

Approximative Methods for Monotone Systems of min-max-Polynomial Equations^{*}

Javier Esparza, Thomas Gawlitza, Stefan Kiefer, and Helmut Seidl

Institut für Informatik
Technische Universität München, Germany
{esparza, gawlitza, kiefer, seidl}@in.tum.de

Abstract. A *monotone system of min-max-polynomial equations* (min-max-MSPE) over the variables X_1, \dots, X_n has for every i exactly one equation of the form $X_i = f_i(X_1, \dots, X_n)$ where each $f_i(X_1, \dots, X_n)$ is an expression built up from polynomials with non-negative coefficients, minimum- and maximum-operators. The question of computing least solutions of min-max-MSPEs arises naturally in the analysis of *recursive stochastic games* [4, 5, 12]. Min-max-MSPEs generalize MSPEs for which convergence speed results of Newton’s method are established in [9, 2]. We present the first methods for approximatively computing least solutions of min-max-MSPEs which converge at least linearly. Whereas the first one converges faster, a single step of the second method is cheaper. Furthermore, we compute ϵ -optimal positional strategies for the player who wants to maximize the outcome in a recursive stochastic game.

1 Introduction

In this paper we study *monotone systems of min-max polynomial equations* (min-max-MSPEs). A min-max-MSPE over the variables X_1, \dots, X_n contains for every $1 \leq i \leq n$ exactly one equation of the form $X_i = f_i(X_1, \dots, X_n)$ where every $f_i(X_1, \dots, X_n)$ is an expression built up from polynomials with non-negative coefficients, minimum- and maximum-operators. An example of such an equation is $X_1 = 3X_1X_2 + 5X_1^2 \wedge 4X_2$ (where \wedge is the minimum-operator). The variables range over non-negative reals. Min-max-MSPEs are called monotone because f_i is a monotone function in all arguments.

Min-max-MSPEs naturally appear in the study of two-player stochastic games and competitive Markov decision processes, in which, broadly speaking, the next move is decided by one of the two players or by tossing a coin, depending on the game’s position (see e.g. [10, 6]). The min and max operators model the competition between the players. The product operator, which leads to non-linear equations, allows to deal with recursive stochastic games [4, 5], a class of games with an infinite number of positions, and having as special case *extinction games*, games in which players influence with their actions the development of a population whose members reproduce and die, and the player’s goals are to extinguish the population or keep it alive (see Section 4).

Min-max-MSPEs generalize several other classes of equation systems. If product is disallowed, we obtain systems of min-max *linear* equations, which appear in classical

^{*} This work was in part supported by the DFG project *Algorithms for Software Model Checking*.

two-person stochastic games with a finite number of game positions. The problem of solving these systems has been thoroughly studied [1, 7, 8]. If both min and max are disallowed, we obtain monotone systems of polynomial equations, which are central to the study of recursive Markov chains and probabilistic pushdown systems, and have been recently studied in [3, 9, 2]. If only one of min or max is disallowed, we obtain a class of systems corresponding to recursive Markov decision processes [4]. All these models have applications in the analysis of probabilistic programs with procedures [12].

In vector form we denote a min-max-MSPE by $\mathbf{X} = \mathbf{f}(\mathbf{X})$ where \mathbf{X} denotes the vector (X_1, \dots, X_n) and \mathbf{f} denotes the vector (f_1, \dots, f_n) . By Kleene’s theorem, if a min-max-MSPE has a solution then it also has a *least* one, denoted by $\mu\mathbf{f}$, which is also the relevant solution for the applications mentioned above. Kleene’s theorem also ensures that the iterative process $\kappa^{(0)} = \mathbf{0}$, $\kappa^{(k+1)} = \mathbf{f}(\kappa^{(k)})$, $k \in \mathbb{N}$, the so-called Kleene sequence, converges to $\mu\mathbf{f}$. However, this procedure can converge very slowly: in the worst case, the number of accurate bits of the approximation grows with the *logarithm* of the number of iterations (cf. [3]). Thus, the goal is to replace the function \mathbf{f} by an operator $G : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that the respective iterative process also converges to $\mu\mathbf{f}$ but faster. In [3, 9, 2] this problem was studied for min-max-MSPEs without the min and max operator. There, G was chosen as one step of the well-known Newton’s method (cf. for instance [11]). This means that, for a given approximate $\mathbf{x}^{(k)}$, the next approximate $\mathbf{x}^{(k+1)} = G(\mathbf{x}^{(k)})$ is determined by the *unique solution* of a linear equation system which is obtained from the first order Taylor approximation of \mathbf{f} at $\mathbf{x}^{(k)}$. It was shown that this choice guarantees *linear convergence*, i.e., the number of accurate bits grows *linearly* in the number of iterations. Notice that when characterizing the convergence behavior the term linear does not refer to the size of \mathbf{f} .

However, this technique no longer works for arbitrary min-max-MSPEs. If we approximate \mathbf{f} at $\mathbf{x}^{(k)}$ through its first order Taylor approximation at $\mathbf{x}^{(k)}$ there is no guarantee that the next approximate still lies below the least solution, and the sequence of approximants may even diverge. For this reason, the PReMo tool [12] uses round-robin iteration for min-max-MSPEs, an optimization of Kleene iteration. Unfortunately, this technique also exhibits “logarithmic” convergence behavior in the worst case.

In this paper we overcome the problem of Newton’s method. Instead of approximating \mathbf{f} (at the current approximate $\mathbf{x}^{(k)}$) by a linear function, both of our methods approximate \mathbf{f} by a *piecewise* linear function. In contrast to the applications of Newton’s method in [3, 9, 2], this approximation may not have a *unique fixpoint*, but it has a *least fixpoint* which we use as the next approximate $\mathbf{x}^{(k+1)} = G(\mathbf{x}^{(k)})$. Our first method uses an approximation of \mathbf{f} at $\mathbf{x}^{(k)}$ whose least fixpoint can be determined using the algorithm for *systems of rational equations* from [8]. The approximation of \mathbf{f} at $\mathbf{x}^{(k)}$ used by our second method allows to use linear programming to compute $\mathbf{x}^{(k+1)}$. Our methods are the first algorithms for approximatively computing $\mu\mathbf{f}$ which converge at least linearly, provided that \mathbf{f} is quadratic, an easily achievable normal form.

The rest of the paper is organized as follows. In Section 2 we introduce basic concepts and state some important facts about min-max-MSPEs. A class of games which can be analyzed using our techniques is presented in Section 4. Our main contribution, the two approximation methods, is presented and analyzed in Sections 5 and 6. In Sec-

tion 7 we study the relation between our two approaches and compare them to previous work. We conclude in Section 8. Missing proofs can be found in a technical report [?].

