \leq j \leq k_i$. The operation of deciding whether P is empty is given by solving a normal LP.

Proof of Theorem 12. To prove the correctness of the algorithm, We need to show that $P \setminus (\bigcup_{D \in \mathcal{A}} D) \neq \emptyset$ if $P \not\subseteq \text{aff}(D)$ for any $D \in \mathcal{A}$. This means that $\text{aff}(P) \not\subseteq \text{aff}(D)$. Observe that

$$\dim(P \cap D) \leq \dim(\text{aff}(P) \cap \text{aff}(D)) < \dim(\text{aff}(P)) = \dim(P)$$

if $\text{aff}(P) \not\subseteq \text{aff}(D)$. So

$$P \setminus \left(\bigcup_{D \in \mathcal{A}} D \right) = P \setminus \left(\bigcup_{D \in \mathcal{A}} (P \cap D) \right)$$

cannot be empty.

Let us analyze the running time for subspace elimination. In the worst case, Algorithm 1 checks whether $P \subset B_i$ for some i for $O(m^2)$ times. During each check, it uses $O(K)$ times of LP with n variables and at most mK constraints. Therefore, the total running time is $O(m^2 K \cdot \text{LP}(n, mK, L))$ where $\text{LP}(a, b, L)$ stands for the running time of linear programming with a variables, b constraints and input numbers bounded by L bits. \square

5.4 Improvement in Running Time

Let us derive the running time for local feasibility of PLP using the time bound from the last section. Suppose the input has n variables, m constraints with degree bounded by D , and the input numbers are bounded by L bits. In this section, the parameters are for the input PLP, not the input of subspace elimination, which is different from Section 5.3. The input dimension for subspace elimination is $O(n^2D)$, the number of subspaces is $O(mnD)$, and K is $O(nD)$. We need to call subspace elimination at most $2nD$ times for different choices of c in Section 5.2. So the final running time is $O(n^4m^2D^4 \cdot \text{LP}(O(n^2D), O(n^2mD^2), L))$. (Recall that $\text{LP}(a, b, L)$ stands for the running time of linear programming with a variables, b constraints and input numbers bounded by L bits.)

However, local feasibility of PLP satisfies special structure that allows us to further reduce the running time. The key observation is that for the subspaces that belong to the same constraint, i.e., subspaces in (9) forms a chain of embeded affine hull. Specifically, let us denote the set on the space of

$$[f_i]_0 = 0, [f_i]_1 = 0, [f_i]_2 = 0, \dots [f_i]_j = 0, \dots [f_i]_{j+1} < 0$$

by $B_{i,j}$, where $0 \leq j \leq d' - 1$. Then we have the following lemma.

Lemma 12. *We have the sequence of inclusions*

$$\mathbb{R}^n = \text{aff}(B_{i,0}) \supset \text{aff}(B_{i,1}) \supset \dots \supset \text{aff}(B_{i,d'-1}).$$

Recall that $D' = O(nD)$ is the bound on the degree of $x(\delta)$.

Thus, if $P \not\subset \text{aff}(B_{i,j})$, then $P \not\subset \text{aff}(B_{i,j'})$ for any $j' > j$. This means that when we are looking for $D_{i,j}$ such that $P \subset \text{aff}(B_{i,j})$, we should keep track of $j_i^* \stackrel{\text{def}}{=} \max_j \{j : P \subset \text{aff}(B_{i,j})\}$ and look at B_{i,j_i^*+1} . If $P \not\subset \text{aff}(B_{i,j})$, there is no need to look at $B_{i,j}$ with bigger index j .

We can further simplify Algorithm 1 to Algorithm 2.

Algorithm 2 Subspace Elimination for Local Feasibility of PLP

```

1: input:  $\mathcal{B} = \{B_{i,j} : i \in [m], 0 \leq j \leq D' - 1\}$  with form
2:  $P \leftarrow \mathbb{R}^n, \mathcal{A} \leftarrow \mathcal{B}$ 
3:  $\forall i \in [m], j_i^* \leftarrow -1$ 
4: while HasUpdate do
5:    $i \leftarrow 1, \text{HasUpdate} \leftarrow \text{False}$ 
6:   while  $i \leq m$  do
7:     if  $P \subseteq \text{aff}(B_{i,j_i^*+1})$  then
8:        $P \leftarrow P \setminus B_{i,j_i^*+1}$ 
9:        $\mathcal{A} \leftarrow \mathcal{A} \setminus \{B_{i,j_i^*+1}\}$ 
10:       $j_i^* \leftarrow j_i^* + 1, \text{HasUpdate} \leftarrow \text{True}$ 
11:    end if
12:     $i \leftarrow i + 1$ 
13:  end while
14: end while
15: if  $P = \emptyset$  then
16:   return False
17: else
18:   return True with a point in the relative interior of  $P$  as feasible solution
19: end if
```

When the algorithm stops, we have either $P = \emptyset$ or for any i , $P \not\subset \text{aff}(D_{i,j_i^*+1})$. Thus $P \not\subset \text{aff}(D_{i,j})$ for any $j > j_i^*$. In other words, for any $D \in \mathcal{A}$, we have $P \not\subset \text{aff}(D)$. So Algorithm 2 is correct.

Compared to directly using Algorithm 1, we only test whether $P \subset \text{aff}(D_{i,j})$ at most once for each i, j rather than potentially $O(nmD)$ times. For each i, j pair, we solve $O(D') = O(nD)$ linear programming instances. There are $O(mD') = O(nmD)$ i, j pairs, so the number of calls to the original LP is $O(nD \cdot mD') = O(mn^2D^2)$. Also, the number of constraints for each original LP is at most $O(mD')$. This means the running time for Algorithm 2 is $O(n^2mD^2 \cdot \text{LP}(O(n^2D), O(nmD), L))$. Recall in Section 5.2, to solve local feasibility of 1-PLP, Algorithm 2 needs to run $O(nD)$ times for different c , so the final running time is $O(n^3mD^3 \cdot \text{LP}(O(n^2d), O(nmD), L))$.

5.5 Outputting a Feasible Solution and its Range

We promised in Theorem 1 and Lemma 6 to give a feasible solution. Note that in Section 5.3, the algorithm only returns whether the instance is feasible, but does not give a feasible solution when the answer is yes. The special structure of subspace elimination (9) also implies the following property which allows us to output a feasible solution with Algorithm 2.

Lemma 13. *When Algorithm 2 returns true,*

$$\begin{aligned} \text{Relint}(P) &= \{(h_{ij})_{ij} | \forall i, [f_i]_0 = 0, [f_i]_1 = 0, \dots, [f_i]_{j_i^*-1} = 0, [f_i]_{j_i^*} > 0\} \\ &\subset \mathbb{R}^n \setminus \left(\bigcup_{i \in [m], 0 \leq j \leq D'-1} B_{i,j} \right) \end{aligned}$$

where Relint stands for the relative interior of a set.

Proof. We will prove this by showing $\text{Relint}(P) \cap B_{i,j} = \emptyset$ for any i, j . First, when the algorithm ends,

$$P = \mathbb{R}^n \setminus \left(\bigcup_{i \in [m], 0 \leq j \leq j_i^*} B_{i,j} \right).$$

So $\text{Relint}(P) \cap B_{i,j} = \emptyset$ for any $i \in [m]$ and $0 \leq j \leq j_i^*$.

Also, $P \subset \text{aff}(B_{i,j_i^*})$ for any i by the algorithm. So based on (9), P can be explicitly expressed as the following set.

$$\forall i, [f_i]_0 = 0, [f_i]_1 = 0, \dots, [f_i]_{j_i^*-1} = 0, [f_i]_{j_i^*} \geq 0$$

Note that when the algorithm terminates, $P \not\subset \{[f_i]_{j_i^*} = 0\}$. So the set $\{[f_i]_{j_i^*} = 0\} \cap P$ is on the relative boundary of P . Therefore, $\text{aff}(P)$ is given by

$$\forall i, [f_i]_0 = 0, [f_i]_1 = 0, \dots, [f_i]_{j_i^*-1} = 0.$$

and $\text{Relint}(P)$ is the following set.

$$\forall i, [f_i]_0 = 0, [f_i]_1 = 0, \dots, [f_i]_{j_i^*-1} = 0, [f_i]_{j_i^*} > 0.$$

By comparing the set with (9), we know that for any $i \in [m]$ and $j > j_i^*$, $B_{i,j} \cap \text{Relint}(P) = \emptyset$. \square

Finding a point in the relative interior of a polyhedron is a standard initialization step for interior point methods. For a concrete algorithm see for example [CG06]. The feasible point for (8) can be used to generate a feasible solution for the original local feasibility of PLP problem using $x(\delta) = \frac{p(\delta)}{\delta^c(1+q_1(\delta))}$.

So far we have generated a rational function solution $x(\delta) = \frac{p(\delta)}{q(\delta)}$ that satisfies the PLP on range $(0, \epsilon)$ for a small enough ϵ (without loss of generality, $q(\delta) > 0$ when δ is a small enough positive number). $x(\delta)$ is feasible if and only if it satisfies $A_i(\delta)p(\delta) - b_i(\delta)q(\delta) \geq 0$ and $q(\delta) > 0$ for any i . The largest possible ϵ such that x is feasible on $(0, \epsilon)$ is given by the smallest positive root among polynomials $A_i(\delta) \cdot p(\delta) - b_i(\delta)q(\delta)$ and $q(\delta)$.

6 Hardness of Polynomial Linear Programming

In this section we prove Theorem 2 and Theorem 6.

6.1 Local Feasibility to Everywhere Feasibility

In this section, we will show that everywhere feasibility (infeasibility) of 1-PLP can be reduced to the local feasibility (infeasibility) of 2-PLP. Note that feasibility and infeasibility are not complements of one another, so we have separate but similar arguments for feasibility and infeasibility. Furthermore, show that everywhere feasibility of 1-PLP does is not easier even if we restrict the polynomials in $A(\delta)$ and $b(\delta)$ to be linear functions.

Lemma 14. *There is a polynomial-time reduction from the everywhere feasibility (infeasibility) of 1-PLP to the local feasibility (infeasibility) of 2-PLP.*

Proof. Let us first show the reduction on feasibility. For any instance of 1-PLP

$$(A(\delta), b(\delta)), \tag{10}$$

let D be the maximum degree of $A(\delta)$ and $b(\delta)$. For the reduction of feasibility, consider an instance of 2-PLP,

$$\left(\delta_2^{D+1} A\left(\frac{\delta_1}{\delta_2}\right), \delta_2^{D+1} b\left(\frac{\delta_1}{\delta_2}\right) \right). \tag{11}$$

We claim that (10) is everywhere feasible if and only if (11) is locally feasible.

