



HAL
open science

Entropy Games and Matrix Multiplication Games

Eugene Asarin, Julien Cervelle, Aldric Degorre, Catalin Dima, Florian Horn,
Victor Kozyakin

► **To cite this version:**

Eugene Asarin, Julien Cervelle, Aldric Degorre, Catalin Dima, Florian Horn, et al.. Entropy Games and Matrix Multiplication Games. 2015. hal-01164086v3

HAL Id: hal-01164086

<https://hal.science/hal-01164086v3>

Preprint submitted on 18 Dec 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Entropy Games and Matrix Multiplication Games*

Eugene Asarin¹, Julien Cervelle², Aldric Degorre¹, Cătălin Dima², Florian Horn¹,
and Victor Kozyakin³

¹LIAFA, University Paris Diderot and CNRS, France

²LACL, University Paris-Est Créteil, France

³IITP, Russian Academy of Science, Russia

December 18, 2015

Abstract

Two intimately related new classes of games are introduced and studied: entropy games (EGs) and matrix multiplication games (MMGs). An EG is played on a finite arena by two-and-a-half players: Despot, Tribune and the non-deterministic People. Despot wants to make the set of possible People's behaviors as small as possible, while Tribune wants to make it as large as possible. An MMG is played by two players that alternately write matrices from some predefined finite sets. One wants to maximize the growth rate of the product, and the other to minimize it. We show that in general MMGs are undecidable in quite a strong sense. On the positive side, EGs correspond to a subclass of MMGs, and we prove that such MMGs and EGs are determined, and that the optimal strategies are simple. The complexity of solving such games is in $\text{NP} \cap \text{coNP}$.

1 Introduction

In recent years, some of us have been working on a new non-probabilistic quantitative approach to classical models in computer science based on the notion of language entropy (growth rate). This approach has produced new insights about timed automata and languages [1] as well as temporal logics [2]. In this article, we apply it to game theory and obtain a new natural class of games that we call *entropy games* (EGs). Such a game is played on a finite arena in a turn-based way, in infinite time, by two-and-a-half¹ players: *Despot*, *Tribune* and the non-deterministic *People*. Whenever Despot and Tribune decide on their strategies σ and τ , it leaves a set $L(\sigma, \tau)$ (an ω -language) of possible behaviors of People. Despot wants $L(\sigma, \tau)$ to be as small as possible, while Tribune wants to make this language as large as possible. Formally the payoff of the game is the entropy of $L(\sigma, \tau)$, with Despot minimizing and Tribune maximizing this value.

Potentially these games can be used to model hidden channel capacity problems in computer security, where the aim of the security policy (Despot) is to minimize the information flow whatever the environment (Tribune) does. EGs can also be rephrased in terms of population dynamics, where one player aims to maximize the population growth rate, while the other minimizes it; applications of this setting to medicine, ecology, and computer security (virus propagation) are still to be explored. On the theoretical side, well-known mean-payoff games on finite graphs can be seen as a subclass of our EGs. However the purpose of this paper is to explore the theoretical

*The support of Agence Nationale de la Recherche under the project EQINOCS (ANR-11-BS02-004) is gratefully acknowledged. The results of Section 4 were obtained at the Institute for Information Transmission Problems, Russian Academy of Science, by V. Kozyakin at the expense of the Russian Science Foundation (project 14-50-00150).

¹Although this term is mostly used for stochastic games, it is also an appropriate description of EGs.

setting of EGs, we therefore leave applications and identification of relevant subclasses of EGs for further work.

The second class of objects studied is that of *matrix multiplication games* (MMGs), which came naturally when analyzing EGs and is, in our opinion, novel and interesting on its own. In such a game, two players, Adam and Eve, each possess a set of matrices, \mathcal{A} and \mathcal{E} , respectively. The game is played in a turn-based way, in infinite time. At every turn, the player writes a matrix from his or her set. Adam wants the norm of the product of matrices $A_1 E_1 A_2 E_2 \dots$ obtained to be as small as possible (in the limit), while Eve wants it to be as large as possible. Formally, the payoff is the growth rate of the norm of the product.

The main interest of MMGs comes from the observation that, in the case when one of the two players is trivial (i.e. his or her set contains only the identity matrix), the game turns into the classical, important, and difficult, problem of computing the joint spectral radius or the joint spectral subradius of a set of matrices, see [14, 21]. Thus, MMGs is a game (or alternating) generalization of this problem. It is thus unsurprising that, in the general case, MMGs are even more difficult to analyze. We prove that several natural problems for MMGs are undecidable, in particular it is impossible to distinguish between games with value 0 and 1 (and thus it is impossible to approximate the value of an MMG).

Fortunately, MMGs have tractable subclasses. We reduce EGs to a particular subclass of MMGs (referred to as IMMGs), when the sets \mathcal{A} and \mathcal{E} are so-called *independent row uncertainty sets* of non-negative matrices [4], and show that for this class the game can be solved: it is determined, and for each player the optimal strategy is to write one and the same matrix at every turn. This result is based on a new, quite technical, *minimax* theorem on the spectral radius of products of the type AB where both A and B belong to sets of matrices with independent row uncertainties. We deduce that EGs are determined, and that the optimal strategies for Despot and Tribune are positional. A careful complexity analysis of the games considered (EGs and IMMGs) allows to prove that comparing their value to a rational constant can be done with complexity $\text{NP} \cap \text{coNP}$.

The article is structured as follows. In Section 2 we recall useful notions from linear algebra and language theory. In Section 3 we formally define the two games and show how they are related, we also prove undecidability of general MMGs. In Section 4 we prove the key technical minimax theorem for matrices. In Section 5 we prove the main properties of EGs and IMMGs: determinacy, existence of simple strategies and complexity bounds. In Section 6 we relate the EGs studied here to classical mean-payoff games and a new kind of population games. We conclude with a discussion on the perspectives. The Appendix contains proofs of all lemmas.

2 Preliminaries

2.1 Some Linear Algebra

Given two vectors $x, y \in \mathbb{R}^N$, we write $x \geq y$, if $x_i \geq y_i$ for each $1 \leq i \leq N$. Similar notation will be applied to matrices. We denote by $\|\cdot\|$ the 1-norm of vectors and matrices. Note that, for non-negative vectors and matrices, $\|x\| = \sum_i x_i$.

Let A be an $(N \times N)$ -matrix. Its *spectral radius* is defined as the maximal modulus of its eigenvalues and denoted by $\rho(A)$. It characterizes the growth rate of A^n for $n \rightarrow \infty$: according to Gelfand's formula $\rho(A) = \lim_{n \rightarrow \infty} \|A^n\|^{1/n}$. The spectral radius depends continuously on the matrix, and is monotone for non-negative matrices [13, Corollary 8.1.19]: $\rho(A) \leq \rho(B)$ when $0 \leq A \leq B$. If $A > 0$, i.e. all the elements of A are positive, then by the Perron-Frobenius theorem, the number $\rho(A)$ is a simple eigenvalue of the matrix A , and all the other eigenvalues of A are strictly less than $\rho(A)$ in modulus. The eigenvector $v = (v_1, v_2, \dots, v_N)^\top$ corresponding to the eigenvalue $\rho(A)$ (normalized, for example, by the equation $\sum v_i = 1$) is uniquely determined and positive.

Following [4], given N sets of M -dimensional rows \mathcal{A}_i we define the *IRU-set* (independent row uncertainty set) \mathcal{A} of $(N \times M)$ -matrices that consists of all matrices of the form $A = (a_{ij})_{\substack{1 \leq i \leq N \\ 1 \leq j \leq M}}$

wherein each of the rows $a_i = [a_{i1}, a_{i2}, \dots, a_{iM}]$ belongs to the respective \mathcal{A}_i . We will need several simple properties of IRU-sets.

Lemma 1. *For an IRU-set \mathcal{A} formed by sets of rows $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_N$ the following holds:*

- (i) *for any matrix B the set $\mathcal{A}B = \{AB \mid A \in \mathcal{A}\}$ is IRU as well;*
- (ii) *the convex hull $\text{conv}(\mathcal{A})$ is the IRU-set formed by the row sets $\text{conv}(\mathcal{A}_1), \dots, \text{conv}(\mathcal{A}_N)$;*
- (iii) *the set \mathcal{A} is compact if and only if so are all the row sets $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_N$.*

2.2 Joint Spectral Radius and Subradius

The *joint spectral radius* [8, 9, 18] of a bounded set \mathcal{A} of $(N \times N)$ -matrices characterizes the maximal growth rate of products of n matrices from the set and admits the following equivalent definitions (where the identity between the upper and the lower formulas constitutes the famous Berger-Wang Theorem [3]):

$$\begin{aligned} \hat{\rho}(\mathcal{A}) &= \lim_{n \rightarrow \infty} \sup \left\{ \|A_1 \cdots A_n\|^{1/n} \mid A_i \in \mathcal{A} \right\} = \inf_{n \geq 1} \sup \left\{ \|A_1 \cdots A_n\|^{1/n} \mid A_i \in \mathcal{A} \right\} \\ &= \lim_{n \rightarrow \infty} \sup \left\{ \rho(A_1 \cdots A_n)^{1/n} \mid A_i \in \mathcal{A} \right\} = \sup_{n \geq 1} \sup \left\{ \rho(A_1 \cdots A_n)^{1/n} \mid A_i \in \mathcal{A} \right\}. \end{aligned} \quad (1)$$

For a compact (closed and bounded) set \mathcal{A} , the suprema in (1) may be replaced by maxima.

