



Mathematisches
Forschungsinstitut
Oberwolfach



Oberwolfach Preprints

OWP 2015 - 20

ENDRE BOROS, KHALED ELBASSIONI, VLADIMIR GURVICH
AND KAZUHISA MAKINO

A Pseudo-Polynomial Algorithm for Mean Payoff
Stochastic Games with Perfect Information and
Few Random Positions

Mathematisches Forschungsinstitut Oberwolfach gGmbH
Oberwolfach Preprints (OWP) ISSN 1864-7596

Oberwolfach Preprints (OWP)

Starting in 2007, the MFO publishes a preprint series which mainly contains research results related to a longer stay in Oberwolfach. In particular, this concerns the Research in Pairs-Programme (RiP) and the Oberwolfach-Leibniz-Fellows (OWLF), but this can also include an Oberwolfach Lecture, for example.

A preprint can have a size from 1 - 200 pages, and the MFO will publish it on its website as well as by hard copy. Every RiP group or Oberwolfach-Leibniz-Fellow may receive on request 30 free hard copies (DIN A4, black and white copy) by surface mail.

Of course, the full copy right is left to the authors. The MFO only needs the right to publish it on its website *www.mfo.de* as a documentation of the research work done at the MFO, which you are accepting by sending us your file.

In case of interest, please send a **pdf file** of your preprint by email to *rip@mfo.de* or *owlf@mfo.de*, respectively. The file should be sent to the MFO within 12 months after your stay as RiP or OWLF at the MFO.

There are no requirements for the format of the preprint, except that the introduction should contain a short appreciation and that the paper size (respectively format) should be DIN A4, "letter" or "article".

On the front page of the hard copies, which contains the logo of the MFO, title and authors, we shall add a running number (20XX - XX).

We cordially invite the researchers within the RiP or OWLF programme to make use of this offer and would like to thank you in advance for your cooperation.

Imprint:

Mathematisches Forschungsinstitut Oberwolfach gGmbH (MFO)
Schwarzwaldstrasse 9-11
77709 Oberwolfach-Walke
Germany

Tel +49 7834 979 50
Fax +49 7834 979 55
Email admin@mfo.de
URL www.mfo.de

The Oberwolfach Preprints (OWP, ISSN 1864-7596) are published by the MFO.
Copyright of the content is held by the authors.

A Pseudo-Polynomial Algorithm for Mean Payoff Stochastic Games with Perfect Information and Few Random Positions*

Endre Boros[†] Khaled Elbassioni[‡] Vladimir Gurvich[§] Kazuhisa Makino[¶]

Abstract

We consider two-person zero-sum stochastic mean payoff games with perfect information, or BWR-games, given by a digraph $G = (V, E)$, with local rewards $r : E \rightarrow \mathbb{Z}$, and three types of positions: black V_B , white V_W , and random V_R forming a partition of V . It is a long-standing open question whether a polynomial time algorithm for BWR-games exists, or not, even when $|V_R| = 0$. In fact, a pseudo-polynomial algorithm for BWR-games would already imply their polynomial solvability. In this paper, we show that BWR-games with a constant number of random positions can be solved in pseudo-polynomial time. More precisely, in any BWR-game with $|V_R| = O(1)$, a saddle point in uniformly optimal pure stationary strategies can be found in time polynomial in $|V_W| + |V_B|$, the maximum absolute local reward, and the common denominator of the transition probabilities.

1 Introduction

1.1 Basic concepts

We consider two-person zero-sum stochastic games with perfect information and mean payoff: Let $G = (V, E)$ be a digraph whose vertex-set V is partitioned into three subsets $V = V_B \cup V_W \cup V_R$ that correspond to black, white, and random positions, controlled respectively, by two players, BLACK - the *minimizer* and WHITE - the *maximizer*, and by nature. We also fix a *local reward* function $r : E \rightarrow \mathbb{Z}$, and probabilities $p(v, u)$ for all arcs (v, u) going out of $v \in V_R$. Vertices $v \in V$ and arcs $e \in E$ are called *positions* and *moves*, respectively. The game begins at time $t = 0$ in the initial position $s_0 = v_0$. In a general step, in time t , we are at position $s_t \in V$. The player who controls s_t chooses an outgoing arc $e_{t+1} = (s_t, v) \in E$, and the game moves to position $s_{t+1} = v$. If $s_t \in V_R$ then an outgoing arc is chosen with the given probability $p(s_t, s_{t+1})$. We assume, in fact without any loss of generality, that every vertex in G has an outgoing arc. (Indeed, if not, one can add loops to terminal vertices.) In general, the strategy of the player is a policy by which (s)he chooses the outgoing arcs from the vertices (s)he controls. This policy may involve the knowledge of the previous steps as well as probabilistic decisions. We call a strategy *stationary* if it does not depend on the history and *pure* if it does not involve probabilistic decisions. For this type of games, it will be enough to consider only such strategies, since these games are known to be (polynomially) equivalent [BEGM13a] to the perfect information stochastic games considered by Gillette [Gil57, LL69].

In the course of this game players generate an infinite sequence of edges $\mathbf{p} = (e_1, e_2, \dots)$ (a *play*) and the corresponding real sequence $r(\mathbf{p}) = (r(e_1), r(e_2), \dots)$ of local rewards. At the end (after infinitely many steps) BLACK pays WHITE $\phi(r(\mathbf{p}))$ amount. Naturally, WHITE's aim is to create a play which maximizes $\phi(r(\mathbf{p}))$, while BLACK tries to minimize it. (Let us note that the local reward function $r : E \rightarrow \mathbb{R}$ may have negative values, and $\phi(r(\mathbf{p}))$ may also be negative, in which case

*This paper combines the results announced in preliminary versions in [BEGM10, BEGM13c]. Part of this research was done at the Mathematisches Forschungsinstitut Oberwolfach during a stay within the "Research in Pairs" Program from July 26, 2015-August 15, 2015.

[†]RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway NJ 08854-8003; (boros@rutcor.rutgers.edu)

[‡]Masdar Institute of Science and Technology, Abu Dhabi, UAE; (kelbassioni@masdar.ac.ae)

[§]RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway NJ 08854-8003; (gurvich@rutcor.rutgers.edu)

[¶]Research Institute for Mathematical Sciences (RIMS) Kyoto University, Kyoto 606-8502, Japan; (makino@kurims.kyoto-u.ac.jp)

WHITE has to pay BLACK. Let us also note that $r(\mathbf{p})$ is a random variable since random transitions occur at positions in V_R .) Here ϕ stands for the *limiting mean payoff*

$$\phi(r(\mathbf{p})) = \liminf_{T \rightarrow \infty} \frac{\sum_{i=0}^T \mathbb{E}[r(e_i)]}{T+1}, \quad (1)$$

where $\mathbb{E}[r(e_i)]$ is the expected reward incurred at step i of the play.

As usual, a pair of (not necessarily pure or stationary) strategies is a saddle point if neither of the players can improve individually by changing her/his strategy. The corresponding $\phi(r(\mathbf{p}))$ is the value of the game with respect to initial position v_0 . Such a pair of strategies are called *optimal*; furthermore, it is called *uniformly optimal* if it provides the value of the game for any initial position.

This class of *BWR-games* was introduced in [GKK88]; see also [CH08]. The special case when $V_R = \emptyset$, *BW-games*, is also known as *cyclic* games. They were introduced for the complete bipartite digraphs in [Mou76], for all (not necessarily complete) bipartite digraphs in [EM79], and for arbitrary digraphs¹ in [GKK88]. A more special case was considered extensively in the literature under the name of *parity games* [BV01a, BV01b, CJH04], [Hal07, Jur98, JPZ06], and later generalized also to include random positions in [CH08]. A BWR-game is reduced to a *minimum mean cycle problem* in case $V_W = V_R = \emptyset$, see, e.g., [Kar78]. If one of the sets V_B or V_W is empty, we obtain a *Markov decision process* (MDP), which can be expressed as a linear program; see, e.g., [MO70]. Finally, if both are empty $V_B = V_W = \emptyset$, we get a *weighted Markov chain*. In the special case when all rewards are zero except at special positions called terminals, each of which only has a single outgoing arc forming a self-loop, we get a *stochastic terminal payoff* game, and when the self-loops have 0/1 payoffs, and every random position has only two outgoing arcs with probability 1/2 each, we obtain the so-called *simple stochastic games* (SSGs), introduced by Condon [Con92, Con93] and considered in several papers [GH08, Hal07]. In the latter games, the objective of WHITE is to maximize the probability of reaching the terminal, while BLACK wants to minimize this probability. Recently, it was shown that Gillette games with perfect information (and hence BWR-games by [BEGM13a]) are equivalent to SSGs under polynomial-time reductions [AM09]. Thus, by results of Halman [Hal07], all these games can be solved in randomized strongly subexponential time $2^{O(\sqrt{n_d} \log n_d)}$, where $n_d = |V_B| + |V_W|$ is the number of *deterministic* vertices. For BW-games several pseudo-polynomial and subexponential algorithms are known [GKK88, KL93, ZP96, Pis99, BV01a, BV01b, HBV04, BV05, BV07, Hal07, Vor08]; see also [JPZ06] for parity games. Besides their many applications (see e.g. [Lit96, Jur00]), all these games are of interest to Complexity Theory: Karzanov and Lebedev [KL93] (see also [ZP96]) proved that the decision problem “whether the value of a BW-game is positive” is in the intersection of NP and co-NP. Yet, no polynomial algorithm is known for these games, see e.g., the survey by Vorobyov [Vor08]. A similar complexity claim can be shown to hold for SSGs and BWR-games, see [AM09, BEGM13a].

1.2 Main result

The problem of developing an efficient algorithm for stochastic games with perfect information was mentioned as an open problem in the survey [RF91]. While there are numerous pseudo-polynomial algorithms known for the BW-case, it is a challenging open question whether a pseudo-polynomial algorithm exists for BWR-games, as the existence of such an algorithm would imply also the polynomial solvability of this class of games [AM09]. Our main result can be viewed as a partial solution of this problem, for the case when the number of random positions is fixed.

For a BWR-game \mathcal{G} let us denote by $n = |V_W| + |V_B| + |V_R|$ the number of positions, by $k = |V_R|$ the number of random positions, and assume that all local rewards are integral with maximum absolute value R , and all transition probabilities are rational with common denominator D . The main result of this paper is as follows.

Theorem 1 *A BWR-game \mathcal{G} can be solved in $(nDk)^{O(k)} R \cdot \text{polylog } R$ time.*

¹In fact, BW-games on arbitrary digraphs can be polynomially reduced to BW-games on bipartite graphs [BEGM13a]; moreover, the latter class can further be reduced to BW-games on complete bipartite graphs [CHKN14].

As we shall see from our proof in Section 4, one can replace $n^{O(k)}$ in the above theorem by $\text{poly}(n)\nu(\mathcal{G})^{O(k)}$, where $\nu(\mathcal{G})$ is defined by a parametrized BW-game obtained from \mathcal{G} . For stochastic terminal payoff games with t terminals, it can be shown that $\nu(\mathcal{G}) \leq t + 1$. Thus, Theorem 1 extends the results in [CdAH09, DG11, GH08, IJM12] and unifies them with the pseudo-polynomial algorithms for deterministic mean payoff games [GKK88, Pis99, ZP96].

It is important to note that the above result requires a new technique to solve BWR-games. Approximating by discounted games cannot give the claimed complexity. The example in [BEGM13b] shows that one needs to choose β exponentially (in n) close to 1 even if we have only a *single* random position.

Theorem 1 combined with the reduction in [BEF⁺11] implies that we can obtain an ε -saddle point for a BWR-game in time $\text{poly}(n, D^k, \frac{1}{\varepsilon})$.

1.3 Main idea of the proof

Our approach for proving Theorem 1 relies heavily on reducing the computation of uniformly optimal strategies for a general BWR-game to the case of *ergodic* BWR-games, i.e., those in which the (optimal) value does not depend on initial position, and showing that this special case can be solved in $(Dk)^{O(k)}R \cdot \text{poly}(n, \log R)$ time.

Our algorithm for the ergodic case is based on potential transformations which change the local reward without changing the normal form of the game; see Section 2.3. We would like to bring the local rewards to such a form that every locally optimal move is also globally optimal. Starting from zero potentials, and depending on the current locally optimal rewards (maximum for WHITE, minimum for BLACK, and average for Random), the algorithm keeps selecting a subset of positions and reducing their potentials until either the locally optimal rewards at different positions become sufficiently close to each other, or a proof of non-ergodicity is obtained in the form of a certain partition of the positions. The upper bound on the running time consists of three technical parts. The first one is to show that if the number of iterations becomes too large, then there is a large enough potential gap to ensure non-ergodicity. In the second part, we show that the range of potentials can be kept sufficiently small throughout the algorithm, namely $\|x^*\|_\infty \leq nRk(2D)^k$, and hence the range of the transformed rewards does not explode. The third part concerns the required accuracy. We show that it is enough to use an accuracy of

$$\varepsilon < (k2^k D^{2k+2})^{-1} \tag{2}$$

in order to guarantee that the algorithm either finds the exact value or discovers non-ergodicity. As we shall see below, this accuracy is also enough to find an optimal strategy in the non-ergodic case. This contrasts with the fact that there is an example with $k = 1$, in which the difference between two distinct values is exponentially small in n ; see [BEGM13b]. We also show the lower bound $D^{\Omega(k)}$ on the running time of the algorithm of Theorem 1 by constructing a series of games, with only random positions (that is, weighted Markov chains).

For a non-ergodic game, the above algorithm can be used to find the classes of positions with the largest and smallest values, but cannot find other classes, since they interact via random positions. We proceed further by following the idea in [GH08] in guessing the order of values of the random positions. However, unlike in the case for SSGs, the resulting game is still very complicated, and in fact, even if we know the actual values of all the positions, it is not clear how to find optimal strategies realizing these values. Nevertheless, we show that if the values are known, then the problem can be reduced to the ergodic case. Then to utilize the algorithm for the ergodic case, we employ the idea of *parametric search* to reduce the search space for all values into a set of k -tuples of rational intervals of cardinality at most $\nu(\mathcal{G})^{O(k)}$. Using such a set of tuples we can iteratively replace random positions by self-loops on which the local reward is a guessed value. In more details:

1. We iterate steps 2, 3 and 4 below over the random positions in the guessed order, keeping only the positions with highest rank (and hence having the same optimal value), and deleting all the other random positions. We iterate until no random positions remain, in which case we solve a BW-game.

2. We consider the situation when all the kept random positions are replaced by a self-loop with local reward parameter x ; we show that the value of any position in the resulting game defines an interval on the real line, as x changes from $-\infty$ to $+\infty$. We identify a set of at most $\nu(\mathcal{G}) + 1$ maximal intervals, in each of which, the values of different positions as functions of x are either constant or equal to x in the entire interval.
3. Since we do not know the real value of x , we guess among the identified intervals one that contains x ; for the guessed interval, we provide optimal strategies for the positions that have values above the lower bound of the interval, assuming our guess is correct.
4. Each of our guesses above yields a pair of strategies that can be verified for optimality by solving two MDPs.

