

COMPUTING LEAST FIXED POINTS OF PROBABILISTIC SYSTEMS OF POLYNOMIALS

JAVIER ESPARZA AND ANDREAS GAISER AND STEFAN KIEFER

Fakultät für Informatik, Technische Universität München, Germany
E-mail address: {esparza, gaiser, kiefer}@model.in.tum.de

ABSTRACT. We study systems of equations of the form $X_1 = f_1(X_1, \dots, X_n), \dots, X_n = f_n(X_1, \dots, X_n)$ where each f_i is a polynomial with nonnegative coefficients that add up to 1. The least nonnegative solution, say μ , of such equation systems is central to problems from various areas, like physics, biology, computational linguistics and probabilistic program verification. We give a simple and strongly polynomial algorithm to decide whether $\mu = (1, \dots, 1)$ holds. Furthermore, we present an algorithm that computes reliable sequences of lower and upper bounds on μ , converging linearly to μ . Our algorithm has these features despite using inexact arithmetic for efficiency. We report on experiments that show the performance of our algorithms.

1. Introduction

We study how to efficiently compute the least nonnegative solution of an equation system of the form

$$X_1 = f_1(X_1, \dots, X_n) \quad \dots \quad X_n = f_n(X_1, \dots, X_n),$$

where, for every $i \in \{1, \dots, n\}$, f_i is a polynomial over X_1, \dots, X_n with positive rational coefficients that *add up to 1*.¹ The solutions are the fixed points of the function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ with $f = (f_1, \dots, f_n)$. We call f a *probabilistic system of polynomials* (short: *PSP*). E.g., the PSP

$$f(X_1, X_2) = \left(\frac{1}{2}X_1X_2 + \frac{1}{2}, \frac{1}{4}X_2X_2 + \frac{1}{4}X_1 + \frac{1}{2} \right)$$

induces the equation system

$$X_1 = \frac{1}{2}X_1X_2 + \frac{1}{2} \quad X_2 = \frac{1}{4}X_2X_2 + \frac{1}{4}X_1 + \frac{1}{2}.$$

Obviously, $\bar{1} = (1, \dots, 1)$ is a fixed point of every PSP. By Kleene's theorem, every PSP has a least nonnegative fixed point (called just least fixed point in what follows), given by the limit of the sequence $\bar{0}, f(\bar{0}), f(f(\bar{0})), \dots$

PSPs are important in different areas of the theory of stochastic processes and computational models. A fundamental result of the theory of branching processes, with numerous applications in physics, chemistry and biology (see e.g. [9, 2]), states that extinction probabilities of species are

1998 ACM Subject Classification: F.2.1 Numerical Algorithms and Problems, G.3 Probability and Statistics.

Key words and phrases: computing fixed points, numerical approximation, stochastic models, branching processes.

¹Later, we allow that the coefficients add up to *at most* 1.

equal to the least fixed point of a PSP. The same result has been recently shown for the probability of termination of certain probabilistic recursive programs [7, 6]. The consistency of stochastic context-free grammars, a problem of interest in statistical natural language processing, also reduces to checking whether the least fixed point of a PSP equals $\bar{1}$ (see e.g. [11]).

Given a PSP f with least fixed point μ_f , we study how to efficiently solve the following two problems: (1) decide whether $\mu_f = \bar{1}$, and (2) given a rational number $\epsilon > 0$, compute $\mathbf{lb}, \mathbf{ub} \in \mathbb{Q}^n$ such that $\mathbf{lb} \leq \mu_f \leq \mathbf{ub}$ and $\mathbf{ub} - \mathbf{lb} \leq \bar{\epsilon}$ (where $\mathbf{u} \leq \mathbf{v}$ for vectors \mathbf{u}, \mathbf{v} means \leq in all components). While the motivation for Problem (2) is clear (compute the probability of extinction with a given accuracy), the motivation for Problem (1) requires perhaps some explanation. In the case study of Section 4.3 we consider a family of PSPs, taken from [9], modelling the neutron branching process in a ball of radioactive material of radius D (the family is parameterized by D). The least fixed point is the probability that a neutron produced through spontaneous fission *does not* generate an infinite “progeny” through successive collisions with atoms of the ball; loosely speaking, this is the probability that the neutron *does not* generate a chain reaction and the ball *does not* explode. Since the number of atoms in the ball is very large, spontaneous fission produces many neutrons per second, and so even if the probability that a given neutron produces a chain reaction is very small, the ball will explode with large probability in a very short time. It is therefore important to determine the largest radius D at which the probability of no chain reaction is still 1 (usually called the *critical radius*). An algorithm for Problem (1) allows to compute the critical radius using binary search. A similar situation appears in the analysis of parameterized probabilistic programs. In [7, 6] it is shown that the question whether a probabilistic program almost surely terminates can be reduced to Problem (1). Using binary search one can find the “critical” value of the parameter for which the program may not terminate any more.

Etessami and Yannakakis show in [7] that Problem (1) can be solved in polynomial time by a reduction to (exact) Linear Programming (LP), which is not known to be strongly polynomial. Our first result reduces Problem (1) to solving a system of linear equations, resulting in a strongly polynomial algorithm for Problem (1). The Maple library offers exact arithmetic solvers for LP and systems of linear equations, which we use to test the performance of our new algorithm. In the neutron branching process discussed above we obtain speed-ups of about one order of magnitude with respect to LP.

The second result of the paper is, to the best of our knowledge, the first practical algorithm for Problem (2). Lower bounds for μ_f can be computed using Newton’s method for approximating a root of the function $f(\bar{X}) - \bar{X}$. This has recently been investigated in detail [7, 10, 5]. However, Newton’s method faces considerable numerical problems. Experiments show that naive use of exact arithmetic is inefficient, while floating-point computation leads to false results even for very small systems. For instance, the PReMo tool [12], which implements Newton’s method with floating-point arithmetic for efficiency, reports $\mu_f \geq \bar{1}$ for a PSP with only 7 variables and small coefficients, although $\mu_f < \bar{1}$ is the case (see Section 3.1).