2 Notations, Basic Concepts and a Fundamental Theorem

As usual, \mathbb{R} and \mathbb{N} denote the set of real and natural numbers. We assume $0 \in \mathbb{N}$. We write $\mathbb{R}_{\geq 0}$ for the set of non-negative real numbers. We use bold letters for vectors, e.g. $\mathbf{x} \in \mathbb{R}^n$. In particular $\mathbf{0}$ denotes the vector $(0, \dots, 0)$. The transpose of a matrix or a vector is indicated by the superscript \top . We assume that the vector $\mathbf{x} \in \mathbb{R}^n$ has the components x_1, \dots, x_n . Similarly, the i -th component of a function $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is denoted by f_i . As in [2], we say that $\mathbf{x} \in \mathbb{R}^n$ has $i \in \mathbb{N}$ valid bits of $\mathbf{y} \in \mathbb{R}^n$ iff $|x_j - y_j| \leq 2^{-i}|y_j|$ for $j = 1, \dots, n$. We identify a linear function from \mathbb{R}^n to \mathbb{R}^m with its representation as a matrix from $\mathbb{R}^{m \times n}$. The identity matrix is denoted by I . The *Jacobian* of a function $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ at $\mathbf{x} \in \mathbb{R}^n$ is the matrix of all first-order partial derivatives of \mathbf{f} at \mathbf{x} , i.e., the $m \times n$ -matrix with the entry $\frac{\partial f_i}{\partial x_j}(\mathbf{x})$ in the i -th row and the j -th column. We denote it by $\mathbf{f}'(\mathbf{x})$.

The partial order \leq on \mathbb{R}^n is defined by setting $\mathbf{x} \leq \mathbf{y}$ iff $x_i \leq y_i$ for all $i = 1, \dots, n$. We write $\mathbf{x} < \mathbf{y}$ iff $\mathbf{x} \leq \mathbf{y}$ and $\mathbf{x} \neq \mathbf{y}$. The operators \wedge and \vee are defined by $x \wedge y := \min\{x, y\}$ and $x \vee y := \max\{x, y\}$ for $x, y \in \mathbb{R}$. These operators are also extended component-wise to \mathbb{R}^n and point-wise to \mathbb{R}^n -valued functions. A function $\mathbf{f} : D \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^m$ is called *monotone* on $M \subseteq D$ iff $\mathbf{f}(\mathbf{x}) \leq \mathbf{f}(\mathbf{y})$ for every $\mathbf{x}, \mathbf{y} \in M$ with $\mathbf{x} \leq \mathbf{y}$. Let $X \subseteq \mathbb{R}^n$ and $\mathbf{f} : X \rightarrow X$. A vector $\mathbf{x} \in X$ is called *fixpoint* of \mathbf{f} iff $\mathbf{x} = \mathbf{f}(\mathbf{x})$. It is the *least fixpoint* of \mathbf{f} iff $\mathbf{y} \geq \mathbf{x}$ for every fixpoint $\mathbf{y} \in X$ of \mathbf{f} . If it exists we denote the least fixpoint of \mathbf{f} by $\mu \mathbf{f}$. We call \mathbf{f} *feasible* iff \mathbf{f} has some fixpoint $\mathbf{x} \in X$.

Let us fix a set $\mathbf{X} = \{X_1, \dots, X_n\}$ of variables. We call a vector $\mathbf{f} = (f_1, \dots, f_m)$ of polynomials f_1, \dots, f_m in the variables X_1, \dots, X_n a *system of polynomials*. \mathbf{f} is called *linear* (resp. *quadratic*) iff the degree of each f_i is at most 1 (resp. 2), i.e., every monomial contains at most one variable (resp. two variables). As usual, we identify \mathbf{f} with its interpretation as a function from \mathbb{R}^n to \mathbb{R}^m . As in [9, 2] we call \mathbf{f} a *monotone system of polynomials* (MSP for short) iff all coefficients are non-negative.

Min-max-MSPs. Given polynomials f_1, \dots, f_k we call $f_1 \wedge \dots \wedge f_k$ a *min-polynomial* and $f_1 \vee \dots \vee f_k$ a *max-polynomial*. A function that is either a min- or a max-polynomial is also called *min-max-polynomial*. We call $\mathbf{f} = (f_1, \dots, f_n)$ a *system of min-polynomials* iff every component f_i is a min-polynomial. The definition of *systems of max-polynomials* and *systems of min-max-polynomials* is analogous. A system of min-max-polynomials is called *linear* (resp. *quadratic*) iff all occurring polynomials are linear (resp. *quadratic*). By introducing auxiliary variables every system of min-max-polynomials can be transformed into a *quadratic* one in time linear in the size of the system (cf. [9]). A system of min-max-polynomials where all coefficients are from $\mathbb{R}_{\geq 0}^n$ is called a *monotone system of min-max-polynomials* (*min-max-MSP*) for short. The terms *min-MSP* and *max-MSP* are defined analogously.

Example 1. $\mathbf{f}(x_1, x_2) = (\frac{1}{2}x_2^2 + \frac{1}{2} \wedge 3, x_1 \vee 2)^\top$ is a quadratic min-max-MSP.

A min-max-MSP $\mathbf{f} = (f_1, \dots, f_n)^\top$ can be considered as a mapping from $\mathbb{R}_{\geq 0}^n$ to $\mathbb{R}_{\geq 0}^n$. The Kleene sequence $(\kappa_{\mathbf{f}}^{(k)})_{k \in \mathbb{N}}$ is defined by $\kappa_{\mathbf{f}}^{(k)} := \mathbf{f}^k(\mathbf{0})$, $k \in \mathbb{N}$. We have:

Lemma 1. *Let $\mathbf{f} : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}^n$ be a min-max-MSP. Then: (1) \mathbf{f} is monotone and continuous on $\mathbb{R}_{\geq 0}^n$; and (2) If \mathbf{f} is feasible (i.e., \mathbf{f} has some fixpoint), then \mathbf{f} has a least fixpoint $\mu\mathbf{f}$ and $\mu\mathbf{f} = \lim_{k \rightarrow \infty} \kappa_{\mathbf{f}}^{(k)}$.*

Strategies. Assume that \mathbf{f} denotes a system of min-max-polynomials. A \vee -strategy σ for \mathbf{f} is a function that maps every max-polynomial $f_i = f_{i,1} \vee \dots \vee f_{i,k_i}$ occurring in \mathbf{f} to one of the $f_{i,j}$'s and every min-polynomial f_i to f_i . We also write f_i^σ for $\sigma(f_i)$. Accordingly, a \wedge -strategy π for \mathbf{f} is a function that maps every min-polynomial $f_i = f_{i,1} \wedge \dots \wedge f_{i,k}$ occurring in \mathbf{f} to one of the $f_{i,j}$'s and every max-polynomial f_i to f_i . We denote the set of \vee -strategies for \mathbf{f} by $\Sigma_{\mathbf{f}}$ and the set of \wedge -strategies for \mathbf{f} by $\Pi_{\mathbf{f}}$. For $s \in \Sigma_{\mathbf{f}} \cup \Pi_{\mathbf{f}}$, we write \mathbf{f}^s for $(f_1^s, \dots, f_n^s)^\top$. We define $\Pi_{\mathbf{f}}^* := \{\pi \in \Pi_{\mathbf{f}} \mid \mathbf{f}^\pi \text{ is feasible}\}$. We drop the subscript whenever it is clear from the context.

Example 2. Consider \mathbf{f} from Example 1. Then $\pi : \frac{1}{2}x_2^2 + \frac{1}{2} \wedge 3 \mapsto 3$, $x_1 \vee 2 \mapsto x_1 \vee 2$ is a \wedge -strategy. The max-MSP \mathbf{f}^π is given by $\mathbf{f}^\pi(x_1, x_2)^\top = (3, x_1 \vee 2)^\top$. \square

We collect some elementary facts concerning strategies.