Note that the number of guesses is bounded by $(k\nu(\mathcal{G}))^{O(k)}$.

Remark 1 *It is interesting to note that in the BW-model the ergodic case is as difficult as the general one [GKK88] while in the BWR-model the non-ergodic case seems more difficult, see Section 4.*

2 Preliminaries

2.1 Markov chains

Let $(G = (V, E), P)$ be a Markov chain, and let $C_1, \dots, C_k \subseteq V$ be the vertex sets of the strongly connected components (classes) of G . For $i \neq j$, let us (standardly) write $C_i \prec C_j$ if there is an arc $(v, u) \in E$ such that $v \in C_i$ and $u \in C_j$. The components C_i , such that there is no C_j with $C_i \prec C_j$ are called the *absorbing* (or *recurrent*) classes, while the other components are called *transient* or *non-recurrent*. Let $J = \{i \mid C_i \text{ is absorbing}\}$, $A = \cup_{i \in J} C_i$, and $T = V \setminus A$. For $X, Y \subseteq V$, a matrix $H \subseteq \mathbb{R}^{V \times V}$, a vector $h \subseteq \mathbb{R}^V$, we denote by $H[X; Y]$ the submatrix of H induced by X as rows and Y as columns, and by $h[X]$ the subvector of h induced by X . Let $I = I[V; V]$ be the $|V| \times |V|$ identity matrix, $e = e[V]$ be the vector of all ones of dimension $|V|$. For simplicity, we drop the indices of $I[\cdot, \cdot]$ and $e[\cdot]$, when they are understood from the context. Then $P[C_i; C_j] = 0$ if $C_j \prec C_i$, and hence in particular, $P[C_i; C_i]e = e$ for all $i \in J$, while $P[T, T]e$ has at least one component of value strictly less than 1.

The following are well-known facts about P^i and the limiting distribution $p_w = e_w P^*$, when the initial distribution is the w th unit vector e_w of dimension $|V|$ (see, e.g., [KS63]):

- (L1) $p_w[A] > 0$ and $p_w[T] = 0$;
- (L2) $\lim_{i \rightarrow \infty} P^i[V; T] = 0$;
- (L3) $\text{rank}(I - P[C_i; C_i]) = |C_i| - 1$ for all $i \in J$, $\text{rank}(I - P[T; T]) = |T|$, and $(I - P[T; T])^{-1} = \sum_{i=0}^{\infty} P[T; T]^i$;
- (L4) the absorption probabilities $y_i \in [0, 1]^V$ into a class C_i , $i \in J$, are given by the unique solution of the linear system: $(I - P[T; T])y_i[T] = P[T; C_i]e$, $y_i[C_i] = e$ and $y_i[C_j] = 0$ for $j \in J$ with $j \neq i$;
- (L5) the limiting distribution $p_w \in [0, 1]^V$ is given by the unique solution of the linear system: $p_w[C_i](I - P[C_i; C_i]) = 0$, $p_w[C_i]e = y_i(w)$, for all $i \in J$, and $p_w[T] = 0$.

2.2 BWR-games, solvability and ergodicity

A BWR-game $\mathcal{G} = (G, p, r)$ is given by a digraph $G = (V, E)$, where $V = V_W \cup V_B \cup V_R$ is a partition of the vertices; G may have loops and multiple arcs, but no terminal positions, i.e., positions of out-degree 0; $p(v, u)$ are probabilities for $v \in V_R$, $(v, u) \in E$ satisfying $\sum_u \mathbb{1}_{(v, u) \in E} p(v, u) = 1$ for all $v \in V_R$; and $r : E \rightarrow \mathbb{R}$ is a local reward function. For convenience we will also assume that $p(v, u) > 0$ whenever $(v, u) \in E$ and $v \in V_R$, and set $p(v, u) = 0$ for $(v, u) \notin E$.

Standardly, we define a strategy $s_W \in S_W$ (respectively, $s_B \in S_B$) as a mapping that assigns a position $u \in V$, such that $(v, u) \in E$, to each position $v \in V_W$ (respectively, $v \in V_B$). A pair of strategies $s = (s_W, s_B)$ is called a *situation*. Given a BWR-game $\mathcal{G} = (G, p, r)$ and situation $s = (s_B, s_W)$, we obtain a weighted Markov chain $\mathcal{G}_s = (P_s, r)$ with transition matrix P_s in the obvious way:

$$p_s(v, u) = \begin{cases} 1 & \text{if } (v \in V_W \text{ and } u = s_W(v)) \text{ or } (v \in V_B \text{ and } u = s_B(v)); \\ 0 & \text{if } (v \in V_W \text{ and } u \neq s_W(v)) \text{ or } (v \in V_B \text{ and } u \neq s_B(v)); \\ p(v, u) & \text{if } v \in V_R. \end{cases}$$

In the obtained Markov chain $\mathcal{G}_s = (P_s, r)$, we define the limiting (mean) effective payoff $\mu_{\mathcal{G}_s}(v) = \mu_s(v)$ as

$$\mu_s(v) = \sum_{w \in V} p_s^*(v, w) \sum_u p_s(w, u) r(w, u), \quad (3)$$

where $p_s^*(v, w)$ is the limit probability in \mathcal{G}_s to be at position w when the initial position is v (see Section 2.1 for more details). It is known [Gil57, LL69] that every such game has a pair of (uniformly optimal) pure stationary strategies (s_W^*, s_B^*) such that for any other pair of stationary strategies (s_W, s_B) and for every initial position v , the following hold:

$$\mu_{(s_W, s_B^*)}(v) \leq \mu_{(s_W^*, s_B^*)}(v) \leq \mu_{(s_W^*, s_B)}(v).$$

The quantity $\mu_{(s_W^*, s_B^*)}(v)$ is called the *value of the game* starting from position v , and will be denoted by $\mu_{\mathcal{G}}(v)$, or simply by $\mu(v)$ if the game is clear from the context. The value $\mu_{\mathcal{G}}(v)$ may depend on v . The BWR-game $\mathcal{G} = (G, p, r)$ is called *ergodic* if the value $\mu(v)$ is the same for all initial positions $v \in V$.

2.3 Potential transformations and canonical forms

Given a BWR-game $\mathcal{G} = (G, p, r)$, let us introduce a mapping $x : V \rightarrow \mathbb{R}$, whose values $x(v)$ will be called *potentials*, and define the transformed reward function $r_x : E \rightarrow \mathbb{R}$ as:

$$r_x(v, u) = r(v, u) + x(v) - x(u), \quad \text{where } (v, u) \in E. \quad (4)$$

It is not difficult to verify that the two normal forms of the obtained game \mathcal{G}^x and the original game \mathcal{G} , are the same, and hence the games \mathcal{G}^x and \mathcal{G} are equivalent (see [BEGM13a]). In particular, their optimal (pure stationary) strategies coincide, and their value functions also coincide: $\mu_{\mathcal{G}^x} = \mu_{\mathcal{G}}$.

It is known that for BW-games there exists a potential transformation such that, in the obtained game the locally optimal strategies are globally optimal, and hence, the value and optimal strategies become obvious [GKK88]. This result was extended for the more general class of BWR-games in [BEGM13a]: in the transformed game, the equilibrium value $\mu_{\mathcal{G}}(v) = \mu_{\mathcal{G}^x}(v)$ is given simply by the maximum local reward for $v \in V_W$, the minimum local reward for $v \in V_B$, and the average local reward for $v \in V_R$. In this case we say that the transformed game is in *canonical* form. To define this more formally, let us use the following notation throughout this section: Given functions $f : E \rightarrow \mathbb{R}$ and $g : V \rightarrow \mathbb{R}$, we define the functions $\bar{M}[f], \bar{M}[g] : V \rightarrow \mathbb{R}$.

$$\bar{M}[f](v) = \begin{cases} \max_{u|(v,u) \in E} f(v, u), & \text{for } v \in V_W, \\ \min_{u|(v,u) \in E} f(v, u), & \text{for } v \in V_B, \\ \sum_{u|(v,u) \in E} p(v, u) f(v, u), & \text{for } v \in V_R. \end{cases}$$

$$\bar{M}[g](v) = \begin{cases} \max_{u|(v,u) \in E} g(u), & \text{for } v \in V_W, \\ \min_{u|(v,u) \in E} g(u), & \text{for } v \in V_B, \\ \sum_{u|(v,u) \in E} p(v, u) g(u), & \text{for } v \in V_R. \end{cases}$$

We say that a BWR-game \mathcal{G} is in canonical form if there exist vectors $\mu, x \in \mathbb{R}^V$ such that

$$(C1) \quad \mu = \bar{M}[\mu] = \bar{M}[r_x] \text{ and,}$$

(C2) for every $v \in V_W \cup V_B$, every move $(v, u) \in E$ such that $\mu(v) = r_x(v, u)$ must also have $\mu(v) = \mu(u)$, or in other words, every locally optimal move (v, u) is globally optimal.

Canonical forms were defined for BW-games in [GKK88], and extended to BWR-games and other more classes of stochastic games in [BEGM13a]. It was shown in [GKK88] that there always exists a potential transformation x such that the optimal local rewards $\bar{M}[r_x](v)$ in a BW-game are equal to the game's value $\mu(v)$ at each vertex v . This result was extended in [BEGM13a] to the BWR-case.

Theorem 2 ([BEGM13a]) *For each BWR-game \mathcal{G} , there is a potential transformation (4) bringing \mathcal{G} to canonical form. Furthermore, in a game in canonical form we have $\mu_{\mathcal{G}} = \bar{M}[r]$.*

In this paper we will provide an algorithm for finding such a potential transformation in the ergodic case.

Proposition 1 *If in a BWR-game if there exists a constant m such that $\bar{M}[r] = m$ for all $v \in V$, then (i) every locally optimal strategy is optimal and (ii) the game is ergodic: $m = \mu(v)$ is its value for every initial position $v \in V$.*

Proof Indeed, if WHITE (BLACK) applies a locally optimal strategy then after every own move (s)he will get (pay) m , while for each move of the opponent the local reward will be at least (at most) m , and finally, for each random position the expected local reward is m . Thus, every locally optimal strategy of a player is optimal. Furthermore, if both players choose their optimal strategies then the expected local reward $\mathbb{E}[r(e_i)]$ equals m for every step i . Hence, the value of the game $\lim_{T \rightarrow \infty} \frac{1}{T+1} \sum_{i=0}^T b_i$ equals m . \square

2.4 Sufficient conditions for ergodicity of BWR-games

A digraph $G = (V = V_W \cup V_B \cup V_R, E)$ is called *ergodic* if all BWR-games $\mathcal{G} = (G, p, r)$ on G are ergodic. We will give a simple characterization of ergodic digraphs, which, obviously, provides a sufficient condition for ergodicity of BWR-games. For BW-games (that is, in case $R = \emptyset$) such characterization was given in [GL89].

In addition to partition $V = V_W \cup V_B \cup V_R$, let us consider another partition $\Pi : V = V^+ \cup V^- \cup V^0$ with the following properties:

- (i) Sets V^+ and V^- are not empty (while V^0 might be empty).
- (ii) There is no arc $(v, u) \in E$ such that either $v \in (V_W \cup V_R) \cap V^-$ and $u \notin V^-$, or $v \in (V_B \cup V_R) \cap V^+$ and $u \notin V^+$. In other words, WHITE cannot leave V^- , BLACK cannot leave V^+ , and there are no random moves leaving V^+ or V^- .
- (iii) For each $v \in V_W \cap V^+$ (respectively, $v \in V_B \cap V^-$) there is a move $(v, u) \in E$ such that $u \in V^+$ (respectively, $u \in V^-$). In other words, WHITE (BLACK) cannot be forced to leave V^+ (respectively, V^-).

In particular, the properties above imply that the induced subgraphs $G[V^+]$ and $G[V^-]$ have no terminal vertices.

A partition $\Pi : V = V^+ \cup V^- \cup V^0$ satisfying (i), (ii), and (iii) will be called a *contra-ergodic* partition for digraph $G = (V_W \cup V_B \cup V_R, E)$.

Theorem 3 *A digraph G is ergodic iff it has no contra-ergodic partition.*

Proof “Only if part”. Let $\Pi : V = V^+ \cup V^- \cup V^0$ be a contra-ergodic partition of G . Let us assign arbitrary positive probabilities to random moves such that $\sum_{u \mid (v, u) \in E} p(v, u) = 1$ for all $v \in V_R$. We still have to assign a local reward $r(v, u)$ to each move $(v, u) \in E$. Let us define $r(v, u) = 1$ whenever $v, u \in V^+$, $r(v, u) = -1$ whenever $v, u \in V^-$, and $r(v, u) = 0$ otherwise. Clearly, if the initial position is in V^+ (respectively, in V^-) then the value of the obtained game is 1 (respectively, -1). Hence, the corresponding game is not ergodic.

“If part”. Given a non-ergodic BWR-game $\mathcal{G} = (G, P, r)$, the value function $\mu_{\mathcal{G}}$ is not constant. Let μ_W and μ_B denote the maximum and minimum values, respectively. Then, let us set $V^+ = \{v \in V \mid \mu(v) = \mu_W\}$, $V^- = \{v \in V \mid \mu(v) = \mu_B\}$, and $V^0 = V \setminus (V^+ \cup V^-)$. It is not difficult to verify that the obtained partition $\Pi : V = V^+ \cup V^- \cup V^0$ is contra-ergodic. \square

The “only if part” can be strengthened as follows.

A *contra-ergodic decomposition* of \mathcal{G} is a contra-ergodic partition $\Pi : V = V^+ \cup V^- \cup V^0$ such that $\bar{M}[r](v) > \bar{M}[r](u)$ for every $v \in V^+$ and $u \in V^-$.

Proposition 2 *Given a BWR-game \mathcal{G} whose graph has a contra-ergodic partition, if $\bar{M}[r](v) > \bar{M}[r](u)$ for every $v \in V^+, u \in V^-$ then $\mu(v) > \mu(u)$ for every $v \in V^+, u \in V^-$. In particular, \mathcal{G} is not ergodic.*

Proof Let us choose a number μ such that $\bar{M}[r](v) > \mu > \bar{M}[r](u)$ for every $v \in V^+$ and $u \in V^-$; it exists, because set V of positions is finite. Obviously, properties (i), (ii), and (iii) imply that WHITE (BLACK) can guarantee more (less) than μ for every initial position $v \in V^+$ (respectively, $v \in V^-$). Hence, $\mu(v) > \mu > \mu(u)$ for every $v \in V^+$ and $u \in V^-$. \square

For example, no contra-ergodic partition can exist if $G = (V_W \cup V_B \cup V_R, E)$ is a complete tripartite digraph.

3 Ergodic BWR-games

3.1 Description of the pumping algorithm

Given a BWR-game $\mathcal{G} = (G, p, r)$, let us compute $m(v) = \bar{M}[r](v)$ for all $v \in V$. The algorithm proceeds in phases. In each phase we iteratively change the potentials such that either the range of m is reduced by a factor of $3/4$, or we discover a contra-ergodic partition. The starting local reward function in the next phase is obtained by transforming the initial rewards by the potential function obtained in the current phase. Once the range becomes small enough we stop repeating these phases, which allows us to conclude that the game is ergodic.