Our algorithm produces a sequence of guaranteed lower and upper bounds, both of which converge linearly to μ_f . Linear convergence means that, loosely speaking, the number of accurate bits of the bound is a linear function of the position of the bound in the sequence. The algorithm is based on the following idea. Newton’s method is an iterative procedure that, given a current lower bound \mathbf{lb} on μ_f , applies a certain operator \mathcal{N} to it, yielding a new, more precise lower bound $\mathcal{N}(\mathbf{lb})$. Instead of computing $\mathcal{N}(\mathbf{lb})$ using exact arithmetic, our algorithm computes *two* consecutive Newton steps, i.e., $\mathcal{N}(\mathcal{N}(\mathbf{lb}))$, using *inexact* arithmetic. Then it checks if the result satisfies a carefully chosen condition. If so, the result is taken as the next lower bound. If not, then the precision is increased, and the computation redone. The condition is eventually satisfied, assuming the results

of computing with increased precision converge to the exact result. Usually, the repeated inexact computation is much faster than the exact one. At the same time, a careful (and rather delicate) analysis shows that the sequence of lower bounds converges linearly to μ_f .

Computing *upper* bounds is harder, and seemingly has not been considered in the literature before. Similarly to the case of lower bounds, we apply f twice to \mathbf{ub} , i.e., we compute $f(f(\mathbf{ub}))$ with increasing precision until a condition holds. The sequence so obtained may not even converge to μ_f . So we need to introduce a further operation, after which we can then prove linear convergence.

We test our algorithm on the neutron branching process. The time needed to obtain lower and upper bounds on the probability of no explosion with $\epsilon = 0.0001$ lies below the time needed to check, using exact LP, whether this probability is 1 or smaller than one. That is, in this case study our algorithm is faster, and provides more information.

The rest of the paper is structured as follows. We give preliminary definitions and facts in Section 2. Sections 3 and 4 present our algorithms for solving Problems (1) and (2), and report on their performance on some case studies. Section 5 contains our conclusions. The full version of the paper, including all proofs, can be found in [4].

2. Preliminaries

Vectors and matrices. We use bold letters for designating (column) vectors, e.g. $\mathbf{v} \in \mathbb{R}^n$. We write \bar{s} with $s \in \mathbb{R}$ for the vector $(s, \dots, s)^\top \in \mathbb{R}^n$ (where $^\top$ indicates transpose), if the dimension n is clear from the context. The i -th component of $\mathbf{v} \in \mathbb{R}^n$ will be denoted by v_i . We write $\mathbf{x} = \mathbf{y}$ (resp. $\mathbf{x} \leq \mathbf{y}$ resp. $\mathbf{x} < \mathbf{y}$) if $x_i = y_i$ (resp. $x_i \leq y_i$ resp. $x_i < y_i$) holds for all $i \in \{1, \dots, n\}$. By $\mathbf{x} < \mathbf{y}$ we mean $\mathbf{x} \leq \mathbf{y}$ and $\mathbf{x} \neq \mathbf{y}$. By $\mathbb{R}^{m \times n}$ we denote the set of real matrices with m rows and n columns. We write Id for the identity matrix. For a square matrix A , we denote by $\rho(A)$ the *spectral radius* of A , i.e., the maximum of the absolute values of the eigenvalues. A matrix is *nonnegative* if all its entries are nonnegative. A nonnegative matrix $A \in \mathbb{R}^{n \times n}$ is *irreducible* if for every $k, l \in \{1, \dots, n\}$ there exists an $i \in \mathbb{N}$ so that $(A^i)_{kl} \neq 0$.

Probabilistic Systems of Polynomials. We investigate equation systems of the form

$$X_1 = f_1(X_1, \dots, X_n) \quad \dots \quad X_n = f_n(X_1, \dots, X_n),$$

where the f_i are polynomials in the variables X_1, \dots, X_n with positive real coefficients, and for every polynomial f_i the sum of its coefficients is *at most* 1. The vector $f := (f_1, \dots, f_n)^\top$ is called a *probabilistic system of polynomials* (PSP for short) and is identified with its induced function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$. If X_1, \dots, X_n are the formal variables of f , we define $\bar{X} := (X_1, \dots, X_n)^\top$ and $\text{Var}(f) := \{X_1, \dots, X_n\}$. We assume that f is represented as a list of polynomials, and each polynomial is a list of its monomials. If $S \subseteq \{X_1, \dots, X_n\}$, then f_S denotes the result of removing the polynomial $f_i(X_1, \dots, X_n)$ from f for every $x_i \notin S$; further, given $\mathbf{x} \in \mathbb{R}^n$ and $B \in \mathbb{R}^{n \times n}$, we denote by \mathbf{x}_S and B_{SS} the vector and the matrix obtained from \mathbf{x} and B by removing the entries with indices i such that $X_i \notin S$. The coefficients are represented as fractions of positive integers. The *size* of f is the size of that representation. The *degree* of f is the maximum of the degrees of f_1, \dots, f_n . PSPs of degree 0 (resp. 1 resp. >1) are called *constant* (resp. *linear* resp. *superlinear*). PSPs f where the degree of each f_i is at least 2 are called *purely superlinear*. We write f' for the *Jacobian* of f , i.e., the matrix of first partial derivatives of f .

Given a PSP f , a variable X_i *depends directly* on a variable X_j if X_j ‘‘occurs’’ in f_i , more formally if $\frac{\partial f_i}{\partial X_j}$ is not the constant 0. A variable X_i *depends* on X_j if X_i depends directly on X_j

or there is a variable X_k such that X_i depends directly on X_k and X_k depends on X_j . We often consider the *strongly connected components* (or SCCs for short) of the dependence relation. The SCCs of a PSP can be computed in linear time using e.g. Tarjan's algorithm. An SCC S of a PSP f is *constant* resp. *linear* resp. *superlinear* resp. *purely superlinear* if the PSP \tilde{f} has the respective property, where \tilde{f} is obtained by restricting f to the S -components and replacing all variables not in S by the constant 1. A PSP is an *scPSP* if it is not constant and consists of only one SCC. Notice that a PSP f is an scPSP if and only if $f'(\bar{1})$ is irreducible.

A fixed point of a PSP f is a vector $\mathbf{x} \geq \bar{0}$ with $f(\mathbf{x}) = \mathbf{x}$. By Kleene's theorem, there exists a least fixed point μ_f of f , i.e., $\mu_f \leq \mathbf{x}$ holds for every fixed point \mathbf{x} . Moreover, the sequence $\bar{0}, f(\bar{0}), f(f(\bar{0})), \dots$ converges to μ_f . Vectors \mathbf{x} with $\mathbf{x} \leq f(\mathbf{x})$ (resp. $\mathbf{x} \geq f(\mathbf{x})$) are called *pre-fixed* (resp. *post-fixed*) points. Notice that the vector $\bar{1}$ is always a post-fixed point of a PSP f , due to our assumption on the coefficients of a PSP. By Knaster-Tarski's theorem, μ_f is the least post-fixed point, so we always have $\bar{0} \leq \mu_f \leq \bar{1}$. It is easy to detect and remove all components i with $(\mu_f)_i = 0$ by a simple round-robin method (see e.g. [5]), which needs linear time in the size of f . We therefore assume in the following that $\mu_f \succ \bar{0}$.

