

Games through Nested Fixpoints

Thomas Martin Gawlitza and Helmut Seidl

TU München, Institut für Informatik, I2, 85748 München, Germany

Abstract. In this paper we consider two-player zero-sum payoff games on finite graphs, both in the deterministic as well as in the stochastic setting. In the deterministic setting, we consider total-payoff games which have been introduced as a refinement of mean-payoff games [10, 18]. In the stochastic setting, our class is a turn-based variant of liminf-payoff games [4, 15, 16]. In both settings, we provide a non-trivial characterization of the values through nested fixpoint equations. The characterization of the values of liminf-payoff games moreover shows that solving liminf-payoff games is polynomial-time reducible to solving stochastic parity games. We construct practical algorithms for solving the occurring nested fixpoint equations based on strategy iteration. As a corollary we obtain that solving deterministic total-payoff games and solving stochastic liminf-payoff games is in $\text{UP} \cap \text{co-UP}$.

1 Introduction

We consider two-player zero-sum payoff games on finite graphs without sinks. The goal of the maximizer (called \vee -player) is to maximize the outcome (his payoff) of the play whereas the goal of the minimizer (called \wedge -player) is to minimize this value (his losing). We consider both deterministic and stochastic games. In the stochastic setting there additionally exists a probabilistic player who makes decisions randomly.

In the deterministic setting we consider *total-payoff* games, where the payoff is the long-term minimal value of the sum of all immediate rewards. This means that, if $r_1, r_2, r_3 \dots$ denotes the sequence of the immediate rewards of a play, then the total reward is $\liminf_{k \rightarrow \infty} \sum_{i=1}^k r_i$. This natural criterion has been suggested by Gimbert and Zielonka in [10], where they prove that both players have positional optimal strategies.

Total-payoff games can be considered as a refinement of *mean-payoff games* (cf. e.g. [6, 19]). Mean-payoff games are payoff games where the outcome of a play is given as the *limiting average* of the immediate rewards. For the same game, the mean-payoff is positive iff the total-payoff is ∞ ; likewise, the mean-payoff is negative iff the total-payoff is $-\infty$. The refinement concerns games where the mean-payoff is 0. For these, the total-payoff is finite and thus allows to differentiate more precisely between games.

Many algorithms have been proposed for solving mean-payoff games (e.g. [1, 11, 17, 19]). While it is still not known whether mean-payoff games can be solved in polynomial time, it at least has been shown that the corresponding decision problem, i.e., deciding whether the \vee -player can enforce a positive outcome, is in $\text{UP} \cap \text{co-UP}$. Much less is known about total-payoff games.

One contribution of this paper therefore is a method for computing the values and positional optimal strategies for total-payoff games. For this, we provide a polynomial-time reduction to *hierarchical systems of simple integer equations* as introduced in [8].

The hierarchical system utilized for the reduction consists of a greatest fixpoint subsystem which is nested inside a least fixpoint system. This reduction allows to compute the values of total-payoff games by means of the strategy improvement algorithm for the corresponding hierarchical system as presented in [8].

Deterministic total-payoff games can be generalized to stochastic total-payoff games, where the objective is to maximize the expected long-term minimal value of the sum of all immediate rewards. The immediate generalization of deterministic total-payoff games to stochastic total-payoff games is problematic, if both $-\infty$ and ∞ are possible rewards with positive probabilities. Here it is possible to restrict the considerations to stochastic total-payoff games, where the payoff for every play is finite. This allows to model the realistic situation, where each player has an initial amount of money and can go bankrupt. These games can be reduced to *stochastic liminf-payoff games* by encoding the currently accumulated rewards into the positions (cf. [15]).

In stochastic liminf-payoff games an immediate reward is assigned to each position and the payoff for a play is the least immediate reward that occurs infinitely often. These games are more general than they appear at first sight. Stochastic liminf-payoff games have been studied by Maitra and Sudderth in the more general concurrent setting with infinitely many positions [15, 16]. In particular, they show that concurrent stochastic liminf-payoff games have a value and that this value can be expressed through a transfinite iteration. Recently, Chatterjee and Henzinger showed that stochastic liminf-payoff games can be solved in $\text{NP} \cap \text{co-NP}$ [4]. This result is obtained by providing polynomial-time algorithms for both 1-player cases, using the fact that these games are positionally determined. A practical algorithm for the 2-player case is not provided.

In this paper we present a practical method for computing the values and positional optimal strategies for liminf-payoff games which consists of a reduction to a class of nested fixpoint equation systems. The resulting algorithm is a strategy improvement algorithm. As a corollary we obtain that solving liminf-payoff games is in fact in $\text{UP} \cap \text{co-UP}$. As a consequence of the provided characterization we moreover get that liminf-payoff games are polynomial-time reducible to stochastic parity games.

The technical contribution of this paper consists of a practical algorithm for solving *hierarchical systems of stochastic equations*. These are generalizations of hierarchical systems of simple integer equations as considered in [8]. The latter hierarchical systems are equivalent the quantitative μ -calculus introduced in [7]. In a hierarchical equation system least and greatest fixpoint computations are nested as in the μ -calculus. In this paper we use the nesting in order to describe the values of both deterministic total-payoff games and stochastic liminf-payoff games.

Omitted proofs can be found in the corresponding technical report.

2 Definitions

As usual, \mathbb{N} , \mathbb{Z} , \mathbb{Q} and \mathbb{R} denote the set of natural numbers, the set of integers, the set of rational numbers and the set of real numbers, respectively. We assume $0 \in \mathbb{N}$ and we denote $\mathbb{N} \setminus \{0\}$ by $\mathbb{N}_{>0}$. We set $\overline{\mathbb{Z}} := \mathbb{Z} \cup \{-\infty, \infty\}$, $\overline{\mathbb{Q}} := \mathbb{Q} \cup \{-\infty, \infty\}$ and $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, \infty\}$. We use a uniform cost measure where we count arithmetic operations and memory accesses for $\mathcal{O}(1)$. The size of a data structure S in the uniform

cost measure will be denoted by $\|S\|$. An algorithm is called *uniform* iff its running time w.r.t. the uniform cost measure only depends on the size of the input in the uniform cost measure, i.e., the running time does not depend on the sizes of the occurring numbers. Given a binary relation $R \subseteq A \times B$ and a subset $A' \subseteq A$, we write $A'R$ for the set $\{b \in B \mid \exists a \in A' : (a, b) \in R\}$.

Game Graphs. A (stochastic) game graph is a tuple $G = (P_\vee, P_\wedge, P_+, \delta, E)$ with the following properties: The sets P_\vee, P_\wedge and P_+ are pair-wise disjoint sets of positions which are controlled by the \vee -player, the \wedge -player, and the probabilistic player, respectively. We always write P for the set $P_\vee \cup P_\wedge \cup P_+$ of all positions. The set $E \subseteq P^2$ is the set of moves, where we claim that $\{p\}E \neq \emptyset$ holds for all $p \in P$, i.e., the graph (P, E) has no sinks. The mapping $\delta : P_+ \rightarrow P \rightarrow [0, 1]$ assigns a probability distribution $\delta(p)$ with $\sum_{p' \in \{p\}E} \delta(p)(p') = 1$ to each probabilistic position $p \in P_+$, i.e., $\delta(p)(p')$ denotes the probability for p' being the next position assuming that p is the current position. For convenience, we assume that $\{p' \mid \delta(p)(p') > 0\} = \{p\}E$ holds for every $p \in P_+$. A game graph G is called deterministic iff $|\{p\}E| = 1$ holds for all $p \in P_+$. It is called 0-player game graph iff $|\{p\}E| = 1$ holds for all $p \in P_\vee \cup P_\wedge$.

Plays and Histories. An infinite word $\mathbf{p} = p_1p_2p_3\cdots$ over the alphabet P with p_iEp_{i+1} for all $i \in \mathbb{N}_{>0}$ is called a play in G . The set of all plays in G is denoted by $\mathbf{P}^\omega(G)$. The set of all plays in G starting at position $p \in P$ is denoted by $\mathbf{P}^\omega(G, p)$. A non-empty, finite word \mathbf{h} is called play prefix (or history) of $\mathbf{p} \in \mathbf{P}^\omega(G)$ iff there exists some $\mathbf{p}' \in \mathbf{P}^\omega(G)$ such that $\mathbf{p} = \mathbf{h}\mathbf{p}'$ holds. The set of all play prefixes of \mathbf{p} is denoted by $\mathbf{P}^*(\mathbf{p})$ and we set $\mathbf{P}^*(G) := \bigcup_{\mathbf{p} \in \mathbf{P}^\omega(G)} \mathbf{P}^*(\mathbf{p})$ and $\mathbf{P}^*(G, p) := \bigcup_{\mathbf{p} \in \mathbf{P}^\omega(G, p)} \mathbf{P}^*(\mathbf{p})$.