Now we describe informally the steps in one phase. The general step of a phase, called *pumping*, consists of reducing all potentials of the positions with m -values in the upper half of the range by the same constant δ . It will be not difficult to show that the m -values of all positions in this upper half can only decrease and by at most δ , while the m -values of all other positions can only increase and also by at most δ . In each step we will choose δ as the largest constant such that no m -value leaves the middle half (that is, second and third quarter) of the range. It can happen that $\delta = +\infty$; in this case, we can construct a contra-ergodic decomposition proving that the game is non-ergodic.

Next, we give a more formal description of phase $h = 0, 1, \dots$. Let us denote by x^h the potential at the end of phase $h-1$, where we set $x^0 = 0$. Given a function $f : V \rightarrow \mathbb{R}$, let us denote by f^+ and f^- its maximum and minimum values respectively. Throughout, we will denote by $[m] = [m^-, m^+]$ and $[r] = [r^-, r^+]$ the range of functions m and r , respectively, and let $M = m^+ - m^-$ and $R = r^+ - r^-$. Given a potential function $x : V \rightarrow \mathbb{R}$, we will denote by m_x, M_x , etc., the functions m, M , etc., with r is replaced by the transformed reward r_x in the above definitions. Given a subset $I \subseteq [m]$, let $V_x(I) = \{v \in V \mid m_x(v) \in I\} \subseteq V$. In the following algorithm, set I will always be a closed or semi-closed interval within $[m]$ (and not within $[m_x]$). For convenience, we will write $(\cdot)_{x^h}$ as $(\cdot)_h$, where (\cdot) could be m, r, r^+ , etc (e.g., $m_h^- = m_{x^h}^-, m_h^+ = m_{x^h}^+$).

Let $m_h^- = t_0 < t_1 < t_2 < t_3 < t_4 = m_h^+$ be thresholds defined by

$$t_i = m_h^- + \frac{i}{4}M_h, \quad i = 0, 1, 2, 3, 4, \quad \text{where } M_h = m_h^+ - m_h^-. \quad (5)$$

Let us introduce $\hat{x}(v) = x(v) - \delta$ for $v \in V_x[t_2, t_4]$ and $\hat{x}(v) = x(v)$ for $v \in V_x[t_0, t_2]$. Let us then introduce the notation m_x^δ for $m_{\hat{x}}$. This notation may look complicated but is necessary since $m_{\hat{x}}$ depends on x and δ in a complicated way.

It is clear that δ can be computed in linear time, and we have $m_x^\delta(v) \geq t_1$ for all $v \in V_x[t_2, t_4]$ and $m_x^\delta(v) \leq t_3$ for all $v \in V_x[t_0, t_2)$, where $m_x^\delta(v)$ is the new value of $m_x(v)$ after all potentials in $V_x[t_2, t_4]$ have been reduced by δ . The value of m_x^δ is given by the following formula:

$$\begin{aligned}
& \max \left\{ \begin{array}{l} \max_{\substack{(v,u) \in E, \\ u \in V_x[t_2, t_4]}} \{r_x(v, u)\}, \max_{\substack{(v,u) \in E, \\ u \in V_x[t_0, t_2]}} \{r_x(v, u)\} - \delta \end{array} \right\} \quad \text{for } v \in V_W \cap V_x[t_2, t_4], \\
& \min \left\{ \begin{array}{l} \min_{\substack{(v,u) \in E, \\ u \in V_x[t_2, t_4]}} \{r_x(v, u)\}, \min_{\substack{(v,u) \in E, \\ u \in V_x[t_0, t_2]}} \{r_x(v, u)\} - \delta \end{array} \right\} \quad \text{for } v \in V_B \cap V_x[t_2, t_4], \\
& \sum_{\substack{(v,u) \in E, \\ u \in V}} p(v, u) r_x(v, u) - \delta \quad \sum_{\substack{(v,u) \in E, \\ u \in V_x[t_0, t_2]}} p(v, u) \quad \text{for } v \in V_R \cap V_x[t_2, t_4], \\
& \max \left\{ \begin{array}{l} \max_{\substack{(v,u) \in E, \\ u \in V_x[t_2, t_4]}} \{r_x(v, u)\} + \delta, \max_{\substack{(v,u) \in E, \\ u \in V_x[t_0, t_2]}} \{r_x(v, u)\} \end{array} \right\} \quad \text{for } v \in V_W \cap V_x[t_0, t_2), \\
& \min \left\{ \begin{array}{l} \min_{\substack{(v,u) \in E, \\ u \in V_x[t_2, t_4]}} \{r_x(v, u)\} + \delta, \min_{\substack{(v,u) \in E, \\ u \in V_x[t_0, t_2]}} \{r_x(v, u)\} \end{array} \right\} \quad \text{for } v \in V_B \cap V_x[t_0, t_2), \\
& \sum_{\substack{(v,u) \in E, \\ u \in V}} p(v, u) r_x(v, u) + \delta \quad \sum_{\substack{(v,u) \in E, \\ u \in V_x[t_2, t_4]}} p(v, u) \quad \text{for } v \in V_R \cap V_x[t_0, t_2).
\end{aligned}$$

Note that $|m_x^\delta(v) - m_x(v)| \leq \delta$. It is also important to note that $\delta \geq M_h/4$. Indeed, the m -values of positions from $V_x[t_2, t_4]$ cannot increase, while those of positions from $V_x[t_0, t_2)$ cannot decrease. Each of them would have to traverse a distance of at least $M_h/4$ before it can reach the border of the interval $V_x[t_1, t_3]$. Moreover, if after some iteration one of the sets $V_x[t_0, t_1)$ or $V_x(t_3, t_4]$ becomes empty then the range of m_x is reduced at least by 25%.

Procedure $\text{PUMP}(\mathcal{G}, \varepsilon)$ below tries to reduce any BWR-game \mathcal{G} by a potential transformation x into one in which $M_x \leq \varepsilon$. Two subroutines are used in the procedure. $\text{REDUCE-POTENTIALS}(\mathcal{G}, x)$ replaces the current potential x with another potential with a sufficiently small norm; see Lemma 4 in Subsection 3.4. This reduction is needed since without it the potentials and, hence, the transformed local rewards too, may grow exponentially; see the analysis of R_h and N_h in the proof of Lemma 5. The second routine $\text{FIND-PARTITION}(\mathcal{G}, x)$ uses the current potential vector x to construct a contra-ergodic decomposition of \mathcal{G} (cf. line 19 of the algorithm below), see Subsection 3.3 for the details.

Note the algorithm can terminate in 3 ways:

1. $\delta = +\infty$ in some iteration. In this case, $(V^+ = V_x[t_0, t_2), V^- = V_x[t_2, t_4], V^0 = \emptyset)$ is a contra-ergodic partition as can be checked from the definition of m_x^δ .
2. The number of pumping iterations performed in some phase h reaches

$$N_h = \left\lceil \frac{8n^2 R_h D^k}{M_h} \right\rceil + 1, \quad (6)$$

where $R_h = r_h^+ - r_h^-$, and yet the range of m_h is not reduced. In this case, we can find a contra-ergodic decomposition by the second subroutine $\text{FIND-PARTITION}(\mathcal{G}, x)$; see Lemma 3.

3. $M_x \leq \varepsilon$ and ε satisfies (2). In this case we can prove that a pair of locally optimal strategies with respect to r_x is optimal in the game \mathcal{G} ; we show this in Section 3.5.

To describe the main idea of the analysis, we will first argue in Section 3.2 that the algorithm terminates in finite time if the considered BWR-game is ergodic, even if we set $N_0 = +\infty$. In the following section, this argument will be made quantitative with the precise bound on the running time. In Section 3.6, we will show that our bound on the running time is tight since we can construct a weighted Markov chain (that is, an R-game) providing an exponential lower bound in k .

3.2 Proof of finiteness for the ergodic case

To simplify notation, we can assume without loss of generality (by shifting and scaling the local rewards) that the range of m is $[0, 1]$, that is, $t_i = \frac{i}{4}$, for $i = 0, \dots, 4$, and that the initial potential $x^0 = 0$. Suppose indirectly that phase 0 of the algorithm does not terminate, that is, $V[0, \frac{1}{4})$ and $V(\frac{3}{4}, 1]$ never become empty and the m -range never gets smaller than $\frac{1}{2}$.

Consider the infinite sequence of iterations and denote by $V^- \subseteq V$ (respectively, by $V^+ \subseteq V$) the set of vertices that were pumped just finitely many times (respectively, always but finitely many times); in other words, $m_x(v) \in [\frac{1}{2}, 1]$ if $v \in V^+$ (respectively, $m_x(v) \in [0, \frac{1}{2})$ if $v \in V^-$) for all but finitely many iterations. By the above assumption, we have

$$\emptyset \neq V[0, \frac{1}{4}) \subseteq V^- \subseteq V[0, \frac{1}{2}), \quad (7)$$

$$\emptyset \neq V(\frac{3}{4}, 1] \subseteq V^+ \subseteq V[\frac{1}{2}, 1]. \quad (8)$$

Proposition 3 *The partition $\Pi : V = V^+ \cup V^- \cup V^0$, where $V^0 = V \setminus (V^+ \cup V^-)$, is a contra-ergodic decomposition.*

Proof First, let us check properties (i), (ii), (iii) of Section 2.4. First, (i) follows by (7) and (8). Let us observe that the transformed local reward on any arc leaving V^+ (resp., V^-) are $= +\infty$ (resp., $= -\infty$). This implies (ii) and (iii).

Finally, it follows from (7) and (8) also that $m_x(v) > \frac{1}{2}$ for all $v \in V^+$, while $m_x(v) \leq \frac{1}{2}$ for all $v \in V^-$. Hence the claim follows by Proposition 2. \square

In other words, our algorithm is finite for the ergodic BWR-games. Below we shall give an upper bound for the number of iterations a vertex can “oscillate” in $[0, \frac{1}{4})$ or $(\frac{3}{4}, 1]$ before it finally enters $[\frac{1}{4}, \frac{3}{4}]$ (to stay there forever).

Remark 2 *For general stochastic games (with imperfect information), the value might not exist [Gil57]. Nevertheless, the pumping algorithm can be extended [BEGM14]. This algorithm either (i) certifies that the game is 24ε -ergodic, that is, it finds a potential transformation x such that all local values $m_x(v)$ are within an interval of length 24ε , or (ii) presents a contra-ergodic partition similar to the one described in Section 2.4; in particular, we obtain two positions u and v such that $|\mu(u) - \mu(v)| > \varepsilon$.*

3.3 Finding a contra-ergodic decomposition: FIND-PARTITION(\mathcal{G}, x)

We assume throughout this section that we are inside phase h of the algorithm, which started with a potential x^h , and we proceed to step 19. For simplicity, we assume that the phase starts with local reward function $r = r_h$ and hence² $x^h = 0$. Given a potential vector x , we use the following notation:

$$\text{EXT}_x = \{(v, u) \in E \mid v \in V_B \cup V_W \text{ and } r_x(v, u) = m_x(v)\},$$

and recall that $x^- = \min\{x(v) \mid v \in V\}$. Let $t_l \leq 0$ be the largest value satisfying the following conditions:

- (i) there are no arcs $(v, u) \in E$ with $v \in V_W \cup V_R$, $x(v) \geq t_l$ and $x(u) < t_l$;
- (ii) there are no arcs $(v, u) \in \text{EXT}_x$ with $v \in V_B$, $x(v) \geq t_l$ and $x(u) < t_l$.

Let $X = \{v \in V \mid x(v) \geq t_l\}$. In words, X is the set of positions with potential as close to 0 as possible, such that no white or random position in X has an arc crossing to $V \setminus X$, and no black position has an extremal arc crossing to $V \setminus X$. Similarly, define $t_u \geq x^-$ to be the smallest value satisfying the following conditions:

²in particular, note that $r_x(v, u)$ and $m_x(v)$ are used, for simplicity of notation, to actually mean $r_{x+x^h}(v, u)$ and $m_{x+x^h}(v)$, respectively.

Algorithm 1 PUMP(\mathcal{G}, ε)

Input: A BWR-game $\mathcal{G} = (G = (V, E), P, r)$ and a desired accuracy ε

Output: Either a potential $x : V \rightarrow \mathbb{R}$ s.t. $|m_x(v) - m_x(u)| \leq \varepsilon$ for all $u, v \in V$, or a contra-ergodic decomposition

```
1: let  $h := 0; x = x^h := 0; i := 1$ 
2: let  $t_0, t_1, \dots, t_4$ , and  $N_h$  be as given by (5) and (6)
3: while  $i < N_h$  do
4:   if  $M_x \leq \varepsilon$  then
5:     return  $x$  and halt
6:   end if
7:    $\delta := \max\{\delta' \mid m_x^{\delta'}(v) \geq t_1 \text{ for all } v \in V_x[t_2, t_4] \text{ and } m_x^{\delta'}(v) \leq t_3 \text{ for all } v \in V_x[t_0, t_2]\}$ 
8:   if  $\delta = \infty$  then
9:     return the contra-ergodic partition ( $V^+ = V_x[t_0, t_2]$ ,  $V^- = V_x[t_2, t_4]$ ,  $V^0 = \emptyset$ )
10:  end if
11:   $x(v) := x(v) - \delta$  for all  $v \in V_x[t_2, t_4]$ 
12:  if  $V_x[t_0, t_1] = \emptyset$  or  $V_x[t_3, t_4] = \emptyset$  then
13:     $x^{h+1} := x := \text{REDUCE-POTENTIALS}(\mathcal{G}, x); h = h + 1; i := 1$ 
14:    recompute the thresholds  $t_0, t_1, \dots, t_4$  and  $N_h$  using (5) and (6)
15:  else
16:     $i := i + 1;$ 
17:  end if
18: end while
19:  $(V^+, V^-, V^0) := \text{FIND-PARTITION}(\mathcal{G}, x)$ 
20: return the contra-ergodic partition ( $V^+, V^-, V^0$ )
```

(iii) there are no arcs $(v, u) \in E$ with $v \in V_B \cup V_R$, $x(v) \leq t_u$ and $x(u) > t_u$;

(iv) there are no arcs $(v, u) \in \text{EXT}_x$ with $v \in V_W$, $x(v) \leq t_u$ and $x(u) > t_u$,

and let $Y = \{v \in V \mid x(v) \leq t_u\}$. Note that both t_l and t_u trivially exist. Note also that the sets X and Y can be computed in $O(|V| \log |V| + |E|)$ time. Define $\theta = \min\{p(v, u) \mid (v, u) \in E\}$ to be the smallest positive transition probability. Clearly, $\theta \geq \frac{1}{D}$.

Lemma 1 *It holds that $\max\{-t_l, t_u - x^-\} \leq nR_h \left(\frac{1}{\theta}\right)^k$.*

To prove Lemma 1, we need the following lemma.

