

A Note on the Approximation of Mean-Payoff Games

Raffaella Gentilini¹

¹University of Perugia, Italy

July 29, 2011

Abstract. We consider the problem of designing approximation schemes for the values of mean-payoff games. It was recently shown that (1) mean-payoff with rational weights scaled on $[-1, 1]$ admit additive fully-polynomial approximation schemes, and (2) mean-payoff games with positive weights admit relative fully-polynomial approximation schemes. We show that the problem of designing additive/relative approximation schemes for general mean-payoff games (i.e. with no constraint on their edge-weights) is P-time equivalent to determining their exact solution.

1 Introduction

Two-player mean-payoff games are played on weighted graphs¹ with two types of vertices: in player-0 vertices, player 0 chooses the successor vertex from the set of outgoing edges; in player-1 vertices, player 1 chooses the successor vertex from the set of outgoing edges. The game results in an infinite path through the graph. The long-run average of the edge-weights along this path, called the *value* of the play, is won by player 0 and lost by player 1.

The *decision problem* for mean-payoff games asks, given a vertex v and a threshold $\nu \in \mathbb{Q}$, if player 0 has a strategy to win a value at least ν when the game starts in v . The *value problem* consists in computing the maximal (rational) value that player 0 can achieve from each vertex v of the game. The associated (*optimal*) *strategy synthesis problem* is to construct a strategy for player 0 that secures the maximal value.

Mean-payoff games have been first studied by Ehrenfeucht and Mycielski in [1], where it is shown that memoryless (or positional) strategies suffice to achieve the optimal value. This result entails that the decision problem for these games lies in $\text{NP} \cap \text{coNP}$ [2, 18], and it was later shown to belong to² $\text{UP} \cap \text{coUP}$ [10]. Despite many efforts [19, 18, 13, 5, 6, 20, 9, 12], no polynomial-time algorithm for the mean-payoff game problems is known so far.

Beside such a theoretically engaging complexity status, mean-payoff games have plenty of applications, especially in the synthesis, analysis and verification of reactive (non-terminating) systems. Many natural models of such systems include quantitative information, and the corresponding question requires the solution of quantitative games, like mean-payoff games. Concrete examples of applications include various kinds of scheduling, finite-window online string matching, or more generally, analysis of online problems and algorithms, as well as selection with limited storage [18]. Mean-payoff games can even be used for solving the max-plus algebra $Ax = Bx$ problem, which in turn has further applications [6]. Beside their applicability to the modeling of quantitative problems, mean-payoff games have tight connections with important problems in game theory and logic. For instance, parity games [8] and the model-checking problem for the modal mu-calculus [11] are poly-time reducible to mean-payoff games [7], and it is a long-standing open question to know whether these problems are in P.

Table 1 summarize the complexity of the main algorithms for solving mean-payoff games in the literature. In particular, all *deterministic* algorithms for mean-payoff games are either pseudopolynomial (i.e., polynomial in the number of vertices $|V|$, the number of edges $|E|$, and the maximal

¹ in which every edge has a positive/negative (rational) weight

² The complexity class UP is the class of problems recognizable by unambiguous polynomial time nondeterministic Turing machines [14]. Obviously $\text{P} \subseteq \text{UP} \cap \text{coUP} \subseteq \text{NP} \cap \text{coNP}$.

Problems			
Algorithms	Decision Problem	Value Problem	Note
[12]	$\mathcal{O}(E \cdot V \cdot W)$	$\mathcal{O}(E \cdot V ^2 \cdot W \cdot (\log V + \log W))$	<i>Deterministic</i>
[18]	$\Theta(E \cdot V ^2 \cdot W)$	$\Theta(E \cdot V ^3 \cdot W)$	<i>Deterministic</i>
[20]	$\mathcal{O}(E \cdot V \cdot 2^{ V })$	$\mathcal{O}(E \cdot V \cdot 2^{ V } \cdot \log W)$	<i>Deterministic</i>
[9]	$\min(\mathcal{O}(E \cdot V ^2 \cdot W), 2^{\mathcal{O}(\sqrt{ V \cdot \log V })} \cdot \log W)$	$\min(\mathcal{O}(E \cdot V ^3 \cdot W \cdot (\log V + \log W)), 2^{\mathcal{O}(\sqrt{ V \cdot \log V })} \cdot \log W)$	<i>Randomized</i>

Table 1. Complexity of the main algorithms to solve mean-payoff games.

absolute weight W , rather than in the binary representation of W) or exponential [19, 18, 13, 12, 20, 17]. The works in [9, 3] define a *randomized* algorithm which is both subexponential and pseudopolynomial. Recently, the authors of [15, 4] show that the pseudopolynomial procedures in [18, 13, 12] can be used to design (fully) polynomial value approximation schemes for certain classes of mean-payoff games: namely, mean-payoff games with positive (integer) weights or rational weights with absolute value less or equal to 1. In this paper, we consider the problem of extending such positive approximation results for general mean-payoff games, i.e. mean-payoff games with weights arbitrary shifted/scaled on the line of rational numbers.