If (10) is infeasible at $\delta = a$, then (11) is infeasible when $\delta_1/\delta_2 = a$ since they only differ by a multiplication of constant on both sides. So (11) cannot be locally feasible.

Conversely, if (10) is feasible for all $\delta \in \mathbb{R}$, then similarly (11) is feasible for any $\delta_1/\delta_2 \in \mathbb{R}$. When $\delta_2 = 0$, (11) becomes $(0, 0)$, which is feasible. So (11) is feasible for any $(\delta_1, \delta_2) \in \mathbb{R}^2$.

To show the reduction for infeasibility, let us consider the following instance of 2-PLP with $n+1$ variables and $m+1$ constraints

$$\begin{cases} \delta_2^D A\left(\frac{\delta_1}{\delta_2}\right) x \geq \delta_2^D b\left(\frac{\delta_1}{\delta_2}\right) \\ x' \delta_2 = 1 \end{cases} \tag{12}$$

where x' is a variable different from x . We claim that (10) is everywhere infeasible if and only if (12) is locally infeasible. The argument is similar as feasibility. If (10) is feasible at $\delta = a$, then (11) is feasible when $\delta_1/\delta_2 = a$.

If (10) is infeasible for any $\delta \in \mathbb{R}$, then (12) is infeasible for any $\delta_1/\delta_2 \in \mathbb{R}$. And (12) is always infeasible when $\delta_2 = 0$ since the last equality cannot be satisfied. So (12) is locally infeasible. \square

Now we show that everywhere feasibility is equally as hard as the degree-1 case.

Definition 6.1 (Degree-1 PLP). A *degree-1 PLP* problem is a PLP problem where $A(\delta)$ and $b(\delta)$ are linear functions of δ .

Lemma 15. *There is a polynomial-time reduction from the everywhere feasibility of 1-PLP to the everywhere feasibility of degree-1 1-PLP.*

Proof. For a PLP instance $A(\delta)$ and $b(\delta)$, suppose the maximum degree is d and write

$$A_{ij} = \sum_{k=0}^d a_{ijk} \delta^k, \quad b_i = \sum_{k=0}^d b_{ik} \delta^k.$$

We design the following degree-1 PLP instance so that it is equivalent to the original one. There are $n(d+1) + d$ variables, $x_i^{(j)}$ for $i \in [n]$, $j = [d] \cup \{0\}$ and y^i for $i \in [d]$. First, let

$$x_i^{(j)} = \delta x_i^{(j-1)}$$

for any $i \in [n]$ and $j \in [d]$, so $x_i^{(j)}$ corresponds to $x_i \delta^j$ in the original PLP instance. Next, let

$$y_1 = \delta, \quad y_i = \delta y_{i-1}$$

for $i \geq 2$, so $y_i = \delta^i$. Finally, we add

$$\sum_{j=1}^n \sum_{k=0}^d a_{ijk} x_i^{(k)} \geq \sum_{k=0}^d b_{ik} y_i$$

for any $i \in [m]$. It is straightforward to check that by construction, the feasibility of the degree-1 PLP instance is the same as the original PLP instance, i.e., each feasible point in one problem corresponds to a feasible point in the other. \square

Next we show that the everywhere feasibility (infeasibility) of 1-PLP problem is NP-hard, which implies that both the everywhere feasibility of degree-1 1-PLP problem and the local feasibility (infeasibility) of 2-PLP problem are also NP-hard.

6.2 Hardness of Everywhere Feasibility of PLP

In this section we only focus on 1-PLP, so any PLP instance is 1-PLP by default.

Theorem 13. *The everywhere feasibility of 1-PLP problem is NP-hard.*

The proof has 3 steps. In the first step, we take the dual and negation of the original PLP to get a form of the problem that is easy to analyze. In the second step, we reduce the independent set problem to this form of PLP by considering a quadratic instance of the form. In the third step, we show that the instance has size bounded by polynomial.

Step 1: Dual and negation. We will first take the dual and negation of the program, then use a construction from [PV91]. From linear programming duality,

$$\forall \delta, \text{ the program } A(\delta)x \geq b(\delta) \text{ is feasible}$$

is equivalent to

$$\forall \delta, y, \text{ s.t. } y \geq 0 \text{ and } A^\top(\delta)y = 0, \text{ it holds that } b^\top(\delta)y \leq 0.$$

Negation of this is

$$\exists \delta, y \text{ satisfying } y \geq 0 \text{ and } A^\top(\delta)y = 0 \text{ with } b^\top(\delta)y > 0. \quad (13)$$

(13) is easier to analyze than the original form of problem because it asks the existence of a solution, while the original problem asks to verify a property for arbitrary δ .

Lemma 16. Fix $A \in \mathbb{R}^{m \times n}$ and $c \in \mathbb{R}^n$ such that there is no $x^* \in \mathbb{R}^n$ simultaneously satisfying

$$\begin{cases} c^\top x^* > 0 \\ Ax^* \geq 0. \end{cases} \quad (14)$$

Then for any $b \in \mathbb{R}^m$ and $d \in \mathbb{R}$ there exists a solution $x \in \mathbb{R}^n$ of

$$\begin{cases} c^\top x > d \\ Ax \geq b \end{cases} \quad (15)$$

if and only if there exists a solution $x_1, x_2 \in \mathbb{R}^n$, $w \in \mathbb{R}$, $s \in \mathbb{R}^m$ of

$$\begin{cases} c^\top (x_1 - x_2) - dw > 0 \\ A(x_1 - x_2) = bw + s \\ x_1, x_2, w, s \geq 0. \end{cases} \quad (16)$$

Proof. Suppose program (15) has a feasible solution x . Then letting $x_1 = \max\{x, 0\}$, $x_2 = \max\{-x, 0\}$, $w = 1$, $s = Ax - b$, we get a feasible point for program (16).

Conversely, suppose program (16) has a feasible solution x_1, x_2, w, s . We claim that $w \neq 0$. Otherwise, by taking $x^* = x_1 - x_2$, we get $c^\top x^* > 0$ and $Ax^* = s \geq 0$, i.e., a feasible solution for (14) – a contradiction. So we can assume $w \neq 0$ and let $x = (x_1 - x_2)/w$ so that $c^\top x > d$ and $Ax - b = s/w \geq 0$. This means program (15) is feasible. \square

Note that (16) is a special case of (13), we have the following corollary.

Corollary 3. Suppose there is a polynomial time algorithm for testing everywhere feasibility of any PLP. Then there is a poly-time algorithm for solving the following problem. Given $m \times n$ matrix polynomial $A(\delta)$, m -dimensional vector polynomial $b(\delta)$, n -dimensional vector polynomial $c(\delta)$ and polynomial $d(\delta)$, decide whether

$$\exists x, \delta \quad \begin{cases} c^\top(\delta)x > d(\delta) \\ A(\delta)x \geq b(\delta) \end{cases} \quad (17)$$

where $A(\delta), c(\delta)$ satisfies for any $\delta \in \mathbb{R}$, the following program has no solution:

$$\begin{cases} c^\top(\delta)x^* > 0 \\ A(\delta)x^* \geq 0. \end{cases} \quad (18)$$

Step 2: Independent set to PLP. Our next step is to take any graph G and produce an instance $A(\delta), b(\delta), c(\delta), d(\delta)$ satisfying (18), such that (17) is feasible if and only if there exists an independent set of size k in G .

Lemma 17 ([PV91]). There exists a polynomial-size linear program P ,

$$\begin{aligned} 0 &\leq x_i \leq 1, 1 \leq i \leq n \\ y_{ij} &\geq 0, y_{ij} \geq x_i + x_j - 1, 1 \leq i < j \leq n \\ \delta &= \sum_{i=1}^n 4^i x_i \\ z &= \sum_{i=1}^n 4^{2i} x_i + \sum_{1 \leq i < j \leq n} 2 \cdot 4^{i+j} y_{ij}, \end{aligned} \quad (P)$$

such that for any fixed x ,

$$\max_{y,z,\delta: (\delta,x,y,z) \in P} \delta^2 - z = 0 \text{ if } x \in \{0,1\}^n,$$

and

$$\max_{y,z,\delta: (\delta,x,y,z) \in P} \delta^2 - z < 0 \text{ otherwise.}$$

We can add linear constraints for an arbitrary instance of k -independent set problem P' ,

$$\begin{aligned} 0 &\leq x_i \leq 1, \forall i \in [n] \\ x_i + x_j &\leq 1, \forall (i,j) \in E \\ \sum_i x_i &\geq k. \end{aligned} \tag{P'}$$

So for a small enough constant ϵ , there exists a feasible point in $P \cap P'$ satisfying $\delta^2 - z + \epsilon > 0$ if and only if the answer to the k -independent set problem is yes. The final instance P^* is $P \cap P'$ with a constraint $\delta^2 - z + \epsilon > 0$. Note that this decision problem satisfies the form in (17).

We can check that (18) is infeasible for P^* . Substituting all constant terms with 0, we get

$$\begin{aligned} x_i &= 0, 1 \leq i \leq n \\ y_{ij} &\geq 0, 1 \leq i < j \leq n \\ z &= \sum_{i=1}^n 4^{2i} x_i + \sum_{1 \leq i < j \leq n} 2 \cdot 4^{i+j} y_{ij} \\ -z &> 0, \end{aligned}$$

which is clearly infeasible.

Remark 6. Here we chose the independent set problem for concreteness, but in fact we can use any NP-hard problem that can be written as an integer linear program on $\{0,1\}^n$ with polynomially bounded parameters.

Step 3: Bounding the size of reduction instance. Next we only need to show that ϵ can be chosen to be represented by polynomial number of bits. From Lemma 18 and Lemma 19, we can set ϵ to be any constant between 0 and $\left(2k^{n^3+n}n!\right)^{-2}$ so that P' contains a feasible integer solution if and only if $\max_{x,y,z,\delta \in P} \delta^2 - z + \epsilon > 0$. Indeed, if P' contains an integer solution x_0 , then fix x_0 , we have $\max_{y,z,\delta \in P} \delta^2 - z + \epsilon = \epsilon > 0$ by Lemma 17. If P' does not contain an integer solution, by Lemma 18 and Lemma 19, fixing any $x \in P'$, we have

$$\max_{y,z,\delta \in P} \delta^2 - z + \epsilon \leq \left(2k^{n^3+n}n!\right)^{-2} + \epsilon < 0.$$

The denominator for ϵ is $O(\exp(\text{poly}(n)))$ so it can be represented by polynomially many of bits.