Lemma 2 *Consider any move $(v, u) \in E$ and let x be the current potential. Then*

$$x(u) \geq \begin{cases} x(v) - (m_h^+ - r_h^-) & \text{if either } (v \in V_W \text{ and } (v, u) \in E) \text{ or } (v \in V_B \text{ and } (v, u) \in \text{EXT}_x) \\ \frac{x(v) - (m_h^+ - r_h^-)}{\theta} & \text{if } v \in V_R \text{ and } (v, u) \in E, \end{cases}$$

and

$$x(u) \leq \begin{cases} x(v) + r_h^+ - m_h^- & \text{if either } (v \in V_B \text{ and } (v, u) \in E) \text{ or } (v \in V_W \text{ and } (v, u) \in \text{EXT}_x) \\ \frac{x(v) + r_h^+ - m_h^- - (1-\theta)x^-}{\theta} & \text{if } v \in V_R \text{ and } (v, u) \in E. \end{cases}$$

Proof We only consider the case for $v \in V_R$, as the other claims are obvious from the definitions. For the first claim, assume that $x(v) \geq x(u)$, since otherwise there is nothing to prove. Then from $m_x(v) \leq m_h^+$, it follows that

$$\begin{aligned} m_h^+ - r_h^- &\geq m_h^+ - \sum_{u'} p(v, u') r_h(v, u') \\ &\geq p(v, u)(x(v) - x(u)) + \sum_{u' \neq u} p(v, u')(x(v) - x(u')) \\ &\geq \theta(x(v) - x(u)) + x(v)(1 - \theta), \end{aligned}$$

from which our claim follows. The other claim can be proved by a similar argument (by replacing $x(\cdot)$ by $x^- - x(\cdot)$ and $m_h^+ - r_h^-$ by $r_h^+ - m_h^-$). \square

Proof of Lemma 1. By definition of X , for every position $v \in X$ there must exist (not necessarily distinct) positions $v_0, v_1, \dots, v_{2j} = v \in X$, $j \leq |X|$, such that $x(v_0) = 0$, and for $i = 1, 2, \dots, j$, $x(v_{2i}) \geq x(v_{2i-1})$, and either $((v_{2i-2}, v_{2i-1}) \in E$ and $v_{2i-2} \in V_W \cup V_R$) or $((v_{2i-2}, v_{2i-1}) \in \text{EXT}_x$ and $v_{2i-2} \in V_B$). Among the even-numbered positions, let $v_{2i_1-2}, \dots, v_{2i_l-2}$ be the ones belonging to V_R , and assume without loss of generality that $l > 0$ and $i_1 < i_2 < \dots < i_l$. Using Lemma 2, we obtain the following inequality by a telescoping sum:

$$x(v_{2i_{q+1}-2}) \geq x(v_{2i_q-2}) - (i_{q+1} - i_q - 1)(m_h^+ - r_h^-), \text{ for } q = 1, \dots, l-1, \quad (9)$$

and $x(v_{2i_1-2}) \geq -(i_1 - 1)(m_h^+ - r_h^-)$.

Now applying Lemma 2 to the pair $v_{2i_q-2} \in V_R$ and $v_{2i_{q-1}-2}$, for $q = 1, \dots, l-1$, and using (9) we obtain:

$$x_{q+1} \geq \frac{x_q}{\theta} - \left(\frac{1}{\theta} + i_{q+1} - i_q - 1 \right) (m_h^+ - r_h^-), \quad x_1 \geq -(i_1 - 1)(m_h^+ - r_h^-), \quad (10)$$

where we write, for convenience, $x_q = x(v_{2i_q-2})$, for $q = 1, \dots, l$. Iterating, we obtain:

$$x_l \geq - \left(\frac{i_1 - 1}{\theta^{l-1}} + \sum_{q=2}^l \frac{1}{\theta^{l-q}} \left(\frac{1}{\theta} + i_q - i_{q-1} - 1 \right) \right) (m_h^+ - r_h^-).$$

Combining this with the inequality $x(v) \geq \frac{x_l}{\theta} - (\frac{1}{\theta} + j - i_l)(m_h^+ - r_h^-)$ and using $\theta < 1$, we get

$$x(v) \geq -\frac{j}{\theta^l} (m_h^+ - r_h^-) \geq -\frac{1}{\theta^k} |X| (m_h^+ - r_h^-).$$

Similarly, one can prove for any $v \in Y$ that $x(v) \leq x^- + \frac{1}{\theta^k} |Y| (r_h^+ - m_h^-)$, and the lemma follows. \square

The correctness of the algorithm follows from the following lemma.

Lemma 3 *Suppose that pumping is performed for $N_h \geq 2nT_h + 1$ iterations, where $T_h = \frac{4nR_h}{M_h\theta^k}$, and neither the set $V_x[t_0, t_1]$ nor $V_x[t_3, t_4]$ becomes empty. Let $V^- = X$ and $V^+ = Y$ be the sets constructed as above, and $V^0 = V \setminus (X \cup Y)$. Then $V^+ \cup V^- \cup V^0$ is a contra-ergodic decomposition.*

Proof We pump in each iteration by $\delta \geq \frac{M_h}{4}$. Furthermore, our formula for δ implies that once a vertex enters the region $V_x[t_1, t_3]$, it never leaves this region. In particular, there are vertices $v_0 \in X \cap V_x[t_0, t_1]$ and $v_n \in Y \cap V_x[t_3, t_4]$ with $x(v_0) = 0$ and $x(v_n) = x^-$.

For a vertex $v \in V$, let $N(v)$ denote the number of times the vertex was pumped. Then $N(v_0) = 0$ and $N(v_n) = N_h$.

We claim that $N(v) \leq T_h$ for any $v \in X$, and $N(v) \geq N_h - T_h$ for any $v \in Y$ (i.e., every vertex in X was *not* pumped in *all* steps but at most T_h , and every vertex in Y was pumped in *all* steps but at most T_h). Indeed, if $v \in X$ (respectively, $v \in Y$) was pumped greater than (respectively, less than) T_h times then $x(v) - x(v_0) \leq -\frac{nR_h}{\theta^k}$ (respectively, $x(v_n) - x(v) \leq -\frac{nR_h}{\theta^k}$), in contradiction to Lemma 1.

Since $N_h > 2T_h$, it follows that $X \cap Y = \emptyset$. Furthermore, among the first $2nT_h + 1$ iterations, in at most nT_h iterations some vertex $v \in X$ was pumped, and in at most nT_h iterations some vertex in Y was not pumped. Thus, there must exist an iteration at which every vertex $v \in X$ was not pumped and every vertex $v \in Y$ was pumped. At that particular iteration, we must have $X \subseteq V_x[t_0, t_2]$ and $Y \subseteq V_x[t_2, t_4]$, and hence $m_x(v) < t_2$ for every $v \in X$ and $m_x(v) \geq t_2$ for every $v \in Y$. By the way the sets X and Y were constructed, we can easily see that X and Y will continue to have this property till the end of the N_h iterations, and hence they induce a contra-ergodic partition. The lemma follows. \square

3.4 Potential reduction: REDUCE-POTENTIALS(\mathcal{G}, x)

One problem that arises during the pumping procedure is that the potentials can increase exponentially in the number of phases, making our bounds on the number of iterations per phase also exponential in n . For the BW-case Pisaruk [Pis99] solved this problem by giving a procedure that reduces the range of the potentials after each round, while keeping all its desired properties needed for the running time analysis.

Pisaruk's potential reduction procedure can be thought of as a combinatorial procedure for finding an *extreme point* of a polyhedron, given a point in it. Indeed, given a BWR-game and a potential x , let us assume without loss of generality, by shifting the potentials if necessary, that $x \geq 0$, and let $E' = \{(v, u) \in E \mid r_x(v, u) \in [m_x^-, m_x^+], v \in V_B \cup V_W\}$, where r is the *original* local reward function. Then the following polyhedron

$$\Gamma_x = \left\{ x' \in \mathbb{R}^V \mid \begin{array}{ll} m_x^- \leq r(v, u) + x'(v) - x'(u) \leq m_x^+, & \forall (v, u) \in E' \\ r(v, u) + x'(v) - x'(u) \leq m_x^+, & \forall v \in V_W, (v, u) \in E \setminus E' \\ m_x^- \leq r(v, u) + x'(v) - x'(u), & \forall v \in V_B, (v, u) \in E \setminus E' \\ m_x^- \leq \sum_{u \in V} p(v, u)(r(v, u) + x'(v) - x'(u)) \leq m_x^+, & \forall v \in V_R \\ x'(v) \geq 0 & \forall v \in V \end{array} \right\}.$$

is non-empty, since $x \in \Gamma_x$. Moreover, Γ_x is pointed, and hence, it must have an extreme point. Let us remark that given a feasible point x , an extreme point can be computed in $O(n^2|E|)$ time (see, e.g., [Sch03]).

Lemma 4 *Consider a BWR-game in which all rewards are integral with range $R = r^+ - r^-$, and probabilities $p(v, u)$ are rational with common denominator D , and let $k = |V_R|$. Then any extreme point x^* of Γ_x satisfies $\|x^*\|_\infty \leq nRk(2D)^k$.*

Proof Consider such an extreme point x^* . Then x^* is uniquely determined by a system of n linearly independent equations chosen from the given inequalities. Thus there exist subsets $V' \subseteq V$, $V'_R \subseteq V_R$ and $E'' \subseteq E$ such that $|V'| + |V'_R| + |E''| = n$, x^* is the unique solution of the subsystem $x'(v) = 0$ for all $v \in V'$, $x'(v) - x'(u) = m_x^* - r(v, u)$ for $(v, u) \in E''$, and $x'(v) - \sum_{u \in V} p(v, u)x'(u) = m_x^* - \sum_{u \in V} p(v, u)r(v, u)$ for $v \in V'_R$, where m_x^* stands for either m_x^- or m_x^+ .

Note that all variables $x'(v)$ must appear in this subsystem, and that the underlying undirected graph of the digraph $G' = (V, E'')$ must be a forest (otherwise the subsystem does not uniquely fix x^* , or it is not linearly independent).

Consider first the case $V_R = \emptyset$. For $i \geq 0$, let V_i be the set of vertices of V at (undirected) distance i from V' (observe that i is finite for every vertex). Then we claim by induction on i that $x^*(v) \leq i\gamma$ for all $v \in V_i$, where $\gamma = \max\{m_x^+ - r^-, r^+ - m_x^-\}$. This is trivially true for $i = 0$. So let us assume that it is also true for some $i > 0$. For any $v \in V_{i+1}$, there must exist either an arc (v, u) or an arc (u, v) where $u \in V_{i-1}$. In the former case, we have $x^*(v) = x^*(u) + m_x^* - r(v, u) \leq i\gamma + m_x^+ - r^- \leq (i+1)\gamma$. In the latter case, we have $x^*(v) = x^*(u) - (m_x^* - r(u, v)) \leq i\gamma + r^+ - m_x^- \leq (i+1)\gamma$.

Now suppose that $|V_R| > 0$. For each connected component D_l in the forest G' , let us fix a position u_l from V' if $D_l \cap V' \neq \emptyset$; otherwise, v_l is chosen arbitrarily. For every position $v \in D_l$ let \mathcal{P}_v be a (not necessary directed) path from v to v_l . Thus, we can write $x'(v)$ uniquely as

$$x'(v) = x'(v_l) + \ell_{v,1}m_x^+ + \ell_{v,2}m_x^- + \sum_{(u', u'') \in \mathcal{P}_v} \ell_{v, u', u''} r(u', u''), \quad (11)$$

for some $\ell_{v,1}, \ell_{v,2} \in \mathbb{Z}$, and $\ell_{v, u', u''} \in \{-1, 1\}$. Thus if $x^*(v_l) = 0$ for some component D_l , then by a similar argument as above, $x^*(v) \leq \gamma|D_l|$ for every $v \in D_l$.

Note that, up to this point, we have used all equations corresponding to arcs in G' and to vertices in V' . The remaining set of $|V'_R|$ equations should uniquely determine the values of the variables in any component which has no position in V' . Substituting the values of $x'(v)$ from (11), for the positions in any such component, we end-up with a linearly independent system on $k' = |V'_R|$

variables $Ax = b$, where A is a $k' \times k'$ matrix in which each entry is at most 1 in absolute value and the sum of each row is at most 2 in absolute value, and $\|b\|_\infty \leq n(R + M_x) \leq 2nR$.

The rest of the argument follows (in a standard way) by Cramer's rule. Indeed, the value of each component in the solution is given by Δ'/Δ , where Δ is the determinant of A and Δ' is the determinant of a matrix obtained by replacing one column of A by b . We upper bound Δ' by $k'\|b\|_\infty\Delta_{max}$, where Δ_{max} is the maximum absolute value of a subdeterminant of A of size $k' - 1$. To bound Δ_{max} , let us consider such a subdeterminant with rows $a_1, \dots, a_{k'-1}$, and use Hadamard's inequality:

$$\Delta' \leq \prod_{i=1}^{k'-1} \|a_i\| \leq 2^{k'-1},$$

since $\|a_i\|_1 \leq 2$, for all i . To lower bound Δ , we note that $D^{k'}\Delta$ is a non-zero integer, and hence has absolute value at least 1. Combining the above inequalities, the lemma follows. \square

Note that any point $x' \in \Gamma_x$ satisfies $M_{x'} \subseteq M_x$, and hence replacing x by x^* does not increase the range of m_x .

3.5 Running time analysis

From Lemmas 3 and 4, we can conclude the following bound.

Lemma 5 *Procedure PUMP(\mathcal{G}, ε) terminates in $O(n^2|E|R(nk(2D))^k \frac{1}{\varepsilon} + \log R)$ time.*

Proof We note the following:

1. By (6), the number of iterations per phase h is at most $N_h = \frac{8n^2R_h}{M_h\theta^k} + 1$.
2. Each iteration requires $O(|E|)$ time, and the end of a phase we need additional $O(n^2|E|)$ time (which is required for REDUCE-POTENTIALS).
3. By Lemma 4, for any $(v, u) \in E$, we have $r_x(v, u) = r(v, u) + x(v) - x(u) \leq r(v, u) + 2nk(2D)^kR$, and similarly, $r_x(v, u) \geq r(v, u) - 2nk(2D)^kR$. In particular, $R_h \leq (1 + 4nk(2D)^k)R$ at the beginning of each phase h in the procedure.

Since $M_h \leq \frac{3}{4}M_{h-1}$ for $h = 1, 2, \dots$, the maximum number of such phases until we reach the required accuracy is at most $H = \log_{4/3}(\frac{M_0}{\varepsilon})$. Putting all the above together, we get that the total running time is at most

$$(1 + 4nk(2D)^k)R \sum_{h=0}^H \frac{8n^2|E|}{M_h} + O(n^2|E|)H.$$

Noting that $M_0 \leq R$ and $M_H \geq \varepsilon$, the lemma follows. \square

We now give an upper bound on the required accuracy ε .