The *joint spectral subradius* [12], or *lower spectral radius*, corresponds to the minimal growth rate of products of matrices:

$$\begin{aligned} \check{\rho}(\mathcal{A}) &= \lim_{n \rightarrow \infty} \inf \left\{ \|A_1 \cdots A_n\|^{1/n} \mid A_i \in \mathcal{A} \right\} = \inf_{n \geq 1} \inf \left\{ \|A_1 \cdots A_n\|^{1/n} \mid A_i \in \mathcal{A} \right\} \\ &= \lim_{n \rightarrow \infty} \inf \left\{ \rho(A_1 \cdots A_n)^{1/n} \mid A_i \in \mathcal{A} \right\} = \inf_{n \geq 1} \inf \left\{ \rho(A_1 \cdots A_n)^{1/n} \mid A_i \in \mathcal{A} \right\}. \end{aligned}$$

The equivalence of the characterizations based on norms and on spectral radii is established in [12, Theorem B1] for finite sets \mathcal{A} , and in [20, Lemma 1.12] and [7, Theorem 1] for arbitrary sets \mathcal{A} . Calculating the joint and lower spectral radii is a challenging problem, and only in exceptional cases these characteristics may be found explicitly, see, e.g., [14, 15] and the bibliography therein. The case of compact IRU-sets of non-negative matrices is such an exception, for which $\hat{\rho}$ and $\check{\rho}$ admit a simple characterization: as stated in [16, Theorem 2], for such a set \mathcal{A} the following equalities hold:

$$\hat{\rho}(\mathcal{A}) = \max_{A \in \mathcal{A}} \rho(A), \quad \check{\rho}(\mathcal{A}) = \min_{A \in \mathcal{A}} \rho(A). \quad (2)$$

Compact IRU-sets of non-negative matrices and their convex hulls have another useful property: as is shown in [16, Corollary 1],

$$\max_{A \in \mathcal{A}} \rho(A) = \max_{A \in \text{conv}(\mathcal{A})} \rho(A), \quad \min_{A \in \mathcal{A}} \rho(A) = \min_{A \in \text{conv}(\mathcal{A})} \rho(A), \quad (3)$$

and hence $\hat{\rho}(\mathcal{A}) = \hat{\rho}(\text{conv}(\mathcal{A}))$, $\check{\rho}(\mathcal{A}) = \check{\rho}(\text{conv}(\mathcal{A}))$.

2.3 Entropy of an ω -Language

The notion of entropy of a language and methods for computing it in the case of regular languages were introduced in [6] for finite words and in [19] for infinite ones. We will use the latter definition. The entropy of an ω -language $L \subseteq \Sigma^\omega$ is defined as

$$H(L) = \limsup_{n \rightarrow \infty} \frac{\log |\text{pref}_n(L)|}{n}$$

(all the logarithms here are in base 2), where $\text{pref}_n(L)$ is the set of prefixes of length n of infinite words in L . Intuitively, $H(L)$ is the information content (“bandwidth”), measured in bits per symbol, in typical words of the language. In particular, $H(\Sigma^\omega) = \log |\Sigma|$.

For a regular $L \subseteq \Sigma^\omega$ accepted by a given Büchi automaton, its entropy can be effectively computed as follows: compute the (finite) automaton recognizing $\text{pref}(L)$, determinize it, and compute the entropy as the logarithm of the spectral radius of the adjacency matrix of the automaton obtained.

3 The Two Games

3.1 Entropy Games

Consider the arena (D, T, Σ, Δ) where D and T are disjoint finite sets of vertices (of two players), Σ a finite alphabet of actions and $\Delta \subseteq T \times \Sigma \times D \cup D \times \Sigma \times T$ is a transition relation. Given such an arena, we define a game with two-and-a-half players: Despot, Tribune and People. The latter plays non-deterministically and counts for half a player. People chooses the initial state in D . When the game is in a state d of D , Despot plays an action $a \in \Sigma$ and the game changes to some $t \in T$ (chosen by People) such that $(d, a, t) \in \Delta$. Then, Tribune plays an action $b \in \Sigma$ and the game changes its state to $d' \in D$, again chosen by People and such that $(t, b, d') \in \Delta$. It is again Despot’s turn. The players must not block the game: they always choose an action that has a corresponding transition $(d, a, \cdot) \in \Delta$, or $(t, b, \cdot) \in \Delta$, respectively. We assume that the arena is non-blocking: at every state there is at least one such transition. Figure 1 shows an example of such an arena, which we will use as a running example in this paper.

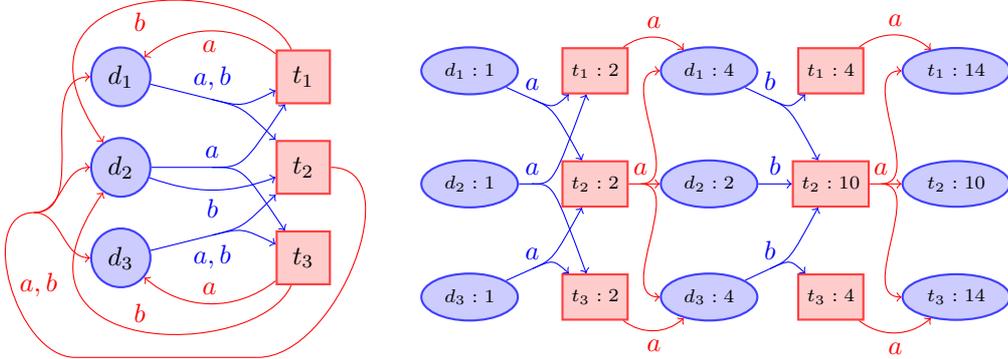


Figure 1: *Left.* Arena of our running example of an entropy game. Circles are states of the Despot while squares are states of the Tribune. At each move, the player has to choose between actions a and b , the outcome of which may sometimes be non-deterministic (e.g. when Despot plays a in state d_2 , the next state may non-deterministically be either t_1 or t_3). *Right.* A finite play on this arena. Despot plays ab (whatever his opponent does) while Tribune plays aa . We only give, for each step, the number of words that end up in each state controlled by the active player.

A *play* of the EG is a finite or infinite sequence $\pi \in (D \cdot \Sigma \cdot T \cdot \Sigma)^\infty$ compatible with the transition relation Δ . Note that four letters in a row correspond to one turn of the game. A *strategy* σ for Despot is a function $(D \cdot \Sigma \cdot T \cdot \Sigma)^* \cdot D \rightarrow \Sigma$ that, given any finite play ending in a D state, outputs an action taken by Despot. The strategy is positional if it only depends on the current state of the game, i.e. it can be expressed just as $\sigma(d)$. A *strategy* τ for Tribune is a function $(D \cdot \Sigma \cdot T \cdot \Sigma)^* \cdot D \cdot \Sigma \cdot T \rightarrow \Sigma$ which, given any finite play ending in a T state, outputs the action taken by Tribune. The strategy is positional if it only depends on the current state of the game. In a natural way we define plays compatible with a Despot’s strategy σ , or with

a Tribune's strategy τ . Then, given σ and τ , we have an ω -language $L(\sigma, \tau)$ containing all the plays compatible with σ and τ . In other words, $L(\sigma, \tau)$ is the set of runs that People can choose if Despot and Tribune commit themselves to σ and τ . What makes EGs different from other games (parity/mean-payoff etc.) is that the payoff does not depend on a single run of the game, but on the whole set of possible runs. More precisely, the payoff (the amount that Despot pays to Tribune) is defined as

$$P(\sigma, \tau) = \limsup_{n \rightarrow \infty} |\text{pref}_{4n}(L(\sigma, \tau))|^{1/n},$$

that is the growth rate (w.r.t. the number of turns) of the number of plays available to the People under the strategies σ and τ . Note that the payoff is a monotone function of the entropy of $L(\sigma, \tau)$, indeed

$$P(\sigma, \tau) = 2^{4H(L(\sigma, \tau))},$$

i.e. Despot tries to diminish the entropy while Tribune aims to augment it.

3.2 Matrix Multiplication Games

Let \mathcal{A} be a set of $M \times N$ -matrices and \mathcal{E} of $N \times M$ -matrices. The MMG between two players, Adam and Eve, is played as follows: in turn, for every $i \in \mathbb{N}$, Adam writes a matrix $A_i \in \mathcal{A}$ and then Eve writes a matrix $E_i \in \mathcal{E}$. Formally, we define a *play* as an infinite sequence $A_1 E_1 A_2 E_2 \dots A_i E_i \dots$ with $A_i \in \mathcal{A}$ and $E_i \in \mathcal{E}$. A strategy for Adam is a function $\sigma : (\mathcal{A} \cdot \mathcal{E})^* \rightarrow \mathcal{A}$ that maps any finite history (which is a sequence of matrices) to Adam's next move. Similarly, a strategy for Eve is a mapping $\tau : (\mathcal{A} \cdot \mathcal{E})^* \cdot \mathcal{A} \rightarrow \mathcal{E}$. A strategy is called *constant* if it does not depend on the history, i.e. is given by just one matrix: $\sigma = A \in \mathcal{A}$ or $\tau = E \in \mathcal{E}$. We define a play compatible with a strategy σ (or τ) in a natural way. Note that, given a strategy σ for Adam and a strategy τ for Eve, there exists a unique play $\pi(\sigma, \tau)$ compatible with both of them. The payoff of a play $\pi = A_1 E_1 A_2 E_2 \dots A_i E_i \dots$ (that is, the amount that Adam pays to Eve) is the growth rate of the norm of the infinite product of matrices:

$$P(\pi) = P(\sigma, \tau) = \limsup_{k \rightarrow \infty} \left\| \prod_{i=1}^k A_i E_i \right\|^{1/k}.$$

3.3 General Matrix Multiplication Games are Undecidable

The difficulty of general MMGs should be compared with results on the difficulty of JSR (joint spectral radius) computation. Thus, as proved in [5, Thm 2], given a finite set \mathcal{E} of non-negative matrices with rational elements, it is undecidable whether $\hat{\rho}(\mathcal{E}) \leq 1$. The decidability status of the problem $\hat{\rho}(\mathcal{E}) < 1$ is unknown. Finally, it is immediate from the characterization (1) that, given a precision $\varepsilon > 0$, it is possible to compute ε -approximation of $\hat{\rho}(\mathcal{E})$ (in other words $\hat{\rho}(\mathcal{E})$ is computable as function of \mathcal{E} in the sense of computable analysis, see [23]).