3. An algorithm for consistency of PSPs

Recall that for applications like the neutron branching process it is crucial to know exactly whether $\mu_f = \bar{1}$ holds. We say a PSP f is *consistent* if $\mu_f = \bar{1}$; otherwise it is *inconsistent*. Similarly, we call a component i consistent if $(\mu_f)_i = 1$. We present a new algorithm for the consistency problem, i.e., the problem to check a PSP for consistency.

It was proved in [7] that consistency is checkable in polynomial time by reduction to Linear Programming (LP). We first observe that consistency of general PSPs can be reduced to consistency of scPSPs by computing the DAG of SCCs, and checking consistency SCC-wise [7]: Take any bottom SCC S , and check the consistency of f_S . (Notice that f_S is either constant or an scPSP; if constant, f_S is consistent iff $f_S = 1$, if an scPSP, we can check its consistency by assumption.) If f_S is inconsistent, then so is f , and we are done. If f_S is consistent, then we remove every f_i from f such that $x_i \in S$, replace all variables of S in the remaining polynomials by the constant 1, and iterate (choose a new bottom SCC, etc.). Note that this algorithm processes each polynomial at most once, as every variable belongs to exactly one SCC.

It remains to reduce the consistency problem for scPSPs to LP. The first step is:

Proposition 3.1. [9, 7] *An scPSP f is consistent iff $\rho(f'(\bar{1})) \leq 1$ (i.e., iff the spectral radius of the Jacobi matrix f' evaluated at the vector $\bar{1}$ is at most 1).*

The second step consists of observing that the matrix $f'(\bar{1})$ of an scPSP f is irreducible and non-negative. It is shown in [7] that $\rho(A) \leq 1$ holds for an irreducible and nonnegative matrix A iff the system of inequalities

$$A\mathbf{x} \geq \mathbf{x} + \bar{1}, \mathbf{x} \geq \bar{0} \tag{3.1}$$

is infeasible. However, no strongly polynomial algorithm for LP is known, and we are not aware that (3.1) falls within any subclass solvable in strongly polynomial time [8].

We provide a very simple, strongly polynomial time algorithm to check whether $\rho(f'(\bar{1})) \leq 1$ holds. We need some results from Perron-Frobenius theory (see e.g. [3]).

Lemma 3.2. *Let $A \in \mathbb{R}^{n \times n}$ be nonnegative and irreducible.*

- (1) $\rho(A)$ is a simple eigenvalue of A .
- (2) There exists an eigenvector $\mathbf{v} \succ \bar{0}$ with $\rho(A)$ as eigenvalue.

- (3) Every eigenvector $\mathbf{v} \succ \bar{\mathbf{0}}$ has $\rho(A)$ as eigenvalue.
 (4) For all $\alpha, \beta \in \mathbb{R} \setminus \{0\}$ and $\mathbf{v} > \bar{\mathbf{0}}$: if $\alpha\mathbf{v} < A\mathbf{v} < \beta\mathbf{v}$, then $\alpha < \rho(A) < \beta$.

The following lemma is the key to the algorithm:

Lemma 3.3. *Let $A \in \mathbb{R}^{n \times n}$ be nonnegative and irreducible.*

- (a) *Assume there is $\mathbf{v} \in \mathbb{R}^n \setminus \{\bar{\mathbf{0}}\}$ such that $(Id - A)\mathbf{v} = \bar{\mathbf{0}}$. Then $\rho(A) \leq 1$ iff $\mathbf{v} \succ \bar{\mathbf{0}}$ or $\mathbf{v} \prec \bar{\mathbf{0}}$.*
 (b) *Assume $\mathbf{v} = \bar{\mathbf{0}}$ is the only solution of $(Id - A)\mathbf{v} = \bar{\mathbf{0}}$. Then there exists a unique $\mathbf{x} \in \mathbb{R}^n$ such that $(Id - A)\mathbf{x} = \bar{\mathbf{1}}$, and $\rho(A) \leq 1$ iff $\mathbf{x} \geq \bar{\mathbf{1}}$ and $A\mathbf{x} < \mathbf{x}$.*

Proof.

- (a) From $(Id - A)\mathbf{v} = \bar{\mathbf{0}}$ it follows $A\mathbf{v} = \mathbf{v}$. We see that \mathbf{v} is an eigenvector of A with eigenvalue 1. So $\rho(A) \geq 1$.
 (\Leftarrow): As both \mathbf{v} and $-\mathbf{v}$ are eigenvectors of A with eigenvalue 1, we can assume w.l.o.g. that $\mathbf{v} \succ \bar{\mathbf{0}}$. By Lemma 3.2(3), $\rho(A)$ is the eigenvalue of \mathbf{v} , and so $\rho(A) = 1$.
 (\Rightarrow): Since $\rho(A) \leq 1$ and $\rho(A) \geq 1$, it follows that $\rho(A) = 1$. By Lemma 3.2(1) and (2), the eigenspace of the eigenvalue 1 is one-dimensional and contains a vector $\mathbf{x} \succ \bar{\mathbf{0}}$. So $\mathbf{v} = \alpha \cdot \mathbf{x}$ for some $\alpha \in \mathbb{R}$, $\alpha \neq 0$. If $\alpha > 0$, we have $\mathbf{v} \succ \bar{\mathbf{0}}$, otherwise $\mathbf{v} \prec \bar{\mathbf{0}}$.
 (b) With the assumption and basic facts from linear algebra it follows that $(Id - A)$ has full rank and therefore $(Id - A)\mathbf{x} = \bar{\mathbf{1}}$ has a unique solution \mathbf{x} . We still have to prove the second part of the conjunction:
 (\Leftarrow): Follows directly from Lemma 3.2(4).
 (\Rightarrow): Let $\rho(A) \leq 1$. Assume for a contradiction that $\rho(A) = 1$. Then, by Lemma 3.2(1), the matrix A would have an eigenvector $\mathbf{v} \neq \bar{\mathbf{0}}$ with eigenvalue 1, so $(Id - A)\mathbf{v} = \bar{\mathbf{0}}$, contradicting the assumption. So we have, in fact, $\rho(A) < 1$. By standard matrix facts (see e.g. [3]), this implies that $(Id - A)^{-1} = A^* = \sum_{i=0}^{\infty} A^i$ exists, and so we have $\mathbf{x} = (Id - A)^{-1}\bar{\mathbf{1}} = A^*\bar{\mathbf{1}} \geq \bar{\mathbf{1}}$. Furthermore, $A\mathbf{x} = \sum_{i=1}^{\infty} A^i\bar{\mathbf{1}} < \sum_{i=0}^{\infty} A^i\bar{\mathbf{1}} = \mathbf{x}$. ■