Strategies. A function $\sigma : \mathbf{P}^*(G) \cap P^*P_\vee \rightarrow P \rightarrow [0, 1]$ is called (mixed) \vee -strategy for G iff $\sigma(\mathbf{h}p)$ is a probability distribution with $\{p' \in P \mid \sigma(\mathbf{h}p)(p') > 0\} \subseteq \{p\}E$ for all $\mathbf{h}p \in \mathbf{P}^*(G) \cap P^*P_\vee$. Mixed \wedge -strategies are defined dually. We denote the set of all mixed \vee -strategies and all mixed \wedge -strategies by $\overline{\Sigma}_G$ and $\overline{\Pi}_G$, respectively. A \vee -strategy σ is called pure iff, for every $\mathbf{h}p \in \mathbf{P}^*(G) \cap P^*P_\vee$, $\sigma(\mathbf{h}p)$ is a Dirac measure, i.e., it holds $|\{p' \in P \mid \sigma(\mathbf{h}p)(p') > 0\}| = 1$. Pure \wedge -strategies are defined dually. A pure \vee -strategy σ for G is called (pure and) positional iff $\sigma(\mathbf{p}p) = \sigma(\mathbf{p}'p)$ holds for all $\mathbf{p}p, \mathbf{p}'p \in \mathbf{P}^*(G) \cap P^*P_\vee$. We identify a positional \vee -strategy σ for G with a function which maps a position $p \in P_\vee$ to some position $\sigma(p) \in \{p\}E$. Positional \wedge -strategies are defined dually. We denote the set of all positional \vee -strategies (resp. positional \wedge -strategies) for G by Σ_G (resp. Π_G). We omit the subscript G , whenever it is clear from the context. For a positional \vee -strategy $\sigma \in \Sigma$, we write $G(\sigma)$ for the game graph that is obtained from G by removing all moves that are not used whenever the \vee -player plays consistent with σ . For a positional \wedge -strategy $\pi \in \Pi$, $G(\pi)$ is defined dually.

After fixing a starting position p , a \vee -strategy σ and a \wedge -strategy π , a probability measure $\mathbb{P}_{G,p,\sigma,\pi}$ on $\mathbf{P}^\omega(G, p)$ is defined as usual, i.e., for a given measurable set A of plays, $\mathbb{P}_{G,p,\sigma,\pi}(A)$ denotes the probability that a play which starts at position p and is played according to the strategies σ and π is in the set A . In a deterministic game graph, after fixing a starting position p , a pure \vee -strategy σ , and a pure \wedge -strategy π , there exists exactly one play which can be played from p according to the strategies σ and π . We denote this play by $\mathbf{play}(G, p, \sigma, \pi)$.

Payoff functions. We assign a *payoff* $U(\mathbf{p})$ to each play $\mathbf{p} \in \mathbf{P}^\omega(G)$ due to a measurable *payoff function* U . A pair $\mathcal{G} = (G, U)$ is called a *payoff game*. We set $\mathcal{G}(\sigma) := (G(\sigma), U)$ for all $\sigma \in \Sigma$ and $\mathcal{G}(\pi) := (G(\pi), U)$ for all $\pi \in \Pi$.

Values. For a payoff game $\mathcal{G} = (G, U)$ we denote the expectation of U under the probability measure $\mathbb{P}_{G,p,\sigma,\pi}$ by $\langle\langle p \rangle\rangle_{\mathcal{G},\sigma,\pi}$. For $p \in P$, $\sigma \in \overline{\Sigma}$, and $\pi \in \overline{\Pi}$ we set:

$$\begin{aligned} \langle\langle p \rangle\rangle_{\mathcal{G},\sigma} &:= \inf_{\pi \in \overline{\Pi}} \langle\langle p \rangle\rangle_{\mathcal{G},\sigma,\pi} & \langle\langle p \rangle\rangle_{\mathcal{G}}^{\vee} &:= \sup_{\sigma \in \overline{\Sigma}} \langle\langle p \rangle\rangle_{\mathcal{G},\sigma} \\ \langle\langle p \rangle\rangle_{\mathcal{G},\pi} &:= \sup_{\sigma \in \overline{\Sigma}} \langle\langle p \rangle\rangle_{\mathcal{G},\sigma,\pi} & \langle\langle p \rangle\rangle_{\mathcal{G}}^{\wedge} &:= \inf_{\pi \in \overline{\Pi}} \langle\langle p \rangle\rangle_{\mathcal{G},\pi} \end{aligned}$$

$\langle\langle p \rangle\rangle_{\mathcal{G}}^{\vee}$ and $\langle\langle p \rangle\rangle_{\mathcal{G}}^{\wedge}$ are called the *lower* and the *upper value* of \mathcal{G} at position p , respectively. The payoff game \mathcal{G} is called *determined* iff $\langle\langle p \rangle\rangle_{\mathcal{G}}^{\vee} = \langle\langle p \rangle\rangle_{\mathcal{G}}^{\wedge}$ holds for all positions p . In this case we simply denote the *value* $\langle\langle p \rangle\rangle_{\mathcal{G}}^{\vee} = \langle\langle p \rangle\rangle_{\mathcal{G}}^{\wedge}$ of \mathcal{G} at p by $\langle\langle p \rangle\rangle_{\mathcal{G}}$. We omit the subscript \mathcal{G} , whenever it is clear from the context. Let \mathcal{C} be a class of payoff games. Deciding whether, for a given payoff game $\mathcal{G} \in \mathcal{C}$, a given position p , and a given value $x \in \overline{\mathbb{R}}$, $\langle\langle p \rangle\rangle_{\mathcal{G}} \geq x$ holds, is called the *quantitative decision problem* for \mathcal{C} .

Optimal Strategies. A \vee -strategy σ is called *optimal* iff $\langle\langle p \rangle\rangle_{\mathcal{G},\sigma} = \langle\langle p \rangle\rangle_{\mathcal{G}}^{\vee}$ holds for all positions p . Accordingly, a \wedge -strategy π is called *optimal* iff $\langle\langle p \rangle\rangle_{\mathcal{G},\pi} = \langle\langle p \rangle\rangle_{\mathcal{G}}^{\wedge}$ holds for all positions p . A payoff game \mathcal{G} is called *positionally determined* iff it is determined and there exist positional optimal strategies for both players.

Mean-payoff function. Let $c : E \rightarrow \mathbb{R}$ be a real-valued function that assigns an immediate payoff to every move. The mean-payoff function U_c^{mp} is defined by

$$U_c^{\text{mp}}(\mathbf{p}) := \liminf_{k \rightarrow \infty} \frac{1}{k-1} \sum_{i=1}^{k-1} c(p_i, p_{i+1}), \quad \mathbf{p} = p_1 p_2 \cdots \in \mathbf{P}^\omega(G)$$

The pair $\mathcal{G} = (G, U_c^{\text{mp}})$ is called *mean-payoff game* and is for simplicity identified with the pair (G, c) . Mean-payoff games are widely studied (e.g. [6, 19]).

Example 1. For the mean-payoff game $\mathcal{G} = (G, c)$ in figure 1 it holds $U_c^{\text{mp}}(a(bc)^\omega) = U_c^{\text{mp}}(a(de)^\omega) = 0$. Moreover, it holds $U_c^{\text{mp}}((bc)^\omega) = U_c^{\text{mp}}((cb)^\omega) = 0$. \square

Total-payoff function. Let $c : E \rightarrow \mathbb{R}$ be a real-valued function that assigns an immediate payoff to every move. The total-payoff function U_c^{tp} is defined by

$$U_c^{\text{tp}}(\mathbf{p}) := \liminf_{k \rightarrow \infty} \sum_{i=1}^{k-1} c(p_i, p_{i+1}), \quad \mathbf{p} = p_1 p_2 \cdots \in \mathbf{P}^\omega(G)$$

The pair $\mathcal{G} = (G, U_c^{\text{tp}})$ is called *total-payoff game* and is for simplicity identified with the pair (G, c) .

We only consider the total-payoff function in the deterministic context, because otherwise the expectation may be undefined. This is the case, whenever the payoff is $-\infty$ with some positive probability as well as ∞ with some positive probability. Motivated by applications in economics [18], Gimbert and Zielonka introduced this total-payoff function in [10]. This payoff function is very natural, since it basically just sums up all occurring immediate payoffs. Since this sum may be undefined, it associates to a play \mathbf{p} the most pessimistic long-term total reward of \mathbf{p} from the \vee -player’s perspective. The following example shows that the total-payoff function allows to differentiate more precisely between plays, whenever the mean-payoff is 0.




Fig. 1. A Payoff Game

Example 2. For the total-payoff game $\mathcal{G} = (G, c)$ in figure 1 it hold $U_c^{\text{tp}}(a(bc)^\omega) = 1$ and $U_c^{\text{tp}}(a(de)^\omega) = 2$. Thus, the \vee -player should prefer the play $a(de)^\omega$. Moreover, the payoff may depend on the starting position. In the example we have $U_c^{\text{tp}}((bc)^\omega) = 0$ and $U_c^{\text{tp}}((cb)^\omega) = -4$. \square

Liminf-payoff function. Let $r : P \rightarrow \mathbb{R}$ be a real-valued function that assigns an immediate payoff to every position. The liminf-payoff function $U_r^{\text{lim inf}}$ is defined by

$$U_r^{\text{lim inf}}(\mathbf{p}) := \liminf_{i \rightarrow \infty} r(p_i), \quad \mathbf{p} = p_1 p_2 \dots \in \mathbf{P}^\omega(G)$$

The pair $\mathcal{G} = (G, U_r^{\text{lim inf}})$ is called *liminf-payoff game* and is for simplicity identified with the pair (G, r) . Limsup-payoff games are defined dually. We only consider liminf-payoff games, since every limsup-payoff game can be transformed in an equivalent liminf-payoff game by negating the immediate payoff function r and switching the roles of the two players (cf. [4]).