Lemma 18 ([FO85], Lemma 1). *Recall that P' is the linear relaxation of the k -independent set problem for graph $G = (V, E)$ as defined above. If $P' \neq \emptyset$ but there is no integer solution, i.e., $P' \cap \{0,1\}^n = \emptyset$, then $d_\infty(P', \{0,1\}^n) \geq \left(2k^{n^3+n}n!\right)^{-1}$.*

Lemma 19. Under the linear program P given in Lemma 17, for a fixed $x \in [0, 1]^n$,

$$\max_{y, z, \delta: (\delta, x, y, z) \in P} \delta^2 - z \leq -d_\infty^2(x, \{0, 1\}^n).$$

Proof. To reach the maximum of $\delta^2 - z$, y_{ij} are at their minimum, $\max\{0, x_i + x_j - 1\} = (x_i + x_j - 1)_+$, so

$$\max_y \delta^2 - z = \sum_{i=1}^n 4^{2i} (x_i^2 - x_i) + \sum_{1 \leq i < j \leq n} 2 \cdot 4^{i+j} (x_i x_j - (x_i + x_j - 1)_+)$$

Note that $x_i x_j - (x_i + x_j - 1)_+ \leq x_j(1 - x_j)$, because when $x_i \leq 1 - x_j$, $x_i x_j - (x_i + x_j - 1)_+ = x_i x_j \leq x_j(1 - x_j)$, otherwise $x_i > 1 - x_j$, $x_i x_j - (x_i + x_j - 1)_+ = (1 - x_i)(1 - x_j) \leq x_j(1 - x_j)$. Let $t_i = x_i(1 - x_i)$, then

$$\begin{aligned} \max_y \delta^2 - z &\leq -\sum_{i=1}^n 4^{2i} t_i + \sum_{1 \leq i < j \leq n} 2 \cdot 4^{i+j} t_j \\ &= \sum_{i=1}^n \left(-4^{2i} + 2 \cdot 4^i \sum_{j=1}^{i-1} 4^j \right) t_i \\ &\leq \sum_{i=1}^n -\frac{4^{2i}}{3} t_i \\ &\leq -\min_i t_i. \end{aligned}$$

Let $d = d_\infty(x, \{0, 1\}^n)$. Because $d \leq 1/2$, we have

$$\min_i t_i = d(1 - d) \geq d^2.$$

□

Remark 7. Note that Lemma 19 implies Lemma 17, so we give an alternate shorter and more intuitive proof of the main theorem in [PV91].

6.3 Hardness of Everywhere Infeasibility of PLP

In this section we prove hardness of everywhere infeasibility of 1-PLP. The proof is similar to feasibility and we again reduce from the independent set problem.

Theorem 14. *The everywhere infeasibility of 1-PLP is NP-hard.*

Proof. The everywhere infeasibility of 1-PLP can be written as follows:

$$\forall \delta \in \mathbb{R}, \forall x, A(\delta)x \geq b(\delta) \text{ is not satisfied.}$$

First, we take the negation, which is

$$\exists \delta, x \text{ such that } A(\delta)x \geq b(\delta)$$

Note that (P) is in such form. Therefore, we can take the following program.

$$\begin{cases} P \\ \delta^2 - z = 0 \\ P' \end{cases}$$

By Lemma 17, the first two constraints restrict x on $\{0, 1\}^n$. This would solve the independent set problem combined with (P') . □

6.4 Everywhere Feasibility is in co-NP

In this section we prove that everywhere feasibility of 1-PLP is in co-NP. This implies that the problem is co-NP complete.

Theorem 15. *Everywhere feasibility problem of 1-PLP is in co-NP, i.e., the following problem is in NP. Given an instance of 1-PLP $(A(\delta), b(\delta))$, decide whether exists $\delta \in \mathbb{R}$ such that there is no x satisfying $A(\delta)x \geq b(\delta)$.*

By linear programming duality, the above decision problem is exactly (13), which is equivalent to the following:

$$\exists \delta, y \text{ satisfying } \begin{cases} y \geq 0 \\ A^\top(\delta)y = 0 \\ b^\top(\delta)y = 1 \end{cases} \quad (19)$$

Intuitively, the value of δ where $A(\delta)x \geq b(\delta)$ is infeasible should be used as the certificate. The difficulty is that δ may not even be a rational number. Our proof has two steps. The first step shows that if answer to the above problem is yes, we can find a δ_0 that is a root of a polynomial. In the second step, we show that y can be chosen to be a rational function of δ_0 and the conditions in (19) can be checked within polynomial time.

Lemma 20. *Suppose the answer to the problem in Theorem 15 is yes. There exists δ_0 satisfying the condition and δ_0 is a root of a polynomial with size $\text{poly}(n, m, d, L)$.*

Proof. By Corollary 2, we can divide \mathbb{R} into a finite number of intervals such that $A(\delta)x \geq b(\delta)$ is always feasible or infeasible when δ is within an interval.

If $A(\delta)x \geq b(\delta)$ is infeasible except at a finite number of points on \mathbb{R} , then we can easily choose a rational δ_0 . Otherwise, there exists two consecutive intervals such that $A(\delta)x \geq b(\delta)$ is feasible on the first interval and infeasible on the second interval. Suppose the two intervals have common endpoint δ_1 , i.e., $A(\delta)x \geq b(\delta)$ is feasible on an open interval $(\delta_1 - \epsilon, \delta_1)$ but infeasible at δ_1 or feasible on $(\delta_1 - \epsilon, \delta_1]$ but infeasible on an open interval (δ_1, δ_2) .

Now we prove that δ_1 is root of a polynomial with size $\text{poly}(n, m, d, L)$. Based on Lemma 8, we can pick ϵ small enough so that there exists $J \subset [n]$ that $A_J^\top b_J(\delta)$ is a solution for x on $(\delta_1 - \epsilon, \delta_1)$. If $A_J^\top b_J(\delta)$ is not continuous at δ_1 , denominator of one of the coordinate becomes zero at δ_1 , which means δ_1 is a root of the denominator. Now suppose $A_J^\top b_J(\delta)$ is continuous around δ_1 , but is not a solution in the right neighborhood of δ_1 . Suppose it violates a constraint $a_i^\top x \geq b_i$ in the right neighborhood of δ_1 . So $a_i^\top A_J^\top b_J(\delta) \geq b_i$ before δ_1 but $a_i^\top A_J^\top b_J(\delta) < b_i$ after δ_1 . Since the function is continuous, we know $a_i^\top A_J^\top b_J(\delta_1) - b_i = 0$. After multiplying any denominator on both sides, we get that δ_1 is root to a polynomial with size $\text{poly}(n, m, d, L)$.

So the proof is completed if $A(\delta)x \geq b(\delta)$ is infeasible at δ_1 . Now assume $A(\delta)x \geq b(\delta)$ is feasible at δ_1 but infeasible at (δ_1, δ_2) . Then by the above analysis, δ_1 and δ_2 are roots of two polynomially sized polynomial p_1 and p_2 . They are both roots of $p_1 p_2$. So by Lemma 4, $\delta_2 - \delta_1 \geq 2^{-\text{poly}(n, m, d, L)}$. Therefore, we can pick a rational number between δ_1 and δ_2 with polynomial number of bits. \square

Lemma 21. *Suppose the answer to the problem in Theorem 15 is yes. Let δ_0 be the number chosen in Lemma 20. There exists $y_0 = f(\delta_0)$ satisfying (19) where f is a rational function of δ_0 with size $\text{poly}(n, m, d, L)$.*

Proof. Note that (19) itself can be viewed as a PLP with variable y . The lemma follows as a corollary of Lemma 7 and Lemma 8. \square

All the ingredients are in place to prove Theorem 15.

Proof of Theorem 15. The certificate of the problem are δ_0 in Lemma 20 and y_0 in Lemma 21. By Lemma 5, we can represent them by irreducible polynomials and signs of their derivatives. Also, (19) can be verified in polynomial time by Lemma 5. \square

7 Application to Probabilistic Cellular Automata

In this section we formally introduce the potential method developed in [HMM19] and discuss how to use the algorithm for the local feasibility of PLP to find a potential.

A one-dimensional discrete-time probabilistic cellular automaton (PCA) is a function on \mathbb{Z} that evolves with time. The PCA has alphabet A so at any time t the configuration of the PCA is denoted by $\eta_t \in A^{\mathbb{Z}}$ where $\eta_t(n)$ is the value on site $n \in \mathbb{Z}$ at time t .

The configuration η_t is Markov chain as t changes. Given η_t , the next configuration η_{t+1} is obtained by updating each site $n \in \mathbb{Z}$ independently. The update rule is local, random, and homogeneous. That is to say,

$$\eta_{t+1}(n) = f(\eta_t(n-a), \eta_t(n-a+1), \dots, \eta_t(n+a)),$$

where f is a random function that follows a fixed distribution independent of t and n , a is a constant. Let F denote the transition of the Markov chain.

The running example is the following. (A similar analysis could be applied to other PCA as well, but as discussed in Section 3, we focus on the one elementary symmetric PCA which ergodicity is unknown.) The PCA has alphabet $\{0, 1\}$ and noise parameter $p \in [0, 1/2]$. The state is updated according to

$$\eta_{t+1}(n) = \text{BSC}_p(\text{NAND}(\eta_t(n-1), \eta_t(n)))$$

where BSC_p is binary symmetric channel with crossover probability p . The Markov transition operator is denoted by A_p . The PCA can be view as a NAND function that has probability $2p$ of turning into a random bit.

The problem of interest is the ergodicity of this PCA, as defined in Definition 3.1. The ergodicity is known to be equivalent to having positive probability of drawing in *percolation game* [HMM19]. When $p = 0$, A_0 is noiseless and it easy to see that η_t may not converge. However, it is proven to be ergodic for any $p > 0$ [HMM19]. An alternate version of noisy NAND function is NAND with edge noise, where the PCA is defined by

$$\eta_{t+1}(n) = \text{NAND}(\text{BSC}_p(\eta_t(n-1)), \text{BSC}_p(\eta_t(n))).$$

The Markov transition operator is denoted by A'_p . It is shown in [DO77] that edge noise is a more general model than vertex noise.