In order to check whether $\rho(A) \leq 1$, we first solve the system $(Id - A)\mathbf{v} = \bar{\mathbf{0}}$ using Gaussian elimination. If we find a vector $\mathbf{v} \neq \bar{\mathbf{0}}$ such that $(Id - A)\mathbf{v} = \bar{\mathbf{0}}$, we apply Lemma 3.3(a). If $\mathbf{v} = \bar{\mathbf{0}}$ is the only solution of $(Id - A)\mathbf{v} = \bar{\mathbf{0}}$, we solve $(Id - A)\mathbf{v} = \bar{\mathbf{1}}$ using Gaussian elimination again, and apply Lemma 3.3(b). Since Gaussian elimination of a rational n -dimensional linear equation system can be carried out in strongly polynomial time using $O(n^3)$ arithmetic operations (see e.g. [8]), we obtain:

Proposition 3.4. *Given a nonnegative irreducible matrix $A \in \mathbb{R}^{n \times n}$, one can decide in strongly polynomial time, using $O(n^3)$ arithmetic operations, whether $\rho(A) \leq 1$.*

Combining Propositions 3.1 and 3.4 directly yields an algorithm for checking the consistency of scPSPs. Extending it to multiple SCCs as above, we get:

Theorem 3.5. *Let $f(X_1, \dots, X_n)$ be a PSP. There is a strongly polynomial time algorithm that uses $O(n^3)$ arithmetic operations and determines the consistency of f .*

3.1. Case study: A family of “almost consistent” PSPs

In this section, we illustrate some issues faced by algorithms that solve the consistency problem. Consider the following family $h^{(n)}$ of scPSPs, $n \geq 2$:

$$h^{(n)} = \left(0.5X_1^2 + 0.1X_n^2 + 0.4, 0.01X_1^2 + 0.5X_2 + 0.49, \dots, 0.01X_{n-1}^2 + 0.5X_n + 0.49 \right)^\top.$$

	$n = 25$	$n = 100$	$n = 200$	$n = 400$	$n = 600$	$n = 1000$
Exact LP	< 1 sec	2 sec	8 sec	67 sec	208 sec	> 2h
Our algorithm	< 1 sec	< 1 sec	1 sec	4 sec	10 sec	29 sec

Table 1: Consistency checks for $h^{(n)}$ -systems: Runtimes of different approaches.

It is not hard to show that $h^{(n)}(\mathbf{p}) \prec \mathbf{p}$ holds for $\mathbf{p} = (1 - 0.02^n, \dots, 1 - 0.02^{2n-1})^\top$, so we have $\mu_{h^{(n)}} \prec \bar{1}$ by Proposition 4.4, i.e., the $h^{(n)}$ are inconsistent.

The tool PReMo [12] relies on Java’s floating-point arithmetic to compute approximations of the least fixed point of a PSP. We invoked PReMo for computing approximants of $\mu_{h^{(n)}}$ for different values of n between 5 and 100. Due to its fixed precision, PReMo’s approximations for $\mu_{h^{(n)}}$ are ≥ 1 in all components if $n \geq 7$. This might lead to the wrong conclusion that $h^{(n)}$ is consistent.

Recall that the consistency problem can be solved by checking the feasibility of the system (3.1) with $A = f'(\bar{1})$. We checked it with `lp_solve`, a well-known LP tool using hardware floating-point arithmetic. The tool wrongly states that (3.1) has no solution for $h^{(n)}$ -systems with $n > 10$. This is due to the fact that the solutions cannot be represented adequately using machine number precision.² Finally, we also checked feasibility with Maple’s Simplex package, which uses exact arithmetic, and compared its performance with the implementation, also in Maple, of our consistency algorithm. Table 1 shows the results. Our algorithm clearly outperforms the LP approach. For more experiments see Section 4.3.

4. Approximating μ_f with inexact arithmetic

It is shown in [7] that μ_f may not be representable by roots, so one can only approximate μ_f . In this section we present an algorithm that computes two sequences, $(\mathbf{lb}^{(i)})_i$ and $(\mathbf{ub}^{(i)})_i$, such that $\mathbf{lb}^{(i)} \leq \mu_f \leq \mathbf{ub}^{(i)}$ and $\lim_{i \rightarrow \infty} \mathbf{ub}^{(i)} - \mathbf{lb}^{(i)} = \bar{0}$. In words: $\mathbf{lb}^{(i)}$ and $\mathbf{ub}^{(i)}$ are lower and upper bounds on μ_f , respectively, and the sequences converge to μ_f . Moreover, they converge linearly, meaning that the *number of accurate bits* of $\mathbf{lb}^{(i)}$ and $\mathbf{ub}^{(i)}$ are linear functions of i . (The number of accurate bits of a vector \mathbf{x} is defined as the greatest number k such that $|(\mu_f - \mathbf{x})_j| / |(\mu_f)_j| \leq 2^{-k}$ holds for all $j \in \{1, \dots, n\}$.) These properties are guaranteed even though our algorithm uses inexact arithmetic: Our algorithm detects numerical problems due to rounding errors, recovers from them, and increases the precision of the arithmetic as needed. Increasing the precision dynamically is, e.g., supported by the GMP library [1].

Let us make precise what we mean by increasing the precision. Consider an elementary operation g , like multiplication, subtraction, etc., that operates on two input numbers x and y . We can *compute $g(x, y)$ with increasing precision* if there is a procedure that on input x, y outputs a sequence $g^{(1)}(x, y), g^{(2)}(x, y), \dots$ that converges to $g(x, y)$. Note that there are no requirements on the convergence speed of this procedure — in particular, we do not require that there is an i with $g^{(i)}(x, y) = g(x, y)$. This procedure, which we assume exists, allows to implement *floating assignments* of the form

$$z \leftarrow g(x, y) \text{ such that } \phi(z)$$

with the following semantics: z is assigned the value $g^{(i)}(x, y)$, where $i \geq 1$ is the smallest index such that $\phi(g^{(i)}(x, y))$ holds. We say that the assignment is *valid* if $\phi(g(x, y))$ holds and ϕ involves

²The mentioned problems of PReMo and `lp_solve` are not due to the fact that the coefficients of $h^{(n)}$ cannot be properly represented using basis 2: The problems persist if one replaces the coefficients of $h^{(n)}$ by similar numbers exactly representable by machine numbers.

only continuous functions and strict inequalities. Our assumption on the arithmetic guarantees that (the computation underlying) a valid floating assignment terminates. As “syntactic sugar”, more complex operations (e.g., linear equation solving) are also allowed in floating assignments, because they can be decomposed into elementary operations.