Liminf-payoff games are more general than it may appear at first sight. Some classes of games on graphs where the decision at a current position only depends on the value achieved so far can be encoded into liminf-payoff games, provided that only finitely different values are possible (cf. [15]). In particular variants of stochastic total-payoff games, where each player starts with an initial amount of money and can go bankrupt, can be expressed using liminf-payoff games.

Parity-payoff function. Let $r : P \rightarrow \mathbb{N}_{>0}$ be a function that assigns a rank to every position. The parity-payoff function U_r^{parity} is defined by

$$U_r^{\text{parity}}(\mathbf{p}) := \begin{cases} 0 & \text{if } \liminf_{i \rightarrow \infty} r(p_i) \text{ is odd} \\ 1 & \text{if } \liminf_{i \rightarrow \infty} r(p_i) \text{ is even} \end{cases}, \quad \mathbf{p} = p_1 p_2 \dots \in \mathbf{P}^\omega(G)$$

The pair $\mathcal{G} = (G, U_r^{\text{parity}})$ is called (stochastic) *parity game* and is for simplicity identified with the pair (G, r) . Recently, Chatterjee et al. showed that stochastic parity games can be solved using a strategy improvement algorithm [5, 2]. Moreover, computing the values of stochastic parity games is polynomial-time reducible to computing the values of stochastic mean-payoff games [3].

Lemma 1 ([2,4,6,9,10,14]). *Liminf-payoff games, Mean-payoff games, parity games, and deterministic total-payoff games are positionally determined.* \square

3 Hierarchical Systems of Stochastic Equations

In this section we introduce *hierarchical systems of stochastic equations*, discuss elementary properties, and show how to solve these hierarchical systems using a strategy

improvement algorithm. Hierarchical systems of stochastic equations are a generalization of hierarchical systems of simple integer equations (studied in [8]).

Assume that a fixed set \mathbf{X} of variables is given. *Rational expressions* are specified by the abstract grammar

$$e ::= q \mid \mathbf{x} \mid e_1 + e_2 \mid t \cdot e \mid e_1 \vee e_2 \mid e_1 \wedge e_2,$$

where $q \in \overline{\mathbb{Q}}$, $t \in \mathbb{Q}_{>0}$, $\mathbf{x} \in \mathbf{X}$, and e, e_1, e_2 are rational expressions. The operator $t \cdot$ has highest precedence, followed by $+$, \wedge and finally \vee which has lowest precedence. We also use the operators $+$, \vee and \wedge as k -ary operators. An element $q \in \overline{\mathbb{Q}}$ is identified with the nullary function that returns q . An equation $\mathbf{x} = e$ is called *rational equation* iff e is a rational expression.

A *system* \mathcal{E} of equations is a finite set $\{\mathbf{x}_1 = e_1, \dots, \mathbf{x}_n = e_n\}$ of equations where $\mathbf{x}_1, \dots, \mathbf{x}_n$ are pairwise distinct. We denote the set $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ of variables occurring in \mathcal{E} by $\mathbf{X}_{\mathcal{E}}$. We drop the subscript, whenever it is clear from the context. The set of subexpressions occurring in right-hand sides of \mathcal{E} is denoted by $\mathcal{S}(\mathcal{E})$ and $\|\mathcal{E}\|$ denotes $|\mathcal{S}(\mathcal{E})|$. An expression e or an equation $\mathbf{x} = e$ is called *disjunctive* (resp. *conjunctive*) iff e does not contain \wedge -operators (resp. \vee -operators). It is called *basic* iff e does neither contain \wedge - nor \vee -operators.

A \vee -*strategy* σ (resp. \wedge -*strategy* π) for \mathcal{E} is a function that maps every expression $e_1 \vee \dots \vee e_k$ (resp. $e_1 \wedge \dots \wedge e_k$) occurring in \mathcal{E} to one of the immediate subexpressions e_j . We denote the set of all \vee -strategies (resp. \wedge -strategies) for \mathcal{E} by $\Sigma_{\mathcal{E}}$ (resp. $\Pi_{\mathcal{E}}$). We drop subscripts, whenever they are clear from the context. For $\sigma \in \Sigma$ the expression $e\sigma$ is defined by

$$(e_1 \vee \dots \vee e_k)\sigma := (\sigma(e_1 \vee \dots \vee e_k))\sigma, \quad (f(e_1, \dots, e_k))\sigma := f(e_1\sigma, \dots, e_k\sigma),$$

where $f \neq \vee$ is some operator. Finally, we set $\mathcal{E}(\sigma) := \{\mathbf{x} = e\sigma \mid \mathbf{x} = e \in \mathcal{E}\}$. The definitions of $e\pi$ and $\mathcal{E}(\pi)$ for a \wedge -strategy π are dual.

The set $\mathbf{X} \rightarrow \overline{\mathbb{R}}$ of all *variable assignments* is a complete lattice. We denote the least upper and the greatest lower bound of a set X by $\bigvee X$ and $\bigwedge X$, respectively. For a *variable assignment* ρ an expression e is mapped to a value $\llbracket e \rrbracket \rho$ by setting $\llbracket \mathbf{x} \rrbracket \rho := \rho(\mathbf{x})$ and $\llbracket f(e_1, \dots, e_k) \rrbracket \rho := f(\llbracket e_1 \rrbracket \rho, \dots, \llbracket e_k \rrbracket \rho)$, where $\mathbf{x} \in \mathbf{X}$, f is a k -ary operator, for instance $+$, and e_1, \dots, e_k are expressions.

A *basic* rational expression e is called *stochastic* iff it is of the form

$$q + t_1 \cdot \mathbf{x}_1 + \dots + t_k \cdot \mathbf{x}_k,$$

where $q \in \overline{\mathbb{Q}}$, $k \in \mathbb{N}$, $\mathbf{x}_1, \dots, \mathbf{x}_k$ are pair-wise distinct variables, $t_1, \dots, t_k \in \mathbb{Q}_{>0}$, and $t_1 + \dots + t_k \leq 1$. A rational expression e is called *stochastic* iff $e\sigma\pi$ is stochastic for every \vee -strategy σ and every \wedge -strategy π . An equation $\mathbf{x} = e$ is called *stochastic* iff e is stochastic.

We define the unary operator $\llbracket \mathcal{E} \rrbracket$ by setting $(\llbracket \mathcal{E} \rrbracket \rho)(\mathbf{x}) := \llbracket e \rrbracket \rho$ for $\mathbf{x} = e \in \mathcal{E}$. A *solution* (resp. *pre-solution*, resp. *post-solution*) is a variable assignment ρ such that $\rho = \llbracket \mathcal{E} \rrbracket \rho$ (resp. $\rho \leq \llbracket \mathcal{E} \rrbracket \rho$, resp. $\rho \geq \llbracket \mathcal{E} \rrbracket \rho$) holds. The set of solutions (resp. pre-solutions, resp. post-solutions) is denoted by $\text{Sol}(\mathcal{E})$ (resp. $\text{PreSol}(\mathcal{E})$, resp. $\text{PostSol}(\mathcal{E})$).

A *hierarchical system* $\mathcal{H} = (\mathcal{E}, r)$ of stochastic equations consists of a system \mathcal{E} of stochastic equations and a rank function r mapping the variables \mathbf{x} of \mathcal{E} to natural

numbers $r(\mathbf{x}) \in \{1, \dots, d\}$, $d \in \mathbb{N}_{>0}$. The canonical solution $\mathbf{CanSol}(\mathcal{H})$ of \mathcal{H} is a particular solution of \mathcal{E} which is characterized by r . In order to define it formally, let, for $j \in \{1, \dots, d\}$ and $\bowtie \in \{=, <, >, \leq, \geq\}$, $\mathbf{X}_{\bowtie j} := \{\mathbf{x} \in \mathbf{X} \mid r(\mathbf{x}) \bowtie j\}$. Then the equations $\mathbf{x} = e$ of \mathcal{E} with $\mathbf{x} \in \mathbf{X}_{\geq j}$ define a monotonic mapping $\llbracket \mathcal{E}, r \rrbracket_j$ from the set $\mathbf{X}_{< j} \rightarrow \overline{\mathbb{R}}$ of variable assignments with domain $\mathbf{X}_{< j}$ into the set $\mathbf{X}_{\geq j} \rightarrow \overline{\mathbb{R}}$ of variable assignments with domain $\mathbf{X}_{\geq j}$ as follows: We set $\llbracket \mathcal{E}, r \rrbracket_{d+1} \rho := \{\}$ for all ρ . Assume that j is odd. Given the mapping $\llbracket \mathcal{E}, r \rrbracket_{j+1}$, the mapping $\llbracket \mathcal{E}, r \rrbracket_j$ is inductively defined by

$$\llbracket \mathcal{E}, r \rrbracket_j \rho := \rho^* + \llbracket \mathcal{E}, r \rrbracket_{j+1}(\rho + \rho^*),$$