Theorem 2. *Given a determined MMG with finite sets of non-negative matrices with rational elements and $\alpha \in \mathbb{Q}_+$, the decision problem for its value $V \leq \alpha$ is undecidable.*

Proof. Let $\mathcal{A} = \{Id\}$ (Adam is trivial) and \mathcal{E} be a finite set of non-negative matrices with rational elements. The corresponding MMG is determined with value $V = \hat{\rho}(\mathcal{E})$ and thus the decision problem $V \leq 1$ is undecidable due to [5, Thm 2], cited above. \square

To prove stronger undecidability results for MMGs without direct counterparts for the JSR, we need a couple of simulation lemmas: for arbitrary matrices and for non-negative ones.

Lemma 3. *Given a two-counter machine M , one can construct two finite sets of integer matrices \mathcal{A} and \mathcal{E} such that the corresponding MMG is determined and its value V satisfies: "if M halts (starting with counters containing 0) then $V = 0$, else $V = 1$."*

Lemma 4. *Given a two-counter machine M , one can construct two finite sets of non-negative integer matrices \mathcal{A} and \mathcal{E} such that the corresponding MMG satisfies: “if M halts then Adam can ensure payoff < 2 , otherwise Eve can ensure payoff ≥ 2 .”*

In both cases the construction, inspired by [10], follows the same principle: Eve tries to simulate the machine M ; if she cheats, then Adam detects this and “resets” the product. Since the halting problem is undecidable, we obtain immediately the following two theorems.

Theorem 5. *Given a determined MMG with finite sets of matrices with integer elements*

- *its value V is not computable from the matrices;*
- *it is not computable even knowing a priori that $V \in \{0, 1\}$.*

Hence the MMG value cannot be approximated and is not computable (as function of \mathcal{A} and \mathcal{E}) in the sense of computable analysis.

Theorem 6. *Given an MMG with finite sets of non-negative matrices with integer elements, it is undecidable whether the maximal payoff that Eve can ensure is < 2 .*

3.4 Relations Between the Two Kinds of Games

Fortunately, as will be shown below, the subclass of MMGs with IRU-sets of non-negative matrices is much easier to solve. In this section, we relate EGs to such MMGs.

Let $\mathbf{A} = (D, T, \Sigma, \Delta)$ be an arena with $D = \{d_1, \dots, d_M\}$ and $T = \{t_1, \dots, t_N\}$. We define matrix sets \mathcal{A}, \mathcal{E} as follows. For each Despot’s vertex $d_i \in D$, and action $a \in \Sigma$ we define the row $c_{ia} = [c_{ia,1}, \dots, c_{ia,N}]$ where $c_{ia,j} = 1$ if $(d_i, a, t_j) \in \Delta$ and $c_{ia,j} = 0$ otherwise. Next we define the row set $\mathcal{A}_i = \{c_{ia} \neq 0 \mid a \in \Sigma\}$ (non-zero rows correspond to non-blocking actions). Row sets $\mathcal{A}_1, \dots, \mathcal{A}_M$ determine an IRU-set of matrices \mathcal{A} . The IRU-set \mathcal{E} corresponding to Tribune’s actions is defined similarly. In the running example in Figure 1, for instance, the row sets are the following: $\mathcal{A}_1 = \{[1, 1, 0]\}$, $\mathcal{A}_2 = \{[0, 1, 0], [1, 0, 1]\}$, $\mathcal{A}_3 = \{[0, 1, 1]\}$, $\mathcal{E}_1 = \{[0, 1, 0], [1, 0, 0]\}$, $\mathcal{E}_2 = \{[1, 1, 1]\}$, $\mathcal{E}_3 = \{[0, 1, 0], [0, 0, 1]\}$.

Note first that there is a natural bijection between the positional strategies of Despot and the set \mathcal{A} : any positional strategy $\sigma : D \rightarrow \Sigma$ corresponds to the matrix $A_\sigma \in \mathcal{A}$ with i -th row $c_{i,\sigma(d_i)}$ for Adam. Similarly, a positional strategy of Tribune τ corresponds to Eve’s matrix $E_\tau \in \mathcal{E}$. The following lemma generalizes this observation to any type of strategies:

Lemma 7. *Let \mathbf{A} be an arena and \mathcal{A}, \mathcal{E} the corresponding IRU matrix sets. Then for every pair of strategies (σ, τ) of Despot and Tribune in the EG on \mathbf{A} there exists a pair of strategies (ς, θ) of Adam and Eve in the MMG $(\text{conv}(\mathcal{A}), \text{conv}(\mathcal{E}))$ with exactly the same payoff. Moreover, if σ is positional, then ς is constant and permanently chooses A_σ . The case of positional τ is similar.*

Note that Lemma 7 provides a rather weak relation between two games and does not mean, by itself, that the two games have the same value. However, we will show later (cf. Lemma 15) that optimal **constant** strategies in the MMG that belong to \mathcal{A} and \mathcal{E} are in bijection with optimal positional strategies in the EG.

4 Minimax Theorem for IRU-Sets of Matrices

In this section, we prove the key theorem of this article.

Theorem 8. *Let \mathcal{A} be a compact IRU-set of non-negative $(N \times M)$ -matrices and \mathcal{B} be a compact IRU-set of non-negative $(M \times N)$ -matrices. Then*

$$\min_{A \in \mathcal{A}} \max_{B \in \mathcal{B}} \rho(AB) = \max_{B \in \mathcal{B}} \min_{A \in \mathcal{A}} \rho(AB). \quad (4)$$

In the rest of the article we will denote this minimax by $\text{mm}(\mathcal{A}, \mathcal{B})$. The study of minimax relations will be based on the following well-known fact:

Lemma 9 (see [22, Section 13.4]). *Let $f(x, y)$ be a continuous function on the product of compact spaces $X \times Y$. Then*

$$\min_x \max_y f(x, y) \geq \max_y \min_x f(x, y).$$

The exact equality holds if and only if there exists a saddle point, i.e. a point (x_0, y_0) satisfying the inequalities

$$f(x_0, y) \leq f(x_0, y_0) \leq f(x, y_0)$$

for all $x \in X, y \in Y$.

We will also use two lemmas on matrices. The first one provides spectral radius bounds and is quite standard in Perron-Frobenius theory; as usual in this theory it relates global characteristics of a non-negative matrix (such as spectral radius) with its behavior on one non-negative vector.

Lemma 10. *Let A be a non-negative $(N \times N)$ -matrix; then the following properties hold:*

- (i) *if $Au \leq \rho u$ for some vector $u > 0$, then $\rho \geq 0$ and $\rho(A) \leq \rho$;*
- (ii) *if furthermore $A > 0$ and $Au \neq \rho u$, then $\rho(A) < \rho$;*
- (iii) *if $Au \geq \rho u$ for some non-zero vector $u \geq 0$ and some number $\rho \geq 0$, then $\rho(A) \geq \rho$;*
- (iv) *if furthermore $Au \neq \rho u$, then $\rho(A) > \rho$.*

The next lemma concerning IRU-sets of matrices is new and can be explained as follows. For an IRU-set of matrices and two vectors u and v we imagine that the sets $B_l = \{x : x \leq v\}$ and $B_u = \{x : v \leq x\}$ form the lower and upper bulbs of an hourglass with the neck at the point v . The lemma asserts that either all the grains Au (for all matrices A in the set) fill one of the bulbs, or there remains at least one grain in the other bulb. Clearly this alternative does not hold for general sets of matrices.

Lemma 11 (hourglass alternative). *Let \mathcal{A} be an IRU-set of $(N \times M)$ -matrices and let $\tilde{A}u = v$ for some matrix $\tilde{A} \in \mathcal{A}$ and vectors u, v . Then the following holds:*

- (i) *either $Au \geq v$ for all $A \in \mathcal{A}$ or exists a matrix $\bar{A} \in \mathcal{A}$ such that $\bar{A}u \leq v$ and $\bar{A}u \neq v$;*
- (ii) *either $Au \leq v$ for all $A \in \mathcal{A}$ or exists a matrix $\bar{A} \in \mathcal{A}$ such that $\bar{A}u \geq v$ and $\bar{A}u \neq v$.*

We are ready to prove the minimax theorem.