2 Preliminaries and Definitions

Game graphs A *game graph* is a tuple $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ where $G^\Gamma = (V, E, w)$ is a weighted graph and $\langle V_0, V_1 \rangle$ is a partition of V into the set V_0 of player-0 vertices and the set V_1 of player-1 vertices. An *infinite game* on Γ is played for infinitely many rounds by two players moving a pebble along the edges of the weighted graph G^Γ . In the first round, the pebble is on some vertex $v \in V$. In each round, if the pebble is on a vertex $v \in V_i$ ($i = 0, 1$), then player i chooses an edge $(v, v') \in E$ and the next round starts with the pebble on v' . A *play* in the game graph Γ is an infinite sequence $p = v_0 v_1 \dots v_n \dots$ such that $(v_i, v_{i+1}) \in E$ for all $i \geq 0$. A *strategy* for player i ($i = 0, 1$) is a function $\sigma : V^* \cdot V_i \rightarrow V$, such that for all finite paths $v_0 v_1 \dots v_n$ with $v_n \in V_i$, we have $(v_n, \sigma(v_0 v_1 \dots v_n)) \in E$. A *strategy-profile* is a pair of strategies $\langle \sigma_0, \sigma_1 \rangle$, where σ_0 (resp. σ_1) is a strategy for player 0 (resp. player 1). We denote by Σ_i ($i = 0, 1$) the set of strategies for player i . A strategy σ for player i is *memoryless* if $\sigma(p) = \sigma(p')$ for all sequences $p = v_0 v_1 \dots v_n$ and $p' = v'_0 v'_1 \dots v'_m$ such that $v_n = v'_m$. We denote by Σ_i^M the set of memoryless strategies of player i . A play $v_0 v_1 \dots v_n \dots$ is *consistent* with a strategy σ for player i if $v_{j+1} = \sigma(v_0 v_1 \dots v_j)$ for all positions $j \geq 0$ such that $v_j \in V_i$. Given an initial vertex $v \in V$, the *outcome* of the *strategy profile* $\langle \sigma_0, \sigma_1 \rangle$ in v is the (unique) play outcome $^\Gamma(v, \sigma_0, \sigma_1)$ that starts in v and is consistent with both σ_0 and σ_1 . Given a memoryless strategy π_i for player i in the game Γ , we denote by $G^\Gamma(\pi_i) = (V, E_{\pi_i}, w)$ the weighted graph obtained by removing from G^Γ all edges (v, v') such that $v \in V_i$ and $v' \neq \pi_i(v)$.

Mean-Payoff Games A *mean-payoff game* (MPG) [1] is an infinite game played on a game graph Γ where player 0 wins a payoff value defined as the long-run average weights of the play, while player 1 loses that value. Formally, the payoff value of a play $v_0 v_1 \dots v_n \dots$ in Γ is

$$\text{MP}(v_0 v_1 \dots v_n \dots) = \liminf_{n \rightarrow \infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} w(v_i, v_{i+1}).$$

The value *secured* by a strategy $\sigma_0 \in \Sigma_0$ in a vertex v is

$$\text{val}^{\sigma_0}(v) = \inf_{\sigma_1 \in \Sigma_1} \text{MP}(\text{outcome}^\Gamma(v, \sigma_0, \sigma_1))$$

and the (*optimal*) value of a vertex v in a mean-payoff game Γ is

$$\text{val}^\Gamma(v) = \sup_{\sigma_0 \in \Sigma_0} \inf_{\sigma_1 \in \Sigma_1} \text{MP}(\text{outcome}^\Gamma(v, \sigma_0, \sigma_1)).$$

We say that σ_0 is *optimal* if $\text{val}^{\sigma_0}(v) = \text{val}^\Gamma(v)$ for all $v \in V$. Secured value and optimality are defined analogously for strategies of player 1. Ehrenfeucht and Mycielski [1] show that mean-payoff games are *memoryless determined*, i.e., memoryless strategies are sufficient for optimality and the optimal (maximum) value that player 0 can secure is equal to the optimal (minimum) value that player 1 can achieve.