We will prove the ergodicity of both PCA for small enough p under the help of the algorithm for the local feasibility of PLP.

Theorem 16. *There exists $\epsilon > 0$ such that A_p is ergodic for any $p \in (0, \epsilon)$.*

Theorem 17. *There exists $\epsilon > 0$ such that A'_p is ergodic for any $p \in (0, \epsilon)$.*

Note that the ergodicity of A'_p is not known prior to this work as we know, but Theorem 16 is a corollary of the result in [HMM19]. Our contribution is to propose a new procedure of automatically finding potential functions by reducing to it to feasibility of PLP.




Figure 4: Update rule of F_p . Here $*$ means arbitrary symbol from $\{0, 1, ?\}$.

7.1 Coupling and Auxiliary Chain

A typical approach to proving the ergodicity is by coupling. Let η_t^+ and η_t^- be two versions of the Markov chain that start at different initial configurations. If we can show that there is a coupling so that for any two initial configuration the probability of having same configuration converges to 1 as t goes to infinity, then by linearity any distribution on configuration would converge to the same stationary distribution. Thus ergodicity is proven by Definition 3.1.

The coupled Markov chain (η_t^+, η_t^-) is given by the following update rule. On any site n , η^+ and η^- either both apply the NAND function with probability $1 - 2p$, or take the same random bit with probability $2p$. The coupled chain has alphabet $\{(0, 0), (1, 1), (0, 1), (1, 0)\}$. If we ignore the difference between $(1, 0)$ and $(0, 1)$ and denote them by $?$, we get a new chain F_p with alphabet $\{0, 1, ?\}$. The update rule of F_p is illustrated in Figure 4 and the update rule of F_p' is shown in Table 2. We can define F_p' for A_p' in the same way. Note that the chain F_p and F_p' are still Markov.

By definition, restricting F_p to not contain $?$ is exactly A_p . So if F_p is ergodic then A_p is ergodic. In fact, we can further reduce the ergodicity of F_p [HMM19].

Lemma 22 ([HMM19]). *F_p is ergodic if any stationary distribution of F_p has 0 probability of symbol $?$.*

The same statement can be proven for F_p' , which can be found in Appendix B for completeness.

Lemma 23. *F_p' is ergodic if any stationary distribution of F_p has 0 probability of symbol $?$.*

7.2 Potential Function Method

We say the distribution of a configuration $\eta(n)$ is *shift-invariant* if $\dots, \eta(-1), \eta(0), \eta(1), \dots$ and $\dots, \eta(-1+k), \eta(0+k), \eta(1+k), \dots$ have the same distribution for any $k \in \mathbb{Z}$. It is without loss of generality to only consider stationary distributions that are shift-invariant. For $s \in \{0, 1, ?\}^n$, and a shift-invariant distribution μ on configurations η , let $\mu(s) = \mu(\eta : (\eta(0), \eta(1), \dots, \eta(n)) = s)$. A potential function is a weighted sum of the probability of strings, defined formally as follows.

Definition 7.1 (Potential function). A *potential function* w with length ℓ is a vector in \mathbb{R}^{3^ℓ} , where the coordinates are labeled by $\{0, 1, ?\}^\ell$. We write $\{s\}$ for the s th standard basis vector, and

$$w = \sum_{s \in S} c_s \{s\}.$$

where $S \subset \{0, 1, ?\}^\ell$. For a shift-invariant distribution μ , define

$$\{s\}(\mu) = \mu(s),$$

and

$$w(\mu) = \sum_{s \in S} c_s \mu(s).$$

The potential function method can be summarised as follows:

Lemma 24 ([HMM19]). *To show that stationary distribution of F_p has 0 probability of symbol $?$, it suffices to design a potential function w such that for any shift-invariant distribution μ ,*

$$w(F_p \mu) \leq w(\mu) - \mu(?s_0) \quad (20)$$

where s_0 is an arbitrary string with alphabet $\{0, 1\}$.

Proof. Because we can assume without loss of generality that the stationary distribution μ is shift-invariant, $w(\mu) = w(F_p \mu) \leq w(\mu) - \mu(?s_0)$ means that $\mu(?s_0) = 0$. Since $\mu(?s_0) = F_p \mu(?s_0) \geq p^{|s_0|} F_p \mu(?) = p^{|s_0|} \mu(?)$, we can conclude that $\mu(?) = 0$. \square

7.3 Connection with PLP

In [HMM19], a carefully designed potential function was introduced. In contrast, we will show how to find a potential function that satisfies (20) with local feasibility of PLP. This connection between finding a potential function and PLP was introduced in [MMP21], but [MMP21] only analyzed the PLP for specific choices of parameter p , which are normal LP problems.

First let us focus on the relation between $w(\mu)$ and $w(F_p \mu)$. Suppose $w = \sum_{s \in S} c_s \{s\}$ and without loss of generality any $s \in S$ has length ℓ . For simplicity we say the *length* of w is ℓ . As stated in Definition 7.1, w can be viewed as a vector with coordinates labeled by $\{0, 1, ?\}^\ell$, also denoted by w . We will show $w(F_p \mu) = w'(\mu)$ where w' is a potential function with length $\ell + 1$. This can be described by the following lemma.

Lemma 25. *For any potential function w with length ℓ , and any PCA, F_p , with alphabet \mathcal{A} and the following form:*

$$\eta_{t+1}(n) = f(\eta_t(n-a), \eta_t(n-a+1), \dots, \eta_t(n+b)),$$

there exists a potential function w' such that $w(F_p \mu) = w'(\mu)$. Further, define matrix $C_\ell \in \mathbb{R}^{|\mathcal{A}|^{\ell+a+b}} \times \mathbb{R}^{|\mathcal{A}|^\ell}$ that for $s' \in \mathcal{A}^{\ell+a+b}$ and $s \in \mathcal{A}^\ell$,

$$\begin{aligned} C_\ell(s', s) &= \Pr[(\eta_{t+1}(1), \dots, \eta_{t+1}(\ell)) = s | (\eta_t(1-a), \dots, \eta_t(\ell+b)) = s'] \\ &= \prod_{n=1}^{\ell} \Pr[f(s'(n-a), s'(n-a+1), \dots, s'(n+b)) = s(n)], \end{aligned}$$

we have

$$w' = C_\ell w.$$

Proof.

$$\begin{aligned} w(F_p \mu) &= \sum_{s \in S} c_s (F_p \mu)(s) \\ &= \sum_{s \in S} c_s \sum_{s' \in \{0, 1, ?\}^{\ell+1}} \mu(s') \Pr[(\eta_{t+1}(1), \dots, \eta_{t+1}(\ell)) = s | (\eta_t(1-a), \dots, \eta_t(\ell+b)) = s'] \\ &= \sum_{s \in S} \sum_{s' \in \{0, 1, ?\}^{\ell+1}} C_\ell(s, s') c_s \mu(s') = C_\ell w(\mu) \end{aligned}$$

\square

Note that in the case of F_p , we have $a = 1, b = 0$ and $\mathcal{A} = \{0, 1, ?\}$. So the size of C_l is $3^{\ell+1} \times 3^\ell$. Also, C_ℓ is a polynomial of p , see Appendix A for the exact expression of C_ℓ for F_p .

For a periodic configuration η that repeats string y and a potential function $w = \sum_{s \in S} c_s \{s\}$, we use $w(y)$ to denote the value of the potential.

Definition 7.2. For a potential function $w = \sum_{s \in S} c_s \{s\}$ with length l and a string $y \in \cup_{n \geq \ell} \{0, 1, ?\}^n$, define

$$w(y) = \sum_{s \in S} c_s \sum_{i=0}^{|y|} \mathbf{I}\{(s_i \bmod |y|, s_{i+1} \bmod |y|, \dots, s_{i+|s|} \bmod |y|) = s\}$$

Lemma 26. A sufficient condition for (20) is that for any string y with alphabet $\{0, 1, ?\}$,

$$C_\ell w(y) \leq w(y) - \{?s_0\}(y). \quad (21)$$

Proof. First, let us show that for any shift-invariant measure μ , the probability of any aperiodic configuration is 0. Suppose for contrary that this is not true. Because there are countably many aperiodic configurations, there must exist one with positive probability. Suppose $\mu(\eta(n)) > 0$ where η is an aperiodic configuration. Since μ is shift-invariant, $\mu(\eta(n)) = \mu(\eta(n+k))$ for any $k \in \mathbb{Z}$. This leads to a contradiction because $\mu(\{\eta(n+k) : k \in \mathbb{Z}\})$ would be $+\infty$.

Therefore, it suffices to show that (21) leads to (20) on μ that has probability 1 on periodic configurations. Suppose L is the period of a periodic configuration η , so $\eta(n+L) = \eta(n)$, let $\mu_\eta = 1/L \sum_{i=1}^L \delta_{\eta_{n+i}}$. Here $\delta_{\eta_{n+i}}$ stands for the delta distribution on η_{n+i} . Any shift-invariant distribution would be a convex combination of μ_η for different η . Let $y = (\eta(1), \eta(2), \dots, \eta(L))$. For any potential function $w = \sum_{s \in S} c_s \mu(s)$, we can conclude by Definition 7.2,

$$w(\mu) = 1/L \sum_{i=1}^L \mathbb{I}[(\eta(i), \eta(1+i), \dots, \eta(|s|+i)) = s] = w(y).$$

Also, by Lemma 25 we have

$$w(F_p \mu) = C_\ell w(\mu) = C_\ell w(y).$$

So (20) is equivalent to

$$C_\ell w(y) \leq w(y) - \{?s_0\}(y).$$

□

Note that $\{s\} = \{s_0\} + \{s_1\} + \{s?\}$. We can extend a potential function to any longer length. Use $P_\ell(w)$ to denote the length- ℓ version of w if w has length not greater than ℓ . P_ℓ is a linear transformation. Specifically, when w is expressed as a vector with length ℓ_0 , P_ℓ is matrix $P_\ell = (1, 1, 1)^{\top \otimes \ell - \ell_0} \otimes I_{\ell_0}$, where \otimes stands for Kronecker product. Therefore, if $|s_0| \leq \ell$ (21) is equivalent to that for any string y ,

$$(P_{\ell+1}(w) - C_\ell w - P_{\ell+1}(\{?s_0\}))(y) \geq 0. \quad (22)$$

This leads to the problem of what potential function w is non-negative over any string. Our goal is to introduce a sufficient condition for (21) expressed by a PLP. A trivial sufficient condition would be that for $w = \sum_{s \in S} c_s \{s\}$, every coefficient $c_s \geq 0$. However, the condition is too strong that no potential would satisfy. To build intuition, let us construct a directed weighted graph G with length l strings as vertices and there is an edge between two vertices s_0 and s_1 if $(s_0(1), s_0(2), \dots, s_0(l-1)) = (s_1(0), s_1(1), \dots, s_1(l-2))$. Each vertex s has weight c_s if $s \in S$ and 0 otherwise. Then by Definition 7.2, $w(y)$ corresponds to the weight of a cycle on the graph. So $w(y) \geq 0$ for any y is equivalent to having no negative cycles on G .