Proof of Theorem 8. According to Lemma 9, the minimax equality (4) may occur if and only if some matrices $\tilde{A} \in \mathcal{A}$ and $\tilde{B} \in \mathcal{B}$ satisfy the inequalities

$$\rho(\tilde{A}\tilde{B}) \leq \rho(\tilde{A}\tilde{B}) \quad \text{for all } B \in \mathcal{B}; \quad (5)$$

$$\rho(\tilde{A}\tilde{B}) \leq \rho(\tilde{A}\tilde{B}) \quad \text{for all } A \in \mathcal{A}. \quad (6)$$

Consider first the case when all the matrices in \mathcal{A} and \mathcal{B} are **positive**. To construct the matrices $\tilde{A} \in \mathcal{A}$ and $\tilde{B} \in \mathcal{B}$ we proceed as follows. For each $B \in \mathcal{B}$ let $A_B \in \mathcal{A}$ be a matrix that minimizes (in A) the quantity $\rho(AB)$. Such a matrix A_B exists due to compactness of the set \mathcal{A} and continuity of the function $\rho(AB)$ in A and B . Then, for each matrix $B \in \mathcal{B}$, the relations

$$\rho(A_B B) = \min_{A \in \mathcal{A}} \rho(AB) \leq \rho(AB)$$

hold for all $A \in \mathcal{A}$. Let \tilde{B} be the matrix maximizing $\min_{A \in \mathcal{A}} \rho(AB)$ over the set \mathcal{B} , and let $\tilde{A} = A_{\tilde{B}}$. In this case

$$\max_{B \in \mathcal{B}} \rho(A_B B) = \max_{B \in \mathcal{B}} \min_{A \in \mathcal{A}} \rho(AB) = \min_{A \in \mathcal{A}} \rho(A\tilde{B}) = \rho(A_{\tilde{B}}\tilde{B}) = \rho(\tilde{A}\tilde{B}), \quad (7)$$

which implies inequality (6) for all $A \in \mathcal{A}$, and it remains to prove (5) for all $B \in \mathcal{B}$.

Let $v = (v_1, v_2, \dots, v_N)^\top$ be the positive eigenvector of the $(N \times N)$ -matrix $\tilde{A}\tilde{B}$ corresponding to the eigenvalue $\tilde{\rho} = \rho(\tilde{A}\tilde{B})$. By denoting

$$w = \tilde{B}v \in \mathbb{R}^M$$

we obtain that $\tilde{\rho}v = \tilde{A}w$. Let us show that in this case

$$\tilde{\rho}v \leq Aw \quad \text{for all } A \in \mathcal{A}. \quad (8)$$

Otherwise, by Lemma 11(i) there would exist a matrix $\bar{A} \in \mathcal{A}$ such that $\tilde{\rho}v \geq \bar{A}w$ and $\tilde{\rho}v \neq \bar{A}w$, which implies, by the definition of the vector w , that $\tilde{\rho}v \geq \bar{A}\tilde{B}v$ and $\tilde{\rho}v \neq \bar{A}\tilde{B}v$. Then by Lemma 10

$$\rho(\bar{A}\tilde{B}) < \tilde{\rho} = \rho(\tilde{A}\tilde{B}),$$

which contradicts (6). This contradiction completes the proof of inequality (8). Similarly, now we show that

$$w \geq Bv \quad \text{for all } B \in \mathcal{B}. \quad (9)$$

Again, assuming the contrary, by Lemma 11(ii) there exists a matrix $\bar{B} \in \mathcal{B}$ such that $w \leq \bar{B}v$ and $w \neq \bar{B}v$. This last inequality, together with (8) applied to the matrix $A_{\bar{B}}$, yields $\tilde{\rho}v \leq A_{\bar{B}}\bar{B}v$ and $\tilde{\rho}v \neq A_{\bar{B}}\bar{B}v$. Then by Lemma 10

$$\tilde{\rho} < \rho(A_{\bar{B}}\bar{B}),$$

which contradicts (7) asserting that $\tilde{\rho} = \rho(\tilde{A}\tilde{B})$ is the maximum value of the function $\rho(A_B B)$ over all $B \in \mathcal{B}$. This contradiction completes the proof of inequality (9).

From $\tilde{\rho}v = \tilde{A}w$ and (9) we obtain the inequality $\tilde{\rho}v \geq \tilde{A}Bv$ valid for all $B \in \mathcal{B}$, which by Lemma 10 implies the relations

$$\rho(\tilde{A}\tilde{B}) = \tilde{\rho} \geq \rho(\tilde{A}B)$$

valid for all $B \in \mathcal{B}$, or, which is the same, inequality (5). The theorem is proved for positive matrices.

Consider now the general case of compact IRU-sets of **non-negative matrices** \mathcal{A} and \mathcal{B} . If the set \mathcal{A} is determined by some sets of M -rows \mathcal{A}_i , $i = 1, 2, \dots, N$, then choose an arbitrary $\varepsilon > 0$ and consider the sets of rows

$$\mathcal{A}_i^{(\varepsilon)} = \{a^{(\varepsilon)} \mid a^{(\varepsilon)} = a + \varepsilon[1, 1, \dots, 1], a \in \mathcal{A}_i\},$$

where $i = 1, 2, \dots, N$. In this case the IRU-set of matrices $\mathcal{A}^{(\varepsilon)}$ consists of positive matrices $A + \varepsilon\mathbf{1}$, where $A \in \mathcal{A}$ and $\mathbf{1}$ is the matrix with all elements equal to 1. Define similarly the IRU-set of matrices $\mathcal{B}^{(\varepsilon)}$.

By the result just proved, for each $\varepsilon > 0$ the minimax equality holds for positive matrices:

$$\min_{A \in \mathcal{A}^{(\varepsilon)}} \max_{B \in \mathcal{B}^{(\varepsilon)}} \rho(AB) = \max_{B \in \mathcal{B}^{(\varepsilon)}} \min_{A \in \mathcal{A}^{(\varepsilon)}} \rho(AB),$$

which by Lemma 9 is equivalent to the existence of $\tilde{A}_\varepsilon \in \mathcal{A}$ and $\tilde{B}_\varepsilon \in \mathcal{B}$ such that

$$\rho((\tilde{A}_\varepsilon + \varepsilon\mathbf{1})(B + \varepsilon\mathbf{1})) \leq \rho((\tilde{A}_\varepsilon + \varepsilon\mathbf{1})(\tilde{B}_\varepsilon + \varepsilon\mathbf{1})) \leq \rho((A + \varepsilon\mathbf{1})(\tilde{B}_\varepsilon + \varepsilon\mathbf{1}))$$

for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$. Taking here $\varepsilon = \varepsilon_n$, where $\{\varepsilon_n\}$ is an arbitrary sequence of positive numbers converging to zero, we get

$$\rho((\tilde{A}_{\varepsilon_n} + \varepsilon_n\mathbf{1})(B + \varepsilon_n\mathbf{1})) \leq \rho((\tilde{A}_{\varepsilon_n} + \varepsilon_n\mathbf{1})(\tilde{B}_{\varepsilon_n} + \varepsilon_n\mathbf{1})) \leq \rho((A + \varepsilon_n\mathbf{1})(\tilde{B}_{\varepsilon_n} + \varepsilon_n\mathbf{1})) \quad (10)$$

for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$. Without loss of generality, in view of the compactness of the sets \mathcal{A} and \mathcal{B} , we may assume the existence of matrices \tilde{A} and \tilde{B} such that $\tilde{A}_{\varepsilon_n} \rightarrow \tilde{A} \in \mathcal{A}$ and $\tilde{B}_{\varepsilon_n} \rightarrow \tilde{B} \in \mathcal{B}$ as $n \rightarrow \infty$. Then turning to the limit in (10), we obtain the inequalities

$$\rho(\tilde{A}B) \leq \rho(\tilde{A}\tilde{B}) \leq \rho(A\tilde{B})$$

for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$, which are equivalent to (5) and (6). This concludes the proof. \square

Corollary 12. *For IRU-sets \mathcal{A} and \mathcal{B} of non-negative matrices it holds that*

$$\text{mm}(\text{conv}(\mathcal{A}), \text{conv}(\mathcal{B})) = \text{mm}(\mathcal{A}, \mathcal{B}).$$

5 Solving the Games

5.1 Solving Matrix Multiplication Games for IRU-Sets

Theorem 13. *Let \mathcal{A} and \mathcal{E} be compact IRU-sets of non-negative matrices. Then the corresponding MMG is determined, and moreover Adam and Eve possess constant optimal strategies.*

Proof. Let us apply Theorem 8 to matrix sets \mathcal{A} and \mathcal{E} . Define V , E_0 and A_0 such that

$$\min_{E \in \mathcal{E}} \rho(EA_0) = \max_{A \in \mathcal{A}} \min_{E \in \mathcal{E}} \rho(EA) = \min_{E \in \mathcal{E}} \max_{A \in \mathcal{A}} \rho(EA) = \max_{A \in \mathcal{A}} \rho(E_0A) = V. \quad (11)$$

Let Adam only play A_0 . Take any compatible play $\pi = A_0E_1A_0E_2 \cdots$ and put $C_i = A_0E_i$. Denote $\mathcal{C} = \{EA_0 | E \in \mathcal{E}\}$; it is an IRU-set by Lemma 1. The payoff P for π yields

$$\begin{aligned} P &= \limsup_{n \rightarrow \infty} \|A_0C_1 \cdots C_{n-1}E_n\|^{1/n} \leq \limsup_{n \rightarrow \infty} (\|A_0\| \cdot \|C_1 \cdots C_{n-1}\| \cdot \|E_n\|)^{1/n} \\ &\leq \lim_{n \rightarrow \infty} K^{\frac{2}{n}} \limsup_{n \rightarrow \infty} \|C_1 \cdots C_{n-1}\|^{\frac{1}{n-1}} \leq \hat{\rho}(\mathcal{C}) \stackrel{1}{=} \max_{C \in \mathcal{C}} \rho(C) = \max_{E \in \mathcal{E}} \rho(EA_0) \stackrel{2}{=} V, \end{aligned}$$

where the constant K is an upper bound for the norms of the matrices in \mathcal{A} and \mathcal{E} , equality 1 comes from the first equality (2) and equality 2 comes from (11).