where $\rho : \mathbf{X}_{< j} \rightarrow \overline{\mathbb{R}}$ is a variable assignment and $\rho^* : \mathbf{X}_{=j} \rightarrow \overline{\mathbb{R}}$ is the *least* variable assignment such that

$$\rho^*(\mathbf{x}) = \llbracket e \rrbracket(\rho + \rho^* + \llbracket \mathcal{E}, r \rrbracket_{j+1}(\rho + \rho^*))$$

holds for all equations $\mathbf{x} = e \in \mathcal{E}$ with $\mathbf{x} \in \mathbf{X}_{=j}$. Here, the operator $+$ denotes combination of two variable assignments with disjoint domains. The case where j is even, i.e., corresponds to a greatest solution is defined dually. Finally, the *canonical solution* is defined by

$$\mathbf{CanSol}(\mathcal{E}, r) := \llbracket \mathcal{E}, r \rrbracket_1 \{\},$$

where $\{\}$ denotes the variable assignment with empty domain. The existence of a canonical solution is ensured by Knaster-Tarski's fixpoint theorem, since all used operators are monotone. For convenience, we denote an equation $\mathbf{x} = e$ with $r(\mathbf{x})$ odd by $\mu : \mathbf{x} = e$ and with $r(\mathbf{x})$ even by $\nu : \mathbf{x} = e$. Equations of lower ranks will be written above equations of higher rank.

Example 3. Consider the following hierarchical system $\mathcal{H} = (\mathcal{E}, r)$:

$$\begin{array}{ll} \mu : \mathbf{b}_1 = 0 \vee (2 + \mathbf{x}_2 \vee 1 + \mathbf{x}_3) & \mu : \mathbf{b}_2 = 0 \vee \frac{1}{3} \cdot (-2 + \mathbf{x}_3) + \frac{2}{3} \cdot (1 + \mathbf{x}_4) \\ \mu : \mathbf{b}_3 = 0 \vee 2 + \mathbf{x}_2 & \mu : \mathbf{b}_4 = 0 \vee -1 + \mathbf{x}_2 \\ \nu : \mathbf{x}_1 = \mathbf{b}_1 \wedge (2 + \mathbf{x}_2 \vee 1 + \mathbf{x}_3) & \nu : \mathbf{x}_2 = \mathbf{b}_2 \wedge \frac{1}{3} \cdot (-2 + \mathbf{x}_3) + \frac{2}{3} \cdot (1 + \mathbf{x}_4) \\ \nu : \mathbf{x}_3 = \mathbf{b}_3 \wedge 2 + \mathbf{x}_2 & \nu : \mathbf{x}_4 = \mathbf{b}_4 \wedge -1 + \mathbf{x}_2 \end{array}$$

Thereby, r is given by $r(\mathbf{b}_i) = 1$ and $r(\mathbf{x}_i) = 2$ for all i . It holds $\mathbf{CanSol}(\mathcal{H}) = \{\mathbf{b}_1 \mapsto 1, \mathbf{b}_2 \mapsto 0, \mathbf{b}_3 \mapsto 0, \mathbf{b}_4 \mapsto 0, \mathbf{x}_1 \mapsto 1, \mathbf{x}_2 \mapsto -2, \mathbf{x}_3 \mapsto 0, \mathbf{x}_4 \mapsto -3\}$. \square

We only consider hierarchical systems $\mathcal{H} = (\mathcal{E}, r)$ with *finite* canonical solutions. A variable assignment $\rho : \mathbf{X} \rightarrow \overline{\mathbb{R}}$ is called *finite* iff $-\infty < \rho(\mathbf{x}) < \infty$ holds for all $\mathbf{x} \in \mathbf{X}$. In order to make finite solutions *unique*, we introduce the lattice

$$\mathbb{R}_d := \mathbb{R} \times (-1)^1 \cdot \mathbb{R}_{\geq 0} \times (-1)^2 \cdot \mathbb{R}_{\geq 0} \times \dots \times (-1)^d \cdot \mathbb{R}_{\geq 0},$$

which is lexicographically linear ordered (for instance: $(2, -3, 5) < (2, -2, 1)$). The same idea is also used in [8]. Finally, we set $\overline{\mathbb{R}}_d := \mathbb{R}_d \cup \{-\infty, \infty\}$, where $-\infty$ and ∞ is the least and the greatest element, respectively. We define:

$$\begin{aligned} (x, y_1, \dots, y_d) + (x', y'_1, \dots, y'_d) &:= (x + x', y_1 + y'_1, \dots, y_d + y'_d) \\ t \cdot (x, y_1, \dots, y_d) &:= (t \cdot x, t \cdot y_1, \dots, t \cdot y_d) \end{aligned}$$

where $(x, y_1, \dots, y_d), (x', y'_1, \dots, y'_d) \in \mathbb{R}_d$ and $t \in \mathbb{R}$. As usual, we set $-\infty + \infty := -\infty$. With these definitions the usual commutativity, associativity and distributivity laws hold, i.e., for instance $t \cdot (x_1 \vee x_2) = t \cdot x_1 \vee t \cdot x_2$ holds for all $t \in \mathbb{R}_{>0}, x_1, x_2 \in \mathbb{R}_d$.

We call the analogon to stochastic equations (resp. stochastic expressions), where constants are from $\overline{\mathbb{R}}_d$, *stochastic equations over $\overline{\mathbb{R}}_d$* (resp. stochastic expressions over $\overline{\mathbb{R}}_d$). Since $\overline{\mathbb{R}}_d$ is not a complete lattice, Knaster-Tarski's fixpoint theorem cannot be applied. Nevertheless, every system \mathcal{E} of stochastic equations over $\overline{\mathbb{R}}_d$ has a least and a greatest solution, which we denote by $\mu[\mathcal{E}]$ and $\nu[\mathcal{E}]$, respectively.

Lemma 2. *Every system \mathcal{E} of stochastic equations over $\overline{\mathbb{R}}_d$ has a least solution $\mu[\mathcal{E}] = \bigwedge \text{PostSol}(\mathcal{E})$ and a greatest solution $\nu[\mathcal{E}] = \bigvee \text{PreSol}(\mathcal{E})$. \square*

Given a stochastic expression e , we define the stochastic expression $e^\#$ over $\overline{\mathbb{R}}_d$ as the expression obtained from e by replacing every constant $x \in \mathbb{R}$ occurring in e with the constant $(x, 0, \dots, 0) \in \mathbb{R}_d$, i.e., $x^\# := (x, 0, \dots, 0)$, $-\infty^\# := -\infty^\#$, $\infty^\# := \infty^\#$, $\mathbf{x}^\# := \mathbf{x}$, $(e_1 + e_2)^\# := e_1^\# + e_2^\#$, $(t \cdot e)^\# := t \cdot e^\#$, $(e_1 \wedge e_2)^\# := e_1^\# \wedge e_2^\#$, and $(e_1 \vee e_2)^\# := e_1^\# \vee e_2^\#$, where $x \in \mathbb{Q}$, $\mathbf{x} \in \mathbf{X}$, $t \in \mathbb{Q}_{>0}$, and e, e_1, e_2 are expressions. For $j \in \{1, \dots, d\}$ we set $\mathbf{1}_j := (0, \delta_{1j} \cdot (-1)^1, \delta_{2j} \cdot (-1)^2, \dots, \delta_{dj} \cdot (-1)^d)$. Thereby δ denotes the Kronecker delta, i.e., δ_{ij} equals 1 if $i = j$ and 0 otherwise. For a hierarchical system $\mathcal{H} = (\mathcal{E}, r)$ of stochastic equations we define the system $\mathcal{E}^{r\#}$ of stochastic equations over $\overline{\mathbb{R}}_d$ by $\mathcal{E}^{r\#} := \{\mathbf{x} = \mathbf{1}_{r(\mathbf{x})} + e^\# \mid \mathbf{x} = e \in \mathcal{E}\}$.