[MMP21] proved a sufficient condition for this expressed by linear constraints.

Lemma 27 ([MMP21], Proposition 10). *There exists A_ℓ, B_ℓ of size $3^{\ell+1} \times 3^\ell$ that only depends on ℓ such that for a potential function w with length $\ell + 1$, if there exists z such that*

$$A_\ell z \leq B_\ell w,$$

we have that for any string y ,

$$w(y) \geq 0.$$

The concrete expression of A_ℓ and B_ℓ is included in Appendix A.

Proof of Theorem 16 and Theorem 17. From Lemma 25, the coefficients in (22) are polynomials of the noise parameter p . So by Lemma 27, a sufficient condition for Theorem 16 is to check the local feasibility of the following PLP for a certain length ℓ and a certain s_0 where $|s_0| \leq \ell$:

$$A_{\ell+1}z \leq B_{\ell+1} (P_{\ell+1}(w) - C_\ell w - P_{\ell+1}(\{?s_0\})). \quad (23)$$

Here z and w are variable vectors.

Finding a potential. The rest of proof is carried out using a computer. To use local feasibility of PLP to find a potential function satisfying (23), we can search in the space of possible length ℓ and choice of s_0 , using the algorithm in Section 1. The local feasibility of (23) for any choice of ℓ and s_0 would lead to ergodicity. Note in Lemma 25, the size of C_ℓ grows exponentially with length of the potential function. So only small ℓ yields a PLP that is computationally efficient in practice. We tried $\ell = 2$ or 3 and all possible s_0 with length no more than ℓ . In the end, we find that (23) is locally feasible for F_p with $\ell = 3$ and $s_0 = 01$. It is locally feasible for F'_p with $\ell = 3$ and $s_0 = 10$. This finishes the proof for Theorem 16 and Theorem 17. See the potential functions in Appendix C. □

7.4 Attempt on Soldier's Rule

Let us first recall the definition of Soldier's Rule. In Soldier's Rule, the alphabet is $\{-1, 1\}$, indicating the direction a soldier is facing. At each time step, each soldier sets their new direction based the majority vote of the direction of itself and the direction of the first neighbor and the third neighbor in their direction. The result then goes through a BSC channel. Thus,

$$\eta_{t+1}(n) = \text{BSC}_p(\text{Maj}(\eta_t(n), \eta_t(n + \eta_t(n)), \eta_t(n + 3\eta_t(n)))) .$$

The same argument can be applied to Soldier's Rule to show that a sufficient condition of ergodicity is (23).

However, by Lemma 25, here the size of C_ℓ is $3^{r+6} \times 3^r$. Therefore, even for potential with length 3, the final PLP would have 3^{10} number of constraints, which makes the algorithm inefficient in practice. We leave it as an open question whether the ergodicity of soldier's rule under BSC noise can be solved efficiently with our approach. This is potentially possible either by improving the running time for Theorem 1 in general, or exploring special structure of (23) so that the algorithm in Theorem 1 can be accelerated.

8 Application in Broadcasting of Information on 2D Grid

In information theory, we are interested in how information broadcasts on a infinite DAG with noise. One important case is when the DAG is a 2D grid. To keep a consistent notation with section 7, the nodes on grid are indexed by (t, i) where $t \geq 0$, $0 \leq i \leq t$. The bit at coordinate (t, i) is denoted by $\eta_t(i)$. At time 0, there is one bit of information at coordinate 0,0, either 0 or 1. At any time $t > 0$, node at $(t-1, i)$ receives information sent by previous layer and pass a bit to (t, i) and $(t, i+1)$. The edge between two nodes is noisy, specifically, a BSC channel with noise parameter p .

For nodes on the boundary, i.e., nodes with coordinates (t, t) and $(t, 0)$, they only receive one bit of input. So we assume that they pass on whatever bit they receive. For other nodes, they receive two inputs and output to two other nodes. The output is a boolean function of the two-bit input, and different functions f corresponds to different Markov chains.

$$\eta_t(i) = f(BSC_p(\eta_{t-1}(i-1)), BSC_p(\eta_{t-1}(i))) \quad (24)$$

We will use η_t^+ and η_t^- to denote the Markov chain that starts with bit 1 and bit 0 respectively. Our goal is to prove that for any function f , η_t^+ and η_t^- converges to the same distribution, i.e.,

$$\lim_{t \rightarrow \infty} \|P_{\eta_t^+}, P_{\eta_t^-}\|_{TV} = 0.$$

The information of the initial bit is lost at infinity if this is true, so we say the recovery is not possible.

This problem can also be approached by potential function method similar to PCA and then be reduced to the local feasibility of a PLP. Results in this section are shown in [MMP21]. We include the results here to show its connection with our work.

The nontrivial choices for f are AND, XOR, NAND and IMP where IMP is the following function.

| x_1 | x_2 | IMP(x_1, x_2) |
|-------|-------|-------------------|
| 0 | 0 | 1 |
| 0 | 1 | 1 |
| 1 | 0 | 0 |
| 1 | 1 | 1 |

Other non-constant functions are equivalent to these four functions, or only depends on one of the inputs.

It is conjectured that it is impossible to recover for any function. In [MMP21], the impossibility of recovery was proven for AND and XOR for any p . They also proved it for NAND at many specific p , which corresponds to PLP at a fixed parameter. Here we will complete the conjecture for small error probability for NAND and IMP and prove the following theorem.

Theorem 18. *There exists $\epsilon > 0$ such that for any $p \in (0, \epsilon)$, the recovery is not possible for any binary function f .*

Intuitively, if a Markov chain converges for small error probability, it should also converge for larger error probability. However, we are not able to prove the monotonicity of convergence with respect to p formally.

8.1 Auxiliary Chain and Potential Function

Similar to PCA, we can define a coupled chain between η_t^+ and η_t^- so that at any edge, the two chain becomes the same random bit with probability $2p$ and remains the initial bit with probability $1 - 2p$. The coupled chain can be simplified with alphabet $\{0, 1, ?\}$ where $?$ means that the two chain are not equal at that site, see Table 2. We always denote the configuration by η_t for simplicity because the same analysis holds for NAND and IMP.

Therefore, it suffices prove that as t tends to infinity, η_t consists of 0 and 1 with probability converging to 1. In other words,

$$\lim_{t \rightarrow \infty} \mathbb{E}[\text{number of } ? \text{ in } \eta_t] = 0. \quad (25)$$

To this end, we introduce the potential function in the setting of broadcasting on grid.

Definition 8.1 (potential function). A potential function w of length ℓ is written as a vector

$$w = \sum_{s \in \{0, 1, ?\}^\ell} c_s \{s\},$$

its value on a string η with length at least ℓ is defined as

$$w(\eta) = \sum_{s \in \{0, 1, ?\}^\ell} c_s \sum_{i=1}^{|\eta|-\ell+1} \mathbf{I}\{(\eta_i, \eta_{i+1} \cdots, \eta_{i+|s|-1}) = s\}$$

For a distribution μ on string η ,

$$w(\mu) = \mathbb{E}_{\eta \sim \mu}[w(\eta)].$$

Therefore, the statement (25) is equivalent to

$$\lim_{t \rightarrow \infty} \{?\}(\eta_t) = 0.$$

The following lemma describes the potential function we need for the condition above to hold.

Lemma 28 (Proposition 4, [MMP21]). *A sufficient condition for $\lim_{t \rightarrow \infty} \{?\}(\eta_t) = 0$ is to find a function w such that*

1. For any η_t , $w(\eta_{t+1}) \leq w(\eta_t)$
2. For any string y , $w(y) \geq \{?\}(y)$.

8.2 Connection with PLP

The goal is to again reduce sufficient condition in Lemma 28 to local feasibility of a PLP.

Intuitively we want to use similar conclusion as Lemma 25 and Lemma 27. The main difference between PCA and broadcasting on grid is that grid has a boundary while PCA is infinite. To circumvent the problem, let us introduce the following concepts.

Definition 8.2. A potential function $w = \sum_{s \in S} c_s \{s\}$ ($c_s \neq 0$) is said to be $?$ -only if any $s \in S$ contains at least one $?$ symbol.

Definition 8.3. An l -boundary-coupled string is a string with alphabet $0, 1, ?$ such that the first l and last l symbols do not contain $?$.

It is observed in [MMP21] that we can only consider l -boundary-coupled strings as for any l , η_t will become l -boundary-coupled with probability converging to 1 when t goes to infinity.

The following lemma is proven in [MMP21]. The proof is similar to Lemma 25.

Lemma 25'. *For a ?-only potential function w with length l , there exists a potential function w' such that $w(\eta_{t+1}) = w'(\eta_t)$. Further, define matrix $C_\ell \in \mathbb{R}^{3^{\ell+1}} \times \mathbb{R}^{3^\ell}$ that for any l -boundary-coupled string $s \in \{0, 1, ?\}^{\ell+1}$ and $s' \in \{0, 1, ?\}^\ell$,*

$$C_\ell(s, s') = \Pr[(\eta_{t+1}(1), \dots, \eta_{t+1}(\ell)) = s' | (\eta_t(0), \dots, \eta_t(\ell+1)) = s],$$

we have

$$w' = C_\ell w.$$

Observe that for ?-only potential functions, Definition 8.1 and Definition 7.2 are equivalent on any l -boundary-coupled strings. Therefore, we have the following corollary of Lemma 27:

Lemma 27'. *There exists A_ℓ, B_ℓ of size $3^{\ell+1} \times 3^\ell$ that only depends on ℓ such that for a ?-only potential function w with length $\ell + 1$, if there exists z such that*

$$A_\ell z \leq B_\ell w,$$

we have that for any l -boundary-coupled string y ,

$$w(y) \geq 0.$$

Proof of Theorem 18. Using Lemma 25' we can rewrite the conditions in Lemma 28 for a ?-only potential function.

$$\begin{aligned} \text{For any } l\text{-boundary-coupled string } y, \text{ we have} \\ (w - C_\ell w)(y) \geq 0, (w - \{?\})(y) \geq 0 \end{aligned} \quad (26)$$

Then we can use Lemma 27' to get a sufficient condition of (26). That is, (26) holds if the following linear program is feasible.

$$\begin{aligned} A_{l+1}z &\leq B_{l+1}(P_{l+1}(w) - C_\ell w) \\ A_{l+1}z &\leq B_{l+1}(P_{l+1}(w) - P_{l+1}(\{?\})) \\ w &\text{ is ?-only} \end{aligned} \quad (27)$$

The explicit expression for A_ℓ , B_ℓ and C_ℓ can be found in Appendix A. Because the linear program is parameterized by polynomials of p , this can be solved by our algorithm. We find a potential function with length 3 for both IMP and NAND.