Let Eve only play E_0 . Take any compatible play $\pi' = A_1E_0A_2E_0 \cdots$. Let us write $D_i = A_iE_0$. Denote $\mathcal{D} = \{AE_0, A \in \mathcal{A}\}$; it is an IRU-set. The payoff P' for π' is such that

$$P' = \limsup_{n \rightarrow \infty} \|C_1 \cdots C_n\|^{1/n} \geq \liminf_{n \rightarrow \infty} \|C_1 \cdots C_n\|^{1/n} \geq \check{\rho}(\mathcal{D}) \stackrel{1}{=} \min_{D \in \mathcal{D}} \rho(D) = \min_{A \in \mathcal{A}} \rho(AE_0) \stackrel{2}{=} V,$$

where equality 1 comes from the second equality (2) and equality 2 from (11) using the equality $\rho(EA_0) = \rho(A_0E)$.

We have proved that Adam (by constantly playing A_0) can ensure payoff $\leq V$ whatever Eve plays; and that Eve (by constantly playing E_0) can ensure payoff $\geq V$ whatever Adam plays. This concludes the proof. \square

Corollary 14. *Let \mathcal{A} and \mathcal{E} be compact IRU-sets of non-negative matrices. In the MMG on $\text{conv}(\mathcal{A})$, $\text{conv}(\mathcal{E})$, the constant optimal strategies can be chosen from sets \mathcal{A} and \mathcal{E} .*

This follows immediately from the proof of the theorem and Corollary 12.

5.2 Solving Entropy Games

In this section, we consider an EG on an arena \mathbf{A} and the corresponding matrix sets \mathcal{A} and \mathcal{E} , as defined in Section 3.4.

Lemma 15. *Let (σ, τ) be two positional strategies in the EG. Then, if corresponding constant strategies A_σ and E_τ are optimal for their respective players in the MMG with matrix sets $\text{conv}(\mathcal{A})$ and $\text{conv}(\mathcal{E})$, then so are σ and τ .*

Theorem 16. *Every EG is determined, and Despot and Tribune possess positional optimal strategies.*

Proof. From Theorem 13, we know that for the MMG $(\text{conv}(\mathcal{A}), \text{conv}(\mathcal{E}))$ both Adam and Eve possess optimal strategies, which consist in constantly playing some matrices A and E . From Corollary 14, the matrices A and E can be chosen from sets \mathcal{A} and \mathcal{E} , respectively. Then, there exist positional strategies σ and τ on \mathbf{A} such that $A = A_\sigma$ and $E = E_\tau$. By Lemma 15, strategies σ and τ are optimal in the EG. \square

Back to the running example. Here a quick exploration of the combinations of rows shows that the matrices realizing the minimax over the two IRU-sets defined by row sets $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ are $A = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$ for Adam/Despot and $E = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}$ for Eve/Tribune. These matrices describe both the optimal constant strategy of the MMG and the optimal positional strategy of the EG induced by this arena. The value of both games is the spectral radius $\rho(AE) = \rho\left(\begin{bmatrix} 2 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 2 \end{bmatrix}\right) = (\sqrt{17} + 3)/2 \simeq 3.562$.

5.3 Complexity Issues

We will analyze the complexity of solving matrix multiplication (and hence entropy) game. We start with necessary and sufficient conditions for inequalities on joint spectral radii and subradii of IRU-sets (recall also (2) relating them to maximal and minimal spectral radii).

Lemma 17. *For any compact IRU-set of positive matrices \mathcal{A} and $\alpha \in \mathbb{Q}_+$ the following equivalences hold:*

$$\hat{\rho}(\mathcal{A}) < \alpha \Leftrightarrow \exists v > 0 \forall A \in \mathcal{A} (Av < \alpha v); \quad (12)$$

$$\hat{\rho}(\mathcal{A}) \leq \alpha \Leftrightarrow \exists v > 0 \forall A \in \mathcal{A} (Av \leq \alpha v); \quad (13)$$

$$\check{\rho}(\mathcal{A}) > \alpha \Leftrightarrow \exists v > 0 \forall A \in \mathcal{A} (Av > \alpha v); \quad (14)$$

$$\check{\rho}(\mathcal{A}) \geq \alpha \Leftrightarrow \exists v > 0 \forall A \in \mathcal{A} (Av \geq \alpha v). \quad (15)$$

If the matrices are only non-negative, the equivalences (12) above and (16) below hold:

$$\check{\rho}(\mathcal{A}) \geq \alpha \Leftrightarrow \exists (v \geq 0, v \neq 0) \forall A \in \mathcal{A} (Av \geq \alpha v). \quad (16)$$

The computational aspects of calculating the values $\hat{\rho}(\mathcal{A})$ and $\check{\rho}(\mathcal{A})$ for IRU-sets of non-negative matrices, based on relations (2), are discussed in [4, 16, 17]. These articles provide polynomial algorithms for approximation of the minimal and maximal spectral radii, as well as a variant of the simplex method for these problems. In the next theorem we prove a complexity result in a form suitable for game analysis.

Theorem 18. *Given a finite IRU-set of nonnegative matrices \mathcal{A} with rational elements (represented by row sets $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_N$), and a number $\alpha \in \mathbb{Q}_+$, the decision problems whether $\hat{\rho}(\mathcal{A}) < \alpha$ and whether $\check{\rho}(\mathcal{A}) \geq \alpha$ belong to the complexity class P. Moreover, if the matrices are positive, then the decision problems $\hat{\rho}(\mathcal{A}) \leq \alpha$ and $\check{\rho}(\mathcal{A}) > \alpha$ are also in P.*

Proof. The polynomial algorithms are based on the previous lemma. Consider the problem of deciding $\hat{\rho}(\mathcal{A}) < \alpha$, which can be rewritten using (12) as $\exists v > 0 \forall A \in \mathcal{A} (Av < \alpha v)$. We will not test all the matrices $A \in \mathcal{A}$ (there are exponentially many of them); instead, we will treat each row separately. The condition $\forall A \in \mathcal{A} (Av < \alpha v)$ can be rewritten as a system of linear inequalities: for each i and for each row $[c_1, c_2, \dots, c_N] \in \mathcal{A}_i$ require that

$$c_1 v_1 + c_2 v_2 + \dots + c_N v_N < \alpha v_i.$$

The condition $v > 0$ can be written as N inequalities $v_i > 0$: one for each coordinate. Using a polynomial algorithm for linear programming we can decide whether a solution v satisfying all these linear inequalities exists.

All other decision procedures, based on (13)–(16), are similar. The condition $v \geq 0, v \neq 0$ can be represented as a disjunction of N linear systems $v_j > 0 \wedge \bigwedge_{i=1}^N v_i \geq 0$. \square

Theorem 19. *Given two finite IRU-sets of nonnegative matrices \mathcal{A} and \mathcal{B} with rational elements, and a number $\alpha \in \mathbb{Q}_+$, the decision problem of whether $\text{mm}(\mathcal{A}, \mathcal{B}) < \alpha$ belongs to $\text{NP} \cap \text{coNP}$. Moreover, if the matrices are positive, then the problem of whether $\text{mm}(\mathcal{A}, \mathcal{B}) \leq \alpha$ is also in $\text{NP} \cap \text{coNP}$.*

Proof. Consider the problem of deciding whether $\text{mm}(\mathcal{A}, \mathcal{B}) < \alpha$, which can be rewritten as

$$\min_{A \in \mathcal{A}} \max_{B \in \mathcal{B}} \rho(BA) < \alpha,$$

or equivalently

$$\exists A_0 \in \mathcal{A} (\hat{\rho}(\mathcal{B}A_0) < \alpha).$$

The nondeterministic polynomial algorithm proceeds as follows:

- guess non-deterministically a matrix $A_0 \in \mathcal{A}$;
- compute the representation of $\mathcal{B}A_0$ as an IRU-set generated by the row sets $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_N$;
- check the inequality $\hat{\rho}(\mathcal{B}A_0) < \alpha$ in polynomial time using Theorem 18.

We conclude that the problem $\text{mm}(\mathcal{A}, \mathcal{B}) < \alpha$ is in NP. The complementary problem $\text{mm}(\mathcal{A}, \mathcal{B}) \geq \alpha$ is also in NP, as it can be rewritten as

$$\max_{B \in \mathcal{B}} \min_{A \in \mathcal{A}} \rho(AB) \geq \alpha,$$

or equivalently

$$\exists B_0 \in \mathcal{B} (\check{\rho}(\mathcal{A}B_0) \geq \alpha),$$

and decided by a non-deterministic polynomial algorithm similarly. We conclude that the two problems belong to $\text{NP} \cap \text{coNP}$.

For positive matrices, the proof for the other decision problem based on the second statement of Theorem 18 is similar. \square

Our main complexity result follows immediately.