Lemma 6 *Consider a Markov chain $\mathcal{M} = (G, P)$ with k positions in which all the entries of the transition matrix P are rational numbers with common denominator D . Assume that \mathcal{M} has only one absorbing class. Let p^* be the limiting distribution starting from any position. Then $p^*(v)$, for all positions v , are rational numbers with common denominator at most $\sqrt{k}2^{k/2}D^k$.*

Proof Consider the system of equations defining p^* : $p^*(I - P) = 0$ and $\mathbf{1}^T p^* = 1$. Let us multiply by D all the equations of this system and note that all coefficients of the resulting system of equations are integers in $[-D, D]$. The value of each component in the solution is given by Δ'/Δ , where Δ is the determinant of a non-singular $k \times k$ submatrix of the resulting system, and Δ' is the determinant of a $k \times k$ integer matrix. It follows by Hadamard's inequality that $\Delta \leq \sqrt{k}2^{k/2}D^k$, as the sum of the entries of each row of P is one. \square

Theorem 4 *When procedure $PUMP(\mathcal{G}, \varepsilon)$ is run with ε as in (2), it either outputs a potential vector x such $m_x(v)$ is constant for all $v \in V$, or finds a contra-ergodic partition. The total running time is $\text{poly}(n)(2D)^{O(k)} R \log R$.*

Proof Suppose that the game is not ergodic and let us fix an optimal situation (pair of optimal strategies) s . Then there must exist at least two absorbing classes in the obtained weighted Markov chain which have different values. Consider such an absorbing class and contract all arcs (v, u) where $v \in V_W \cup V_B$. We obtain an absorbing Markov chain on a subset of V_R . We now can apply Lemma 6 to conclude that the value of any position in this class is a rational number with denominator at most $\sqrt{k}2^{k/2}D^{k+1}$. Consequently, the difference between any two different values of absorbing classes of \mathcal{G}_s is at least $(k2^k D^{2k+2})^{-1}$. If $PUMP(\mathcal{G}, \varepsilon)$ terminates with $M_x < \varepsilon$ then we conclude by Lemma 6, analogously to the previous arguments, that all $m_x(v)$ values are the same. This implies by Proposition 1 that this is the optimal value and the locally optimal moves are globally optimal. The running time bound follows now from Lemma 5. \square

3.6 Lower bound for ergodic games

Note that if $|V_R| = 0$, there is a simple example with only one player showing that the running time of the pumping algorithm can be linear in R ; see Figure 1.

We show now that the execution time of the algorithm, in the worst case, can be exponential in the number of random positions k , already for weighted Markov chains, that is, for R-games. Consider the following example; see Figure 2. Let $G = (V, E)$ be a digraph on $k = 2l + 1$ vertices $u_l, \dots, u_1, u_0 = v_0, v_1, \dots, v_l$, and with the following set of arcs:

$$E = \{(u_l, u_l), (v_l, v_l)\} \cup \{(u_{i-1}, u_i), (u_i, u_{i-1}), (v_{i-1}, v_i), (v_i, v_{i-1}) \mid i = 1, \dots, l\}.$$

Let $D \geq 1$ be an integer. All positions are random with the following transition probabilities: $p(u_l, u_l) = p(v_l, v_l) = 1 - \frac{1}{D+1}$, $p(u_0, u_1) = p(u_0, v_1) = \frac{1}{2}$, $p(u_{i-1}, u_i) = p(v_{i-1}, v_i) = 1 - \frac{1}{D+1}$, for $i = 2, \dots, l$, and $p(u_i, u_{i-1}) = p(v_i, v_{i-1}) = \frac{1}{D+1}$, for $i = 1, \dots, l$. The local rewards are zero on every arc, except for $r(u_l, u_l) = -r(v_l, v_l) = 1$ (See Figure 2 for $l = 4$). Clearly this Markov chain consists of a single recurrent class, and it is easy to verify that the limiting distribution p^* is as follows:

$$p^*(u_0) = \frac{D-1}{(D+1)D^l - 2}, \quad p^*(u_i) = p^*(v_i) = \frac{D^{i-1}(D^2-1)}{2((D+1)D^l - 2)} \text{ for } i = 1, \dots, l.$$

The optimal expected reward at each vertex is

$$\mu(u_i) = \mu(v_i) = -1 \cdot \left(1 - \frac{1}{D+1}\right) p^*(u_l) + 1 \cdot \left(1 - \frac{1}{D+1}\right) p^*(u_l) = 0,$$

for $i = 0, \dots, l$. Up to a shift, there is a unique set of potentials x that transform the Markov chain into the canonical form, and they satisfy the following system of equations:

$$\begin{aligned} 0 &= -\frac{1}{2}\Delta_1 - \frac{1}{2}\Delta'_1, \\ 0 &= -\left(1 - \frac{1}{D+1}\right)\Delta_{i+1} + \frac{1}{D+1}\Delta_i, \text{ for } i = 1, \dots, k-1 \\ 0 &= -\left(1 - \frac{1}{D+1}\right)\Delta'_{i+1} + \frac{1}{D+1}\Delta'_i, \text{ for } i = 1, \dots, k-1 \\ 0 &= -\left(1 - \frac{1}{D+1}\right) + \frac{1}{D+1}\Delta_i, \\ 0 &= 1 - \frac{1}{D+1} + \frac{1}{D+1}\Delta'_l, \end{aligned}$$

where $\Delta_i = x(u_i) - x(u_{i-1})$ and $\Delta'_i = x(v_i) - x(v_{i-1})$; by solving the system, we get $\Delta_i = -\Delta'_i = D^{k-i+1}$, for $i = 1, \dots, l$.

Lower bound on pumping algorithms. Any pumping algorithm that starts with 0 potentials and modifies the potentials in each iteration by at most γ will not have a number of iterations less than $\frac{D^{l-1}}{2\gamma}$ on the above example. In particular, the algorithm in Section 3 has $\gamma \leq 1/\min\{p(v, u) \mid (v, u) \in E, p(v, u) \neq 0\}$, which is $\Omega(D)$ in our example. We conclude that the running time of the algorithm is $\Omega(D^{l-2}) = D^{\Omega(k)}$ on this example.




Figure 1: A lower bound example for pumping algorithm for B-games. In this example positions u_2 and u_3 will oscillate $2R$ times before converging to value 0.




Figure 2: An exponential time example.

4 Non-ergodic BWR-games

The main difficulty in solving non-ergodic BWR-games is the fact that random positions may introduce transient classes. In fact, our algorithm developed for the ergodic case can be used to find the top/bottom classes, and if no random node enters these classes from outside, then we could call the same procedure (called Find-Top(\cdot)/Find-Bottom(\cdot)) recursively for the rest of the positions and find an optimum solution. However, this is not always the case. To handle this we need to introduce parametrized games.

4.1 Preliminary results and basic lemmas

Let $\mathcal{G} = (G = (V = V_B \cup V_W \cup V_R, E), p, r)$ be a BWR-game. In what follows we will use the following notation. For a $\theta \in \mathbb{R}$, let $S(\theta) := \{v \in V \mid \mu_{\mathcal{G}}(v) = \theta\}$. If $S(\theta) \neq \emptyset$, we refer to it as an *ergodic class*. If $S(\theta) = V$ the game \mathcal{G} is said to be ergodic. Every BWR-game \mathcal{G} has a unique sequence of numbers $\theta_1 < \theta_2 < \dots < \theta_\ell$ such that $S(\theta_i) \neq \emptyset$ for all $i \in [\ell]$ and $\bigcup_{i=1}^{\ell} S(\theta_i) = V$.

Proposition 4 *Ergodic classes necessarily satisfy the following properties.*

- (i) *There exists no arc $(v, u) \in E$ such that $v \in V_W \cap S(\theta_i)$, $u \in S(\theta_j)$, and $j > i$;*
- (ii) *there exists no arc $(v, u) \in E$ such that $v \in V_B \cap S(\theta_i)$, $u \in S(\theta_j)$, and $j < i$;*
- (iii) *for every $v \in V_W \cap S(\theta_i)$, there exists an arc $(v, u) \in E$ such that $u \in S(\theta_i)$;*
- (iv) *for every $v \in V_B \cap S(\theta_i)$, there exists an arc $(v, u) \in E$ such that $u \in S(\theta_i)$;*
- (v) *there exists no arc $(v, u) \in E$ such that $v \in V_R \cap S(\theta_1)$, and $u \notin S(\theta_1)$;*

(vi) there exists no arc $(v, u) \in E$ such that $v \in V_R \cap S(\theta_\ell)$, and $u \notin S(\theta_\ell)$.

Proof All claims follow from the existence of canonical form for \mathcal{G} by Theorem 2, since the existence of arcs forbidden by (i), (ii), (v), or (vi), or the non-existence of arcs required by (iii) and (iv), would violate the value equations (C1). \square

For a set of positions $S \subseteq V$, we define the *black closure* $\text{cl}_B(S)$ (respectively, the *black semi-closure* $\text{cl}'_B(S)$) of S to be the set of positions which is recursively obtained from S by adding

- (1) a position $v \in V_B$ (respectively, $v \in V_B \cup V_R$), if some arc $(v, u) \in E$ satisfies $u \in S$, or
- (2) a position $v \in V_W \cup V_R$ (respectively, $v \in V_W$), if all arcs $(v, u) \in E$ satisfy $u \in S$.

In words, $\text{cl}_B(S)$ (respectively, $\text{cl}'_B(S)$) is the set of positions to which BLACK can force a move with probability 1 (respectively, with some positive probability).

The white closure and semi-closure of S , $\text{cl}_W(S)$ and $\text{cl}'_W(S)$, are defined analogously. Associated with a black closure (respectively white closure $\text{cl}_W(S)$) is a (partial) strategy $s_B(\text{cl}_B(S))$ (respectively, $s_W(\text{cl}_W(S))$) which guarantees BLACK (respectively, WHITE) a move into S . Similar strategies are defined with respect to semi-closures.

Observation 1 *If $\mu_{\mathcal{G}}(v) < z$ for all $v \in S$ then the same holds for all $v \in \text{cl}_B(S)$. Furthermore, if $z = \theta_\ell$ then the same holds for all $v \in \text{cl}'_B(S) \supseteq \text{cl}_B(S)$.*

We now introduce parametrized games.

To a BWR-game \mathcal{G} and a set $Y \subseteq V$ such that every $v \in V_W \cup V_B$ has some arc (v, u) with $u \in Y$, we associate the restriction $\mathcal{G}[Y]$ obtained by deleting all positions in $V \setminus Y$ and all arcs (v, u) where $v \notin Y$ or $u \notin Y$. Note that $\mathcal{G}[Y]$ is not a BWR-game in general. Given such a restriction and $x \in \mathbb{R}$, we further define a BWR-game $\mathcal{G}[Y](x)$ and a BW-game $\widehat{\mathcal{G}[Y]}(x)$ as follows:

$\mathcal{G}[Y](x)$: we add a new deterministic position $w = w(Y)$ with a self-loop (w, w) with local reward x , and for every $v \in V_R \cap Y$ add an arc from v to $w(Y)$ with local reward 0 and transition probability $p(v, w(Y)) := 1 - \sum_{u \in Y} p(v, u)$ if $p(v, w(Y)) > 0$;

$\widehat{\mathcal{G}[Y]}(x)$: we remove all arcs leaving $v \in V_R \cap Y$ and contract the set $v \in V_R \cap Y$ into a deterministic position $w(Y)$ (black or white, arbitrarily) with a self-loop $(w(Y), w(Y))$ having local reward value x .

We call such games *parametrized BWR-/BW-games* (with parameter x).

For a situation $s = (s_W, s_B)$ such that $s(v) \in Y$ for all $v \in Y$, we denote by the $s[Y] := (s_W[Y], s_B[Y])$ the restriction of s on Y . For brevity in what follows, we will call (v, u) a black, white, or random arc, depending on whether $v \in V_B$, $v \in V_W$, or $v \in V_R$, respectively.

Lemma 7 (i) *Given a BWR-game $\mathcal{G} = (G = (V = V_B \cup V_W \cup V_R, E), P, r)$, let $\widehat{\mathcal{G}}$ be the BW-game obtained from \mathcal{G} by replacing each random position $v \in V_R$ with a terminal deterministic position (black or white, arbitrarily) with a local reward $\mu_{\mathcal{G}}(v)$ on the self-loop (v, v) . Then $\mu_{\widehat{\mathcal{G}}}(v) = \mu_{\mathcal{G}}(v)$ for all $v \in V$.*

(ii) *Let $\widehat{\mathcal{G}}$ be as above and $U \subseteq V$ be such that $\mu_{\mathcal{G}}(v) \neq \mu_{\mathcal{G}}(u)$ for all $v \in U$ and $u \in V \setminus U$. Then $\mu_{\widehat{\mathcal{G}}[U]}(v) = \mu_{\mathcal{G}}(v)$ for all $v \in U$.*

Proof (i) By Theorem 2, there is a potential $x : V \rightarrow \mathbb{R}$ transforming \mathcal{G} to canonical form and certifying for all positions $v \in V$ that the value of v in \mathcal{G} is $\mu_{\mathcal{G}}(v)$. It is obvious that the same potential gives a canonical form for $\widehat{\mathcal{G}}$ (namely, given by the canonical form equations for \mathcal{G} , with the equations for the random positions dropped), and hence, by Theorem 2 certifies that the value of v in $\widehat{\mathcal{G}}$ is also $\mu_{\mathcal{G}}(v)$. The proof of (ii) follows also from this argument, since U contains some complete ergodic classes of \mathcal{G} . \square

We will write $V(\theta_i) := S(\theta_i) \cup \{w(S(\theta_i))\}$ and denote by $\mathbf{1}$ the vector of all ones with appropriate dimension.

For $i \in [\ell]$, define $\mathcal{G}[\theta_i]$ to be the game $\mathcal{G}[S(\theta_i)](\theta_i)$. Proposition 4 guarantees that the game $\mathcal{G}[\theta_i]$ is well-defined, that is, for every $v \in S(\theta_i)$ there is at least one arc going out of v in $\mathcal{G}[\theta_i]$.

The following two lemmas state that if we identify ergodic classes together with their values, then we can find an optimal strategy in the whole game by solving each ergodic class independently.