For NAND function, the potential function found by the algorithm is the following:

$$\begin{aligned} w = & -\frac{\delta}{1+509\delta/216}\{00?\} + \frac{2+3\delta}{1+509\delta/216}\{01?\} + \frac{4\delta}{1+509\delta/216}\{0?0\} + \frac{2+2\delta}{1+509\delta/216}\{0?1\} \\ & + \frac{2+2\delta}{1+509\delta/216}\{0??\} + \frac{\delta}{1+509\delta/216}\{10?\} + \frac{1+\delta}{1+509\delta/216}\{11?\} + \frac{1-\delta}{1+509\delta/216}\{1?0\} \\ & + \frac{2}{1+509\delta/216}\{1?1\} + \frac{2+3\delta/2}{1+509\delta/216}\{1??\} + \frac{1+223\delta/108}{1+509\delta/216}\{?00\} + \frac{1+439\delta/108}{1+509\delta/216}\{?01\} \\ & + \frac{1+2\delta}{1+509\delta/216}\{?0?\} + \frac{2-4\delta}{1+509\delta/216}\{?10\} + \frac{1+73\delta/432}{1+509\delta/216}\{?11\} + \frac{2-\delta}{1+509\delta/216}\{?1?\} \\ & + \frac{1-4\delta}{1+509\delta/216}\{??0\} + \frac{2-29\delta/48}{1+509\delta/216}\{??1\} + \frac{2-\delta}{1+509\delta/216}\{???\} \end{aligned} \quad (28)$$

For IMP, the potential function is

$$\begin{aligned}
w = & (40 + 6175\delta/38)\{00?\} + (80 + 611\delta/19)\{01?\} + 40\{0?0\} + 80\{0?1\} \\
& + (80 + 200\delta/19)\{0??\} + (80 - 1611\delta/19)\{10?\} + (40 - \delta)\{11?\} + 80\{1?0\} \\
& + (80 + 7521\delta/38)\{1?1\} + (120 + 11807\delta/38)\{1??\} + 40\{?00\} + 80\{?01\} + 80\{?0?\} \\
& + (-40 + 1481\delta/38)\{?10\} + (-40 - 2941\delta/38)\{?11\} + (-20 - 747\delta/19)\{?1?\} \\
& - 19\{??0\} + (10 + 2373\delta/19)\{??1\} + (40 + 5138\delta/19)\{???\}
\end{aligned} \tag{29}$$

□

9 Efficient Solution with Only Equality Constraints

In this section we will prove Theorem 7. Although we showed that the problem of local feasibility of 2-PLP is not tractable, it turns out that the equality version is solvable. For a linear system

$$A(\delta)x = b(\delta), \tag{30}$$

where A and b are polynomials parameterized by $\delta \in \mathbb{R}^d$, we are interested in whether there exists $x \in \mathbb{R}^n$ satisfying the equalities when δ is around the origin.

Theorem 19. *The following decision problem can be solved in polynomial time assuming d is constant. Input an $m \times n$ matrix A and an m -dimensional vector b as polynomials of $\delta \in \mathbb{R}^d$. Is there an $\epsilon > 0$ such that for all $\delta \in B_\epsilon(0)$, there exists x satisfying $A(\delta)x = b(\delta)$?*

Note that this is different than the problem of solving a linear system in the quotient field of $\mathbb{Q}[\delta]$. A general solution $x(\delta)$ in the field that satisfies (30) may not be defined on some value of δ . The most simple example would be $\delta x = 1$ where δ is 1-dimensional. A solution exists in the field $x = 1/\delta$ but the equation is still infeasible when $\delta = 0$. Also, a general solution not being defined at δ_0 does not mean that the system is infeasible at δ_0 , with $\delta^2 x = \delta$ as an example.

The idea is to divide the space of δ into polynomially many regions such that at each region x can be expressed as a rational function of δ . The key is to find regions where the pseudoinverse of A is continuous.

Definition 9.1. A real algebraic set in \mathbb{R}^d is a set of form

$$V(F) = \{\delta \in \mathbb{R}^d : p(\delta) = 0, \forall p \in F\}$$

where F is a set of polynomials. This is called the variety of F . Use $V(p)$ to denote the variety of a single polynomial, $\{\delta \in \mathbb{R}^d : p(\delta) = 0\}$.

9.1 Preliminaries in Algebraic Geometry

Definition 9.2. An irreducible algebraic set S is an algebraic set that cannot be the union of two algebraic sets S_1 and S_2 such that both of them is neither \emptyset nor S .

An irreducible component of an algebraic set S is a subset of S that is an irreducible algebraic set.

Lemma 29. *For any d irreducible polynomials p_1, p_2, \dots, p_d in ring $\mathbb{Q}[\delta_1, \delta_2, \dots, \delta_d]$, any irreducible component of $\cap_{i=1}^d V(p_i) = V(\{p_i : i \in [d]\})$ is 0-dimensional, i.e., a single point on \mathbb{R}^d .*

Proof. First choose any irreducible branch of $V(\{p_i : i \in [d]\})$ and suppose its ideal is I . From Theorem 1.8A in [Har13], $\dim(\mathbb{Q}[\delta_1, \delta_2, \dots, \delta_d]/I) = d - \text{ht}(I)$ where ht stands for height of an ideal. We have $(0) \subsetneq (p_1) \subsetneq (p_1, p_2) \subsetneq \dots (p_1, \dots, p_d) \subseteq I$. Also, $\text{ht}(I) \leq d$. So the height of I is d , which means

$$\dim(\mathbb{Q}[\delta_1, \delta_2, \dots, \delta_d]/I) = 0.$$

The extension from $\mathbb{Q}[\delta_1, \delta_2, \dots, \delta_d]$ to $\mathbb{C}[\delta_1, \delta_2, \dots, \delta_d]$ preserves dimension, so $\mathbb{C}[\delta]/I$ is 0-dimensional. This means $\mathbb{R}[\delta]/I$ can only be 0-dimensional \square

9.2 Algorithm

Recall that $Ax = b$ is feasible if and only if $AA^+b - b = 0$ (Lemma 2). So a direct approach is to calculate A^+ . The main issue is that $A_J^+(\delta_1, \delta_2, \dots, \delta_d)$ is not necessarily a continuous function on \mathbb{R}^d . But we will show that this is true if we further divide \mathbb{R}^d into polynomially many regions where A_J^+ is a rational function on each of them.

Lemma 30. *There exists an $\epsilon > 0$ such that $B_\epsilon(0)$ can be divided into polynomially many regions and on each region $A^+(\delta)$ is a continuous rational function. A^+ on each region can be computed in polynomial time (given as a rational function).*

Proof. To this end, we use the following expression for A_J^+ again.

$$\begin{aligned} A_J^+ &= \lim_{y \rightarrow 0} (A^T A + yI)^{-1} A_J^T \\ &= \lim_{y \rightarrow 0} \frac{1}{\det(A^T A + yI)} \text{adj}(A^T A + yI) A_J^T, \end{aligned}$$

where $\text{adj}(B)$ denotes the adjugate of a matrix B , equal to the transpose of the cofactor matrix.

Note that $\det(A^T A + yI)$ and $\text{adj}(A^T A + yI) A_J^T$ are polynomials in δ_1, δ_2 and y . Now assume

$$\begin{aligned} \det(A^T A + yI) &= q_0(\delta) + yq_1(\delta) + y^2q_2(\delta) + \dots y^kq_k(\delta) \\ \text{adj}(A^T A + yI) A_J^T &= p_0(\delta) + yp_1(\delta) + y^2p_2(\delta) + \dots y^kp_k(\delta), \end{aligned}$$

where p_i are matrices and $k \leq m$. An alternative expression for A_J^+ is the following:

$$A_J^+(\delta) = \begin{cases} \frac{p_0(\delta)}{q_0(\delta)}, & \text{when } q_0(\delta) \neq 0 \\ \frac{p_1(\delta)}{q_1(\delta)}, & \text{when } q_0(\delta) = 0, q_1(\delta) \neq 0 \\ \dots & \\ \frac{p_k(\delta)}{q_k(\delta)}, & \text{when } q_0(\delta) = q_1(\delta) = \dots = q_{k-1}(\delta) = 0 \end{cases} \quad (31)$$

An important observation is that Suppose on a certain region $R \subset \mathbb{R}^d$, for any $1 \leq i \leq k$, whether $q_i = 0$ is fixed. Then $A^+(\delta)$ can be expressed as a rational function on R , given by the expression above.

Now take all irreducible factors of non-zero polynomials from q_0, q_1, \dots, q_k , suppose they are f_1, \dots, f_r . The degree of q_i is bounded by md , so $r \leq kmd \leq m^2d$. There is a line of research for polynomial-time algorithms for factorizing in $\mathbb{Q}[\delta_1, \delta_2, \dots, \delta_d]$, a survey can be found in [Kal92]. So the running time of this step is bounded by a polynomial.