We define mappings α and β by $\alpha(-\infty) := -\infty$, $\alpha(\infty) := \infty$, $\alpha(x, y_1, \dots, y_d) := x$, $\beta(-\infty) := \beta(\infty) := (0, \dots, 0)$, and $\beta(x, y_1, \dots, y_d) := (y_1, \dots, y_d)$ for all $(x, y_1, \dots, y_d) \in \mathbb{R}_d$. For a basic stochastic expression e over $\overline{\mathbb{R}}_d$, we define $W(e)$ by $W(x) := \beta(x)$, $W(\mathbf{x}) := (0, \dots, 0)$, $W(e_1 + e_2) := W(e_1) + W(e_2)$, and $W(t \cdot e) := t \cdot W(e)$, where $x \in \overline{\mathbb{R}}_d$, $\mathbf{x} \in \mathbf{X}$, $t \in \mathbb{R}_{>0}$, and e, e_1, e_2 are expressions. We call e *non-zero* iff $\text{Vars}(e) = \emptyset$ or $W(e) \neq (0, \dots, 0)$. A stochastic expression e (resp. stochastic equation $\mathbf{x} = e$) over $\overline{\mathbb{R}}_d$ is finally called *non-zero* iff $e\sigma\pi$ is non-zero for every \vee -strategy σ and every \wedge -strategy π . A non-zero expression e remains non-zero if one rewrites it using commutativity, associativity, and distributivity. Note that $\mathcal{E}^{r\#}$ is non-zero. We have:

Lemma 3. *1. Every system \mathcal{E} of non-zero stochastic equations over $\overline{\mathbb{R}}_d$ has at most one finite solution ρ . If ρ exists, then its size is polynomial in the size of \mathcal{E} .*

2. Let $\mathcal{H} = (\mathcal{E}, r)$ be a hierarchical system of stochastic equations. Assume that the canonical solution $\text{CanSol}(\mathcal{H})$ of \mathcal{H} is finite. The non-zero system $\mathcal{E}^{r\#}$ has exactly one finite solution ρ and it holds $\alpha \circ \rho = \text{CanSol}(\mathcal{H})$. \square

Example 4. Consider again the hierarchical system $\mathcal{H} = (\mathcal{E}, r)$ of stochastic equations from example 3. The system $\mathcal{E}^{r\#}$ of non-zero stochastic equations over $\overline{\mathbb{R}}_d$ is given by:

$$\begin{aligned} \mathbf{b}_1 &= (0, -1, 0) + ((0, 0, 0) \vee ((2, 0, 0) + \mathbf{x}_2 \vee (1, 0, 0) + \mathbf{x}_3)) \\ \mathbf{b}_2 &= (0, -1, 0) + ((0, 0, 0) \vee \frac{1}{3} \cdot ((-2, 0, 0) + \mathbf{x}_3) + \frac{2}{3} \cdot ((1, 0, 0) + \mathbf{x}_4)) \\ \mathbf{b}_3 &= (0, -1, 0) + ((0, 0, 0) \vee (2, 0, 0) + \mathbf{x}_2) \\ \mathbf{b}_4 &= (0, -1, 0) + ((0, 0, 0) \vee (-1, 0, 0) + \mathbf{x}_2) \\ \mathbf{x}_1 &= (0, 0, 1) + (\mathbf{b}_1 \wedge ((2, 0, 0) + \mathbf{x}_2 \vee (1, 0, 0) + \mathbf{x}_3)) \\ \mathbf{x}_2 &= (0, 0, 1) + (\mathbf{b}_2 \wedge \frac{1}{3} \cdot ((-2, 0, 0) + \mathbf{x}_3) + \frac{2}{3} \cdot ((1, 0, 0) + \mathbf{x}_4)) \\ \mathbf{x}_3 &= (0, 0, 1) + (\mathbf{b}_3 \wedge (2, 0, 0) + \mathbf{x}_2) \\ \mathbf{x}_4 &= (0, 0, 1) + (\mathbf{b}_4 \wedge (-1, 0, 0) + \mathbf{x}_2) \end{aligned}$$

The unique finite solution ρ of \mathcal{E}^\sharp is given by $\rho(\mathbf{b}_1) = (1, -2, 1)$, $\rho(\mathbf{b}_2) = (0, -1, 0)$, $\rho(\mathbf{b}_3) = (0, -1, 0)$, $\rho(\mathbf{b}_4) = (0, -1, 0)$, $\rho(\mathbf{x}_1) = (1, -2, 2)$, $\rho(\mathbf{x}_2) = (-2, -1, 6)$, $\rho(\mathbf{x}_3) = (0, -1, 1)$, $\rho(\mathbf{x}_4) = (-3, -1, 7)$. Lemma 3 implies $\text{CanSol}(\mathcal{H}) = \alpha \circ \rho$. \square

An equation $\mathbf{x} = e$ is called *bounded* iff e is of the form $e' \wedge b \vee a$ with $-\infty < a, b < \infty$. A system \mathcal{E} of equations is called *bounded* iff every equation of \mathcal{E} is bounded. A bounded system \mathcal{E} of non-zero stochastic equations over $\overline{\mathbb{R}}_d$ has exactly one solution ρ^\sharp . The problem of deciding, whether, for a given bounded system \mathcal{E} of non-zero stochastic equations over $\overline{\mathbb{R}}_d$, a given variable \mathbf{x} , and a given $x \in \overline{\mathbb{R}}_d$, $\rho^\sharp(\mathbf{x}) \geq x$ holds, is called the quantitative decision problem for bounded systems of non-zero stochastic equations over $\overline{\mathbb{R}}_d$. Thereby ρ^\sharp denotes the unique solution of \mathcal{E} . Since by lemma 3 the size of ρ^\sharp is polynomial in the size of \mathcal{E} , the uniqueness of ρ^\sharp implies that the quantitative decision problem is in $\text{UP} \cap \text{co-UP}$.

We use a strategy improvement algorithm for computing the unique solution ρ^\sharp of a bounded system \mathcal{E} of non-zero stochastic equations over $\overline{\mathbb{R}}_d$. Given a \vee -strategy σ and a pre-solution ρ , we say that a \vee -strategy σ' is an *improvement* of σ w.r.t. ρ iff $\llbracket \mathcal{E}(\sigma') \rrbracket \rho > \llbracket \mathcal{E}(\sigma) \rrbracket \rho$ holds. Assuming that $\rho \in \text{PreSol}(\mathcal{E}) \setminus \text{Sol}(\mathcal{E})$ holds, an improvement σ' of σ w.r.t. ρ can be computed in time $\mathcal{O}(\|\mathcal{E}\|)$ (w.r.t. the uniform cost measure). For instance one can choose σ' such that, for every expression $e = e_1 \vee \dots \vee e_k$ occurring in \mathcal{E} , $\sigma'(e) = e_j$ only holds, if $\llbracket e_j \rrbracket \rho = \llbracket e \rrbracket \rho$ holds (known as *all profitable switches*). The definitions for \wedge -strategies are dual.

Our method for computing ρ^\sharp starts with an arbitrary \vee -strategy σ_0 such that $\mathcal{E}(\sigma_0)$ has a unique finite solution ρ_0 . The \vee -strategy which select a for every right-hand side $e \wedge b \vee a$ can be used for this purpose. Assume that, after the i -th step, σ_i is a \vee -strategy such that $\mathcal{E}(\sigma_i)$ has a unique finite solution $\rho_i < \rho^\sharp$. Let σ_{i+1} be an improvement of σ_i w.r.t. ρ_i . By Lemmas 2 and 3, it follows that $\mathcal{E}(\sigma_{i+1})$ has a unique finite solution ρ_{i+1} and it holds $\rho_i \leq \rho_{i+1} \leq \rho^\sharp$. Since ρ_i is not a solution of $\mathcal{E}(\sigma_{i+1})$, it follows $\rho_i < \rho_{i+1} \leq \rho^\sharp$. Thus ρ_0, ρ_1, \dots is strictly increasing until the unique solution ρ^\sharp of \mathcal{E} is reached. This must happen at the latest after iterating over all \vee -strategies.

Analogously, the unique finite solution ρ_i of the non-zero system $\mathcal{E}(\sigma_i)$ of stochastic equations over $\overline{\mathbb{R}}_d$ can also be computed by iterating over at most all \wedge -strategies for $\mathcal{E}(\sigma_i)$. Every \wedge -strategy π considered during this process leads to a non-zero system of *basic* stochastic equations over $\overline{\mathbb{R}}_d$ with a unique finite solution, which, according to the following lemma, can be computed using an arbitrary algorithm for solving linear equation systems.

Lemma 4. *Let \mathcal{E} be a system of basic non-zero stochastic equations over $\overline{\mathbb{R}}_d$ with a unique finite solution ρ . For $i \in \{1, \dots, d\}$ we set $\beta_i(x, y_1, \dots, y_d) := y_i$ for $(x, y_1, \dots, y_d) \in \overline{\mathbb{R}}_d$ and $\beta_i(c) = c$ for $c \in \{-\infty, \infty\}$. For $\gamma \in \{\alpha, \beta_1, \dots, \beta_d\}$ we denote the system of stochastic equations obtained from \mathcal{E} by replacing every constant c occurring in \mathcal{E} with $\gamma(c)$ by $\gamma(\mathcal{E})$. The following assertions hold:*

1. *The system $\gamma(\mathcal{E})$ has a unique finite solution ρ_γ for all $\gamma \in \{\alpha, \beta_1, \dots, \beta_d\}$.*
2. *It holds $\rho(\mathbf{x}) = (\rho_\alpha(\mathbf{x}), \rho_{\beta_1}(\mathbf{x}), \dots, \rho_{\beta_d}(\mathbf{x}))$ for all $\mathbf{x} \in \mathbf{X}$.* \square

As a consequence of the above lemma, the unique solution of a system \mathcal{E} of basic non-zero stochastic equations over $\overline{\mathbb{R}}_d$, provided that it exists, can be computed in time $\mathcal{O}(d \cdot (|\mathbf{X}|^3 + \|\mathcal{E}\|))$ (w.r.t. the uniform cost measure). Moreover, if \mathcal{E} is a system of

basic non-zero *simple integer* equations over $\overline{\mathbb{R}}_d$ (see [8]), then this can be done in time $\mathcal{O}(d \cdot (|\mathbf{X}| + \|\mathcal{E}\|))$ (w.r.t. the uniform cost measure), since there are no cyclic variable dependencies.