We feel that any implementation of arbitrary precision arithmetic should satisfy our requirement that the computed values converge to the exact result. For instance, the documentation of the GMP library [1] states: “Each function is defined to calculate with ‘infinite precision’ followed by a truncation to the destination precision, but of course the work done is only what’s needed to determine a result under that definition.”

To approximate the least fixed point of a PSP, we first transform it into a certain normal form. A purely superlinear PSP f is called *perfectly superlinear* if every variable depends directly on itself and every superlinear SCC is purely superlinear. The following proposition states that any PSP f can be made perfectly superlinear.

Proposition 4.1. *Let f be a PSP of size s . We can compute in time $O(n \cdot s)$ a perfectly superlinear PSP \tilde{f} with $\text{Var}(\tilde{f}) = \text{Var}(f) \cup \{\tilde{X}\}$ of size $O(n \cdot s)$ such that $\mu_f = (\mu_{\tilde{f}})_{\text{Var}(f)}$.*

4.1. The algorithm

The algorithm receives as input a perfectly superlinear PSP f and an error bound $\epsilon > 0$, and returns vectors \mathbf{lb}, \mathbf{ub} such that $\mathbf{lb} \leq \mu_f \leq \mathbf{ub}$ and $\mathbf{ub} - \mathbf{lb} \leq \bar{\epsilon}$. A first initialization step requires to compute a vector \mathbf{x} with $\bar{0} \prec \mathbf{x} \prec f(\mathbf{x})$, i.e., a “strict” pre-fixed point. This is done in Section 4.1.1. The algorithm itself is described in Section 4.1.2.

4.1.1. *Computing a strict pre-fixed point.* Algorithm 1 computes a strict pre-fixed point:

Algorithm 1: Procedure computeStrictPrefix

Input: perfectly superlinear PSP f
Output: \mathbf{x} with $\bar{0} \prec \mathbf{x} \prec f(\mathbf{x}) \prec \bar{1}$
 $\mathbf{x} \leftarrow \bar{0}$;
while $\bar{0} \not\prec \mathbf{x}$ **do**
 $Z \leftarrow \{i \mid 1 \leq i \leq n, f_i(\mathbf{x}) = 0\}$;
 $P \leftarrow \{i \mid 1 \leq i \leq n, f_i(\mathbf{x}) > 0\}$;
 $\mathbf{y}_Z \leftarrow \bar{0}$;
 $\mathbf{y}_P \leftarrow f_P(\mathbf{x})$ **such that** $\bar{0} \prec \mathbf{y}_P \prec f_P(\mathbf{y}) \prec \bar{1}$;
 $\mathbf{x} \leftarrow \mathbf{y}$;

Proposition 4.2. *Algorithm 1 is correct and terminates after at most n iterations.*

The reader may wonder why Algorithm 1 uses a floating assignment $\mathbf{y}_P \leftarrow f_P(\mathbf{x})$, given that it must also perform exact comparisons to obtain the sets Z and P and to decide exactly whether $\mathbf{y}_P \prec f_P(\mathbf{y})$ holds in the **such that** clause of the floating assignment. The reason is that, while we perform such operations exactly, we do not want to use the *result* of exact computations as input for other computations, as this easily leads to an explosion in the required precision. For instance, the size of the exact result of $f_P(\mathbf{y})$ may be larger than the size of \mathbf{y} , while an approximation of smaller size may already satisfy the **such that** clause. In order to emphasize this, we *never* store the result of an exact numerical computation in a variable.

4.1.2. *Computing lower and upper bounds.* Algorithm 1 uses Kleene iteration $\bar{0}, f(\bar{0}), f(f(\bar{0})), \dots$ to compute a strict pre-fixed point. One could, in principle, use the same scheme to compute lower bounds of μ_f , as this sequence converges to μ_f from below by Kleene's theorem. However, convergence of Kleene iteration is generally slow. It is shown in [7] that for the 1-dimensional PSP f with $f(X) = 0.5X^2 + 0.5$ we have $\mu_f = 1$, and the i -th Kleene approximant $\kappa^{(i)}$ satisfies $\kappa^{(i)} \leq 1 - \frac{1}{2^i}$. Hence, Kleene iteration may converge only logarithmically, i.e., the number of accurate bits is a logarithmic function of the number of iterations.

In [7] it was suggested to use Newton's method for faster convergence. In order to see how Newton's method can be used, observe that instead of computing μ_f , one can equivalently compute the least nonnegative zero of $f(\bar{X}) - \bar{X}$. Given an approximant \mathbf{x} of μ_f , Newton's method first computes $g^{(\mathbf{x})}(\bar{X})$, the first-order linearization of f at the point \mathbf{x} :

$$g^{(\mathbf{x})}(\bar{X}) = f(\mathbf{x}) + f'(\mathbf{x})(\bar{X} - \mathbf{x})$$

The next Newton approximant \mathbf{y} is obtained by solving $\bar{X} = g^{(\mathbf{x})}(\bar{X})$, i.e.,

$$\mathbf{y} = \mathbf{x} + (Id - f'(\mathbf{x}))^{-1}(f(\mathbf{x}) - \mathbf{x}).$$

We write $\mathcal{N}_f(\mathbf{x}) := \mathbf{x} + (Id - f'(\mathbf{x}))^{-1}(f(\mathbf{x}) - \mathbf{x})$, and usually drop the subscript of \mathcal{N}_f . If $\nu^{(0)} \leq \mu_f$ is any pre-fixed point of f , for instance $\nu^{(0)} = \bar{0}$, we can define a *Newton sequence* $(\nu^{(i)})_i$ by setting $\nu^{(i+1)} = \mathcal{N}(\nu^{(i)})$ for $i \geq 0$. It has been shown in [7, 10, 5] that Newton sequences converge at least linearly to μ_f . Moreover, we have $\bar{0} \leq \nu^{(i)} \leq f(\nu^{(i)}) \leq \mu_f$ for all i .