Theorem 1 ([1]). *For all MPG $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ and for all vertices $v \in V$, we have*

$$\text{val}^\Gamma(v) = \sup_{\sigma_0 \in \Sigma_0} \inf_{\sigma_1 \in \Sigma_1} \text{MP}(\text{outcome}^\Gamma(v, \sigma_0, \sigma_1)) = \inf_{\sigma_1 \in \Sigma_1} \sup_{\sigma_0 \in \Sigma_0} \text{MP}(\text{outcome}^\Gamma(v, \sigma_0, \sigma_1)),$$

and there exist two memoryless strategies $\pi_0 \in \Sigma_0^M$ and $\pi_1 \in \Sigma_1^M$ such that

$$\text{val}^\Gamma(v) = \text{val}^{\pi_0}(v) = \text{val}^{\pi_1}(v).$$

Moreover, *uniform* optimal strategies exist for both players, i.e. there exists a strategy profile $\langle \sigma_0, \sigma_1 \rangle$ that can be used to secure the optimal value independently of the initial vertex [1]. Such a strategy profile is said the *optimal strategy profile*.

The following lemma characterizes the shape of MPG values in a MPG $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ with integer weights in $\{-W, \dots, W\}$. Note that solving MPG with rational weights is P-time reducible to solving MPG with integer weights [20, 18].

Lemma 1 ([1, 20]). *Let $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ be a MPG with integer weights and let $W = \max_{(v, v') \in E} |w(v, v')|$. For each vertex $v \in V$, the optimal value $\text{val}^\Gamma(v)$ is a rational number $\frac{n}{d}$ such that $1 \leq d \leq |V|$ and $|n| \leq d \cdot W$.*

We consider the following three classical problems [18, 9] for a MPG $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$:

1. *Decision Problem.* Given a threshold $\nu \in \mathbb{Q}$ and a vertex $v \in V$, decide if $\text{val}^\Gamma(v) \geq \nu$.
2. *Value Problem.* Compute for each vertex $v \in V$ the value $\text{val}^\Gamma(v)$.
3. *(Optimal) Strategy Problem.* Given an MPG Γ , compute an (optimal) strategy profile for Γ .

Approximate Solutions for MPG

Dealing with approximate MPG solutions, we can take into consideration either absolute or relative error measures, and define the notions of *additive* and *relative MPG approximate value*.

Definition 1 (MPG additive ε -value). *Let $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ be a MPG, let $v \in V$ and consider $\varepsilon \geq 0$. The value $\widetilde{\text{val}} \in \mathbb{Q}$ is said an additive ε -value on v if and only if:*

$$|\widetilde{\text{val}} - \text{val}^\Gamma(v)| \leq \varepsilon$$

Definition 2 (MPG relative ε -value). Let $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ be a MPG, let $v \in V$ and consider $\varepsilon \geq 0$. The value $\widetilde{val} \in \mathbb{Q}$ is said a relative ε -value on v if and only if:

$$\frac{|\widetilde{val} - val^\Gamma(v)|}{|val^\Gamma(v)|} \leq \varepsilon$$

Note that additive MPG ε -values are shift-invariant. More precisely, if \widetilde{val} is an additive approximate ε -value on the vertex v in $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$, then $\widetilde{val} + k$ is an additive approximate ε -value in the MPG $\Gamma' = (V, E, w + k, \langle V_0, V_1 \rangle)$, where all the weights are shifted by k . On the contrary, additive MPG ε -values are not scale-invariant. In fact, if \widetilde{val} is a relative ε -value for v in the MPG $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$, then $k \cdot \widetilde{val}$ is a relative $\varepsilon \cdot k$ -value for v in the MPG $\Gamma' = (V, E, w \cdot k, \langle V_0, V_1 \rangle)$, where all the weights are multiplied by k . In other words, the additive error on Γ is amplified by a factor k in the scaled version of the game, Γ' . Conversely, relative MPG ε -values are scale invariant but not shift invariant.

The notions of (fully) polynomial approximation schemes w.r.t relative and additive errors are formally defined below.