We have shown that ρ^\sharp can be computed using a strategy improvement algorithm which needs at most $\mathcal{O}(|\Sigma| \cdot |\Pi| \cdot d \cdot (|\mathbf{X}|^3 + \|\mathcal{E}\|))$ arithmetic operations. The numbers $|\Sigma|$ and $|\Pi|$ are exponential in the size of \mathcal{E} . However, it is not clear, whether there exist instances for which super-polynomial many strategy improvement steps are necessary.

Observe that all steps of the above procedure can be done in polynomial time. Summarizing, we have the following main result:

Lemma 5. *The unique solution of a bounded system \mathcal{E} of non-zero stochastic equations over $\overline{\mathbb{R}}_d$ can be computed in time $|\Sigma| \cdot |\Pi| \cdot \mathcal{O}(d \cdot (|\mathbf{X}|^3 + \|\mathcal{E}\|))$ (w.r.t. the uniform cost measure) using a strategy improvement algorithm. The quantitative decision problem for bounded systems of non-zero stochastic equations over $\overline{\mathbb{R}}_d$ is in $\mathbf{UP} \cap \mathbf{co-UP}$. \square*

A hierarchical system $\mathcal{H} = (\mathcal{E}, r)$ is called *bounded* iff \mathcal{E} is bounded. The problem of deciding, whether, for a given bounded hierarchical system of stochastic equations $\mathcal{H} = (\mathcal{E}, r)$, a given $\mathbf{x} \in \mathbf{X}$, and a given $x \in \overline{\mathbb{R}}$, $\mathbf{CanSol}(\mathcal{H})(\mathbf{x}) \geq x$ holds, is called the quantitative decision problem for bounded hierarchical systems of stochastic equations.

We finally want to compute the canonical solution $\mathbf{CanSol}(\mathcal{H})$ of a bounded hierarchical system $\mathcal{H} = (\mathcal{E}, r)$ of stochastic equations. By Lemma 3, it suffices to compute the *unique* solution of the bounded system \mathcal{E}^{r^\sharp} which can be done according to Lemma 5. Thus, we obtain:

Theorem 1. *The canonical solution of a bounded hierarchical system $\mathcal{H} = (\mathcal{E}, r)$ of stochastic equations can be computed in time $|\Sigma| \cdot |\Pi| \cdot \mathcal{O}(d \cdot (|\mathbf{X}|^3 + \|\mathcal{E}\|))$ (w.r.t. the uniform cost measure) using a strategy improvement algorithm, where $d = \max r(\mathbf{X})$ is the highest occurring rank. The quantitative decision problem for bounded hierarchical systems of stochastic equations is in $\mathbf{UP} \cap \mathbf{co-UP}$. \square*

If a hierarchical system $\mathcal{H} = (\mathcal{E}, r)$ of stochastic equations is not bounded but has a finite canonical solution $\mathbf{CanSol}(\mathcal{H})$, then it is nevertheless possible to compute $\mathbf{CanSol}(\mathcal{H})$ using the developed techniques. For that, one first computes a bound $B(\mathcal{E})$ such that $|\mathbf{CanSol}(\mathcal{H})(\mathbf{x})| \leq B(\mathcal{E})$ holds for all $\mathbf{x} \in \mathbf{X}$. Then one solves the hierarchical system $(\mathcal{E} \wedge B(\mathcal{E}) \vee -B(\mathcal{E}), r)$ instead of \mathcal{H} , where $\mathcal{E} \wedge B(\mathcal{E}) \vee -B(\mathcal{E}) := \{\mathbf{x} = e \wedge B(\mathcal{E}) \vee -B(\mathcal{E}) \mid \mathbf{x} = e \in \mathcal{E}\}$. An adequate bound $B(\mathcal{E})$ can be computed in polynomial time. Moreover, there will be obvious bounds to the absolute values of the components of the canonical solutions in the applications in the following sections.

4 Solving Deterministic Total-Payoff Games

A lower complexity bound. In a first step we establish a lower complexity bound for deterministic total-payoff games by reducing deterministic mean-payoff games to deterministic total-payoff games. The quantitative decision problem for mean-payoff games is in $\mathbf{UP} \cap \mathbf{co-UP}$ [13], but, despite of many attempts for finding polynomial algorithms, not known to be in \mathbf{P} . Let $\mathcal{G} = (G, c)$ be a deterministic mean-payoff game, where we w.l.o.g. assume that all immediate rewards are from \mathbb{Z} , i.e., it holds $c(E) \subseteq \mathbb{Z}$.

Lemma 6. For all positions $p \in P$ and all $x \in \mathbb{R}$ it holds:

$$\langle\langle p \rangle\rangle_{(G, U_c^{\text{mp}})} \begin{cases} < x & \text{if } \langle\langle p \rangle\rangle_{(G, U_{c-p}^{\text{tp}})} = -\infty \\ = x & \text{if } \langle\langle p \rangle\rangle_{(G, U_{c-p}^{\text{tp}})} \in \mathbb{R} \\ > x & \text{if } \langle\langle p \rangle\rangle_{(G, U_{c-p}^{\text{tp}})} = \infty. \end{cases} \quad \square$$

By lemma 6 the quantitative decision problem for mean-payoff games is polynomial-time reducible to the problem of computing the values of total-payoff games. Since the problem of computing the values of mean-payoff games is polynomial-time reducible to the quantitative decision problem for mean-payoff games we finally get:

Theorem 2. Computing the values of deterministic mean-payoff games is polynomial-time reducible to computing the values of deterministic total-payoff games. \square

Computing the values. Our next goal is to compute the values of deterministic total-payoff games. We do this through *hierarchical systems of stochastic equations*. Since stochastic expressions like $\frac{1}{2}e_1 + \frac{1}{2}e_2$ are ruled out, we only need *hierarchical systems of simple integer equations* as considered in [8].

Let $\mathcal{G} = (G, c)$ be a deterministic total-payoff game where we w.l.o.g. assume that $P_+ = \emptyset$ holds. For every Position $p \in P$ we define the expression $e_{\mathcal{G},p}$ by

$$e_{\mathcal{G},p} := \begin{cases} \bigvee_{p' \in \{p\}E} c(p, p') + \mathbf{x}_{p'} & \text{if } p \in P_V \\ \bigwedge_{p' \in \{p\}E} c(p, p') + \mathbf{x}_{p'} & \text{if } p \in P_\wedge. \end{cases}$$

Finally, we define, for $b \in \mathbb{N}$, the hierarchical system $\mathcal{H}_{\mathcal{G},b} = (\mathcal{E}_{\mathcal{G},b}, r_{\mathcal{G}})$ by

$$\mathcal{E}_{\mathcal{G},b} := \{\mathbf{b}_p = (0 \vee e_{\mathcal{G},p}) \wedge 2b + 1, \mathbf{x}_p = (\mathbf{b}_p \wedge e_{\mathcal{G},p}) \vee -2b - 1 \mid p \in P\}$$

and $r_{\mathcal{G}} := \{\mathbf{b}_p \mapsto 1, \mathbf{x}_p \mapsto 2 \mid p \in P\}$. We set $B(\mathcal{G}) := \sum_{m \in E} |c(m)|$, $\mathcal{E}_{\mathcal{G}} := \mathcal{E}_{\mathcal{G},B(\mathcal{G})}$, and $\mathcal{H}_{\mathcal{G}} := (\mathcal{E}_{\mathcal{G}}, r_{\mathcal{G}})$.

Thus $\mathcal{H}_{\mathcal{G}}$ is a hierarchical system, where the variables \mathbf{x}_p ($p \in P$) live within the scope of the variables \mathbf{b}_p ($p \in P$). Moreover, the \mathbf{x}_p 's ($p \in P$) are “greatest” whereas the \mathbf{b}_p 's ($p \in P$) are “least fixpoint variables”. Intuitively, we are searching for the greatest values for the \mathbf{x}_p 's ($p \in P$) under the least possible bounds \mathbf{b}_p 's ($p \in P$) that are greater than or equal to 0. We illustrate the reduction by an example:

Example 5 (The Reduction). Let \mathcal{G} be the deterministic total-payoff game shown in figure 1. The system $\mathcal{E}_{\mathcal{G},11}$ of simple integer equations is given by:

$$\begin{array}{ll} \mathbf{b}_a = (0 \vee (1 + \mathbf{x}_b \vee 2 + \mathbf{x}_d)) \wedge 23 & \mathbf{x}_a = (\mathbf{b}_a \wedge (1 + \mathbf{x}_b \vee 2 + \mathbf{x}_d)) \vee -23 \\ \mathbf{b}_b = (0 \vee 4 + \mathbf{x}_c) \wedge 23 & \mathbf{x}_b = (\mathbf{b}_b \wedge 4 + \mathbf{x}_c) \vee -23 \\ \mathbf{b}_c = (0 \vee -4 + \mathbf{x}_b) \wedge 23 & \mathbf{x}_c = (\mathbf{b}_c \wedge -4 + \mathbf{x}_b) \vee -23 \\ \mathbf{b}_d = (0 \vee \mathbf{x}_e) \wedge 23 & \mathbf{x}_d = (\mathbf{b}_d \wedge \mathbf{x}_e) \vee -23 \\ \mathbf{b}_e = (0 \vee \mathbf{x}_d) \wedge 23 & \mathbf{x}_e = (\mathbf{b}_e \wedge \mathbf{x}_d) \vee -23. \end{array}$$