Lemma 8 *For all $i \in [\ell]$ and $v \in V(\theta_i)$, it holds that $\mu_{\mathcal{G}[\theta_i]}(v) = \theta_i$.*

Proof Consider a potential transformation $x : V \rightarrow \mathbb{R}$ bringing \mathcal{G} to canonical form. Let $x' : V(\theta_i) \rightarrow \mathbb{R}$ be the vector of potentials defined as follows: $x'(v) := x(v)$ if $v \in S(\theta_i)$, and

$$x'(w(S(\theta_i))) := x(v) + \frac{1}{p(v, w(S(\theta_i)))} \left(\sum_{u \in S(\theta_i)} p(v, u)(r(v, u) + x(v) - x(u)) - \theta_i \right), \quad (12)$$

for $v \in V_R \cap S(\theta_i)$. Then it is immediate that the pair $(x', \theta_i \mathbf{1})$ satisfies the canonical form conditions (C1) and (C2) at any deterministic position $v \in V(\theta_i)$. Furthermore for $v \in V_R \cap S(\theta_i)$, $(x', \theta_i \mathbf{1})$ also satisfies (C1) since, trivially, $\theta_i = \sum_{u \in S(\theta_i)} p(v, u)\theta_i + p(v, w(S(\theta_i)))\theta_i$, and moreover,

$$\theta_i = \sum_{u \in S(\theta_i)} p(v, u)(r(v, u) + x(v) - x(u)) + p(v, w(S(\theta_i)))(0 + x(v) - x'(w(S(\theta_i))))$$

holds by (12). □

Lemma 9 *For $i \in [\ell]$, let $s^i := (s_W^i, s_B^i)$ be a pair of optimal strategies in $\mathcal{G}[\theta_i]$. Then the situation $s^* = (s_W^*, s_B^*)$ obtained by concatenating all these strategies together (that is, $s_W^*(v) := s_W^i(v)$ for $v \in V_W \cap S(\theta_i)$ and $s_B^*(v) := s_B^i(v)$ for $v \in V_B \cap S(\theta_i)$) is optimal in \mathcal{G} .*

Proof For a strategy s denote by $c_s(v)$ the effective payoff (3) starting from position v in the Markov chain \mathcal{G}_s obtained from \mathcal{G} by fixing the arcs determined by s . The lemma follows from the following claims.

Claim 1 *Let $s = (s_W^*, s_B)$ or $s = (s_W, s_B^*)$, where $s_B \in S_B$ and $s_W \in S_W$ are arbitrary strategies of BLACK and WHITE, respectively. Then any absorbing class U in the Markov chain \mathcal{G}_s satisfies: $U \subseteq S(\theta_i)$ for some $i \in [\ell]$.*

Proof Without loss of generality, consider $s = (s_W^*, s_B)$ for some $s_B \in S_B$. Suppose that there an absorbing class U in \mathcal{G}_s such that $U \not\subseteq S(\theta_i)$ for all $i \in [\ell]$. Let i be the largest index such that $U \cap S(\theta_i) \neq \emptyset$. By the strong connectivity of the subgraph induced by U in \mathcal{G}_s , there must exist an arc (v, u) from some $v \in U \cap S(\theta_i)$ to some $u \in U \cap S(\theta_j) \neq \emptyset$, for some $j < i$. By Proposition 4-(ii), $v \notin V_B$, and by the choice of s , $v \notin V_W$. Then $v \in V_R$, and since U is absorbing, there are no arcs in \mathcal{G} from v to some $u \notin U$, and in particular to no $u \in S(\theta_j)$ with $j > i$. We get the following contradiction from (C1):

$$\theta_i = \sum_{u \in S(\theta_i)} p(v, u)\theta_i + \sum_{u \notin S(\theta_i)} p(v, u)\mu_{\mathcal{G}}(u) < \theta_i,$$

since $\mu_{\mathcal{G}}(u) < \theta_i$ for all $u \notin S(\theta_i)$. □

Claim 2 $\mu_{s^*}(v) = \mu_{\mathcal{G}}(v)$ for all $v \in V$.

Proof For non-transient positions, the claim follows from Claim 1, which implies that any absorbing class of \mathcal{G}_{s^*} , included in $S(\theta_i)$, is also an absorbing class in the Markov chain $(\mathcal{G}[\theta_i])_{s^i}$. In particular, the limiting distribution and hence the value of any position v in such an absorbing class are identical to the corresponding ones in $(\mathcal{G}[\theta_i])_{s^i}$, that is, $\mu_{s^*}(v) = \mu_{\mathcal{G}[\theta_i]}(v)$. By Lemma 8, we get $\mu_{s^*}(v) = \theta_i = \mu_{\mathcal{G}}(v)$.

Consider now the set of transient positions T . Using the notation in Section 2.1, let C_1, \dots, C_h be the set of absorbing classes. Then it follows from (L3) and (L4) in Section 2.1 that $\mu_{s^*}[T] := (\mu_{s^*}(v) : v \in T)$ is the unique solution of the equation $A\mathbf{x} = \alpha$, where $A := I - P_{s^*}[T; T]$, $\alpha := \sum_{i=1}^h \mu_i P_{s^*}[T; C_i] \mathbf{1}$, and μ_i is the value $\mu_G(v)$ for $v \in C_i$. Note that this equation is the value equation given in condition (C1) of the canonical form, where the value $\mu(v)$ is set to $\mu_G(v)$ for all positions v in the absorbing classes. Since the vector $\mathbf{x} := (\mu_G(v) : v \in T)$ satisfies this equation, it is the unique solution, implying that $\mu_{s^*}(v) = \mu_G(v)$ for all $v \in T$. \square

Claim 3 *Let $s' = (s_W^*, s_B)$ and $s'' = (s_W, s_B^*)$, where $s_B \in S_B$ and $s_W \in S_W$ are arbitrary strategies of BLACK and WHITE, respectively. Then $\mu_{G_{s'}}(v) \geq \mu_G(v)$ and $\mu_{G_{s''}}(v) \leq \mu_G(v)$.*

Proof Without loss of generality we only prove the claim for s' . Let $S \subseteq V$ be the set of absorbing positions in $\mathcal{G}_{s'}$, and $\xi \in \mathbb{R}^S$ be the vector of corresponding values. Then the vector of values for the set of transient positions $T := V \setminus S$ in $\mathcal{G}_{s'}$ is given by the unique solution of the following equation in \mathbf{y}

$$\mathbf{y} = \bar{B} \begin{pmatrix} \mathbf{y} \\ \xi \end{pmatrix} = B\mathbf{y} + D\xi, \quad (13)$$

where $\bar{B} = [B \mid D]$ is a stochastic matrix with $B := P_{s'}[T; T]$ and $D := P_{s'}[T; S]$. Similarly, the value equation for the set of positions T in \mathcal{G}_{s^*} is given by

$$\mathbf{x} = \bar{A} \begin{pmatrix} \mathbf{x} \\ \eta \end{pmatrix} = A\mathbf{x} + C\eta, \quad (14)$$

where $\bar{A} = [A \mid C]$, $A := P_{s^*}[T; T]$, $D := P_{s^*}[T; S]$, $\mathbf{x} \in \mathbb{R}^T$ and $\eta \in \mathbb{R}^S$ are the vector of values of the positions in T and S , respectively.

By Claim 2, $[\mathbf{x} \mid \eta] = \mu_G$ satisfies (14), and by Claim 1 we have $\xi \geq \eta$, since s_B^i is an optimal BLACK strategy in $\mathcal{G}[\theta_i]$ for all i . Note that \bar{A} and \bar{B} are identical on any row that corresponds to a position $v \in V_W \cup V_R$. Furthermore, Proposition 4-(ii) implies the following *shifting property*, for any $v \in V_B$

$$\bar{A}(v, u) = 1 \implies \bar{B}(v, u') = 1 \quad \text{for some } u' \text{ such that } \mu_G(u') \geq \mu_G(u),$$

which in turn implies that $\bar{A} \begin{pmatrix} \mathbf{x} \\ \eta \end{pmatrix} \leq \bar{B} \begin{pmatrix} \mathbf{x} \\ \eta \end{pmatrix} \leq \bar{B} \begin{pmatrix} \mathbf{x} \\ \xi \end{pmatrix}$, or

$$A\mathbf{x} + C\eta \leq B\mathbf{x} + D\xi. \quad (15)$$

By (L3) in Section 2.1, $(I - B)^{-1}$ exists and is non-negative. Combining (13), (14) and (15), we get

$$\mathbf{x} \leq (I - B)^{-1} D\xi = \mathbf{y}.$$

The claim follows. \square

This completes the proof of Lemma 9. \square

Remark 3 *Lemma 9 states that, if we know the values of all the positions, then we can get uniformly optimal strategies by solving ℓ ergodic different games. It should be noted however that, even if we know those values and the corresponding ergodic classes, a pseudo-polynomial algorithm for the ergodic case, as the one described in the previous section, does not yield in general a pseudo-polynomial algorithm for solving the game. The reason is that in our reduction in the proof of Lemma 9, we introduce local rewards on self-loops (v, v) of value $\mu_G(v)$, which might be exponentially small in n , even for games with a single random position; see, e.g., [BEGM13b]. This is due to the fact that some random positions are transitional, and hence the precision estimate in Lemma 6 does not apply.*

In view of the above remark, we need to analyze the structural dependence of the ergodic classes on the guessed values of the random positions. We achieve this by considering a parametrized BW-game as described in the next lemma.

Lemma 10 *Let v be a position in a parametrized BW-game $\widehat{\mathcal{G}}(x)$, and $x, y \in \mathbb{R}$.*

- (i) *If $\mu_{\widehat{\mathcal{G}}(x)}(v) < x$, then for any $y \geq \mu_{\widehat{\mathcal{G}}(x)}(v)$, $\mu_{\widehat{\mathcal{G}}(y)}(v) = \mu_{\widehat{\mathcal{G}}(x)}(v)$;*
- (ii) *if $\mu_{\widehat{\mathcal{G}}(x)}(v) > x$, then for any $y \leq \mu_{\widehat{\mathcal{G}}(x)}(v)$, $\mu_{\widehat{\mathcal{G}}(y)}(v) = \mu_{\widehat{\mathcal{G}}(x)}(v)$;*
- (iii) *if $\mu_{\widehat{\mathcal{G}}(x)}(v) = x$, then for any $y > x$, $y \geq \mu_{\widehat{\mathcal{G}}(y)}(v) \geq x$;*
- (iv) *if $\mu_{\widehat{\mathcal{G}}(x)}(v) = x$, then for any $y < x$, $y \leq \mu_{\widehat{\mathcal{G}}(y)}(v) \leq x$.*

Proof We prove (i) and (iii); (ii) and (iv) are analogous.

(i) Suppose that $\mu_{\widehat{\mathcal{G}}(x)}(v) < x$ and $y \geq \mu_{\widehat{\mathcal{G}}(x)}(v)$. Let (s_W^*, s_B^*) be an optimal strategy in $\widehat{\mathcal{G}}(x)$, and consider a strategy $s := (s_W^*, s_B)$, for some $s_B \in S_B$. By optimality of s_W^* in $\widehat{\mathcal{G}}(x)$, v either reaches in $\widehat{\mathcal{G}}_s(x)$ the self-loop (w, w) , or reaches another cycle C of mean value (that is, $\sum_{e \in C} r(e)/|C|$) at least $\mu_{\widehat{\mathcal{G}}(x)}(v)$. In both cases, since $y \geq \mu_{\widehat{\mathcal{G}}(x)}(v)$, v reaches in $\widehat{\mathcal{G}}_s(y)$ a cycle of mean value at least $\mu_{\widehat{\mathcal{G}}(x)}(v)$, i.e., $\mu_{\widehat{\mathcal{G}}(y)}(v) \geq \mu_{\widehat{\mathcal{G}}(x)}(v)$. Now consider a strategy $s := (s_W, s_B^*)$, for some $s_W \in S_W$. The optimality of s_B^* in $\widehat{\mathcal{G}}(x)$ implies that v does not reach the self-loop (w, w) in $\widehat{\mathcal{G}}_s(x)$, since $x > \mu_{\widehat{\mathcal{G}}(x)}(v)$. This implies that $\mu_{\widehat{\mathcal{G}}(y)}(v) \leq \mu_{\widehat{\mathcal{G}}(x)}(v)$.

(iii) Suppose that $\mu_{\widehat{\mathcal{G}}(x)}(v) = x < y$. Let (s_W^*, s_B^*) be an optimal strategy in $\widehat{\mathcal{G}}(x)$. In $\widehat{\mathcal{G}}(y)$, let us consider a strategy $s := (s_W^*, s_B)$, for some $s_B \in S_B$. By optimality of s^* in $\widehat{\mathcal{G}}(x)$, v either reaches the self-loop (w, w) , or reaches in $\widehat{\mathcal{G}}_s(x)$ another cycle of mean value at least $\mu_{\widehat{\mathcal{G}}(x)}(v)$. This implies that, in $\widehat{\mathcal{G}}(y)$, we either have $\mu_{\widehat{\mathcal{G}}_s(y)}(v) = y > x$ or $\mu_{\widehat{\mathcal{G}}_s(y)}(v) \geq \mu_{\widehat{\mathcal{G}}(x)}(v) = x$. In both cases, we get $\mu_{\widehat{\mathcal{G}}(y)}(v) \geq x$. Similarly, if $s := (s_W, s_B^*)$, for some $s_W \in S_W$, then v either reaches in $\widehat{\mathcal{G}}_s(x)$ the self-loop (w, w) , or reaches another cycle of mean value at most $\mu_{\widehat{\mathcal{G}}(x)}(v) = x < y$. This implies that, in $\widehat{\mathcal{G}}(y)$, we either have $\mu_{\widehat{\mathcal{G}}_s(y)}(v) = y$ or $\mu_{\widehat{\mathcal{G}}_s(y)}(v) \leq \mu_{\widehat{\mathcal{G}}(x)}(v) < y$. In both cases, we get $\mu_{\widehat{\mathcal{G}}(y)}(v) \leq y$. \square

Corollary 1 *For any position $v \in V$ of a parametrized BW-game $\widehat{\mathcal{G}}(x)$ there is an interval $I(v) := [\lambda_1(v), \lambda_2(v)]$, such that*

- (i) $\mu_{\widehat{\mathcal{G}}(x)}(v) = \lambda_2(v)$ if $x \geq \lambda_2(v)$;
- (ii) $\mu_{\widehat{\mathcal{G}}(x)}(v) = \lambda_1(v)$ if $x \leq \lambda_1(v)$;
- (iii) $\mu_{\widehat{\mathcal{G}}(x)}(v) = x$ if $x \in I(v)$.

Proof Consider a position $v \in V$. Then $\lambda_1(v) := \mu_{\widehat{\mathcal{G}}(-R)}(v)$ and $\lambda_2(v) := \mu_{\widehat{\mathcal{G}}(R)}(v)$ satisfy the conditions of the claim. Indeed, Lemma 10-(i) and (ii) imply, respectively, claims (i) and (ii) of the corollary. Moreover, for any $x \in I(v)$, Claims (iii) and (iv) of Lemma 10 imply respectively that $\lambda_1(v) \leq \mu_{\widehat{\mathcal{G}}(x)} \leq y$ and $y \leq \mu_{\widehat{\mathcal{G}}(x)} \leq \lambda_2(v)$, and hence Claim (iii) of the corollary. \square

Let $\widehat{\mathcal{G}}(x)$ be a parametrized BW-game with a self-loop of reward x . By Corollary 1, the end-points of the intervals $I(v)$, for $v \in V$, partition the range $[-R, R]$ into a set of at most $2\nu(\widehat{\mathcal{G}}) + 1 \leq 2n + 1$ intervals $\mathcal{I} := \mathcal{I}(\widehat{\mathcal{G}}(x))$, such that the "structure" of the game with respect to the dependence of the positions' values on x is fixed over each interval. That is, for each $I = [\lambda_1(I), \lambda_2(I)] \in \mathcal{I}$, there is a uniquely defined partition $S^-(I) \cup S^0(I) \cup S^+(I) = V$, such that

- $\mu_{\widehat{\mathcal{G}}(x)}(v) = \lambda_2(v)$ for all $v \in S^-(I)$, where $\lambda_2(v) < \lambda_1(I)$;
- $\mu_{\widehat{\mathcal{G}}(x)}(v) = x$ for all $v \in S^0(I)$;




Figure 3: In this figure we illustrate Corollary 1. The two piecewise linear functions represent the value function of $\widehat{\mathcal{G}}(x)$ for initial positions u and v , respectively, as functions of x . For the interval $I = [\lambda_1(u), \lambda_2(v)]$ we have $\{u, v\} \subseteq S^0(I)$, while for $I' = [\lambda_1(v), \lambda_1(u)]$ we have $u \in S^-(I')$ and $v \in S^0(I')$.