These facts were shown only for Newton sequences that are computed exactly, i.e., without rounding errors. Unfortunately, Newton approximants are hard to compute exactly: Since each iteration requires to solve a linear equation system whose coefficients depend on the results of the previous iteration, the size of the Newton approximants easily explodes. Therefore, we wish to use inexact arithmetic, but without losing the good properties of Newton's method (reliable lower bounds, linear convergence).

Algorithm 2 accomplishes these goals, and additionally computes post-fixed points \mathbf{ub} of f , which are upper bounds on μ_f . Let us describe the algorithm in some detail. The lower bounds are stored in the variable \mathbf{lb} . The first value of \mathbf{lb} is not simply $\bar{0}$, but is computed by `computeStrictPrefix(f)`, in order to guarantee the validity of the following floating assignments. We use Newton's method for improving the lower bounds because it converges fast (at least linearly) when performed exactly. In each iteration of the algorithm, *two* Newton steps are performed using inexact arithmetic. The intention is that two inexact Newton steps should improve the lower bound at least as much as one exact Newton step. While this may sound like a vague hope for small rounding errors, it can be rigorously proved thanks to the **such that** clause of the floating assignment in line 4. The proof involves two steps. The first step is to prove that $\mathcal{N}(\mathcal{N}(\mathbf{lb}))$ is a (strict) post-fixed point of the function $g(\bar{X}) = f(\mathbf{lb}) + f'(\mathbf{lb})(\bar{X} - \mathbf{lb})$, i.e., $\mathcal{N}(\mathcal{N}(\mathbf{lb}))$ satisfies the first inequality in the **such that** clause. For the second step, recall that $\mathcal{N}(\mathbf{lb})$ is the least fixed point of g . By Knaster-Tarski's theorem, $\mathcal{N}(\mathbf{lb})$ is actually the least post-fixed point of g . So, our value \mathbf{x} , the inexact version of $\mathcal{N}(\mathcal{N}(\mathbf{lb}))$, satisfies $\mathbf{x} \geq \mathcal{N}(\mathbf{lb})$, and hence two inexact Newton steps are in fact at least as "fast" as one exact Newton step. Thus, the \mathbf{lb} converge linearly to μ_f .

The upper bounds \mathbf{ub} are post-fixed points, i.e., $f(\mathbf{ub}) \leq \mathbf{ub}$ is an invariant of the algorithm. The algorithm computes the sets Z and P so that inexact arithmetic is only applied to the components i with $f_i(\mathbf{ub}) < 1$. In the P -components, the function f is applied to \mathbf{ub} in order to improve the upper bound. In fact, f is applied twice in line 9, analogously to applying \mathcal{N} twice in line 4. Here, the **such that** clause makes sure that the progress towards μ_f is at least as fast as the progress of one exact application of f would be. One can show that this leads to linear convergence to μ_f .

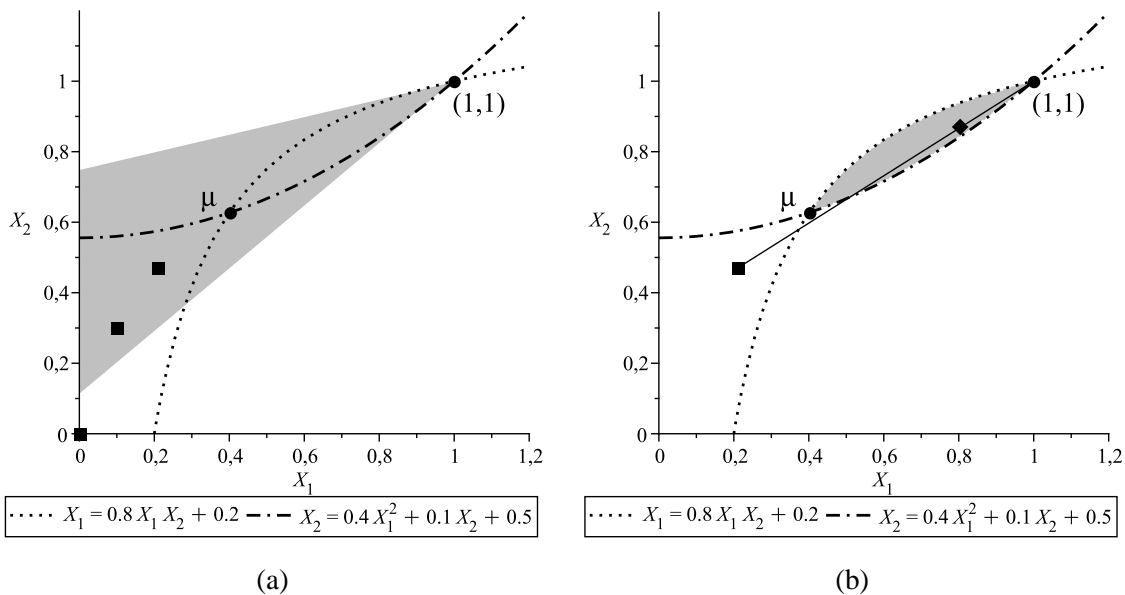
Algorithm 2: Procedure `calcBounds`

Input: perfectly superlinear PSP f , error bound $\epsilon > 0$
Output: vectors \mathbf{lb}, \mathbf{ub} such that $\mathbf{lb} \leq \mu_f \leq \mathbf{ub}$ and $\mathbf{ub} - \mathbf{lb} \leq \bar{\epsilon}$

```

1  $\mathbf{lb} \leftarrow \text{computeStrictPrefix}(f)$ ;
2  $\mathbf{ub} \leftarrow \bar{\mathbf{1}}$ ;
3 while  $\mathbf{ub} - \mathbf{lb} \not\leq \bar{\epsilon}$  do
4    $\mathbf{x} \leftarrow \mathcal{N}(\mathcal{N}(\mathbf{lb}))$  such that  $f(\mathbf{lb}) + f'(\mathbf{lb})(\mathbf{x} - \mathbf{lb}) \prec \mathbf{x} \prec f(\mathbf{x}) \prec \bar{\mathbf{1}}$ ;
5    $\mathbf{lb} \leftarrow \mathbf{x}$ ;
6    $Z \leftarrow \{i \mid 1 \leq i \leq n, f_i(\mathbf{ub}) = 1\}$ ;
7    $P \leftarrow \{i \mid 1 \leq i \leq n, f_i(\mathbf{ub}) < 1\}$ ;
8    $\mathbf{y}_Z \leftarrow \bar{\mathbf{1}}$ ;
9    $\mathbf{y}_P \leftarrow f_P(f(\mathbf{ub}))$  such that  $f_P(\mathbf{y}) \prec \mathbf{y}_P \prec f_P(\mathbf{ub})$ ;
10  forall superlinear SCCs  $S$  of  $f$  with  $\mathbf{y}_S = \bar{\mathbf{1}}$  do
11     $\mathbf{t} \leftarrow \bar{\mathbf{1}} - \mathbf{lb}_S$ ;
12    if  $f'_{SS}(\bar{\mathbf{1}})\mathbf{t} \succ \mathbf{t}$  then
13       $\mathbf{y}_S \leftarrow \bar{\mathbf{1}} - \min \left\{ 1, \frac{\min_{i \in S}(f'_{SS}(\bar{\mathbf{1}})\mathbf{t} - \mathbf{t})_i}{2 \cdot \max_{i \in S}(f_S(\bar{\mathbf{2}}))_i} \right\} \cdot \mathbf{t}$  such that  $f_S(\mathbf{y}) \prec \mathbf{y}_S \prec \bar{\mathbf{1}}$ ;
14   $\mathbf{ub} \leftarrow \mathbf{y}$ ;
```