Note that if $f_i(0) \neq 0$, f_i would be non-zero on $B_\epsilon(0)$ for a small enough ϵ since it is continuous. So we only care about varieties of all f_i where $f_i(0) = 0$. Suppose the varieties of these f_i 's are $V_1, \dots, V_{r'}$. By Lemma 29, the intersection of any d of the V_i 's is either empty or finite number of

points on \mathbb{R}^2 . Therefore, we can choose ϵ small enough such that the intersection of any d -tuple of V_i on $B_\epsilon(0)$ is either empty or equal to $\{0\}$. Let V_S where $S \subset [r]$ denote $\cap_{i \in S} V_i$. Now we can partition $B_\epsilon(0)$ into the following regions:

$$\begin{aligned} & \{0\}, B_\epsilon(0) \setminus (\cup_{|S| \leq d, S \subset [r]} V_S), \\ & B_\epsilon(0) \cap V_S, \forall |S| \leq d-1, S \subset [r]. \end{aligned} \tag{32}$$

f_1, \dots, f_r are either always zero or always non-zero on any region, so the same holds for q_1, \dots, q_k . So on each region A^+ is given by (31). \square

Now we are ready to prove Theorem 7. With Lemma 30, $B_\epsilon(0)$ is partitioned into polynomially many regions in (32). On each region A^+ is a continuous rational function and so is $AA^+b - b$. We can check the feasibility of $Ax = b$ by checking whether the following rational function $AA^+b - b$ is zero on every region. There are three cases:

1. The region is $\{0\}$, A and b are fixed. So the system becomes a linear system without parameter.
2. The region is $B_\epsilon(0) \setminus (\cup_{|S| \leq d, S \subset [r]} V_S)$, assume $AA^+b - b = g(\delta)$. The region is d -dimensional, so $g(\delta) = 0$ on this region if and only if $g = 0$.
3. The region is $B_\epsilon(0) \cap V_S$ for $|S| \leq d-1$ where V_S is the intersection of varieties of V_i for all $i \in S$, let f_{k_i} be the corresponding irreducible polynomial. So we only need to check whether the numerator of $AA^+b - b$, denoted by $g(\delta)$, is zero on V_S . Since f_{k_i} are irreducible, this is equivalent to whether there exists $i \in S$ such that f_{k_i} divides $g(\delta)$.

The overall algorithm can be expressed as the following.

Algorithm 3 The local feasibility of $A(\delta)x = b(\delta)$

- 1: Calculate $\det(A^T A + yI) = \sum_{i=0}^k y^i q_i(\delta)$
 - 2: Factorize every q_i and get a set of irreducible factors in $\mathbb{Q}[\delta_1, \delta_2, \dots, \delta_d]$, f_1, f_2, \dots, f_r
 - 3: Pick factors that are 0 when $\delta = 0$, suppose they are $f_{k_1}, \dots, f_{k_{r'}}$. Denote their varieties by $V_1, \dots, V_{r'}$
 - 4: Calculate the numerator of $AA^+b - b$ on $B_\epsilon(0) \cap V_S$ for $|S| \leq d-1$, on $B_\epsilon(0) \setminus (\cup_{|S| \leq d, S \subset [r]} V_S)$, and on 0 by (31). Denote it by $g(\delta)$.
 - 5: Check the following three conditions:
 - $g(0) = 0$
 - $g(\delta) = 0$ on $B_\epsilon(0) \setminus (\cup_{|S| \leq d, S \subset [r]} V_S)$
 - $\forall S \subset [r'], \exists i \in S$ that f_{k_i} divides $g(\delta)$ on V_S
 - 6: **if** all conditions are met **then**
 - 7: **return Feasible**
 - 8: **else**
 - 9: **return Not feasible**
 - 10: **end if**
-

| | 0 | 1 | ? |
|----|---------|---------|----------|
| 00 | p | $1 - p$ | 0 |
| 01 | p | $1 - p$ | 0 |
| 0? | p | $1 - p$ | 0 |
| 10 | p | $1 - p$ | 0 |
| 11 | $1 - p$ | p | 0 |
| 1? | p | p | $1 - 2p$ |
| ?0 | p | $1 - p$ | 0 |
| ?1 | p | p | $1 - 2p$ |
| ?? | p | p | $1 - 2p$ |

Table 1: Transition matrix for PCA with NAND function and vertex noise

| | 0 | 1 | ? |
|----|----------------|---------------|-----------------|
| 00 | p^2 | $1 - p^2$ | 0 |
| 01 | $p - p^2$ | $1 - p + p^2$ | 0 |
| 0? | p^2 | $1 - p + p^2$ | $p - 2p^2$ |
| 10 | $p - p^2$ | $1 - p + p^2$ | 0 |
| 11 | $1 - 2p + p^2$ | $2p - p^2$ | 0 |
| 1? | $p - p^2$ | $2p - p^2$ | $1 - 3p + 2p^2$ |
| ?0 | p^2 | $1 - p + p^2$ | $p - 2p^2$ |
| ?1 | $p - p^2$ | $2p - p^2$ | $1 - 3p + 2p^2$ |
| ?? | p^2 | $2p - p^2$ | $1 - 2p$ |

Table 2: Transition matrix for PCA with NAND function and vertex noise, and for broadcasting on grid with NAND function

A Details of Lemma 25 and Lemma 27

In this section we describe the exact expression of matrices used in Lemma 25 and Lemma 27 in both application to PCA and application to broadcasting on grid.

In Lemma 25, the matrix C_ℓ is defined to be

$$\begin{aligned}
C_\ell(s, s') &= \Pr[(\eta_{t+1}(1), \dots, \eta_{t+1}(\ell)) = s' | (\eta_t(1 - a), \dots, \eta_t(\ell + b)) = s] \\
&= \prod_{n=a+1}^{\ell+a+1} \Pr[f(s(n - a), s(n - a + 1), \dots, s(n + b)) = s'(n - a)],
\end{aligned}$$

In all our applications, $a = 1$ and $b = 0$, so c_ℓ is a $3^{\ell+1} \times 3^\ell$ matrix. C_1 is simply the transition matrix. For PCA with NAND function and vertex noise, we can get from the definition that C_1 is Table 1. For PCA with NAND function and edge noise, or broadcasting of information on grid with NAND function, C_1 is Table 2.

For broadcasting of information on the grid with IMP function, C_1 is given by Table 3.

For arbitrary ℓ , we can define

$$C_\ell(s, s') = \prod_{i=1}^{\ell} C_1((s_i, s_{i+1}), (s'_i)).$$

| | 0 | 1 | ? |
|----|----------------|---------------|-----------------|
| 00 | $p - p^2$ | $1 - p + p^2$ | 0 |
| 01 | p^2 | $1 - p^2$ | 0 |
| 0? | p^2 | $1 - p + p^2$ | $p - 2p^2$ |
| 10 | $1 - 2p + p^2$ | $2p - p^2$ | 0 |
| 11 | $p - p^2$ | $1 - p + p^2$ | 0 |
| 1? | $p - p^2$ | $2p - p^2$ | $1 - 3p + 2p^2$ |
| ?0 | $p - p^2$ | $2p - p^2$ | $1 - 3p + 2p^2$ |
| ?1 | p^2 | $1 - p + p^2$ | $p - 2p^2$ |
| ?? | p^2 | $2p - p^2$ | $1 - 2p$ |

Table 3: tab:Transition matrix for broadcasting on grid with IMP function

Note that C_ℓ is a matrix polynomial of p in all the settings we discuss.

For Lemma 27, the matrices are defined through a graph that connects patterns that appear consecutively in a string. Let G be a directed weighted graph with length l strings as vertices and there is an edge between two vertices s_0 and s_1 if $(s_0(1), s_0(2), \dots, s_0(l-1)) = (s_1(0), s_1(1), \dots, s_1(l-2))$, where $s(i)$ stands for the i th symbol of s . For a potential functions $w = \sum_{s \in S} c_s \{s\}$ with length ℓ , each edge (s, s') has weight c_s if $s \in S$ and 0 otherwise. This way, any string s would corresponds to a circle on the graph with i th node being $(s_{i \bmod \ell}, s_{i+1 \bmod \ell}, \dots, s_{i+\ell \bmod \ell})$. $w(s)$ is equal to weight of the cycle defined by the sum of weight of all nodes on the cycle. $w(s) \geq 0$ for any s would be equivalent to having no negative cycles.

Let B_{in} and B_{out} be matrices of size $3^\ell \times 3^{\ell+1}$ with vertices of G being rows and edges being columns. Entry of B_{in} at row v and column e is 1 if the corresponding e points to v and 0 otherwise. Similarly, B_{out} has entry 1 if e points out from v .

We can show that to show a directed graph having no negative cycle, it is sufficient to show that

$$(B_{out} - B_{in})^T z \leq B_{out}^T c$$

where c is the vector of weights of all edges. So A_ℓ in Lemma 27 is $(B_{out} - B_{in})^T$ and B_ℓ is B_{out}^T .

B Proof of of Lemma 23

To simplify the discussion, we extend the definition of *NAND* and *BSC* channel on alphabet $\{0, 1, ?\}$ as follows.

| NAND | result |
|------|--------|
| 0* | 1 |
| 11 | 0 |
| 1? | ? |
| ?1 | ? |
| ?? | ? |

| BSC | 0 | 1 | ? |
|-----|---------|---------|----------|
| 0 | $1 - p$ | p | 0 |
| 1 | p | $1 - p$ | 0 |
| ? | p | p | $1 - 2p$ |

So the chain of F'_p can still be written as

$$\eta_{t+1}(n) = \text{NAND}(\text{BSC}(\eta_t(n)), \text{BSC}(\eta_t(n-1)))$$

Here the only randomness comes from the extended BSC, which we refer to as edge noise.

Definition B.1. We define a partial order between symbols $0 < ? > 1$. The order between configurations are done by entry-wise comparison, and the order between distribution of configurations are defined by stochastic domination.

Lemma 31. *For two configurations η and η' such that $\eta \geq \eta'$, the order remains after evolution under the same realization of edge noises.*

For two distributions μ and ν on configurations, if $\mu \geq \nu$, then $F'_p\mu \geq F'_p\nu$.

Proof. The first statement is easy to check based on the definition of NAND and BSC .

For the second statement, since $\mu \geq \nu$, there exists a coupling (η, η') between the distributions such that $\eta \geq \eta'$ always holds. We can couple them using same realization of edge noise, we have $F'_p\delta_\eta \geq F'_p\delta_{\eta'}$. So $F'_p\mu \geq F'_p\nu$. \square

Proof of Lemma 23. Let us first extend the time range of the PCA to \mathbb{Z} . Each BSC has probability p being 0, probability p being 1 and probability $1 - 2p$ being the original symbol. Let τ be the random outcome of extended BSCs at all time and location, and define a random configuration η^* as follows. $\eta^*(i)$ is 1 (or 0) if $\exists t > 0$ such that starting with arbitrary configuration on η_{-t} and let the PCA evolve based on τ , $\eta_0(i)$ would be 1 (or 0). Otherwise $\eta^*(i)$ is ?.