- $\mu_{\widehat{\mathcal{G}}(x)}(v) = \lambda_1(v)$ for all $v \in S^+(I)$, where $\lambda_1(v) > \lambda_2(I)$.

Indeed, $S^0(I)$ is defined as the set S such that $I = \bigcap_{v \in S} I(v)$; see Figure 3 for an illustration.

It is not difficult to see that for each position v , both $\lambda_1(v)$ and $\lambda_2(v)$ are rational numbers with integer denominator not exceeding n and integer numerator not exceeding nR .

Lemma 11 *Given a parametrized BW-game $\widehat{\mathcal{G}}(x)$ and a position v , we can find $\lambda_1(v)$ and $\lambda_2(v)$ in time $O(n^2 R \log(nR))$.*

Proof The smallest possible value of $\mu_{\widehat{\mathcal{G}}(x)}(v)$ is obtained for $x = -R$ and the largest for $x = R$. After this we can do two binary searches to locate λ_1 and λ_2 . Since in each step of the binary search we solve a BW-game with integer coefficients except for x which is rational with denominator at most n , the required precision is only $\frac{1}{n^2}$. Consequently, we have to solve $O(\log(nR))$ BW-games and hence the claimed complexity bound follows by [ZP96]. \square

We can now apply the above lemma for each position $v \in V$ and obtain a set of intervals $\{I(v) \mid v \in V\}$. From these we can obtain the set of intervals $\mathcal{I}(\widehat{\mathcal{G}}(x))$. We shall call this procedure $\text{BW-FindIntervals}(\widehat{\mathcal{G}}(x))$ which finds this set of intervals $\mathcal{I}(\widehat{\mathcal{G}}(x))$ together with the partitions $S^-(I) \cup S^0(I) \cup S^+(I) = V$, for $I \in \mathcal{I}(\widehat{\mathcal{G}}(x))$.

In the algorithm, we will use the above results assuming that all random positions have the same value x , which we do not know exactly. By the above structure we will guess the interval I containing x . Suppose that our guess is correct. It follows that the sets $S^+(I)$ and $S^-(I)$ contain no random positions and hence they provide a set of deterministic ergodic classes for which we obtain the values and optimal strategies by solving BW-games. On the other hand, the set $S^0(I)$ contains random positions for which the BW-game $\widehat{\mathcal{G}}[S^0(I)](x)$ has not enough information to obtain the optimal strategies. For this, we consider the parametrized BWR-game $\mathcal{G}[S^0(I)](x)$, and find the interval of x -values for which this game is ergodic. By Lemmas 8 and 9, the optimal strategies obtained for such ergodic BWR-game will yield the optimal strategies for the corresponding set in \mathcal{G} .

Lemma 12 *For a parametrized BWR-game $\mathcal{G}(x)$, the set $I(\mathcal{G}(x)) := \{x \in [-R, R] \mid \mathcal{G}(x) \text{ is ergodic}\}$ forms a closed (possibly empty) interval in $[-R, R]$.*

Proof Suppose that $\mathcal{G}(x)$ is non-ergodic. Then either

- (i) there is a strategy $s_B^* \in S_B$ such that for all $s_W \in S_W$, there is position v belonging to an absorbing class C_{s_W} in the Markov chain $\mathcal{G}_{(s_W, s_B^*)}(x)$ with value $\mu_{\mathcal{G}_{(s_W, s_B^*)}(x)} < x$, or
- (ii) there is a strategy $s_W^* \in S_W$ such that for all $s_B \in S_B$, there is position v belonging to an absorbing class C_{s_B} in the Markov chain $\mathcal{G}_{(s_W^*, s_B)}(x)$ with value $\mu_{\mathcal{G}_{(s_W^*, s_B)}(x)} > x$.

In case (i), since C_{s_W} does not include w , the same strategy s_B^* guarantees that $\mu_{\mathcal{G}_{(s_W, s_B^*)}(y)} < y$, for any $y > x$. In case (ii), the strategy s_W^* guarantees that $\mu_{\mathcal{G}_{(s_W^*, s_B)}(y)} > y$, for any $y < x$. Thus we conclude that the game $\mathcal{G}(y)$ remains non-ergodic for either all $y \leq x$ or all $y \geq x$.

Suppose that there exist $\tau_1 < \tau_2$ such that $\mathcal{G}(\tau_1)$ and $\mathcal{G}(\tau_2)$ are ergodic. Then the claim above implies that $\mathcal{G}(\tau)$ is ergodic for all $\tau \in [\tau_1, \tau_2]$. \square

Lemma 13 *In a parametrized BWR-game $\mathcal{G}(x)$, let $\tau_1 \leq \tau_2$ be two real numbers in $I(\mathcal{G}(x))$, and $s_W^* \in S_W$ and $s_B^* \in S_B$ be optimal WHITE and BLACK strategies in the games $\mathcal{G}(\tau_2)$ and $\mathcal{G}(\tau_1)$, respectively. Then (s_W^*, s_B^*) is a pair of optimal strategies in $\mathcal{G}(x)$ for all $x \in [\tau_1, \tau_2]$.*

Proof By the definition of $\mathcal{G}(x)$ the position w with the self-loop forms a single-vertex absorbing class with value x . Thus, for all $x \in [\tau_1, \tau_2]$, all other positions have the same value x , since $\mathcal{G}(x)$ is an ergodic game. We claim that for all $v \in V$ and $s_W \in S_W$ we have $\mu_{\mathcal{G}_{(s_W, s_B^*)}(\tau_1)}(v) \leq \tau_1 \leq x$. To see this, consider a strategy $s = (s_W, s_B^*)$ for some $s_W \in S_W$. In the Markov chain $\mathcal{G}_s(\tau_1)$, any position v in an absorbing class, that is not formed by the singleton $\{w\}$, has value $\mu_{\mathcal{G}_s(\tau_1)}(v) \leq \tau_1 \leq x$, and hence it has value $\mu_{\mathcal{G}_s(x)}(v) \leq \tau_1 \leq x$ in $\mathcal{G}_s(x)$. It follows that the value of any transient position v in $\mathcal{G}_s(x)$ is also at most x since it is a convex combination of the values in the absorbing classes.

Analogously, we obtain that for all $v \in V$ and $s_B \in S_B$: $\mu_{\mathcal{G}_{(s_W^*, s_B)}(x)}(v) \geq x$. \square

To be able to compute efficiently $\mathcal{I}(\mathcal{G}(x))$ and the corresponding optimal strategies mentioned in the previous lemma, we need a procedure to find the top class of a given BWR-game. This is provided in the next lemma.

Lemma 14 *Let $\bar{\mathcal{G}} := \mathcal{G}(\frac{p}{q})$ be a BWR-game obtained from the parametrized BWR-game $\mathcal{G}(x)$, where p, q are integers such that $p \leq \sqrt{k}2^{k/2}D^k R$ and $q \leq \sqrt{k}2^{k/2}D^k$. Then, we can determine both the bottom and top ergodic classes, and find a pair of strategies proving this in time $\text{poly}(n)(2D)^{O(k)} R \log R$, using algorithm PUMP($\bar{\mathcal{G}}, \varepsilon$) with ε as in (2).*

Proof We first show how to find the top ergodic class (that is, the set of positions which have the highest value in $\bar{\mathcal{G}}$). We apply Algorithm 2 called BWR-FindTop($\bar{\mathcal{G}}$). It works by calling PUMP($\bar{\mathcal{G}}, \varepsilon$) with ε as in (2). If no contra-ergodic partition is produced then the game is ergodic and its solution has been already found by the selection of ε , in view of Theorem 4. Otherwise, any position $v \in V^-$ and hence in $\text{cl}'_B(V^-)$ has value less than the top value, which can be seen from the value equations (C1), and using induction as positions are added to $\text{cl}'_B(V^-)$. Thus if we remove all these positions we get a well-defined game $\mathcal{G}' := \bar{\mathcal{G}}[V \setminus \text{cl}'_B(V^-)]$ which includes the top class of $\bar{\mathcal{G}}$. By induction, BWR-FindTop(\mathcal{G}') returns the top class $T \subseteq Y := V \setminus \text{cl}'_B(V^-)$ in \mathcal{G}' , and we claim that T is the top class also in $\bar{\mathcal{G}}$. To see this, note that by Proposition 4, there are no black or random arcs from T to $Y \setminus T$, and by the definition of Y , there are no black or random arcs from T to $V \setminus Y$ either. Thus, if $s^* = s^*[T]$ is the situation returned by BWR-FindTop(\mathcal{G}'), then s_W^* guarantees for WHITE in $\bar{\mathcal{G}}$ the same value y guaranteed by s_W^* in \mathcal{G}' . Furthermore, since T is the top class in \mathcal{G}' , there is a BLACK strategy $\bar{s}_B = \bar{s}_B[Y]$ such that $\mu_{\bar{\mathcal{G}}_{(s_W, \bar{s}_B)}} \leq y$ for all WHITE strategies $s_W = s_W[Y]$. By conditions (ii) and (iii) of the contra-ergodic partition, there is strategy $\hat{s}_B = \hat{s}_B[V^-]$ that forces WHITE to stay in V^- and ensures that $\mu_{\bar{\mathcal{G}}_{(s_W[V^-], \hat{s}_B)}}(v) < y$ for all $v \in V^-$ and $s_W \in S_W$. It follows that the strategy $s_B \in S_B$ obtained by concatenating $\bar{s}_B[Y]$, $s_B(\text{cl}'_B(V^-))$ and $\hat{s}_B[V^-]$ satisfies $\mu_{\bar{\mathcal{G}}_s}(v) \leq y$ for all $v \in V$ and $s_W \in S_W$. This proves our claim.

Finding the bottom ergodic class is analogous and can be done by a similar procedure BWR-FindBottom($\bar{\mathcal{G}}$). \square

Algorithm 2 BWR-FindTop(\mathcal{G})

Input: A BWR-game $\mathcal{G} = (G = (V, E), p, r)$

Output: A set $T \subseteq V$ such that for all $v \in V$, $\mu_{\mathcal{G}}(v) = \max_{u \in V} \mu_{\mathcal{G}}(u)$ and a situation s^* solving the game $\mathcal{G}[T]$

- 1: Set ε as in (2)
 - 2: **if** PUMP(\mathcal{G}, ε) returns a contra-ergodic partiton (V^+, V^0, V^-) **then**
 - 3: FindTop($\mathcal{G}[V \setminus \text{cl}'_B(V^-)]$)
 - 4: **else**
 - 5: Let s^* be the situation returned by PUMP(\mathcal{G}, ε)
 - 6: **return** (V, s^*)
 - 7: **end if**
-

Lemma 15 *Let us consider a parametrized BWR-game $\mathcal{G}(x)$ and let $\mathcal{I}(\mathcal{G}(x)) = [\tau_1, \tau_2]$. We can compute τ_1 , τ_2 , and the optimal strategies described in Lemma 13 in time $\text{poly}(n)(2D)^{O(k)} R \log^2 R$.*

Proof We employ binary search, calling in each step the procedure BWR-FindTop($\mathcal{G}(x)$), defined in the previous lemma, with different guessed values of the parameter x . Suppose that we start the search on the interval $[\lambda_1, \lambda_2]$. If the top class does not include all the positions then the game is non-ergodic, and the procedure will return a position $v \in V$, and either a strategy $s_B \in S_B$ certifying that $\mu_{\mathcal{G}(x)}(v) < x$ or a strategy $s_W \in S_W$ certifying that $\mu_{\mathcal{G}(x)}(v) > x$. In the former case, we reduce the search into the interval $[\lambda_1, x]$, and in the latter case, we reduce the search into the interval $[x, \lambda_2]$. Since by Lemma 6, the precision needed is $(k^{1/2} 2^{k/2} D^k)^{-1}$, we need $O(\log R + k \log D)$ many search steps. \square

We call the procedure described in the above proof BWR-FindErgodicityInterval($\mathcal{G}(x)$).

4.2 Description of main algorithm

For a BWR-game \mathcal{G} and a parameter $x \in \mathbb{R}$, define $\widehat{\mathcal{G}}(x)$ to be the BW-game obtained from \mathcal{G} by replacing each random position $v \in V_R$ in \mathcal{G} with a terminal deterministic position (black or white, arbitrarily) with a local reward of value x on the self-loop (v, v) .

Our algorithm uses four auxiliary routines: BWR-FindErgodicityInterval($\mathcal{G}(x)$) and find BW-FindIntervals($\mathcal{G}(x)$) which we described above; BWR-SolveErgodic(\mathcal{G}) which solves a given ergodic BWR-game \mathcal{G} using the pumping algorithm as in Theorem 4; and BW-Solve(\mathcal{G}) which solves a BW-game using e.g. [GKK88, Pis99, ZP96].

For a position $v \in V_R$, we define $\text{rank}(v) = |\{\mu_{\mathcal{G}}(u) \mid u \in V_R, \mu_{\mathcal{G}}(u) > \mu_{\mathcal{G}}(v)\}| + 1$. For each $v \in V_R$, our algorithm guesses its rank as $g(v)$. We remark that there are at most k^k possible guesses.

For each such guess $g : V_R \rightarrow [k]$, we call procedure BWR-Solve($\mathcal{G}, U, g, \ell, s^*[V \setminus U]$) with $U = V$ and $\ell = 1$. At any point in time, U represents the set of positions for which strategies have not been fixed yet.

This procedure keeps constructing complete situations for \mathcal{G} until it finds an optimal one or discovers that our current guess is not correct. Each time we check optimality by solving two MDPs using linear programming (see, e.g., [MO70]). We will prove that for each guess g the procedure will only construct $O(\nu(\mathcal{G})^k)$ complete situations. We will also prove that it always finds an optimal one if our guess is correct.