Theorem 20. *Given an EG or an MMG with finite IRU-sets of non-negative matrices with rational elements and $\alpha \in \mathbb{Q}_+$, the decision problem for its value: $V < \alpha$ is in $\text{NP} \cap \text{coNP}$.*

6 Related Models

6.1 Weighted Entropy Games

Up to now we have considered entropy games with *simple* transitions, but it is straightforward to add multiplicities (weights) to them. A *weighted entropy game* is played on a *weighted arena* $\mathbf{A} = (D, T, \Sigma, \Delta, w)$ with a function $w : \Delta \rightarrow \mathbb{N}_+$ assigning weights to transitions (informally a weight is the number of ways in which a transition can be taken). Strategies and plays are defined as in the unweighted case. Let L be some set of (infinite) plays. For every $u \in \text{pref}(L)$ we define its weight $w(u)$ as the product of weights of all the transitions taken along u . We define $w_n(L) = \sum_{u \in \text{pref}_{4n}(L)} w(u)$, and finally the payoff corresponding to strategies σ and τ of two players is defined as:

$$P = \limsup_{n \rightarrow \infty} (w_n(L(\sigma, \tau)))^{1/n}.$$

Our main results on EGs (Thms 16 and 20) extend straightforwardly to weighted EGs.

6.2 Mean-Payoff Games

Well-known mean-payoff finite-state games (MPG) [11] can be considered as a deterministic subclass of weighted entropy games. A (variant of) MPG is played on arena (D, T, Δ, w) with transition relation $\Delta \subseteq D \times T \cup T \times D$ and weight function $w : \Delta \rightarrow \mathbb{N}$. The play starts in some state $d_0 \in D$, and the two players choose transitions in turn. The resulting play is an infinite word

$\gamma_{d_0} \in (D \cdot T)^\omega$. The mean-payoff corresponding to the play $\gamma_{d_0} = d_0, t_0, d_1, t_1, \dots$ is the limit of the average weight of transitions taken:

$$\text{mp}(\gamma_{d_0}) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (w(d_{i-1}, t_{i-1}) + w(t_{i-1}, d_i)).$$

Finally, player D wants to minimize and player T to maximize the payoff $\max_{d_0 \in D} \text{mp}(\gamma_{d_0})$. As proved in [11], MPGs are determined and their optimal strategies are positional. As for complexity, [24] shows that testing whether the value of an MPG is smaller than a rational α is in $\text{NP} \cap \text{coNP}$ and becomes polynomial for weights presented in the unary system.

An MPG $\mathbf{A} = (D, T, \Delta, w)$ can be transformed into a weighted EG $\mathbf{A}' = (D, T, \Sigma, \Delta', w')$ as follows. The states of both players are the same, Σ is large enough, and for each transition $(p, q) \in \Delta$ there is a corresponding transition $(p, a, q) \in \Delta'$ with some a (occurring only in this transition). Its weight is $w'(p, a, q) = 2^{w(p, q)}$. We notice that the EG obtained is deterministic: due to unique transition labels for any strategies σ and τ , the language $L(\sigma, \tau)$ contains one play for each initial state. Strategies and plays of both games \mathbf{A} and \mathbf{A}' are now in natural bijection and the payoff of \mathbf{A} equals the logarithm of the payoff of \mathbf{A}' .

This way, we obtain the classical results that MPGs are determined and both players have optimal positional strategies. Due to the exponential encoding of payoffs, the complexity obtained using our approach is, however, not as good as using direct algorithms, see [24].

6.3 Population Dynamics

Consider an EG with arena $\mathbf{A} = (D, T, \Sigma, \Delta)$. It can be interpreted as the following population game between two players, Damien and Theo. Elements of D and T correspond to species (forms of viruses, microorganisms, etc.). Initially there is one (or any non-zero number of) organism(s) for each species in D . At his turn Damien chooses an action $a \in \Sigma$ and applies it to each organism. An organism of species d , when subject to action a , turns into the set of organisms of species $\{t \mid (d, a, t) \in \Delta\}$. Theo plays similarly. The aim of Damien is to minimize the growth rate of the population, while Theo wants to maximize it. The value of the game and the optimal (positional) strategies are the same as for the EG.

7 Conclusions

We have introduced two (closely interrelated) families of games: entropy games played on finite arenas (graphs), and matrix multiplication games. The main result is that entropy games are determined and optimal strategies are positional in EG, while MMGs for IRU-sets of non-negative matrices are determined and optimal strategies are constant. These results are based on a new minimax theorem on spectral radii of products of IRU-sets of matrices. The results obtained prove the existence of equilibria in zero-sum games with a new type of limit payoffs, which is neither computed on a single play of the game nor probabilistic. On the other hand, they rely upon and generalize important results on the computability of joint spectral radii and subradii, an important problem in switching dynamic systems.

A presumably straightforward extension would be the “probabilization” of our game models, in that both Despot and Tribune would be allowed to play randomized strategies. The minimax theorem ensures the existence of optimal pure strategies for both players. However the entropy-based payoff of the game needs to be given a proper generalization to this probabilistic setting. We may mention that such a generalization could be seen as entropy games on stochastic branching processes, and provide interesting links with this research domain. Finally, both our games are turn-based games with perfect information. The first generalization to be considered is to go to concurrent games — where perhaps some polynomial-size memory is needed, similarly to the classic case of concurrent games played on graphs in infinite time. The more difficult case is that of games of imperfect information: corresponding matrix games no longer have a simple structure

(independent row uncertainty), and we conjecture that analysis of such games is non-computable. Last but not least, potential applications sketched in the introduction should be addressed.

References

- [1] Eugene Asarin, Nicolas Basset, and Aldric Degorre. Entropy of regular timed languages. *Inform. Comput.*, 241:142–176, 2015.
- [2] Eugene Asarin, Michel Blockelet, Aldric Degorre, Cătălin Dima, and Chunyan Mu. Asymptotic behaviour in temporal logic. In *Proc. CSL-LICS*, pages 10:1–10:9. ACM, 2014.
- [3] Marc A. Berger and Yang Wang. Bounded semigroups of matrices. *Linear Algebra Appl.*, 166:21–27, 1992.
- [4] Vincent D. Blondel and Yurii Nesterov. Polynomial-time computation of the joint spectral radius for some sets of nonnegative matrices. *SIAM J. Matrix Anal. A.*, 31(3):865–876, 2009.
- [5] Vincent D. Blondel and John N. Tsitsiklis. The boundedness of all products of a pair of matrices is undecidable. *Syst. Control Lett.*, 41(2):135–140, 2000.
- [6] Noam Chomsky and George A. Miller. Finite state languages. *Inform. Control*, 1(2):91–112, 1958.
- [7] Adam Czornik. On the generalized spectral subradius. *Linear Algebra Appl.*, 407:242–248, 2005.
- [8] Ingrid Daubechies and Jeffrey C. Lagarias. Sets of matrices all infinite products of which converge. *Linear Algebra Appl.*, 161:227–263, 1992.
- [9] Ingrid Daubechies and Jeffrey C. Lagarias. Corrigendum/addendum to [8]. *Linear Algebra Appl.*, 327(1-3):69–83, 2001.
- [10] Aldric Degorre, Laurent Doyen, Raffaella Gentilini, Jean-François Raskin, and Szymon Torunczyk. Energy and mean-payoff games with imperfect information. In *Proc. CSL*, LNCS 6247, pages 260–274. Springer, 2010.
- [11] Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.
- [12] Leonid Gurvits. Stability of discrete linear inclusion. *Linear Algebra Appl.*, 231:47–85, 1995.
- [13] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2013.
- [14] Raphaël Jungers. *The Joint Spectral Radius: Theory and Applications*. LNCIS 385. Springer, 2009.
- [15] Victor Kozyakin. An annotated bibliography on convergence of matrix products and the theory of joint/generalized spectral radius. Preprint, Institute for Information Transmission Problems, Moscow, 2013. <http://dx.doi.org/10.13140/2.1.4257.5040>.
- [16] Yurii Nesterov and Vladimir Yu. Protasov. Optimizing the spectral radius. *SIAM J. Matrix Anal. A.*, 34(3):999–1013, 2013.
- [17] Vladimir Yu. Protasov. Spectral simplex method. *Math. Program.*, pages 1–27, 2015.
- [18] Gian-Carlo Rota and Gilbert Strang. A note on the joint spectral radius. *Nederl. Akad. Wetensch. Proc. Ser. A 63 = Indag. Math.*, 22:379–381, 1960.
- [19] Ludwig Staiger. Entropy of finite-state omega-languages. *Probl. Control Inform.*, 14(5):383–392, 1985.

- [20] Jacques Theys. *Joint Spectral Radius: Theory and Approximations*. PhD thesis, Université Catholique de Louvain, 2005.
- [21] John N. Tsitsiklis and Vincent D. Blondel. The Lyapunov exponent and joint spectral radius of pairs of matrices are hard — when not impossible — to compute and to approximate. *Math. Control Signals Systems*, 10(1):31–40, 1997.
- [22] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1947.
- [23] Klaus Weihrauch. *Computable Analysis*. Springer, 2000.
- [24] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theor. Comput. Sci.*, 158(1–2):343–359, 1996.