Definition 3 (MPG Fully Polynomial Time Approximation Scheme (FPTAS)). An additive (resp. relative) fully polynomial approximation scheme for the MPG $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ is an algorithm \mathcal{A} such that for all $\varepsilon > 0$, \mathcal{A} computes an additive (resp. relative) ε -value in time polynomial w.r.t. the size³ of Γ and $\frac{1}{\varepsilon}$.

Definition 4 (MPG Polynomial Time Approximation Scheme (PTAS)). An additive (resp. relative) polynomial approximation scheme for the MPG $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ is an algorithm \mathcal{A} such that for all $\varepsilon > 0$, \mathcal{A} computes an additive (resp. relative) ε -value in time polynomial w.r.t. the size of Γ .

3 Mean-Payoff Games and Additive Approximation Schemes

Recently, [15] provides an additive fully polynomial scheme for the MPG value problem on graphs with rational weights in the interval $[-1, +1]$. A natural question is whether we could efficiently approximate the value in MPG with no restrictions on the weights. The next theorem shows that a generalization of the positive approximation result in [15] on MPG with arbitrary (rational) weights would indeed provide a polynomial time *exact* solution to the MPG value problem.

Theorem 2. *The MPG value problem does not admit an additive FPTAS, unless it is in P.*

Proof. We start to consider the MPG problem on graphs with integer weights. Assume that the MPG value problem on graphs with integer weights admits an additive FPTAS. Given a MPG $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$ and a vertex $v \in V$, let $|V| = n$ and $\varepsilon = \frac{1}{2n(n-1)}$. Then, our FPTAS computes an additive ε -value \widetilde{val} on v in time polynomial w.r.t. n . By Lemma 1, $val^\Gamma(v)$ is a rational number with denominator d such that $1 \leq d \leq n$. Two rationals with denominator d for which $1 \leq d \leq n$ have distance at least $\frac{1}{n(n-1)}$. Hence, there is a unique rational with denominator d , $1 \leq d \leq n$, within the interval $I = \{q \in \mathbb{Q} \mid \widetilde{val} - \varepsilon \leq q \leq \widetilde{val} + \varepsilon\}$, where $\varepsilon = \frac{1}{2n(n-1)}$. Such unique rational is $val^\Gamma(v)$ and can be easily found in time logarithmic w.r.t. n [16]. Thus, we have an algorithm \mathcal{A} to solve the value problem on Γ in time polynomial w.r.t. n . The MPG problem on graphs with rational weights can be reduced in polynomial time (w.r.t. the size of Γ) to the MPG on graphs with integer weights by simply resizing the weights in the original graph [20, 12]. \square

³ Given $\Gamma = (V, E, w, \langle V_0, V_1 \rangle)$, $size(\Gamma) = |E| + |V| + \log(W)$, where W is the maximum (absolute value) of a weight in Γ .

In view of the proof of the above theorem, we could still hope to obtain some positive approximation results for general (i.e. arbitrarily scaled) MPG by considering weaker notion of approximations with respect to FPTAS. Unfortunately, the next lemma shows that the following is sufficient to show that the MPG value problem is in P: determining in time polynomial w.r.t. the size of a given MPG Γ a k -approximate value of v , where $v \in V$ and k is an arbitrary constant.

Theorem 3. *For any constant k : If the problem of computing an additive k -approximate MPG value can be solved in polynomial time (w.r.t. the size of the input MPG), then the MPG value problem belongs to P.*

Proof. We start to consider MPG with integer weights. Let v be a vertex in the MPG $\Gamma = (V, E, w : E \mapsto [-W, W], \langle V_0, V_1 \rangle)$ and denote $|V| = n$. If $2k + 1 > (n - 2)!$, then the problem of determining $val^\Gamma(v)$ can be solved in time $\mathcal{O}(k^k) = \mathcal{O}(1)$ by simply enumerating all the strategies available to the players.

Otherwise, assume $2k + 1 \leq (n - 2)!$. Consider the game $\Gamma' = (V, E, w', \langle V_0, V_1 \rangle)$, where $\forall e \in E : w'(e) = w(e) \cdot n!$. By hypothesis, there is an algorithm \mathcal{A} that computes a k -approximate value \widetilde{val} for v on Γ' in time T polynomial w.r.t. the size of Γ' . Since $\log(W \cdot n!) = \mathcal{O}(n \cdot \log(n) + \log(W))$, T is also polynomial w.r.t. the size of Γ . By construction, $val^{\Gamma'}(v)$ is an integer. There are at most $2k + 1$ integers in the interval $[\widetilde{val} - k, \widetilde{val} + k]$, thus we have at most $2k + 1$ candidates $\{\frac{\lfloor \widetilde{val} - k \rfloor}{n!}, \dots, \frac{\lfloor \widetilde{val} + k \rfloor}{n!}\}$ for $val^\Gamma(v)$. Moreover, those candidates lie in an interval of length $L \leq \frac{2k+1}{n!} \leq \frac{(n-2)!}{n!} = \frac{1}{n \cdot (n-1)}$. The minimum distance between two possible candidates for $val^\Gamma(v)$ is $\frac{1}{n \cdot (n-1)}$.