Lemma 2. *Let \mathbf{f} be a feasible min-max-MSP. Then (1) $\mu\mathbf{f}^\sigma \leq \mu\mathbf{f}$ for every $\sigma \in \Sigma$; (2) $\mu\mathbf{f}^\pi \geq \mu\mathbf{f}$ for every $\pi \in \Pi^*$; (3) $\mu\mathbf{f}^\pi = \mu\mathbf{f}$ for some $\pi \in \Pi^*$.*

In [4] the authors consider a subclass of recursive stochastic games for which they prove that a positional optimal strategy exists for the player who wants to maximize the outcome (Theorem 2). The outcome of such a game is the least fixpoint of some min-max-MSP \mathbf{f} . In our setting, Theorem 2 of [4] implies that there exists a \vee -strategy σ such that $\mu\mathbf{f}^\sigma = \mu\mathbf{f}$ — provided that \mathbf{f} is derived from such a recursive stochastic game. Example 4 shows that this property does not hold for arbitrary min-max-MSPs.

Example 3. Consider \mathbf{f} from Example 1. Let $\sigma_1, \sigma_2 \in \Sigma$ be defined by $\sigma_1(x_1 \vee 2) = x_1$ and $\sigma_2(x_1 \vee 2) = 2$. Then $\mu\mathbf{f}^{\sigma_1} = (1, 1)^\top$, $\mu\mathbf{f}^{\sigma_2} = (\frac{5}{2}, 2)^\top$ and $\mu\mathbf{f} = (3, 3)^\top$. \square

The proof of the following fundamental result is inspired by the proof of Theorem 2 in [4]. Although the result looks very natural it is non-trivial to prove.

Theorem 1. *Let \mathbf{f} be a feasible max-MSP. Then $\mu\mathbf{f}^\sigma = \mu\mathbf{f}$ for some $\sigma \in \Sigma$.*

iiiiiii .mine

3 An Example of Application: Extinction Games

=====

4 A Class of Applications: Extinction Games

~~~~~.r1222 In order to illustrate the interest of min-max-MSPs we consider *extinction games*, which are special stochastic games. Consider a world of  $n$  different species  $s_1, \dots, s_n$ . Each species  $s_i$  is controlled by one of two adversarial players. For each  $s_i$  there is a non-empty set  $A_i$  of actions. An action  $a \in A_i$  replaces a single individual of species  $s_i$  by other individuals specified by the action  $a$ . The actions can be probabilistic. E.g., an action could transform an adult rabbit to zero individuals with probability 0.2, to an adult rabbit with probability 0.3 and to an adult and a baby rabbit with probability 0.5. Another action could transform an adult rabbit to a fat rabbit. The max-player (min-player) wants to maximize (minimize) the probability that some initial population is extinguished. During the game each player continuously chooses an individual of a species  $s_i$  controlled by her/him and applies an action from  $A_i$  to it. Note that actions on different species are never in conflict and the execution order is irrelevant. What is the probability that the population is extinguished if the players follow optimal strategies?

To answer those questions we set up a min-max-MSP  $f$  with one min-max-polynomial for each species, thereby following [?,4]. The variables  $X_i$  represent the probability that a population with only a single individual of species  $s_i$  is extinguished. In the rabbit example we have  $X_{\text{adult}} = 0.2 + 0.3X_{\text{adult}} + 0.5X_{\text{adult}}X_{\text{baby}} \vee X_{\text{fat}}$ , assuming that the adult rabbits are controlled by the max-player. The probability that an initial population with  $p_i$  individuals of species  $s_i$  is extinguished is given by  $\prod_{i=1}^n ((\mu f)_i)^{p_i}$ . The stochastic termination games of [4, 5, 12] can be considered as extinction games. In the following we present another instance.

*The primaries game.* In order to illustrate the interest of min-max-MSPs we consider the following stochastic game. Hillary Clinton has to decide her strategy in the primaries. Her team estimates that undecided voters have not yet decided to vote for her for three possible reasons: they consider her (a) cold and calculating, (b) too much part of Washington's establishment, or (c) they listen to Obama's campaign. So the team decides to model those problems as species in an extinction game. The larger the population of a species, the more influenced is an undecided voter by the problem. The goal of Clinton's team is to maximize the extinction probabilities.

~~~~~.mine So the team decides to model the state of an undecided voter by a triple  $(n_a, n_b, n_c)$  of natural numbers indicating the degree to which the voter is influenced by these three problems (the higher the number, the more influenced he or she is). The goal of Clinton's team is to bring as many voters as possible to the state  $(0, 0, 0)$ .

If $n_a > 1$, Clinton can tackle problem (a) by either showing emotions or concentrating on her program. If she shows emotions, her team estimates that the voter's state changes by $(-1, 0, 0)$ with probability 0.3, but with probability 0.7 the action may backfire and change it by $(+1, 0, 0)$. If she concentrates on her program, the state changes by $(-1, 0, 0)$ with probability 0.1, and by $(-1, 0, +1)$ with probability 0.9. Let X_1 (resp. X_2, X_3) denote the probability that the voters's state changes *eventually* from $(1, 0, 0)$ (resp. $(0, 1, 0), (0, 0, 1)$) to $(0, 0, 0)$, assuming Clinton and Obama follow perfect strategies. Then X_1 satisfies equation (1) below. If $n_b > 1$, Clinton can choose between ===== Clinton's possible actions for problem (a) are *showing emotions* or *concentrating on her program*. If she shows emotions, her team estimates that the individual of problem (a) is removed with probability 0.3, but with probability 0.7 the

action backfires and produces yet another individual of (a). This and the effect of concentrating on her program can be read off from Equation (1) below. For problem (b), Clinton can choose between \dots concentrating on her voting record or her statement “I’ll be ready from day 1”. Her team estimates the effect as given in Equation (2). Problem (c) is controlled by Obama, who has the choice between his “change” message, or attacking Clinton for her position on Iraq, see Equation (3).

$$X_a = 0.3 + 0.7X_a^2 \quad \vee \quad 0.1 + 0.9X_c \quad (1)$$

$$X_b = 0.1 + 0.9X_c \quad \vee \quad 0.4X_b + 0.3X_c + 0.3 \quad (2)$$

$$X_c = 0.5X_b + 0.3X_b^2 + 0.2 \quad \wedge \quad 0.5X_a + 0.4X_aX_b + 0.1X_b \quad (3)$$

(*I’ve added some lines here!*)

\dots .mine

What is the probability, assuming perfect strategies, that a voter in state $(1, 0, 0)$ is eventually moved to $(0, 0, 0)$? It is easy to see that it is given by the first component of the least solution of this min-max-MSP. We are also interested in the strategies by Clinton and Obama that lead to this probability. In the next two sections we show how to efficiently solve these problems.