We now describe this procedure BWR-Solve(\cdot). For an integer $\ell \in [k]$, define U^ℓ to be the set of positions obtained from \mathcal{G} by removing all positions in the black closure $\text{cl}_B(\{v \in V_R \mid g(v) > \ell\})$; these are positions for which the values are smaller than the value of the random positions at rank ℓ , assuming that our guess is correct. We first form the game $\widehat{\mathcal{G}}[U^\ell](x)$. Then we find the set of intervals $\mathcal{I}(\widehat{\mathcal{G}}[U^\ell](x))$ using the routine BW-FindIntervals($\widehat{\mathcal{G}}[U^\ell](x)$) described above. Then for each such interval $I = [\lambda_1(I), \lambda_2(I)]$, we consider three subgames, defined by the sets $S^+(I)$, $S^0(I)$, and $S^-(I)$. By the definition of $S^+(I)$, the first subgame $\mathcal{G}[S^+(I)]$ is a BW-game. Hence, the optimal strategy $s^*[S^+(I)]$ in $\mathcal{G}[S^+(I)]$ can be obtained by calling BW-Solve($\mathcal{G}[S^+(I)]$). The positions in the second subgame $\mathcal{G}[S^0(I)]$ have the same value x . Although we do not know what the exact value of

x is, we can find the interval of ergodicity of the BWR-game $\mathcal{G}[S^0(I)](x)$ by calling procedure BWR-FindErgodicityInterval($\mathcal{G}[S^0(I)](x)$) (step 11). Once we determine this interval $[\tau_1, \tau_2]$, we solve the two ergodic games $\mathcal{G}[S^0(I)](\tau_1)$ and $\mathcal{G}[S^0(I)](\tau_2)$ using procedure BWR-IsErgodic(\cdot), and then combine the strategies according to Lemma 13 to obtain an optimal situation for $\mathcal{G}[S^0(I)]$. Finally, the rest of the game is solved by calling the procedure recursively with $\mathcal{G} := \mathcal{G}[U \setminus (S^+(I) \cup S^0(I))]$ and $\ell := \ell + 1$.

Algorithm 3 BWR-Solve($\mathcal{G}, U, g, \ell, s^*[V \setminus U]$)

Input: A BWR-game $\mathcal{G} = (G = (V = (V_B \cup V_W \cup V_R), E), p, r)$, a set of positions $U \subseteq V$, a vector of rank guesses $g : V_R \rightarrow [k]$, an integer ℓ , and a situation $s^*[V \setminus U]$ on the set of positions $V \setminus U$

- 1: **if** $V_R = \emptyset$ **then**
- 2: $s^*[U] := \text{BW-Solve}(\mathcal{G}[U])$
- 3: **if** s^* is optimal in \mathcal{G} **then**
- 4: OUTPUT s^* and HALT
- 5: **end if**
- 6: **else**
- 7: $U^\ell := U \setminus \text{cl}_B(\{v \in V_R \mid g(v) > \ell\})$
- 8: $(\mathcal{I}, S^-, S^0, S^+) := \text{BW-FindIntervals}(\widehat{\mathcal{G}[U^\ell]}(x))$
- 9: **for** each $I = [\lambda_1(I), \lambda_2(I)] \in \mathcal{I}$ **do**
- 10: $s^*[S^+(I)] := \text{BW-Solve}(\mathcal{G}[S^+(I)])$
- 11: $[\tau_1, \tau_2] := \text{BWR-FindErgodicityInterval}(\mathcal{G}[S^0(I)](x))$
- 12: $s^1 := \text{BWR-SolveErgodic}(\mathcal{G}[S^0(I)](\tau_1))$
- 13: $s^2 := \text{BWR-SolveErgodic}(\mathcal{G}[S^0(I)](\tau_2))$
- 14: $s^*[S^0(I)] := (s_W^2, s_B^1)$
- 15: **if** $U = (S^+(I) \cup S^0(I))$ **then**
- 16: **if** s^* is optimal in \mathcal{G} **then**
- 17: OUTPUT s^* and HALT
- 18: **end if**
- 19: **else**
- 20: BWR-Solve($\mathcal{G}, U \setminus (S^+(I) \cup S^0(I)), g, \ell + 1, s^*[(V \setminus U) \cup S^+(I) \cup S^0(I)]$)
- 21: **end if**
- 22: **end for**
- 23: **end if**

The following lemma states that if the guess is correct, then the procedure actually solves the game.

Lemma 16 *Let \mathcal{G} be a BWR-game. If procedure BWR-Solve($\mathcal{G}, U, g, \ell, s^*[V \setminus U]$) is called with $g(v) = \text{rank}(v)$ for all $v \in V_R$, $U = V$ and $\ell = 1$, then it returns an optimal situation s^* .*

Proof Suppose that g is as in the statement of the lemma. We prove by induction (on ℓ) that there is a path in the recursion tree (from the root to a node at depth ℓ) through which the algorithm finds correctly all the classes $S(\theta_i) := \{u \in V \mid \mu_{\mathcal{G}}(u) = \theta_i\}$, for $\theta_i \geq \mu_{\mathcal{G}}(v)$ for all $v \in V_R$ with $\text{rank}(v) = \ell$, and that the obtained situation s^* induces optimal situations in $\mathcal{G}[\theta_i]$ for each such i . This claim together with Lemma 9 would prove the theorem.

For $\ell = 0$ our claim is trivial. Suppose that it is correct up to level $\ell - 1$ and consider the path from the root of the recursion tree to a node \mathcal{N} verifying this. Consider the call to procedure BWR-Solve(\cdot) at this node. Assume that $\mu_{\mathcal{G}}(v) = x$ for all $v \in V_R$ with $\text{rank}(v) = \ell$. Let \mathcal{I} be the set of intervals computed in step 8, and $I \in \mathcal{I}$ be the interval for which $x \in I$; this interval will be eventually chosen in step 9. Since all the positions in $\text{cl}_B(\{v \in V_R \mid g(v) > \ell\})$ have value less than x , Lemma 7-(ii) implies that the values computed for all positions in the set $S^+(I) = \{v \in U^\ell \mid \mu_{\mathcal{G}}(v) > x\}$ are correct.

Furthermore, by definition of $S^0(I)$ the subgame $\mathcal{G}[S^0(I)]$ is ergodic with value x , and hence x belongs to the interval of ergodicity $[\tau_1, \tau_2]$ computed in step 11. By Lemma 13, the situation $s^*[S^0(I)]$ computed in step 14 is optimal in the subgame $\mathcal{G}[S^0(I)]$.

This completes the proof of the induction step. \square

Proof of Theorem 1. The fact the the algorithm returns an optimal situation follows from the previous lemma. For the complexity analysis, note that the depth of the recursion tree is at most $k + 1$. The number of intervals tried at each recursion level is at most $2\nu(\mathcal{G}) + 1$. The running time of the local work at each edge of the recursion tree is limited by $\text{poly}(n)(2D)^{O(k)}R \log^2 R$. We have at most $(2\nu(\mathcal{G}) + 1)^{k+1}$ such edges, and hence the claimed complexity follows. \square

Finally, we remark that our proof of Theorem 1 actually gives the following stronger bound. For a BW-game \mathcal{G}' , let $\nu(\mathcal{G}')$ be the number of distinct optimal values achieved by positions of \mathcal{G}' . For a BWR-game, we extend this by defining $\nu(\mathcal{G}) := \max_{Y \subseteq V, x \in \mathbb{R}} \nu(\widehat{\mathcal{G}[Y]}(x))$. This definition implies the following observation.

Observation 2 *For a BWR-game, $\nu(\mathcal{G}) \leq n$, and for a stochastic terminal payoff game with t terminals of distinct rewards, $\nu(\mathcal{G}) \leq t + 1$.*

Thus, the algorithm of Theorem 1 solves any BWR-game in time $(\nu Dk)^{O(k)}R \cdot \text{poly}(n, \log R)$, implying a running time of $(tDk)^{O(k)}R \cdot \text{poly}(n, \log R)$ for these games, and of $k^{O(k)} \text{poly}(n)$ for simple stochastic games.

Acknowledgements

The first author thanks the National Science Foundation for partial support (Grant IIS-1161476). The third author was partially supported by Grant-in-Aid for Scientific Research from the Ministry of Education, Culture, Sports, Science and Technology of Japan. The authors are also grateful for the support received from the Research-in-Pairs Program of the Mathematisches Forschungsinstitut Oberwolfach.

References

- [AM09] D. Andersson and P. B. Miltersen. The complexity of solving stochastic games on graphs. In *ISAAC*, pages 112–121, 2009.
- [BEF⁺11] E. Boros, K. Elbassioni, M. Fouz, V. Gurvich, K. Makino, and B. Manthey. Stochastic mean payoff games: Smoothed analysis and approximation schemes. In *ICALP (1)*, pages 147–158, 2011.
- [BEGM10] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. A pumping algorithm for ergodic stochastic mean payoff games with perfect information. In *Proc. 14th IPCO*, pages 341–354, 2010.
- [BEGM13a] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. On canonical forms for zero-sum stochastic mean payoff games. *Dynamic Games and Applications*, 3(2):128–161, 2013.
- [BEGM13b] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. On discounted approximations of undiscounted stochastic games and markov decision processes with limited randomness. *Operations Research Letters*, 41(4):357 – 362, 2013.
- [BEGM13c] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. A pseudo-polynomial algorithm for mean payoff stochastic games with perfect information and a few random positions. In Fedor V. Fomin, Rusins Freivalds, Marta Kwiatkowska, and David Peleg, editors, *Automata, Languages, and Programming*, volume 7965 of *Lecture Notes in Computer Science*, pages 220–231. Springer Berlin Heidelberg, 2013.

- [BEGM14] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. A potential reduction algorithm for ergodic two-person zero-sum limiting average payoff stochastic games. In Zhao Zhang, Lidong Wu, Wen Xu, and Ding-Zhu Du, editors, *Combinatorial Optimization and Applications*, volume 8881 of *Lecture Notes in Computer Science*. Springer International Publishing, 2014.
- [BV01a] E. Beffara and S. Vorobyov. Adapting gurvich-karzanov-khachiyan’s algorithm for parity games: Implementation and experimentation. Technical Report 2001-020, Department of Information Technology, Uppsala University, available at: <https://www.it.uu.se/research/reports/#2001>, 2001.
- [BV01b] E. Beffara and S. Vorobyov. Is randomized gurvich-karzanov-khachiyan’s algorithm for parity games polynomial? Technical Report 2001-025, Department of Information Technology, Uppsala University, available at: <https://www.it.uu.se/research/reports/#2001>, 2001.
- [BV05] H. Björklund and S. Vorobyov. Combinatorial structure and randomized sub-exponential algorithm for infinite games. *Theoretical Computer Science*, 349(3):347–360, 2005.
- [BV07] H. Björklund and S. Vorobyov. A combinatorial strongly sub-exponential strategy improvement algorithm for mean payoff games. *Discrete Applied Mathematics*, 155(2):210–229, 2007.
- [CdAH09] K. Chatterjee, L. de Alfaro, and T. A. Henzinger. Termination criteria for solving concurrent safety and reachability games. In *SODA*, pages 197–206, 2009.
- [CH08] K. Chatterjee and T. A. Henzinger. Reduction of stochastic parity to stochastic mean-payoff games. *Inf. Process. Lett.*, 106(1):1–7, 2008.
- [CHKN14] K. Chatterjee, M. Henzinger, S. Krinninger, and D. Nanongkai. Polynomial-time algorithms for energy games with special weight structures. *Algorithmica*, 70(3):457–492, 2014.
- [CJH04] K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Quantitative stochastic parity games. In *SODA ’04*, pages 121–130, Philadelphia, PA, USA, 2004. Society for Industrial and Applied Mathematics.
- [Con92] A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.
- [Con93] A. Condon. An algorithm for simple stochastic games. In *Advances in computational complexity theory, volume 13 of DIMACS series in discrete mathematics and theoretical computer science*, 1993.
- [DG11] D. Dai and R. Ge. Another sub-exponential algorithm for the simple stochastic game. *Algorithmica*, 61:1092–1104, 2011.
- [EM79] A. Eherenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8:109–113, 1979.
- [GH08] H. Gimbert and F. Horn. Simple stochastic games with few random vertices are easy to solve. In *FoSSaCS*, pages 5–19, 2008.
- [Gil57] D. Gillette. Stochastic games with zero stop probabilities. In A.W. Tucker M. Dresher and P. Wolfe, editors, *Contribution to the Theory of Games III, in Annals of Mathematics Studies*, volume 39, pages 179–187. Princeton University Press, 1957.
- [GKK88] V. Gurvich, A. Karzanov, and L. Khachiyan. Cyclic games and an algorithm to find minimax cycle means in directed graphs. *USSR Computational Mathematics and Mathematical Physics*, 28:85–91, 1988.

- [GL89] V. A. Gurvich and V. N. Lebedev. A criterion and verification of ergodicity of cyclic game forms. *Uspehi Mat. Nauk (in Russian)*, 265:193–194, English transl. in *Russian Mathematical Surveys* 44(1) (1989), 243–244, 1989.
- [Hal07] N. Halman. Simple stochastic games, parity games, mean payoff games and discounted payoff games are all LP-type problems. *Algorithmica*, 49(1):37–50, 2007.
- [HBV04] S. Sandberg H. Björklund and S. Vorobyov. A combinatorial strongly sub-exponential strategy improvement algorithm for mean payoff games. DIMACS Technical Report 2004-05, DIMACS, Rutgers University, 2004.
- [IJM12] R. Ibsen-Jensen and P. B. Miltersen. Solving simple stochastic games with few coin toss positions. In *ESA*, pages 636–647, 2012.
- [JPZ06] M. Jurdziński, M. Paterson, and U. Zwick. A deterministic subexponential algorithm for solving parity games. In *SODA '06*, pages 117–123, New York, NY, USA, 2006. ACM.
- [Jur98] M. Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Inf. Process. Lett.*, 68(3):119–124, 1998.
- [Jur00] M. Jurdziński. *Games for Verification: Algorithmic Issues*. PhD thesis, Department of Computer Science, University of Aarhus, Denmark, 2000.
- [Kar78] R. M. Karp. A characterization of the minimum cycle mean in a digraph. *Discrete Math.*, 23:309–311, 1978.
- [KL93] A. V. Karzanov and V. N. Lebedev. Cyclical games with prohibition. *Mathematical Programming*, 60:277–293, 1993.
- [KS63] J. G. Kemeny and J. L. Snell. *Finite Markov chains*. Springer-Verlag, New York, 1963.
- [Lit96] M. L. Littman. *Algorithm for sequential decision making, CS-96-09*. PhD thesis, Dept. of Computer Science, Brown Univ., USA, 1996.
- [LL69] T. M. Liggett and S. A. Lippman. Stochastic games with perfect information and time-average payoff. *SIAM Review*, 4:604–607, 1969.
- [MO70] H. Mine and S. Osaki. *Markovian decision process*. American Elsevier Publishing Co., New York, 1970.
- [Mou76] H. Moulin. Extension of two person zero sum games. *Journal of Mathematical Analysis and Application*, 5(2):490–507, 1976.
- [Pis99] N. N. Pisaruk. Mean cost cyclical games. *Mathematics of Operations Research*, 24(4):817–828, 1999.
- [RF91] T. E. S. Raghavan and J. A. Filar. Algorithms for stochastic games: A survey. *Mathematical Methods of Operations Research*, 35(6):437–472, 1991.
- [Sch03] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency, Algorithms and Combinatorics*, volume 24. Springer, New York, 2003.
- [Vor08] S. Vorobyov. Cyclic games and linear programming. *Discrete Applied Mathematics, Special volume in Memory of Leonid Khachiyan (1952 - 2005)*, 156(11):2195–2231, 2008.
- [ZP96] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996.