The rest of the algorithm (lines 10-13) deals with the problem that, given a post-fixed \mathbf{ub} , the sequence $\mathbf{ub}, f(\mathbf{ub}), f(f(\mathbf{ub})), \dots$ does not necessarily converge to μ_f . For instance, if $f(X) = 0.75X^2 + 0.25$, then $\mu_f = 1/3$, but $1 = f(1) = f(f(1)) = \dots$. Therefore, the if-statement of Algorithm 2 allows to improve the upper bound from $\bar{\mathbf{1}}$ to a post-fixed point less than $\bar{\mathbf{1}}$, by exploiting the lower bounds \mathbf{lb} . This is illustrated in Figure 1 for a 2-dimensional scPSP f . The


 Figure 1: Computation of a post-fixed point less than $\bar{\mathbf{1}}$.

dotted lines indicate the curve of the points (X_1, X_2) satisfying $X_1 = 0.8X_1X_2 + 0.2$ and $X_2 = 0.4X_1^2 + 0.1X_2 + 0.5$. Notice that $\mu_f \prec \bar{1} = f(\bar{1})$. In Figure 1 (a) the shaded area consists of those points \mathbf{lb} where $f'(\bar{1})(\bar{1} - \mathbf{lb}) \succ \bar{1} - \mathbf{lb}$ holds, i.e., the condition of line 12. One can show that μ_f must lie in the shaded area, so by continuity, any sequence converging to μ_f , in particular the sequence of lower bounds \mathbf{lb} , finally reaches the shaded area. In Figure 1 (a) this is indicated by the points with the square shape. Figure 1 (b) shows how to exploit such a point \mathbf{lb} to compute a post-fixed point $\mathbf{ub} \prec \bar{1}$ (post-fixed points are shaded in Figure 1 (b)): The post-fixed point \mathbf{ub} (diamond shape) is obtained by starting at $\bar{1}$ and moving a little bit along the straight line between $\bar{1}$ and \mathbf{lb} , cf. line 13. The sequence $\mathbf{ub}, f(\mathbf{ub}), f(f(\mathbf{ub})), \dots$ now converges linearly to μ_f .

Theorem 4.3. *Algorithm 2 terminates and computes vectors \mathbf{lb}, \mathbf{ub} such that $\mathbf{lb} \leq \mu_f \leq \mathbf{ub}$ and $\mathbf{ub} - \mathbf{lb} \leq \bar{\epsilon}$. Moreover, the sequences of lower and upper bounds computed by the algorithm both converge linearly to μ_f .*

Notice that Theorem 4.3 is about the convergence speed of the approximants, not about the time needed to compute them. To analyse the computation time, one would need stronger requirements on how floating assignments are performed.

The lower and upper bounds computed by Algorithm 2 have a special feature: they satisfy $\mathbf{lb} \prec f(\mathbf{lb})$ and $\mathbf{ub} \geq f(\mathbf{ub})$. The following proposition guarantees that such points are in fact lower and upper bounds.

Proposition 4.4. *Let f be a perfectly superlinear PSP. Let $\bar{0} \leq \mathbf{x} \leq \bar{1}$. If $\mathbf{x} \prec f(\mathbf{x})$, then $\mathbf{x} \prec \mu_f$. If $\mathbf{x} \geq f(\mathbf{x})$, then $\mathbf{x} \geq \mu_f$.*

So a user of Algorithm 2 can immediately verify that the computed bounds are correct. To summarize, Algorithm 2 computes provably and even verifiably correct lower and upper bounds, although exact computation is restricted to detecting numerical problems. See Section 4.3 for experiments.

4.2. Proving consistency using the inexact algorithm

In Section 3 we presented a simple and efficient algorithm to check the consistency of a PSP. Algorithm 2 is aimed at approximating μ_f , but note that it can also prove the inconsistency of a PSP: when the algorithm sets $\mathbf{ub}_i < 1$, we know $(\mu_f)_i < 1$. This raises the question whether Algorithm 2 can also be used for proving consistency. The answer is yes, and the procedure is based on the following proposition.

Proposition 4.5. *Let f be an scPSP. Let $\mathbf{t} \succ \bar{0}$ be a vector with $f'(\bar{1})\mathbf{t} \leq \mathbf{t}$. Then f is consistent.*

Proposition 4.5 can be used to identify consistent components.

Use Algorithm 2 with some (small) ϵ to compute \mathbf{ub} and \mathbf{lb} . Take any bottom SCC S .

- If $f'(\bar{1})(\bar{1} - \mathbf{lb}_S) \leq \bar{1} - \mathbf{lb}_S$, mark all variables in S as consistent and remove the S -components from f . In the remaining components, replace all variables in S with 1.
- Otherwise, remove S and all other variables that depend on S from f .

Repeat with the new bottom SCC until all SCCs are processed.

There is no guarantee that this method detects all i with $(\mu_f)_i = 1$.

D n	2			3			6			10		
	20	50	100	20	50	100	20	50	100	20	50	100
inconsistent (yes/no)	n	n	n	y	y	y	y	y	y	y	y	y
Cons. check (Alg. Sec. 3)	< 1	< 1	2	< 1	< 1	2	< 1	< 1	2	< 1	< 1	2
Cons. check (exact LP)	< 1	20	258	< 1	22	124	< 1	16	168	< 1	37	222
Approx. Q_D ($\epsilon = 10^{-3}$)	< 1	< 1	4	2	8	32	1	5	21	1	4	17
Approx. Q_D ($\epsilon = 10^{-4}$)	< 1	< 1	4	2	8	34	2	7	28	1	6	23

Table 2: Runtime in seconds of various algorithms on different values of D and n .