This is an example of an extinction game. We have three populations of “problem instances” (a state (n_a, n_b, n_c) indicates that the populations have $n_a, n_b,$ and n_c individuals). Individuals can reproduce or die. The goal of one of the players (Clinton in our case) is to get all populations “extinct”, while the other player wants to keep them alive. Another example for an extinction game in appendix ?? . ===== What should Clinton and Obama do? What are the extinction probabilities, assuming perfect strategies? In the next sections we show how to efficiently solve these problems. \dots .r1222

5 The τ -Method

Assume that f denotes a feasible min-max-MSP. In this section we present our first method for computing μf approximatively. We call it τ -method. This method computes, for each approximate $x^{(i)}$, the next approximate $x^{(i+1)}$ as the least fixpoint of a piecewise linear approximation $\mathcal{L}(f, x^{(i)}) \vee x^{(i)}$ (see below) of f at $x^{(i)}$. This approximation is a system of *linear* min-max-polynomials where all coefficients of monomials of degree 1 are non-negative. Here, we call such a system a *monotone linear min-max-system* (*min-max-MLS* for short). Note that a min-max-MLS f is not necessarily a min-max-MSP, since negative coefficients of monomials of degree 0 are allowed, e.g. the min-max-MLS $f(x_1) = x_1 - 1$ is not a min-max-MSP.

In [8] a min-max-MLS f is considered as a system of equations (called *system of rational equations* in [8]) which we denote by $X = f(X)$ in vector form. We identify a min-max-MLS f with its interpretation as a function from $\overline{\mathbb{R}}^n$ to $\overline{\mathbb{R}}^n$ ($\overline{\mathbb{R}}$ denotes the complete lattice $\mathbb{R} \cup \{-\infty, \infty\}$). Since f is monotone on $\overline{\mathbb{R}}^n$, it has a least fixpoint $\mu f \in \overline{\mathbb{R}}^n$ which can be computed using the strategy improvement algorithm from [8].

We now define the min-max-MLS $\mathcal{L}(f, y)$, a piecewise linear approximation of f at y . As a first step, let us consider a monotone polynomial $f : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$. Given some approximate $y \in \mathbb{R}_{\geq 0}^n$, a linear approximation $\mathcal{L}(f, y) : \mathbb{R}^n \rightarrow \mathbb{R}$ of f at y is given by the first order Taylor approximation at y , i.e.,

$$\mathcal{L}(f, y)(x) := f(y) + f'(y)(x - y), \quad x \in \mathbb{R}^n.$$

This is precisely the linear approximation which is used for Newton's method. Now consider a max-polynomial $f = f_1 \vee \dots \vee f_k : \mathbb{R}^n \rightarrow \mathbb{R}$. We define the approximation $\mathcal{L}(f, \mathbf{y}) : \mathbb{R}^n \rightarrow \mathbb{R}$ of f at \mathbf{y} by $\mathcal{L}(f, \mathbf{y}) := \mathcal{L}(f_1, \mathbf{y}) \vee \dots \vee \mathcal{L}(f_k, \mathbf{y})$. We emphasize that in this case, $\mathcal{L}(f, \mathbf{y})$ is in general not a linear function but a linear max-polynomial. Accordingly, for a min-MSP $f = f_1 \wedge \dots \wedge f_k : \mathbb{R}^n \rightarrow \mathbb{R}$, we define $\mathcal{L}(f, \mathbf{y}) := \mathcal{L}(f_1, \mathbf{y}) \wedge \dots \wedge \mathcal{L}(f_k, \mathbf{y})$. In this case $\mathcal{L}(f, \mathbf{y})$ is a linear min-polynomial. Finally, for a min-max-MSP $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, we define the approximation $\mathcal{L}(\mathbf{f}, \mathbf{y}) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ of \mathbf{f} at \mathbf{y} by $\mathcal{L}(\mathbf{f}, \mathbf{y}) := (\mathcal{L}(f_1, \mathbf{y}), \dots, \mathcal{L}(f_n, \mathbf{y}))^\top$ which is a min-max-MLS.

Example 4. Consider the min-max-MSP \mathbf{f} from Example 1. The approximation $\mathcal{L}(\mathbf{f}, (\frac{1}{2}, \frac{1}{2}))$ is given by $\mathcal{L}(\mathbf{f}, (\frac{1}{2}, \frac{1}{2}))(x_1, x_2) = (\frac{1}{2}x_2 + \frac{3}{8} \wedge 3, x_1 \vee 2)$. \square

Using the approximation $\mathcal{L}(\mathbf{f}, \mathbf{x}^{(i)})$ we define the operator $\mathcal{N}_{\mathbf{f}} : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}^n$ which gives us, for an approximate $\mathbf{x}^{(i)}$, the next approximate $\mathbf{x}^{(i+1)}$ by

$$\mathcal{N}_{\mathbf{f}}(\mathbf{x}) := \mu(\mathcal{L}(\mathbf{f}, \mathbf{x}) \vee \mathbf{x}), \quad \mathbf{x} \in \mathbb{R}_{\geq 0}^n.$$

Observe that $\mathcal{L}(\mathbf{f}, \mathbf{x}) \vee \mathbf{x}$ is still a min-max-MLS (at least after introducing auxiliary variables in order to eliminate components which contain \vee - and \wedge -operators).

Example 5. In Example 5 we have: $\mathcal{N}_{\mathbf{f}}(\frac{1}{2}, \frac{1}{2}) = \mu(\mathcal{L}(\mathbf{f}, (\frac{1}{2}, \frac{1}{2})) \vee (\frac{1}{2}, \frac{1}{2})^\top) = (\frac{11}{8}, 2)^\top$.

We collect basic properties of $\mathcal{N}_{\mathbf{f}}$ in the following lemma:

Lemma 3. *Let \mathbf{f} be a feasible min-max-MSP and $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\geq 0}^n$. Then:*

1. $\mathbf{x}, \mathbf{f}(\mathbf{x}) \leq \mathcal{N}_{\mathbf{f}}(\mathbf{x})$;
2. $\mathbf{x} = \mathcal{N}_{\mathbf{f}}(\mathbf{x})$ whenever $\mathbf{x} = \mathbf{f}(\mathbf{x})$;
3. (Monotonicity of $\mathcal{N}_{\mathbf{f}}$) $\mathcal{N}_{\mathbf{f}}(\mathbf{x}) \leq \mathcal{N}_{\mathbf{f}}(\mathbf{y})$ whenever $\mathbf{x} \leq \mathbf{y}$;
4. $\mathcal{N}_{\mathbf{f}}(\mathbf{x}) \leq \mathbf{f}(\mathcal{N}_{\mathbf{f}}(\mathbf{x}))$ whenever $\mathbf{x} \leq \mathbf{f}(\mathbf{x})$;
5. $\mathcal{N}_{\mathbf{f}}(\mathbf{x}) \geq \mathcal{N}_{\mathbf{f}^\sigma}(\mathbf{x})$ for every \vee -strategy $\sigma \in \Sigma$;
6. $\mathcal{N}_{\mathbf{f}}(\mathbf{x}) \leq \mathcal{N}_{\mathbf{f}^\pi}(\mathbf{x})$ for every \wedge -strategy $\pi \in \Pi$;
7. $\mathcal{N}_{\mathbf{f}}(\mathbf{x}) = \mathcal{N}_{\mathbf{f}^\pi}(\mathbf{x})$ for some \wedge -strategy $\pi \in \Pi$.