A Proofs

Proof of Lemma 1

Proof. (i) Let \mathcal{A}_i be the set of admissible i -th rows in \mathcal{A} ,

$$\mathcal{R}_i = \left\{ \left[\sum_{k=1}^n a_k b_{kj} \right]_{1 \leq j \leq n} \mid a \in \mathcal{A}_i \right\},$$

and \mathcal{R} be the IRU-set made from sets \mathcal{R}_i . One has that $\mathcal{A}B = \mathcal{R}$:

- if $M \in \mathcal{R}$ then let $a^{(i)} \in \mathcal{A}_i$ be such that the i -th row of M is $\left[\sum_{k=1}^n a_k^{(i)} b_{kj} \right]_{1 \leq j \leq n}$, then $M = AB$ where A is the matrix made with rows a_i ;
 - conversely, if $A \in \mathcal{A}$ and $a^{(i)}$ is the i -th row of A , then the i -th row of AB equals $\left[\sum_{k=1}^n a_k^{(i)} b_{kj} \right]_{1 \leq j \leq n}$ and belongs to \mathcal{R}_i .
- (ii) The easy direction is \subseteq . Let M be a matrix of $\text{conv}(\mathcal{A})$. Then, there exist matrices $M_1, \dots, M_k \in \mathcal{A}$ and real numbers $\lambda_1, \dots, \lambda_k$ such that

$$M = \sum_{i=1}^k \lambda_i M_i.$$

Let j be an integer in $\{1, \dots, n\}$. For all $i \in \{1, \dots, k\}$, there exists a vector $v_i \in \mathcal{A}_j$ such that row j of M_i is v_i . Then, row j of M being $\sum_{i=1}^k \lambda_i v_i$, it belongs to $\text{conv}(\mathcal{A}_j)$.

For the direction \supseteq , let M be a matrix of the IRU-set formed by $\text{conv}(\mathcal{A}_1), \dots, \text{conv}(\mathcal{A}_n)$. Let u_1, \dots, u_n be the rows of the matrix M . By definition of M , there are integers k_i for $i \in \{1, \dots, n\}$, real numbers $\lambda_j^i \in [0, 1]$ and vectors $v_j^i \in \mathcal{A}_i$ for $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, k_i\}$ such that

$$u_i = \sum_{j=1}^{k_i} \lambda_j^i v_j^i \text{ and } \sum_{j=1}^{k_i} \lambda_j^i = 1.$$

Then, for all $i \in \{1, \dots, n\}$, one has:

$$u_i = \sum_{j_i=1}^{k_i} \lambda_{j_i}^i v_{j_i}^i = \left(\prod_{l=1}^{i-1} \sum_{j_l=1}^{k_l} \lambda_{j_l}^l \right) \left(\sum_{j_i=1}^{k_i} \lambda_{j_i}^i v_{j_i}^i \right) \left(\prod_{l=i+1}^n \sum_{j_l=1}^{k_l} \lambda_{j_l}^l \right) = \sum_{j_1=1}^{k_1} \dots \sum_{j_n=1}^{k_n} \left(\prod_{l=1}^n \lambda_{j_l}^l \right) v_{j_i}^i.$$

Hence

$$M = \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix} = \sum_{j_1=1}^{k_1} \dots \sum_{j_n=1}^{k_n} \left(\prod_{l=1}^n \lambda_{j_l}^l \right) \begin{bmatrix} v_{j_1}^1 \\ \vdots \\ v_{j_n}^n \end{bmatrix},$$

each matrix in the sum being in \mathcal{A} . The proof is finished stating that

$$\sum_{j_1=1}^{k_1} \dots \sum_{j_n=1}^{k_n} \prod_{l=1}^n \lambda_{j_l}^l = \prod_{l=1}^n \sum_{j_l=1}^{k_l} \lambda_{j_l}^l = 1.$$

- (iii) Immediate from the characterization of compact sets (of finite dimension) as bounded and closed. \square

Proof of Lemmas 3 and 4

In both lemmas, we announce that it is possible to reduce the halting problem of a 2-counter Minsky machine (2CMM) to the threshold problem for MMG.

So let us first remind the reader about 2CMMs. Such a machine can be defined as a set of instructions, indexed by a finite set of states Q , with $q_0 \in Q$ marked as initial, operating on two non-negative integer counters x and y . There are three types of instructions ($q...$ are states in Q , and c is either x or y):

1. q_i : increment c then execute q_j ;
2. q_i : if $c = 0$ execute q_j , else decrement c then execute q_k ;
3. q_i : stop.

The computation starts from instruction q_0 , executing it and thus triggering a sequence of instructions, which may be finite, if it eventually reaches a *stop* instruction, or otherwise infinite. Whether or not the execution will be finite is undecidable.

Obviously, both reductions consist in encoding any 2CMM into an MMG, the payoffs of which depend on whether the machine halts or not.

Now let us describe the encoding used in Lemma 3. Here the 2CMM is translated into two sets \mathcal{A} and \mathcal{E} of square matrices of dimension $|Q| + 5$. States of the 2CMM (discrete location and counter values) are encoded, along with some other information, as (row) vectors of the space on which these matrices operate. The $|Q|$ first coordinates of such a vector, labelled with $q_0, \dots, q_{|Q|-1}$, take a non-zero value only for the current state of the simulation. The two next coordinates x and y represent the two counters. Finally, there are three additional coordinates: *One*, *E* and *Neg* (the role of which will be explained later on).

Eve's matrices allow her to simulate the machine execution (as long as it goes on). The set \mathcal{E} consists in exactly one matrix per transition of the 2CMM (warning: instruction 2 consists in two different transitions, depending on the test $c = 0$, while instruction 3 consists in no transition, i.e., a state with *stop* instruction is a deadlock state). For the sake of presentation, we describe them below as sets of assignments of variables, but it is easy to see that all assignments actually are linear operations:

- matrices $I_{qq'c}$ (as *Increment*): $q := q - \text{One}$; $q' := q' + \text{One}$; $c := c + \text{One}$;
- matrices $K_{qq'c}$ (as *Keep* current counter value): $q := q - \text{One}$; $q' := q' + \text{One}$; $c := -c$;
- matrices $D_{qq'c}$ (as *Decrement*): $q := q - \text{One}$; $q' := q' + \text{One}$; $c := c - \text{One}$.

Notice that matrix $K_{qq'c}$ should be normally applied when $c = 0$, and thus the operation $c = -c$ does not harm. On the contrary, if it is applied illegally, for a positive counter value, then it results in a negative c .

Adam has five kinds of matrices, which he can use to detect whenever Eve does not faithfully simulate the machine, and then punish her by forcing a payoff of 0. Here is the set \mathcal{A} :

- the matrix *Init* (initialize the 2CMM): $q_0 := E$; $q_{i \neq 0} := 0$; $x := 0$; $y := 0$; $\text{One} := E$; $\text{Neg} := 0$
- the identity matrix *Id* (do nothing and just let Eve continue playing);
- the matrices F_c (*flash* and take a picture of coordinate c) for c corresponding to a state or a counter: $\text{Neg} := c$;
- the matrix A (*adjust* the value of *Neg*): $\text{Neg} := \text{Neg} + \text{One}$;
- the matrix P (*punish* Eve by assigning 0 to E): $E = E + \text{Neg}$.

Now, in order to prove Lemma 3, it suffices to prove the two following sublemmas:

Lemma A.1. *The MMG obtained by the translation above from a non-halting 2CMM is determined with value 1, i.e. it has the following properties:*

1. *there exists a strategy of Adam σ_0 such that for any strategy τ of Eve, $P(\sigma_0, \tau) \leq 1$;*
2. *there exists a strategy of Eve τ_0 such that for any strategy σ of Adam, $P(\sigma, \tau_0) \geq 1$.*

Lemma A.2. *The MMG obtained by the translation above from a halting 2CMM is determined with value 0, i.e. it has the following properties:*

1. *there exists a strategy of Adam σ_0 such that for any strategy τ of Eve, $P(\sigma_0, \tau) \leq 0$;*
2. *there exists a strategy of Eve τ_0 such that for any strategy σ of Adam, $P(\sigma, \tau_0) \geq 0$.*

Proof sketch of Lemma A.1. 1. Let the strategy σ_0 consist in always playing identity. Then $\max_{\tau} P(\sigma_0, \tau) = \hat{\rho}(\mathcal{E})$. It is easy to see that applying a matrix of \mathcal{E} to a vector only changes the value of coordinates labelled by a state or by a counter, and only modifies it by adding or removing the value of coordinate *One* (its value is left unchanged). Thus, for any vector v and any sequence of matrices of \mathcal{E} : E_1, \dots, E_n , all coordinates of vector $vE_1 \cdots E_n$ are bounded in absolute value by $n \cdot v_{One}$, which means that $\|E_1 \cdots E_n\| \leq k$, and thus $\hat{\rho}(\mathcal{E}) \leq 1$.

2. Assume τ_0 is as follows: Eve always stores in a variable t the last time Adam played *Init* (initially $t = 0$). Then at turn i , she plays the matrix that corresponds to the $(i - t)$ -th transition of the execution of the 2CMM.

We fix a non-negative vector

$$v_0 = (q_0 = 1, q_{i \neq 0} = 0, x = 0, y = 0, E = 1, One = 1, Neg = 0)$$

and prove by induction the following invariant on the vector $v_n = v_0 A_1 E_1 A_2 E_2 \cdots A_n E_n$:

- all coordinates of v_n are non-negative;
- $v_n^E \geq 1$.

Indeed applying a matrix of \mathcal{A} while respecting the rules of the 2CMM ensures that state and counter coordinates remain non-negative, while *One*, *E* and *Neg* remain unchanged.

On the other hand Adam's matrices are all non-negative, implying the first bullet of the invariant, and can only modify *E* by adding the value of *Neg*, which, by the invariant, was non-negative at the previous step.