Note by Lemma 31, the order between configurations remains under same edge noise. Let $\eta^?$ be the result of η_0 with η_{-t} being all ?, and η^a be the result of η_0 with η_{-t} being an arbitrary configuration. We have $\eta^? \geq \eta^a$. So $\eta^?(i)$ is 1 (or 0) imply that $\eta^a(i)$ being 1 (or 0).

Therefore, η^* can be equivalently defined by the following limit. Let η^t be that $\eta^t(i)$ is 1 (or 0) if starting with a configuration with all ? on η_{-t} and let the PCA evolve based on τ , $\eta_0(i)$ would be 1 (or 0). Otherwise $\eta^t(i)$ is ?. And $\eta^* = \lim_{t \rightarrow \infty} \eta^t$. This means starting with all ?, the PCA converges to η^* . So η^* is a stationary distribution and has 0 probability of containing ?.

Suppose we start with arbitrary configuration, by Lemma 31, $\lim_{t \rightarrow \infty} \eta_t \leq \eta^*$. When η^* have no ?, the only configuration η satisfying $\eta \leq \eta^*$ is itself. So we proved that arbitrary configuration would converge to η^* . \square

C Potential Functions for Ergodicity of PCA

We give an explicit solution of the potential function for Theorem 16 and Theorem 17. The following is a feasible potential function that satisfies (23) and proves Theorem 16.

$$\begin{aligned} w = & (-4 + 4\delta)\{000\} + (-8 + 4\delta)\{001\} + (-1 - 2\delta)\{00?\} + (2 + 4\delta)\{010\} \\ & + (-2 + 4\delta)\{011\} + (1 + 5\delta/2)\{01?\} + (-2 + 4\delta)\{0?0\} + (-3 - 13\delta)\{0?1\} \\ & + (-2 - 6\delta)\{0??\} + (-6 + 4\delta)\{100\} + (-10 + 4\delta)\{101\} + 3\{10?\} \\ & + 4\delta\{110\} + (-4 + 4\delta)\{111\} + (-1 - 2\delta)\{11?\} + (-3 - 21\delta)\{1?0\} \\ & - 12\{1?1\} + (-5 + 7\delta/2)\{1??\} + (8 - 8\delta)\{?00\} + 13\{?01\} \\ & + 21\{?0?\} + (8 + 3\delta)\{?10\} + (4 - 4\delta)\{?11\} + (7 - 16\delta)\{?1?\} \\ & + (6 - 26\delta)\{??0\} + (4 - 15\delta)\{??1\} + (6 - 43\delta)\{???\} \end{aligned} \tag{33}$$

The following is a feasible potential function that proves Theorem 17.

$$\begin{aligned}
w = & (-2 + \delta)\{000\} + (-2 + \delta)\{001\} + (-4 - \delta)\{00?\} + (-2 + \delta)\{010\} \\
& + (-2 + \delta)\{011\} + (-9 - 25\delta)\{01?\} - 2\{0?0\} + (5 - 52\delta)\{0?1\} \\
& + (3 - 28\delta)\{0??\} + (-2 + \delta)\{100\} + (-2 + \delta)\{101\} + (-4 + \delta)\{10?\} \\
& + (-2 + \delta)\{110\} + (-2 + \delta)\{111\} + (-10 - 45055\delta/2496)\{11?\} + (6 - 4275059\delta/87360)\{1?0\} \\
& - (12 - 88969\delta/2496)\{1?1\} + (10 - 65239\delta/5824)\{1??\} \\
& + (2 + 374207\delta/33280)\{?00\} + (2 + 440767\delta/33280)\{?01\} \\
& + 284767\delta\{?0?\}/33280 + (-2 - 1877\delta/210)\{?10\} \\
& + (-3 - 28031\delta/8736)\{?11\} + (-12 + 242057\delta/8736)\{?1?\} \\
& + (-4 - 9596961\delta/232960)\{??0\} + (2 - 786109\delta/29120)\{??1\} - 15311\delta\{???\}/4368
\end{aligned} \tag{34}$$

acknowledgements

GB was supported in part by NSF CAREER award CCF-1940205. This material is based upon work supported by the National Science Foundation under Grants No CCF-1717842, CCF-2131115.

References

- [ACM84] Dennis S Arnon, George E Collins, and Scott McCallum. Cylindrical algebraic decomposition I: The basic algorithm. *SIAM Journal on Computing*, 13(4):865–877, 1984.
- [Alb72] Arthur Albert. Regression and the moore-penrose pseudoinverse. Technical report, 1972.
- [BBC11] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM review*, 53(3):464–501, 2011.
- [BD07] Christopher W Brown and James H Davenport. The complexity of quantifier elimination and cylindrical algebraic decomposition. In *Proceedings of the 2007 international symposium on Symbolic and algebraic computation*, pages 54–60, 2007.
- [BM04] Yann Bugeaud and Maurice Mignotte. On the distance between roots of integer polynomials. *Proceedings of the Edinburgh Mathematical Society*, 47(3):553–556, 2004.
- [BPS04] Dimitris Bertsimas, Dessislava Pachamanova, and Melvyn Sim. Robust linear optimization under general norms. *Operations Research Letters*, 32(6):510–516, 2004.
- [BRZ95] Pavel M Bleher, Jean Ruiz, and Valentin A Zagrebnov. On the purity of the limiting gibbs state for the Ising model on the bethe lattice. *Journal of Statistical Physics*, 79(1):473–482, 1995.
- [BS04] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1):35–53, 2004.
- [BTN99] Aharon Ben-Tal and Arkadi Nemirovski. Robust solutions of uncertain linear programs. *Operations research letters*, 25(1):1–13, 1999.
- [CD98] Bastien Chopard and Michel Droz. Cellular automata modeling of physical systems,(1998). *Cambridge University Press*, pages 122–137, 1998.
- [CG06] Coralia Cartis and Nicholas IM Gould. Finding a point in the relative interior of a polyhedron. Technical report, Citeseer, 2006.
- [CR88] Michel Coste and Marie-Françoise Roy. Thom’s lemma, the coding of real algebraic numbers and the computation of the topology of semi-algebraic sets. *Journal of Symbolic Computation*, 5(1-2):121–129, 1988.
- [DH88] James H Davenport and Joos Heintz. Real quantifier elimination is doubly exponential. *Journal of Symbolic Computation*, 5(1-2):29–35, 1988.
- [DO77] Roland L’vovich Dobrushin and SI Ortyukov. Lower bound for the redundancy of self-correcting arrangements of unreliable functional elements. *Problemy Peredachi Informatsii*, 13(1):82–89, 1977.
- [DPSV69] RL Dobrushin, II Pyatetski-Shapiro, and NB Vasilev. Markov processes in an infinite product of discrete spaces. In *Soviet-Japanese Symposium in Probability Theory, Khavarovsk*, 1969.
- [ED16] Matthew England and James H Davenport. The complexity of cylindrical algebraic decomposition with respect to polynomial degree. In *International Workshop on Computer Algebra in Scientific Computing*, pages 172–192. Springer, 2016.

- [EEK93] G Bard Ermentrout and Leah Edelstein-Keshet. Cellular automata approaches to biological modeling. *Journal of theoretical Biology*, 160(1):97–133, 1993.
- [EKPS00] William Evans, Claire Kenyon, Yuval Peres, and Leonard J Schulman. Broadcasting on trees and the Ising model. *Annals of Applied Probability*, pages 410–433, 2000.
- [ES99] William S Evans and Leonard J Schulman. Signal propagation and noisy circuits. *IEEE Transactions on Information Theory*, 45(7):2367–2373, 1999.
- [FO85] Robert M Freund and James B Orlin. On the complexity of four polyhedral set containment problems. *Mathematical programming*, 33(2):139–145, 1985.
- [Gác86] Peter Gács. Reliable computation with cellular automata. *Journal of Computer and System Sciences*, 32(1):15–78, 1986.
- [GWS80] Lawrence Gray, John C Wierman, and RT Smythe. Lower bounds for the critical probability in percolation models with oriented bonds. *Journal of Applied Probability*, 17(4):979–986, 1980.
- [Har13] Robin Hartshorne. *Algebraic geometry*, volume 52. Springer Science & Business Media, 2013.
- [HMM19] Alexander E Holroyd, Irène Marcovici, and James B Martin. Percolation games, probabilistic cellular automata, and the hard-core model. *Probability Theory and Related Fields*, 174(3):1187–1217, 2019.
- [Kal92] Erich Kaltofen. Polynomial factorization 1987–1991. In *Latin American Symposium on Theoretical Informatics*, pages 294–313. Springer, 1992.
- [LN18] Pierre-Yves Louis and Francesca R Nardi. *Probabilistic cellular automata: Theory, applications and future perspectives*, volume 27. Springer, 2018.
- [MM14] Jean Mairesse and Irene Marcovici. Around probabilistic cellular automata. *Theoretical Computer Science*, 559:42–72, 2014.
- [MMP19] Anuran Makur, Elchanan Mossel, and Yury Polyanskiy. Broadcasting on random networks. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 1632–1636. IEEE, 2019.
- [MMP21] Anuran Makur, Elchanan Mossel, and Yury Polyanskiy. Reconstruction on 2d regular grids. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pages 2107–2112. IEEE, 2021.
- [NB66] John von Neumann and Arthur W Burks. *Theory of self-reproducing automata*. University of Illinois Press, 1966.
- [PV91] Panos M Pardalos and Stephen A Vavasis. Quadratic programming with one negative eigenvalue is np-hard. *Journal of Global optimization*, 1(1):15–22, 1991.
- [PW17] Yury Polyanskiy and Yihong Wu. Strong data-processing inequalities for channels and bayesian networks. In *Convexity and Concentration*, pages 211–249. Springer, 2017.
- [Too95] Andrei Toom. Cellular automata with errors: problems for students of probability. *Topics in contemporary probability and its applications*, pages 117–157, 1995.

- [TVS⁺90] AL Toom, NB Vasilyev, ON Stavskaya, LG Mityushin, GL Kurdyumov, and SA Pirogov. Discrete local Markov systems. *Stochastic cellular systems: ergodicity, memory, morphogenesis*, 1, 1990.