In particular Lemma 3 implies that the least fixpoint of $\mathcal{N}_{\mathbf{f}}$ is equal to the least fixpoint of \mathbf{f} . Moreover, iteration based on $\mathcal{N}_{\mathbf{f}}$ is at least as fast as Kleene iteration. We therefore use this operator for computing approximates to the least fixpoint. Formally, we define:

Definition 1. *We call the sequence $(\tau_{\mathbf{f}}^{(k)})$ of approximates defined by $\tau_{\mathbf{f}}^{(k)} := \mathcal{N}_{\mathbf{f}}^k(\mathbf{0})$ for $k \in \mathbb{N}$ the τ -sequence for \mathbf{f} . We drop the subscript if it is clear from the context.*

Proposition 1. *Let \mathbf{f} be a feasible min-max-MSP. The τ -sequence $(\tau^{(k)})$ for \mathbf{f} (see definition 1) is monotonically increasing, bounded from above by $\mu\mathbf{f}$, and converges to $\mu\mathbf{f}$. Moreover, $\kappa^{(k)} \leq \tau^{(k)}$ for all $k \in \mathbb{N}$.*

We now show that the new approximation method converges at least *linearly* to the least fixpoint. Theorem 6.2 of [2] implies the following lemma about the convergence of Newton's method for MSPs, i.e., systems without maxima and minima.

Lemma 4. *Let \mathbf{f} be a feasible quadratic MSP. The sequence $(\tau^{(k)})_{k \in \mathbb{N}}$ converges linearly to $\mu\mathbf{f}$. More precisely, there is a $k_{\mathbf{f}} \in \mathbb{N}$ such that $\tau^{(k_{\mathbf{f}} + i \cdot (n+1) \cdot 2^n)}$ has at least i valid bits of $\mu\mathbf{f}$ for every $i \in \mathbb{N}$.*

We emphasize that linear convergence is the worst case. In many practical examples, in particular if the matrix $I - \mathbf{f}'(\mu\mathbf{f})$ is invertible, Newton's method converges *exponentially*. We mean by this that the number of accurate bits of the approximation grows *exponentially* in the number of iterations.

As a first step towards our main result for this section, we use Lemma 4 to show that our approximation method converges linearly whenever \mathbf{f} is a max-MSPs. In this case we obtain the same convergence speed as for MSPs.

Lemma 5. *Let \mathbf{f} be a feasible max-MSP. Let $M := \{\sigma \in \Sigma \mid \mu\mathbf{f}^\sigma = \mu\mathbf{f}\}$. The set M is non-empty and $\tau_{\mathbf{f}}^{(i)} \geq \tau_{\mathbf{f}^\sigma}^{(i)}$ for all $\sigma \in M$ and $i \in \mathbb{N}$.*

Proof. Theorem 1 implies that there exists a \vee -strategy $\sigma \in \Sigma$ such that $\mu\mathbf{f}^\sigma = \mu\mathbf{f}$. Thus M is non-empty. Let $\sigma \in M$. By induction on k Lemma 3 implies $\tau_{\mathbf{f}}^{(k)} = \mathcal{N}_{\mathbf{f}}^k(\mathbf{0}) \geq \mathcal{N}_{\mathbf{f}^\sigma}^k(\mathbf{0}) = \tau_{\mathbf{f}^\sigma}^{(k)}$ for every $k \in \mathbb{N}$. \square

Combining Lemma 4 and Lemma ?? we get linear convergence for max-MSPs:

Theorem 2. *Let \mathbf{f} be a feasible quadratic max-MSP. The τ -sequence $(\tau_{\mathbf{f}}^{(k)})$ for \mathbf{f} (see definition 1) converges linearly to $\mu\mathbf{f}$. More precisely, there is a $k_{\mathbf{f}} \in \mathbb{N}$ such that $\tau_{\mathbf{f}}^{(k_{\mathbf{f}}+i \cdot (n+1) \cdot 2^n)}$ has at least i valid bits of $\mu\mathbf{f}$ for every $i \in \mathbb{N}$.*

A direct consequence of Lemma ?? is that the τ -sequence $(\tau_{\mathbf{f}}^{(i)})$ converges exponentially if $(\tau_{\mathbf{f}^\sigma}^{(i)})$ converges exponentially for some $\sigma \in \Sigma$ with $\mu\mathbf{f}^\sigma = \mu\mathbf{f}$. This is in particular the case if the matrix $I - (\mathbf{f}^\sigma)'(\mu\mathbf{f})$ is invertible. In order to extend this result to min-max-MSPs we state the following lemma which enables us to relate the sequence $(\tau_{\mathbf{f}}^{(i)})$ to the sequences $(\tau_{\mathbf{f}^\pi}^{(i)})$ where $\mu\mathbf{f}^\pi = \mu\mathbf{f}$.

Lemma 6. *Let \mathbf{f} be a feasible min-max-MSP and m denote the number of strategies $\pi \in \Pi$ with $\mu\mathbf{f} = \mu\mathbf{f}^\pi$. There is a constant $k \in \mathbb{N}$ such that for all $i \in \mathbb{N}$ there exists some strategy $\pi \in \Pi$ with $\mu\mathbf{f} = \mu\mathbf{f}^\pi$ and $\tau_{\mathbf{f}^\pi}^{(i)} \leq \tau_{\mathbf{f}}^{(k+m \cdot i)}$.*

We now present the main result of this section which states that our approximation method converges at least linearly also in the general case, i.e., for min-max-MSPs.

Theorem 3. *Let \mathbf{f} be a feasible quadratic min-max-MSP and m denote the number of strategies $\pi \in \Pi$ with $\mu\mathbf{f} = \mu\mathbf{f}^\pi$. The τ -sequence $(\tau_{\mathbf{f}}^{(k)})$ for \mathbf{f} (see definition 1) converges linearly to $\mu\mathbf{f}$. More precisely, there is a $k_{\mathbf{f}} \in \mathbb{N}$ such that $\tau_{\mathbf{f}}^{(k_{\mathbf{f}}+i \cdot m \cdot (n+1) \cdot 2^n)}$ has at least i valid bits of $\mu\mathbf{f}$ for every $i \in \mathbb{N}$.*

The upper bound on the convergence rate provided by Theorem 2 is by the factor m worse than the upper bound obtained for MSPs. Since m is the number of strategies $\pi \in \Pi$ with $\mu\mathbf{f}^\pi = \mu\mathbf{f}$, m is trivially bounded by $|\Pi|$ but is usually much smaller. The τ -sequence $(\tau_{\mathbf{f}}^{(i)})$ converges exponentially whenever $(\tau_{\mathbf{f}^\pi}^{(i)})$ converges exponentially for every π with $\mu\mathbf{f}^\pi = \mu\mathbf{f}$ (see [?]). The latter condition is typically satisfied (see the discussion after Theorem 2).

In order to determine the approximate $\tau^{(i+1)} = \mathcal{N}_{\mathbf{f}}(\tau^{(i)})$ from $\tau^{(i)}$ we must compute the least fixpoint of the min-max-MLS $\mathcal{L}(\mathbf{f}, \tau^{(i)}) \vee \tau^{(i)}$. This can be done by using the strategy improvement algorithm from [8]. The algorithm iterates over \vee -strategies.

For each strategy it solves a linear program or alternatively iterates over \wedge -strategies. The number of \vee -strategies used by this algorithm is trivially bounded by the number of \vee -strategies for $\mathcal{L}(\mathbf{f}, \tau^{(i)}) \vee \tau^{(i)}$ which is exponentially in the number of \vee -expressions occurring in $\mathcal{L}(\mathbf{f}, \tau^{(i)}) \vee \tau^{(i)}$. However, we do not know an example for which the algorithm considers more than linearly many strategies.