Together with the rank function $r_{\mathcal{G}}$ defined by $r_{\mathcal{G}}(\mathbf{b}_a) = r_{\mathcal{G}}(\mathbf{b}_b) = r_{\mathcal{G}}(\mathbf{b}_c) = r_{\mathcal{G}}(\mathbf{b}_d) = r_{\mathcal{G}}(\mathbf{b}_e) = 1$ and $r_{\mathcal{G}}(\mathbf{x}_a) = r_{\mathcal{G}}(\mathbf{x}_b) = r_{\mathcal{G}}(\mathbf{x}_c) = r_{\mathcal{G}}(\mathbf{x}_d) = r_{\mathcal{G}}(\mathbf{x}_e) = 2$ we have given the hierarchical system $\mathcal{H}_{\mathcal{G},11} = (\mathcal{E}_{\mathcal{G},11}, r_{\mathcal{G}})$. Then $\text{CanSol}(\mathcal{H}_{\mathcal{G},11}) = \{\mathbf{b}_a \mapsto 2, \mathbf{b}_b \mapsto 0, \mathbf{b}_c \mapsto 0, \mathbf{b}_d \mapsto 0, \mathbf{b}_e \mapsto 0, \mathbf{x}_a \mapsto 2, \mathbf{x}_b \mapsto 0, \mathbf{x}_c \mapsto -4, \mathbf{x}_d \mapsto 0, \mathbf{x}_e \mapsto 0\}$. \square

In order to show that the values of a total-payoff game \mathcal{G} are given by the canonical solution of the hierarchical system $\mathcal{H}_{\mathcal{G}}$, we first show this result for the case that no player has choices. In order to simplify notations, for $b \in \mathbb{N}$, we set

$$v_b(x) := \begin{cases} -\infty & \text{if } x < -b \\ x & \text{if } b \leq x \leq b \\ \infty & \text{if } x > b \end{cases}, \quad x \in \overline{\mathbb{R}}.$$

Lemma 7. *Let $\mathcal{G} = (G, c)$ be a 0-player deterministic total-payoff game and $b \geq B(\mathcal{G})$. It holds $\langle\langle p \rangle\rangle = v_b(\mathbf{CanSol}(\mathcal{H}_{\mathcal{G},b})(\mathbf{x}_p))$ for all $p \in P$. \square*

In order to generalize the statement of the above lemma, we show that (in a certain sense) there exist positional optimal strategies for the hierarchical system $\mathcal{H}_{\mathcal{G},b}$:

Lemma 8. *Let \mathcal{G} be a deterministic total-payoff game and $b \in \mathbb{N}$. It hold*

1. $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G},b}) = \max_{\sigma \in \Sigma} \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma),b})$ and
2. $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G},b}) = \min_{\pi \in \Pi} \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\pi),b})$.

Proof. We only show statement 1. Statement 2 can be shown analogously. By monotonicity it holds $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma),b}) \leq \mathbf{CanSol}(\mathcal{H}_{\mathcal{G},b})$ for all $\sigma \in \Sigma$. It remains to show that there exists some $\sigma \in \Sigma$ such that $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma),b}) = \mathbf{CanSol}(\mathcal{H}_{\mathcal{G},b})$ holds. Since by lemma 3 $\mathcal{E}_{\mathcal{G},b}^{r\sigma\sharp}$ has a unique finite solution ρ^\sharp , there exists a $\sigma \in \Sigma$ such that ρ^\sharp is the unique finite solution of the system $\mathcal{E}_{\mathcal{G}(\sigma),b}^{r\sigma\sharp}$. Thus lemma 3 gives us that $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma),b}) = \alpha \circ \rho^\sharp = \mathbf{CanSol}(\mathcal{H}_{\mathcal{G},b})$ holds. \square

We now prove our main result for deterministic total-payoff games:

Lemma 9. *Let $b \geq B(\mathcal{G})$. It holds $\langle\langle p \rangle\rangle = v_b(\mathbf{CanSol}(\mathcal{H}_{\mathcal{G},b})(\mathbf{x}_p))$ for every $p \in P$.*

Proof. For every $p \in P$, it holds

$$\begin{aligned} \langle\langle p \rangle\rangle_{\mathcal{G}} &= \max_{\sigma \in \Sigma} \min_{\pi \in \Pi} \langle\langle p \rangle\rangle_{\mathcal{G}(\sigma)(\pi)} && \text{(Lemma 1)} \\ &= \max_{\sigma \in \Sigma} \min_{\pi \in \Pi} v_b(\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma)(\pi),b})(\mathbf{x}_p)) && \text{(Lemma 7)} \\ &= v_b(\max_{\sigma \in \Sigma} \min_{\pi \in \Pi} (\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma)(\pi),b})(\mathbf{x}_p))) \\ &= v_b(\mathbf{CanSol}(\mathcal{H}_{\mathcal{G},b})(\mathbf{x}_p)) && \text{(Lemma 8). } \square \end{aligned}$$

We additionally want to compute positional optimal strategies. By lemma 3 $\mathcal{E}_{\mathcal{G}}^{r\sigma\sharp}$ has a unique finite solution ρ^\sharp and it holds $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}}) = \alpha \circ \rho^\sharp$. The unique finite solution can be computed using lemma 5. A \vee -strategy σ and a \wedge -strategy π such that ρ^\sharp is a solution and thus the unique finite solution of $\mathcal{E}_{\mathcal{G}(\sigma),B(\mathcal{G})}^{r\sigma\sharp}$ and $\mathcal{E}_{\mathcal{G}(\pi),B(\mathcal{G})}^{r\sigma\sharp}$ can then be computed in linear time. Thus, it holds $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma),B(\mathcal{G})}) = \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}}) = \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\pi),B(\mathcal{G})})$ which (by lemma 9) implies $\langle\langle \cdot \rangle\rangle_{\mathcal{G}(\sigma)} = \langle\langle \cdot \rangle\rangle_{\mathcal{G}} = \langle\langle \cdot \rangle\rangle_{\mathcal{G}(\pi)}$. Thus, σ and π are positional optimal strategies. Summarizing, we have shown:

Theorem 3. *Let \mathcal{G} be a deterministic total-payoff game. The values and positional optimal strategies for both players can be computed using a strategy improvement algorithm. Each inner strategy improvement step can be done in time $\mathcal{O}(\|\mathcal{G}\|)$ (w.r.t. the uniform cost measure). The quantitative decision problem is in $\mathbf{UP} \cap \mathbf{co-UP}$. \square*

Additionally one can show that the strategy improvement algorithm presented in section 3 needs only polynomial time for solving deterministic 1-player total-payoff games.

5 Solving Stochastic Liminf-Payoff Games

In this section we characterize the values of a liminf-payoff game $\mathcal{G} = (G, r)$ by a hierarchical system of stochastic equations. Firstly, observe that the values fulfill the *optimality equations*, i.e., $\langle\langle p \rangle\rangle = \bigvee_{p' \in \{p\}^E} \langle\langle p' \rangle\rangle$ for all $p \in P_\vee$, $\langle\langle p \rangle\rangle = \bigwedge_{p' \in \{p\}^E} \langle\langle p' \rangle\rangle$ for all $p \in P_\wedge$, and $\langle\langle p \rangle\rangle = \sum_{p' \in \{p\}^E} \delta(p)(p') \cdot \langle\langle p' \rangle\rangle$ for all $p \in P_+$. The optimality equations, however, do not have unique solutions. In order to obtain the desired solution, we define the hierarchical system $\mathcal{H}_{\mathcal{G}} = (\mathcal{E}_{\mathcal{G}}, r_{\mathcal{G}})$ by

$$\mathcal{E}_{\mathcal{G}} := \{ \mathbf{b}_p = r(p) \vee e_{\mathcal{G},p}, \mathbf{x}_p = \mathbf{b}_p \wedge e_{\mathcal{G},p} \mid p \in P \}$$

and $r_{\mathcal{G}}(\mathbf{b}_p) := 1$, $r_{\mathcal{G}}(\mathbf{x}_p) := 2$ for all $p \in P$, where

$$e_{\mathcal{G},p} := \begin{cases} \bigvee_{p' \in \{p\}^E} \mathbf{x}_{p'} & \text{if } p \in P_\vee \\ \bigwedge_{p' \in \{p\}^E} \mathbf{x}_{p'} & \text{if } p \in P_\wedge \\ \sum_{p' \in \{p\}^E} \delta(p)(p') \cdot \mathbf{x}_{p'} & \text{if } p \in P_+ \end{cases}, \quad p \in P.$$

Analogously to lemma 8 we find:

Lemma 10. *Let \mathcal{G} be a stochastic liminf-payoff game. It hold*

1. $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}}) = \max_{\sigma \in \Sigma} \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma)})$ and
2. $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}}) = \min_{\pi \in \Pi} \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\pi)})$. □