The exact value $val^\Gamma(v)$ is thus the unique rational number with denominator of size at most n that lies in the interval $[\frac{\lfloor \widetilde{val} - k \rfloor}{n!}, \frac{\lfloor \widetilde{val} + k \rfloor}{n!}]$ and can be easily found in time logarithmic w.r.t. n [16].

The MPG problem on graphs with rational weights can be reduced in polynomial time (w.r.t. the size of Γ) to the MPG on graphs with integer weights by simply resizing the weights in the original graph [20, 12]. \square

A direct consequence of Theorem 3 is that the MPG value problem does not admit a PTAS, unless it is in P. More precisely, Theorem 2 and Theorem 3 entail a result of P-time equivalence between the exact MPG value problem and the three classes of approximations listed in Corollary 1.

Corollary 1. *The following problems are P-time equivalent:*

1. *Solving the MPG value problem.*
2. *Determining an additive FPTAS for the MPG value problem.*
3. *Determining an additive PTAS for the MPG value problem.*
4. *Computing an additive k -approximate MPG value in polynomial time, for any constant k .*

4 Mean-Payoff Games and Relative Approximation Schemes

In the recent work in [4], the authors consider the design approximation schemes for the MPG value problem based on the relative—rather than absolute—error. In particular, they provide a relative FPTAS for the MPG value problem on graphs with *nonnegative* weights. Note that negative weights are necessary to encode parity games and the μ -calculus model checking into MPG games [10]. The following theorem considers the problem of designing (fully) polynomial approximation schemes for the MPG value problem on graphs with arbitrary (positive and negative) rational weights. It shows that solving such a problem would indeed provide an exact solution to the MPG value problem, computable in time polynomial w.r.t. the size of the MPG.

Theorem 4. *The MPG value problem does not admit a relative PTAS, unless it is in P.*

Proof. Let $\Gamma = (V, E, p, \langle V_0, V_1 \rangle)$ be a MPG, let $v \in V$. Assume that MPG admit a relative PTAS and consider $\varepsilon = \frac{1}{2}$. Our assumption entails that we have an algorithm \mathcal{A} that computes a relative $\frac{1}{2}$ -value \widetilde{val} for v in time polynomial w.r.t. the size of Γ . We show that $\widetilde{val} \geq 0$ if and only if $val^\Gamma(v) \geq 0$. In other words, we show that the MPG decision problem is PTIME reducible to the computation of a relative $\frac{1}{2}$ -value. By definition of relative ε -value, for $\varepsilon = \frac{1}{2}$, we have:

$$\frac{|\widetilde{val} - val^\Gamma(v)|}{|val^\Gamma(v)|} \leq \frac{1}{2} \quad (1)$$

We have four cases to consider:

1. In the first case, assume that $\widetilde{val} - val^\Gamma(v) \geq 0$ and $\widetilde{val} \geq 0$. By contradiction, suppose $val^\Gamma(v) < 0$. Then, Disequation implies:

$$\begin{aligned} \widetilde{val} - val^\Gamma(v) &\leq \frac{1}{2} \cdot |val^\Gamma(v)| \quad \Rightarrow \\ \widetilde{val} &\leq val^\Gamma(v) + \frac{1}{2} \cdot |val^\Gamma(v)| < 0 \end{aligned}$$

that contradicts our hypothesis.

2. In the second case, assume that $\widetilde{val} - val^\Gamma(v) \geq 0$ and $\widetilde{val} < 0$. Then, $0 > \widetilde{val} \geq val^\Gamma(v)$. that contradicts our hypothesis.
3. In the third case, assume that $\widetilde{val} - val^\Gamma(v) < 0$ and $\widetilde{val} < 0$. By contradiction, suppose $val^\Gamma(v) > 0$. Then, Disequation implies:

$$\begin{aligned} val^\Gamma(v) - \widetilde{val} &\leq \frac{1}{2} \cdot |val^\Gamma(v)| \quad \Rightarrow \\ \widetilde{val} &\geq val^\Gamma(v) - \frac{1}{2} \cdot |val^\Gamma(v)| \geq 0 \end{aligned}$$

that contradicts our hypothesis.