6 The ν -Method

The τ -method, presented in the previous section, uses strategy iteration over \vee -strategies to compute $\mathcal{N}_{\mathbf{f}}(\mathbf{y})$. This could be expensive, as there may be exponentially many \vee -strategies. Therefore, we derive an alternative generalization of Newton's method that in each step picks the currently most promising \vee -strategy directly, without strategy iteration.

Consider again a fixed feasible min-max-MSP \mathbf{f} whose least fixpoint we want to approximate. Assume that \mathbf{y} is some approximation of $\mu\mathbf{f}$. Instead of applying $\mathcal{N}_{\mathbf{f}}$ to \mathbf{y} , as the τ -method, we now choose a strategy $\sigma \in \Sigma$ such that $\mathbf{f}(\mathbf{y}) = \mathbf{f}^{\sigma}(\mathbf{y})$, and compute $\mathcal{N}_{\mathbf{f}^{\sigma}}(\mathbf{y})$, where $\mathcal{N}_{\mathbf{f}^{\sigma}}$ was defined in Section 5 as $\mathcal{N}_{\mathbf{f}^{\sigma}}(\mathbf{y}) := \mu(\mathcal{L}(\mathbf{f}^{\sigma}, \mathbf{y}) \vee \mathbf{y})$. In the following we write \mathcal{N}_{σ} instead of $\mathcal{N}_{\mathbf{f}^{\sigma}}$ if \mathbf{f} is understood.

Assume for a moment that \mathbf{f} is a max-MSP and that there is a unique $\sigma \in \Sigma$ such that $\mathbf{f}(\mathbf{y}) = \mathbf{f}^{\sigma}(\mathbf{y})$. The approximant $\mathcal{N}_{\sigma}(\mathbf{y})$ is the result of applying one iteration of Newton's method, because $\mathcal{L}(\mathbf{f}^{\sigma}, \mathbf{y})$ is not only a linearization of \mathbf{f}^{σ} , but the first order Taylor approximation of \mathbf{f} at \mathbf{y} . More precisely, $\mathcal{L}(\mathbf{f}^{\sigma}, \mathbf{y})(\mathbf{x}) = \mathbf{f}(\mathbf{y}) + \mathbf{f}'(\mathbf{y}) \cdot (\mathbf{x} - \mathbf{y})$, and $\mathcal{N}_{\sigma}(\mathbf{y})$ is obtained by solving $\mathbf{x} = \mathcal{L}(\mathbf{f}^{\sigma}, \mathbf{y})(\mathbf{x})$. In this sense, the ν -method is a more direct generalization of Newton's method than the τ -method. Formally, we define the ν -method by a sequence of approximates, the ν -sequence.

Definition 2 (ν -sequence). A sequence $(\nu_{\mathbf{f}}^{(k)})_{k \in \mathbb{N}}$ is called ν -sequence of a min-max-MSP \mathbf{f} if $\nu_{\mathbf{f}}^{(0)} = \mathbf{0}$ and for each k there is a strategy $\sigma_{\mathbf{f}}^{(k)} \in \Sigma$ with $\mathbf{f}(\nu_{\mathbf{f}}^{(k)}) = \mathbf{f}^{\sigma_{\mathbf{f}}^{(k)}}(\nu_{\mathbf{f}}^{(k)})$ and $\nu_{\mathbf{f}}^{(k+1)} = \mathcal{N}_{\sigma_{\mathbf{f}}^{(k)}}(\nu_{\mathbf{f}}^{(k)})$. We may drop the subscript if \mathbf{f} is understood.

Notice the nondeterminism here if there is more than one \vee -strategy that attains $\mathbf{f}(\nu_{\mathbf{f}}^{(k)})$. The following proposition is analogous to Proposition 1 and states some basic properties of ν -sequences.

Proposition 2. Let \mathbf{f} be a feasible min-max-MSP. The sequence $(\nu^{(k)})$ is monotonically increasing, bounded from above by $\mu\mathbf{f}$, and converges to $\mu\mathbf{f}$. More precisely, we have $\kappa^{(k)} \leq \nu^{(k)} \leq \mathbf{f}(\nu^{(k)}) \leq \nu^{(k+1)} \leq \mu\mathbf{f}$ for all $k \in \mathbb{N}$.

The goal of this section is again to strengthen Proposition 2 towards quantitative convergence results for ν -sequences. To achieve this goal we again relate the convergence of ν -sequences to the convergence of Newton's method for MSPs. If \mathbf{f} is an MSP, Lemma 4 allows to argue about the Newton operator $\mathcal{N}_{\mathbf{f}}$ when applied to approximates $\mathbf{x} \leq \mu\mathbf{f}$. To transfer this result to min-max-MSPs \mathbf{f} we need an invariant like $\nu^{(k)} \leq \mu\mathbf{f}^{\sigma^{(k)}}$ for ν -sequences. As a first step to such an invariant we further restrict the selection of the $\sigma^{(k)}$. Roughly speaking, the strategy in a component i is only changed when it is immediate that component i has not yet reached its fixpoint.

Algorithm 1 lazy ν -method

```

procedure lazy- $\nu(\mathbf{f}, k)$ 
assumes:  $\mathbf{f}$  is a min-max-MSP
returns:  $\nu^{(k)}, \sigma^{(k)}$  obtained by  $k$  iterations
of the lazy  $\nu$ -method
 $\nu \leftarrow \mathbf{0}$ 
 $\sigma \leftarrow$  any  $\sigma \in \Sigma$  such that  $\mathbf{f}(\mathbf{0}) = \mathbf{f}^\sigma(\mathbf{0})$ 
for  $i$  from 1 to  $k$  do
     $\nu \leftarrow \mathcal{N}_{\mathbf{f}^\sigma}(\nu)$ 
     $\sigma \leftarrow$  lazy strategy update from  $\nu$  and  $\sigma$ 
od
return  $\nu, \sigma$ 

```

```

procedure  $\mathcal{N}_{\mathbf{f}}(\mathbf{y})$ 
assumes:  $\mathbf{f}$  is a min-MSP,  $\mathbf{y} \in \mathbb{R}_{\geq 0}^n$ 
returns:  $\mu(\mathcal{L}(\mathbf{f}, \mathbf{y}) \vee \mathbf{y})$ 
 $\mathbf{g} \leftarrow$  linear min-MSP with
     $\mathbf{g}(\mathbf{d}) = \mathcal{L}(\mathbf{f}, \mathbf{y})(\mathbf{y} + \mathbf{d}) - \mathbf{y}$ 
 $\mathbf{u} \leftarrow \kappa_{\mathbf{g}}^{(n)}$ 
 $\tilde{\mathbf{g}} \leftarrow (\tilde{g}_1, \dots, \tilde{g}_n)^\top$  where
     $\tilde{g}_i = \begin{cases} 0 & \text{if } u_i = 0 \\ g_i & \text{if } u_i > 0 \end{cases}$ 
 $\mathbf{d}^* \leftarrow$  maximize  $x_1 + \dots + x_n$  subject
to  $\mathbf{0} \leq \mathbf{x} \leq \tilde{\mathbf{g}}(\mathbf{x})$  by 1 LP
return  $\mathbf{y} + \mathbf{d}^*$ 

```

Definition 3 (lazy strategy update). Let $\mathbf{x} \leq \mathbf{f}^\sigma(\mathbf{x})$ for a $\sigma \in \Sigma$. We say that $\sigma' \in \Sigma$ is obtained from \mathbf{x} and σ by a lazy strategy update if $\mathbf{f}(\mathbf{x}) = \mathbf{f}^{\sigma'}(\mathbf{x})$ and $\sigma'(f_i) = \sigma(f_i)$ holds for all components i with $f_i(\mathbf{x}) = x_i$. We call a ν -sequence $(\nu^{(k)})_{k \in \mathbb{N}}$ lazy if for all k , the strategy $\sigma^{(k)}$ is obtained from $\nu^{(k)}$ and $\sigma^{(k-1)}$ by a lazy strategy update.