4.3. Case study: A neutron branching process

One of the main applications of the theory of branching processes is the modelling of cascade creation of particles in physics. We study a problem described by Harris in [9]. Consider a ball of fissionable radioactive material of radius D . Spontaneous fission of an atom can liberate a neutron, whose collision with another atom can produce further neutrons etc. If D is very small, most neutrons leave the ball without colliding. If D is very large, then nearly all neutrons eventually collide, and the probability that the neutron's progeny never dies is large. A well-known result shows that, loosely speaking, the population of a process that does not go extinct grows exponentially over time with large probability. Therefore, the neutron's progeny never dying out actually means that after a (very) short time all the material is fissioned, which amounts to a nuclear explosion. The task is to compute the largest value of D for which the probability of extinction of a neutron born at the centre of the ball is still 1 (if the probability is 1 at the centre, then it is 1 everywhere). This is often called the critical radius. Notice that, since the number of atoms that undergo spontaneous fission is large (some hundreds per second for the critical radius of plutonium), if the probability of extinction lies only slightly below 1, there is already a large probability of a chain reaction. Assume that a neutron born at distance ξ from the centre leaves the ball without colliding with probability $l(\xi)$, and collides with an atom at distance η from the centre with probability density $R(\xi, \eta)$. Let further $f(x) = \sum_{i \geq 0} p_i x^i$, where p_i is the probability that a collision generates i neutrons. For a neutron's progeny to go extinct, the neutron must either leave the ball without colliding, or collide at some distance η from the centre, but in such a way that the progeny of all generated neutrons goes extinct. So the extinction probability $Q_D(\xi)$ of a neutron born at distance ξ from the centre is given by [9], p. 86:

$$Q_D(\xi) = l(\xi) + \int_0^D R(\xi, \eta) f(Q_D(\eta)) d\eta$$

Harris takes $f(x) = 0.025 + 0.830x + 0.07x^2 + 0.05x^3 + 0.025x^4$, and gives expressions for both $l(\xi)$ and $R(\xi, \eta)$. By discretizing the interval $[0, D]$ into n segments and replacing the integral by a finite sum we obtain a PSP of dimension $n + 1$ over the variables $\{Q_D(jD/n) \mid 0 \leq j \leq n\}$. Notice that $Q_D(0)$ is the probability that a neutron born in the centre does not cause an explosion.

Results. For our experiments we used three different discretizations $n = 20, 50, 100$. We applied our consistency algorithm from Section 3 and Maple's Simplex to check inconsistency, i.e., to check whether an explosion occurs. The results are given in the first 3 rows of Table 2: Again our algorithm dominates the LP approach, although the polynomials are much denser than in the $h^{(n)}$ -systems.

We also implemented Algorithm 2 using Maple for computing lower and upper bounds on $Q_D(0)$ with two different values of the error bound ϵ . The runtime is given in the last two rows. By setting the *Digits* variable in Maple we controlled the precision of Maple's software floating-point numbers for the floating assignments. In all cases starting with the standard value

of 10, Algorithm 2 increased *Digits* at most twice by 5, resulting in a maximal *Digits* value of 20. We mention that Algorithm 2 computed an upper bound $\prec \bar{1}$, and thus proved inconsistency, after the first few iterations in all investigated cases, almost as fast as the algorithm from Section 3.

Computing approximations for the critical radius. After computing $Q_D(0)$ for various values of D one can suspect that the critical radius, i.e., the smallest value of D for which $Q_D(0) = 1$, lies somewhere between 2.7 and 3. We combined binary search with the consistency algorithm from Section 3 to determine the critical radius up to an error of 0.01. During the binary search, the algorithm from Section 3 has to analyze PSPs that come closer and closer to the verge of (in)consistency. For the last (and most expensive) binary search step that decreases the interval to 0.01, our algorithm took $< 1, 1, 3, 8$ seconds for $n = 20, 50, 100, 150$, respectively. For $n = 150$, we found the critical radius to be in the interval $[2.981, 2.991]$. Harris [9] estimates 2.9.

5. Conclusions

We have presented a new, simple, and efficient algorithm for checking the consistency of PSPs, which outperforms the previously existing LP-based method. We have also described the first algorithm that computes reliable lower and upper bounds on μ_f . The sequence of bounds converges linearly to μ_f . To achieve these properties without sacrificing efficiency, we use a novel combination of exact and inexact (floating-point) arithmetic. Experiments on PSPs from concrete branching processes confirm the practicality of our approach. The results raise the question whether our combination of exact and inexact arithmetic could be transferred to other computational problems.

Acknowledgments. We thank several anonymous referees for pointing out inaccuracies and helping us clarify certain aspects of the paper. The second author was supported by the DFG Graduiertenkolleg 1480 (PUMA). We also thank Andreas Reuss for proofreading the manuscript.

References

- [1] GMP library. <http://gmplib.org>.
- [2] K. B. Athreya and P. E. Ney. *Branching Processes*. Springer, 1972.
- [3] A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. SIAM, 1994.
- [4] J. Esparza, A. Gaiser, and S. Kiefer. Computing least fixed points of probabilistic systems of polynomials. Technical report, Technische Universität München, Institut für Informatik, 2009.
- [5] J. Esparza, S. Kiefer, and M. Luttenberger. Convergence thresholds of Newton’s method for monotone polynomial equations. In *Proceedings of STACS*, pages 289–300, 2008.
- [6] J. Esparza, A. Kučera, and R. Mayr. Model checking probabilistic pushdown automata. In *LICS 2004*, pages 12–21. IEEE Computer Society, 2004.
- [7] K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. *Journal of the ACM*, 56(1):1–66, 2009.
- [8] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer, 1993.
- [9] T. E. Harris. *The theory of branching processes*. Springer, Berlin, 1963.
- [10] S. Kiefer, M. Luttenberger, and J. Esparza. On the convergence of Newton’s method for monotone systems of polynomial equations. In *Proceedings of STOC*, pages 217–226. ACM, 2007.
- [11] C. D. Manning and H. Schuetze. *Foundations of Statistical Natural Language Processing*. MIT Press, June 1999.
- [12] D. Wojtczak and K. Etessami. PReMo: an analyzer for probabilistic recursive models. In *TACAS*, volume 4424 of *Lecture Notes in Computer Science*, pages 66–71. Springer, 2007.