Finally, we can show the correctness of our reduction. Because of Lemma 1 and 10, we only have to consider the 0-player case. There the interesting situation is within an end-component, i.e., a strongly connected component without out-going edges. In such a component every position is visited infinitely often. Thus, the value of every position in such a component is equal to the minimal immediate payoff within the component. By induction one can show that this corresponds to the canonical solution of the hierarchical system $\mathcal{H}_{\mathcal{G}}$. Since the canonical solution solves the optimality equations and there exists exactly one solution of the optimality equations where the values of all positions in end-components are fixed, we finally obtain the following lemma:

Lemma 11. *It holds $\langle\langle p \rangle\rangle = \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}})(\mathbf{x}_p)$ for every position $p \in P$.* □

We want to compute positional optimal strategies. By lemma 3 $\mathcal{E}_{\mathcal{G}}^{r_{\mathcal{G}}\#}$ has a unique finite solution $\rho^\#$ and it holds $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}}) = \alpha \circ \rho^\#$. The unique finite solution can be computed using lemma 5. A \vee -strategy σ and a \wedge -strategy π such that $\rho^\#$ is a solution and thus the unique finite solution of $\mathcal{E}_{\mathcal{G}(\sigma)}^{r_{\mathcal{G}}\#}$ and $\mathcal{E}_{\mathcal{G}(\pi)}^{r_{\mathcal{G}}\#}$ can then be computed in time $\mathcal{O}(\|\mathcal{G}\|)$ (w.r.t. the uniform cost measure). Thus, it holds $\mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\sigma)}) = \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}}) = \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}(\pi)})$ which (by lemma 11) implies $\langle\langle \cdot \rangle\rangle_{\mathcal{G}(\sigma)} = \langle\langle \cdot \rangle\rangle_{\mathcal{G}} = \langle\langle \cdot \rangle\rangle_{\mathcal{G}(\pi)}$. Thus, σ and π are positional optimal strategies. Summarizing, we have shown:

Theorem 4. *Let \mathcal{G} be a stochastic liminf-payoff game. The values and positional optimal strategies for both players can be computed using a strategy improvement algorithm. Each inner strategy improvement step can be done in time $\mathcal{O}(|P|^3)$ (w.r.t. the uniform cost measure). The quantitative decision problem is in $\mathbf{UP} \cap \mathbf{co-UP}$.* □

It is not clear whether the strategy improvement algorithm presented in section 3 needs only polynomial time for solving deterministic 1-player total-payoff games.

6 Solving Stochastic Parity Games

In this section we characterize the values of a stochastic parity game $\mathcal{G} = (G, r)$. For that, we define the hierarchical system $\mathcal{H}_{\mathcal{G}} = (\mathcal{E}_{\mathcal{G}}, r_{\mathcal{G}})$ of stochastic equations by $\mathcal{E}_{\mathcal{G}} := \{\mathbf{x}_p = e_{\mathcal{G},p} \wedge 1 \vee 0 \mid p \in P\}$ and $r_{\mathcal{G}}(\mathbf{x}_p) := r(p)$, $p \in P$. Thereby, the expression $e_{\mathcal{G},p}$ is defined as in section 5. We get:

Lemma 12. *For every position $p \in P$ it holds $\langle\langle p \rangle\rangle = \mathbf{CanSol}(\mathcal{H}_{\mathcal{G}})(\mathbf{x}_p)$.* \square

Since lemma 10 holds also for stochastic parity games (with the definition of $\mathcal{H}_{\mathcal{G}}$ given in this section), one can use the strategy improvement algorithm presented in section 3 for computing the values and positional optimal strategies. This is literally the same as for liminf-payoff games. We get:

Theorem 6.1 *Let \mathcal{G} be a stochastic parity game. The values and positional optimal strategies for both players can be computed using a strategy improvement algorithm. Each inner strategy improvement step can be done in time $\mathcal{O}(|P|^3)$ (w.r.t. the uniform cost measure). The quantitative decision problem is in $\mathbf{UP} \cap \mathbf{co-UP}$.* \square

Finally, we show that computing the values of stochastic parity games is at least as hard as computing the values of stochastic liminf-payoff games. The values of a stochastic liminf-payoff game \mathcal{G} can (by lemma 11) be characterized by the hierarchical system $\mathcal{H}_{\mathcal{G}}$ of equations (see the definition in section 5). Let us assume that $r(P) \subseteq [0, 1]$ holds. This precondition can be achieved in polynomial time by shifting and scaling r . Now one can set up a stochastic parity game $\overline{\mathcal{G}}$ such that the hierarchical system $\mathcal{H}_{\overline{\mathcal{G}}}$ (see the definition in this section) which characterizes the values of $\overline{\mathcal{G}}$, and the hierarchical system $\mathcal{H}_{\mathcal{G}}$ have the same canonical solution (besides additional variables). Thus, using lemma 11 and lemma 12 we obtain the following result:

Theorem 6.2 *Computing the values of stochastic liminf-payoff games is polynomial-time reducible to computing the values of stochastic parity games with at most 2 different ranks.* \square

Stochastic parity games are polynomial-time reducible to stochastic mean-payoff games [3]. The latter are polynomial-time equivalent to simple stochastic games [12]. Since stochastic liminf-payoff games generalize simple stochastic games, computing the values of stochastic liminf-payoff games is in fact polynomial-time equivalent to computing the values of stochastic parity games.

7 Conclusion

We introduced hierarchical systems of stochastic equations and presented a practical algorithm for solving bounded hierarchical systems of stochastic equations. The obtained results could be applied to deterministic total-payoff games and stochastic liminf-payoff games. We showed how to compute the values of as well as optimal strategies for these games through a reduction to hierarchical systems of stochastic equations — which

results in practical strategy improvement algorithms. In the deterministic case, the 1-player games can be solved in polynomial time. We also proved that the quantitative decision problems for these games are in $\mathbf{UP} \cap \mathbf{co-UP}$. As a by-product we obtained that the problem of computing the values of stochastic liminf-payoff games is polynomial-time reducible to the problem of computing the values of stochastic parity games.

References

1. Björklund, H., Vorobyov, S.G.: A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Applied Mathematics* (2007)
2. Chatterjee, K., Henzinger, T.A.: Strategy improvement and randomized subexponential algorithms for stochastic parity games. In: Durand, B., Thomas, W. (eds.) *STACS 2006*. LNCS, vol. 3884, pp. 512–523. Springer, Heidelberg (2006)
3. Chatterjee, K., Henzinger, T.A.: Reduction of stochastic parity to stochastic mean-payoff games. *Inf. Process. Lett.* 106(1), 1–7 (2008)
4. Chatterjee, K., Henzinger, T.A.: Probabilistic systems with limsup and liminf objectives. *ILC* (to appear, 2009)
5. Chatterjee, K., Jurdzinski, M., Henzinger, T.A.: Quantitative stochastic parity games. In: *SODA*, pp. 121–130 (2004)
6. Ehrenfeucht, A., Mycielski, J.: Positional strategies for mean payoff games. *IJGT* 8, 109–113 (1979)
7. Fischer, D., Grädel, E., Kaiser, L.: Model checking games for the quantitative μ -calculus. In: *STACS*, pp. 301–312 (2008)
8. Gawlitza, T., Seidl, H.: Computing game values for crash games. In: Namjoshi, K.S., Yoneda, T., Higashino, T., Okamura, Y. (eds.) *ATVA 2007*. LNCS, vol. 4762, pp. 177–191. Springer, Heidelberg (2007)
9. Gillette, D.: Stochastic games with zero stop probabilities. In: *Contributions to the Theory of Games III*, vol. 39, pp. 179–187. Princeton University Press, Princeton (1957)
10. Gimbert, H., Zielonka, W.: When can you play positionally? In: Fiala, J., Koubek, V., Kratochvíl, J. (eds.) *MFCS 2004*. LNCS, vol. 3153, pp. 686–697. Springer, Heidelberg (2004)
11. Gurvich, V., Karzanov, A., Khachiyan, L.: Cyclic games and an algorithm to find minimax cycle means in directed graphs. *U.S.S.R. Computational Mathematics and Mathematical Physics* 28(5), 85–91 (1988)
12. Gurvich, V., Miltersen, P.B.: On the computational complexity of solving stochastic mean-payoff games. *CoRR*, abs/0812.0486 (2008)
13. Jurdzinski, M.: Deciding the winner in parity games is in $\mathbf{UP} \cap \mathbf{co-UP}$. *Inf. Process. Lett.* 68(3), 119–124 (1998)
14. Liggett, T.M., Lippman, S.A.: Stochastic games with perfect information and time average payoff. *SIAM Review* 11, 604–607 (1969)
15. Maitra, A., Sudderth, W.: An operator solution of stochastic games. *Israel Journal of Mathematics* 78, 33–49 (1992)
16. Maitra, A., Sudderth, W.: Borel stochastic games with limsup payoff. *Annals of Probability* 21, 861–885 (1993)
17. Pizaruk, N.: Mean cost cyclical games. *Mathematics of Operations Research* 24(4), 817–828 (1999)
18. Thuijsman, F., Vrieze, O.: Total reward stochastic games and sensitive average reward strategies. *Journal of Optimization Theory and Applications* 98 (1998)
19. Zwick, U., Paterson, M.: The Complexity of Mean Payoff Games on Graphs. *Theoretical Computer Science (TCS)* 158(1,2), 343–359 (1996)