The key property of lazy ν -sequences is the following non-trivial invariant.

Lemma 7. Let $(\nu^{(k)})_{k \in \mathbb{N}}$ be a lazy ν -sequence. Then $\nu^{(k)} \leq \mu \mathbf{f}^{\sigma^{(k)}} \pi$ holds for all $k \in \mathbb{N}$ and all $\pi \in \Pi^*$.

The following example shows that lazy strategy updates are essential to Lemma 6 even for max-MSPs.

Example 6. Consider the MSP $\mathbf{f}(x, y) = (\frac{1}{2} \vee x, xy + \frac{1}{2})$. Let $\sigma^{(0)}(\frac{1}{2} \vee x) = \frac{1}{2}$ and $\sigma^{(1)}(\frac{1}{2} \vee x) = x$. Then there is a ν -sequence $(\nu^{(k)})$ with $\nu^{(0)} = \mathbf{0}$, $\nu^{(1)} = \mathcal{N}_{\sigma^{(0)}}(\mathbf{0}) = (\frac{1}{2}, 0)$, $\nu^{(2)} = \mathcal{N}_{\sigma^{(1)}}(\nu^{(1)})$. However, the conclusion of Lemma 6 does not hold, because $(\frac{1}{2}, 0) = \nu^{(1)} \not\leq \mu \mathbf{f}^{\sigma^{(1)}} = (0, \frac{1}{2})$. Notice that $\sigma^{(1)}$ is not obtained by a lazy strategy update, as $f_1(\nu^{(1)}) = \nu_1^{(1)}$. \square

Lemma 6 falls short of our subgoal to establish $\nu^{(k)} \leq \mu \mathbf{f}^{\sigma^{(k)}}$, because $\Pi \setminus \Pi^*$ might be non-empty. In fact, we provide an example in [?] showing that $\nu^{(k)} \leq \mu \mathbf{f}^{\sigma^{(k)}} \pi$ does not always hold for all $\pi \in \Pi$, even when $\mathbf{f}^{\sigma^{(k)}} \pi$ is feasible. Luckily, Lemma 6 will suffice for our convergence speed result.

The left procedure of Algorithm 1 summarizes the *lazy ν -method* which works by computing lazy ν -sequences. The following lemma relates the ν -method for min-max-MSPs to Newton's method for MSPs.

Lemma 8. Let \mathbf{f} be a feasible min-max-MSP and $(\nu^{(k)})$ a lazy ν -sequence. Let m be the number of strategy pairs $(\sigma, \pi) \in \Sigma \times \Pi$ with $\mu \mathbf{f} = \mu \mathbf{f}^{\sigma \pi}$. Then $m \geq 1$ and there is a constant $k_{as} \in \mathbb{N}$ such that, for all $k \in \mathbb{N}$, there exist strategies $\sigma \in \Sigma$, $\pi \in \Pi$ with $\mu \mathbf{f} = \mu \mathbf{f}^{\sigma \pi}$ and $\nu_{\mathbf{f}}^{(k_{as} + m \cdot k)} \geq \tau_{\mathbf{f}^{\sigma \pi}}^{(k)}$.

In typical cases, i.e., if $I - (\mathbf{f}^{\sigma \pi})'(\mu \mathbf{f})$ is invertible for all $\sigma \in \Sigma$ and $\pi \in \Pi$ with $\mu \mathbf{f}^{\sigma \pi} = \mu \mathbf{f}$, Newton's method converges exponentially. The following theorem captures the worst-case, in which the lazy ν -method still converges linearly.

Theorem 4. *Let \mathbf{f} be a quadratic feasible min-max-MSP. The lazy ν -sequence $(\nu^{(k)})_{k \in \mathbb{N}}$ converges linearly to $\mu\mathbf{f}$. More precisely, let m be the number of strategy pairs $(\sigma, \pi) \in \Sigma \times \Pi$ with $\mu\mathbf{f} = \mu\mathbf{f}^{\sigma\pi}$. Then there is a $k_{\mathbf{f}} \in \mathbb{N}$ such that $\nu^{(k_{\mathbf{f}} + i \cdot m \cdot (n+1) \cdot 2^n)}$ has at least i valid bits of $\mu\mathbf{f}$ for every $i \in \mathbb{N}$.*

Next we show that $\mathcal{N}_{\mathbf{f}^\sigma}(\mathbf{y})$ can be computed exactly by solving a single LP. The right procedure of Algorithm 1 accomplishes this by taking advantage of the following proposition which states that $\mathcal{N}_{\mathbf{f}^\sigma}(\mathbf{y})$ can be determined by computing the least fixpoint of some linear min-MSP \mathbf{g} .

Proposition 3. *Let $\mathbf{y} \leq \mathbf{f}^\sigma(\mathbf{y}) \leq \mu\mathbf{f}$. Then $\mathcal{N}_{\mathbf{f}^\sigma}(\mathbf{y}) = \mathbf{y} + \mu\mathbf{g}$ for the linear min-MSP \mathbf{g} with $\mathbf{g}(\mathbf{d}) = \mathcal{L}(\mathbf{f}^\sigma, \mathbf{y})(\mathbf{y} + \mathbf{d}) - \mathbf{y}$.*

After having computed the linear min-MSP \mathbf{g} , Algorithm 1 determines the 0-components of $\mu\mathbf{g}$. This can be done by performing n Kleene steps, since $(\mu\mathbf{g})_i = 0$ whenever $(\kappa_{\mathbf{g}}^{(n)})_i = 0$. Let $\tilde{\mathbf{g}}$ be the linear min-MSP obtained from \mathbf{g} by substituting the constant 0 for all components g_i with $(\mu\mathbf{g})_i = 0$. The least fixpoint of $\tilde{\mathbf{g}}$ can be computed by solving a single LP, as implied by the following lemma. The correctness of Algorithm 1 follows.

Lemma 9. *Let \mathbf{g} be a linear min-MSP such that $g_i = 0$ whenever $(\mu\mathbf{g})_i = 0$ for all components i . Then $\mu\mathbf{g}$ is the greatest vector \mathbf{x} with $\mathbf{x} \leq \mathbf{g}(\mathbf{x})$.*

The following theorem is a direct consequence of Lemma 6 for the case where $\Pi = \Pi^*$. It shows the second major advantage of the lazy ν -method, namely, that the strategies $\sigma^{(k)}$ are meaningful in terms of games.