4. The last case to consider is: $\widetilde{val} - val^\Gamma(v) < 0$ and $\widetilde{val} \geq 0$. Then, $val^\Gamma(v) > \widetilde{val} \geq 0$.

Provided a P-time algorithm for deciding whether $val^\Gamma(v) \geq 0$, a dichotomic search can be used to determine $val^\Gamma(v)$ in time polynomial w.r.t. the size of Γ [12, 20]. \square

As a direct consequence of Theorem 4 we obtain the following result of P-time equivalence involving the computation of MPG exact and approximate solutions.

Corollary 2. *The following problems are P-time equivalent:*

1. *Solving the MPG value problem.*
2. *Determining a relative FPTAS for the MPG value problem.*
3. *Determining a relative PTAS for the MPG value problem.*

References

1. A. Ehrenfeucht and J. Mycielski. International journal of game theory. *Positional Strategies for Mean-Payoff Games*, 8:109–113, 1979.
2. A. V. Karzanov and V. N. Lebedev. Cyclical games with prohibitions. *Mathematical Programming*, 60:277–293, 1993.
3. D. Andersson and S. Vorobyov. Fast algorithms for monotonic discounted linear programs with two variables per inequality. Technical Report Preprint NI06019-LAA, Isaac Newton Institute for Mathematical Sciences, Cambridge, UK, 2006.
4. Y. Boros, K. Elbassioni, M. Fouz, V. Gurvich, K. Makino, and B. Manthey. Stochastic mean-payoff games: Smoothed analysis and approximation schemes. In *Proc. of ICALP: Colloquium on Automata, Languages and Programming*, 2011.

5. A. Condon. On algorithms for simple stochastic games. In *Advances in Computational Complexity Theory*, volume 13 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 51–73. American Mathematical Society, 1993.
6. V. Dhingra and S. Gaubert. How to solve large scale deterministic games with mean payoff by policy iteration. In *Proc. Performance evaluation methodologies and tools*, article no. 12. ACM, 2006.
7. E. A. Emerson, C. Jutla, and A. P. Sistal. On model checking for fragments of the μ -calculus. In *Proc. of CAV: Computer Aided Verification*, LNCS 697, pages 385–396. Springer, 1993.
8. Y. Gurevich and L. Harrington. Trees, automata, and games. In *Proc. of STOC: Symposium on Theory of Computing*, pages 60–65. ACM, 1982.
9. H. Bjorklund and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Applied Mathematics*, 155:210–229, 2007.
10. M. Jurdzinski. Deciding the winner in parity games is in $UP \cap coUP$. *Inf. Process. Lett.*, 68(3):119–124, 1998.
11. D. Kozen. Results on the propositional mu-calculus. *Theor. Comput. Sci.*, 27:333–354, 1983.
12. L. Brim, J. Chaloupka, L. Doyen, R. Gentilini, and J-F. Raskin. Faster algorithms for mean payoff games. *Formal Methods in System Design*, 38(2):97–118, 2011.
13. N. Pisanuk. Mathematics of operations research. *Mean Cost Cyclical Games*, 4(24):817–828, 1999.
14. C. M. Papadimitriou. *Computational complexity*. Addison-Wesley, Reading, Massachusetts, 1994.
15. A. Roth, M. Balcan, A. Kalai, and Y. Mansour. On the equilibria of alternating move games. In *Proc. of SODA: Symposium on Discrete Algorithms*, pages 805–816. ACM, 2010.
16. S. Kwek and K. Mehlhorn. Optimal search for rationals. *Information Processing Letters*, 86:23–26, 2003.
17. S. Schewe. From parity and payoff games to linear programming. In *Proceedings of MFCS: Mathematical Foundations of Computer Science*, LNCS 5734, pages 675–686. Springer, 2009.
18. U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158:343–359, 1996.
19. V. A. Gurvich, A. V. Karzanov, and L. G. Kachiyan. Ussr computational mathematics and mathematical physics. *Cyclic Games and an Algorithm to Find Minmax Cycle Means in Directed Graphs*, 5(28):85–91, 1988.
20. Y. Lifshits and D. Pavlov. Potential theory for mean payoff games. *Journal of Mathematical Sciences*, 145(3):4967–4974, 2007.