Theorem 5. *Let $\Pi = \Pi^*$. Let $(\nu^{(k)})_{k \in \mathbb{N}}$ be a lazy ν -sequence. Then $\nu^{(k)} \leq \mu\mathbf{f}^{\sigma^{(k)}}$ holds for all $k \in \mathbb{N}$.*

As $(\nu^{(k)})$ converges to $\mu\mathbf{f}$, the max-strategy $\sigma^{(k)}$ can be considered ϵ -optimal. In terms of games, Theorem 5 states that the strategy $\sigma^{(k)}$ guarantees the max-player an outcome of at least $\nu^{(k)}$. It is open whether an analogous theorem holds for the τ -method.

Application to the primaries example. We solved the equation system of Section 4 approximately by performing 5 iterations of the lazy ν -method. Using Theorem 5 we found that Clinton can extinguish a problem (a) individual with a probability of at least $X_a = 0.492$ by concentrating on her program and her “ready from day 1” message. (More than 70 Kleene iterations would be needed to infer that X_a is at least 0.49.) As $\nu^{(5)}$ seems to solve above equation system quite well in the sense that $\|\mathbf{f}(\nu^{(5)}) - \nu^{(5)}\|$ is small, we are pretty sure about Obama’s optimal strategy: he should talk about Iraq. As $\nu_{X_1}^{(2)} > 0.38$ and $\sigma^{(2)}$ maps f_1 to $0.3 + 0.7X_1^2$, Clinton’s team can use Theorem 5 to infer that $X_a \geq 0.38$ by showing emotions and using her “ready from day 1” message.

7 Discussion

In order to compare our two methods in terms of convergence speed, assume that \mathbf{f} denotes a feasible min-max-MSP. Since $\mathcal{N}_{\mathbf{f}}(\mathbf{x}) \geq \mathcal{N}_{\mathbf{f}^\sigma}$ (Lemma 3.5), it follows

that $\tau_f^{(i)} \geq \nu_f^{(i)}$ holds for all $i \in \mathbb{N}$. This means that the τ -method is at least as fast as the ν -method if one counts the number of approximation steps. Next, we construct an example which shows that the number of approximation steps needed by the lazy ν -method can be much larger than the respective number needed by the τ -method. It is parameterized with an arbitrary $k \in \mathbb{N}$ and given by $f(x_1, x_2) = (x_2 \wedge 2, x_1^2 + 0.25 \vee x_1 + 2^{-2(k+1)})^\top$. Since the constant $2^{-2(k+1)}$ is represented using $\mathcal{O}(k)$ bits, it is of size linear in k . It can be shown (see [?]) that the lazy ν -method needs at least k steps. More precisely, $\nu f - \tau^{(k)} \geq (1.5, 1.95)$. The τ -method needs exactly 2 steps.

We now compare our approaches with the tool PReMo [12]. PReMo employs 4 different techniques to approximate μf for min-max-MSPs f : It uses Newton’s method only for MSPs without min or max. In this case both of our methods coincide with Newton’s method. For min-max-MSPs, PReMo uses Kleene iteration, round-robin iteration (called Gauss-Seidel in [12]), and an “optimistic” variant of Kleene which is not guaranteed to converge. In the following we compare our algorithms only with Kleene iteration, as our algorithms are guaranteed to converge and a round-robin step is not faster than n Kleene steps.

Our methods improve on Kleene iteration in the sense that $\kappa^{(i)} \leq \tau^{(i)}, \nu^{(i)}$ holds for all $i \in \mathbb{N}$, and our methods converge linearly, whereas Kleene iteration does not converge linearly in general. For example, consider the MSP $g(x) = \frac{1}{2}x^2 + \frac{1}{2}$ with $\mu g = 1$. Kleene iteration needs exponentially many iterations for j bits [3], whereas Newton’s method gives exactly 1 bit per iteration. For the slightly modified MSP $\tilde{g}(x) = g(x) \wedge 1$ which has the same fixpoint, PReMo no longer uses Newton’s method, as \tilde{g} contains a minimum. Our algorithms still produce exactly 1 bit per iteration.

In the case of linear min-max systems our methods compute the precise solution and not only an approximation. This applies, for example, to the max-linear system of [12] describing the expected time of termination of a nondeterministic variant of Quicksort. Notice that Kleene iteration does not compute the precise solution (except for trivial instances), even for linear MSPs without min or max.

We implemented our algorithms prototypically in Maple and ran them on the quadratic nonlinear min-max-MSP describing the termination probabilities of a recursive simple stochastic game. This game stems from the example suite of PReMo (`rssg2.c`) and we used PReMo to produce the equations. Both of our algorithms reached the least fixpoint after 2 iterations. So we could compute the precise μf and optimal strategies for both players, whereas PReMo computes only approximations of μf .

8 Conclusion

We have presented the first methods for approximately computing the least fixpoint of min-max-MSPs, which are guaranteed to converge at least linearly. Both of them are generalizations of Newton’s method. Whereas the τ -method converges faster in terms of number of approximation steps, one approximation step of the ν -method is cheaper. Furthermore, we have shown that the ν -method computes ϵ -optimal strategies for games. Whether such a result can also be established for the τ -method is still open. A direction for future research is to evaluate our methods in practice. In particular, the

influence of imprecise computation through floating point arithmetic should be studied. It would also be desirable to find a bound on the “threshold” k_f .

Acknowledgement. We thank the anonymous referees for valuable comments.

References

1. A. Condon. The complexity of stochastic games. *Inf. and Comp.*, 96(2):203–224, 1992.
2. J. Esparza, S. Kiefer, and M. Luttenberger. Convergence thresholds of Newton’s method for monotone polynomial equations. In *Proceedings of STACS 2008*, to appear.
3. K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. In *STACS*, pages 340–352, 2005.
4. K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. In *Proc. ICALP 2005*, volume 3580 of *LNCS*. Springer, 2005.
5. K. Etessami and M. Yannakakis. Efficient qualitative analysis of classes of recursive Markov decision processes and simple stochastic games. In *STACS*, pages 634–645, 2006.
6. J. Filar and K. Vrieze. *Competitive Markov Decision processes*. Springer, 1997.
7. T. Gawlitza and H. Seidl. Precise fixpoint computation through strategy iteration. In *European Symposium on Programming (ESOP)*, LNCS 4421, pages 300–315. Springer, 2007.
8. T. Gawlitza and H. Seidl. Precise relational invariants through strategy iteration. In J. Duparc and T. A. Henzinger, editors, *CSL*, volume 4646 of *LNCS*, pages 23–40. Springer, 2007.
9. S. Kiefer, M. Luttenberger, and J. Esparza. On the convergence of Newton’s method for monotone systems of polynomial equations. In *STOC*, pages 217–226. ACM, 2007.
10. A. Neyman and S. Sorin. *Stochastic Games and Applications*. Kluwer Academic Press, 2003.
11. J. Ortega and W. Rheinboldt. *Iterative solution of nonlinear equations in several variables*. Academic Press, 1970.
12. D. Wojtczak and K. Etessami. PReMo: An analyzer for probabilistic recursive models. In *Proc. of TACAS*, volume 4424 of *LNCS*, pages 66–71. Springer, 2007.