The above proves the invariant which implies the second item of the lemma. □

Proof sketch of Lemma A.2. 1. For σ_0 , we consider the following strategy:

- first play *Init*: Adam initializes a simulation of the 2CMM in his private memory in the form of a vector $v_0 = (q_0 = 1, q_{i \neq 0} = 0, x = 0, y = 0, E = 1, One = 1, Neg = 0)$, on which all subsequent matrices will be applied (yielding v_1, v_2, \dots);
- then play *Id* as long as Eve plays valid transitions of the 2CMM;
- play F_c as soon as Eve plays an invalid move (if Eve lied on a counter value, then c is the name of this counter; if Eve lied on current state, then c is the name of this state; in both cases this corresponds to a negative coordinate);
- play *A* until $v_n^{Neg} = -1$;
- finally play *P* (nulling *E*) and a last time *Init* (nulling the whole vector).

Explanation: the invariant from Lemma A.1 holds as long as Eve simulates the 2CMM. When she stops simulating, the value of *E* is still 1 but some coordinate c is negative. The ending sequence $F_c A \cdots A P$ *Init* forces the final vector to be 0, no move of Eve can then prevent this from happening, as she cannot modify *One*, *E* or *Neg*.

Now remark that, for any vector v such that $v^E = 1$, it holds that $v \cdot \text{Init} = v_0$ and thus $v \cdot \Omega = 0$ where Ω is the product matrix for the whole play until Adam plays a last time *Init*. This proves that as much yields for any initial vector, thus $\Omega = 0$ and therefore $P(\sigma_0, \tau) = 0$.

2. The payoff function of an MMG is always non-negative. \square

Proof of Lemma 3. It directly follows from Lemmas A.1 and A.2. \square

Proof of Lemma 4. General idea: Here, in order to use only non-negative matrices, we introduce a slightly different construction. Indeed, previous encoding relied on the fact that when Eve cheats, a negative coordinate appears that Adam can use to punish her (by nulling the product matrix with clever additions using the negative integers). Now there is no hope to create a negative element in the product matrix (since the matrices of the MMG are non-negative), so the punishment will be less drastic: Adam will just try to obtain a product that grows more slowly than the product of matrices corresponding to a faithful simulation. For this Adam needs to reset the game infinitely often, as to force Eve to cheat as many times, within a bounded horizon.

The encoding uses the following idea: a counter of value k is encoded at time n by a coordinate of value 2^{n+k} . This way, a counter decrement consists in keeping its coordinate value unchanged, while a counter increment consists in multiplying its coordinate by 4. A counter stalling at a given step still sees its coordinate multiplied by 2.

The vector space: Matrices are square of dimension $|Q| + 4$. They act on vectors such that their $|Q|$ first coordinates represent the states of the 2CMM (positive value only in the coordinate corresponding to the active state of the simulation) and the four other coordinates x_+, x_-, y_+, y_- , are the two counters and their opposite (i.e. their value will be 2^{n-c} instead of 2^{n+c}).

Eve's matrices: In this game too, Eve tries to faithfully simulate the 2CMM. Her matrices are the following:

- matrices $I_{qq'x}$: $q' := 2q$; $q := 0$; $x_+ := 4x_+$; $y_+ := 2y_+$; $y_- := 2y_-$;
- matrices $K_{qq'x}$: $q' := 2q$; $q := 0$; $(x_+, x_-) := 2(x_-, x_+)$; $(y_+, y_-) := 2(y_-, y_+)$;
- matrices $D_{qq'x}$: $q' := 2q$; $q := 0$; $x_- := 4x_-$; $y_+ := 2y_+$; $y_- := 2y_-$;
- matrices $I_{qq'y}$, $K_{qq'y}$ and $D_{qq'y}$ are defined likewise.

Notice that the coordinate inversion of the previous construction, for the case of a successful $x = 0$ test, is now translated as a coordinate swap between x_+ and x_- . Thus when Eve cheats on a counter value, be it one way or the other, x_+ has a value smaller than 2^n .

Adam's matrices:

- matrix Id ;
- matrices P_x (and P_y): $q_0 := x_+$; $q_{i \neq 0} := 0$; $x_- = x_+$; $y_+ = x_+$; $y_- = x_+$;
- matrices P_q : let $s = \sum_{q \in Q} q$ then $q_0 := s$; $q_{i \neq 0} := 0$; $x_+ = s$; $x_- = s$; $y_+ = s$; $y_- = s$.

Both $P_{\{x,y\}}$ and P_q reset the simulation, forcing the copied value as the new norm for the product matrix.

Adam's strategy consists in playing Id most of the time; playing P_q whenever Eve cheats on the state, implying a null product; and playing P_c whenever Eve cheats on a counter $c \in \{x, y\}$ value, implying a factor of norm $\leq 2^{f-1}$ since the last time when Adam played a P matrix (factor of length f).

Since Eve needs to cheat with a positive frequency if the run of the 2CMM is finite, then the final payoff will be < 2 (from the product of such factors, which are of bounded length).

If the run is infinite, whether Adam plays Id or a P , there will be some coordinate that remains of magnitude $\geq 2^n$.

Details of the proof are similar to those of Lemma 3. We prove determinacy, with a value of 2, in the case when the 2CMM has an infinite run. For the other case we prove that Adam can ensure a payoff < 2 , but determinacy remains an open problem. \square

Proof of Lemmas 7–17

Proof of Lemma 7. Assume $D = \{d_1, \dots, d_M\}$ and $T = \{t_1, \dots, t_N\}$. Given arbitrary strategies (σ, τ) for the two players in the EG, let us represent the set of all compatible plays as a forest. Its nodes are labeled by elements of D on even levels and elements of T on odd levels, and its edges are labeled by symbols in Σ . The label of a node q is denoted $\ell(q)$; the sequence of labels on the path reaching q from the appropriate root in the forest is referred to as its address $\alpha(q)$. The forest \mathbf{F} is defined inductively as follows:

- \mathbf{F} has M root nodes labeled by d_1, \dots, d_M ;
- all the outgoing edges of a node q labeled $d \in D$ carry the symbol $a = \sigma(\alpha(q))$ and the sons of the node q correspond to (and are labeled by) the elements of $\{t \mid (d, a, t) \in \Delta\}$;
- all the outgoing edges of a node q labeled $t \in T$ carry the symbol $b = \tau(\alpha(q))$ and the sons of the node q correspond to (and are labeled by) the elements of $\{d \mid (t, a, d) \in \Delta\}$.

The payoff of the EG can be characterized in terms of the growth rate of this forest:

$$P(\sigma, \tau) = \limsup_{n \rightarrow \infty} |\mathbf{F}_{2n}|^{1/n},$$

where \mathbf{F}_k denotes the set of nodes of \mathbf{F} at the level k . Indeed $L(\sigma, \tau)$ is the set of labels of infinite paths of \mathbf{F} , hence $\text{pref}(L(\sigma, \tau))$ is the set of addresses of nodes in \mathbf{F} (we use the fact that our strategies are required to be non-blocking). To words of length $4n$ in $\text{pref}(L(\sigma, \tau))$ correspond addresses of nodes of level $2n$, and thus

$$\limsup_{n \rightarrow \infty} |\text{pref}_{4n}(L(\sigma, \tau))|^{1/n} = \limsup_{n \rightarrow \infty} |\mathbf{F}_{2n}|^{1/n}$$

as required.

Let us characterize the number of nodes $|\mathbf{F}_{2n}|$ in terms of matrices. Let the vector $x^{(n)} = (x_1^{(n)}, \dots, x_j^{(n)})$ be such that $x_i^{(n)}$ is the number of nodes labeled by d_i on $2n$ -th level of \mathbf{F} ; similarly let $y^{(n)} = (y_1^{(n)}, \dots, y_N^{(n)})$ be such that $y_j^{(n)}$ is the number of nodes labeled by t_j on $(2n + 1)$ -th level of \mathbf{F} . To relate $y^{(n)}$ to $x^{(n)}$ we observe that

$$y_j^{(n)} = \sum_{i=1}^M \sum_{a \in \Sigma} |\{q \in \mathbf{F}_{2n} \mid \ell(q) = d_i \wedge \sigma(\alpha(q)) = a\}| c_{ia,j}.$$

Indeed, every node on level $2n$ with label d_i and action a generates on the next level a node with label t_j whenever $c_{ia,j} = 1$. Summing up on all i, a and q we obtain the quantity $y_j^{(n)}$. The expression for y can be rewritten as

$$y_j^{(n)} = \sum_{i=1}^M x_i^{(n)} \sum_{a \in \Sigma} \mu_{ia} c_{ia,j} \tag{17}$$

with $\mu_{ia}^{(n)} = |\{q \in \mathbf{F}_{2n} \mid \ell(q) = d_i \wedge \sigma(\alpha(q)) = a\}| / x_i^{(n)}$ (whenever $x_i^{(n)} = 0$, coefficients $\mu_{ia}^{(n)}$ can be chosen arbitrarily, only respecting conditions (18) below). Intuitively, $\mu_{ia}^{(n)}$ is the proportion among the states d_i on level $2n$, of those for which Despot takes the action a . In matrix form (17) can be rewritten as $y^{(n)} = x^{(n)} A_n$ with $A_{n,ij} = \sum_{a \in \Sigma} \mu_{ia}^{(n)} c_{ia,j}$. We notice that

$$\mu_{ia}^{(n)} \geq 0 \text{ and } \sum_{a \in \Sigma} \mu_{ia}^{(n)} = 1, \tag{18}$$

thus i -th row of A_n belongs to $\text{conv}(\mathcal{A}_i)$, hence $A_n \in \text{conv}(\mathcal{A})$. Similarly, $x^{(n+1)} = y^{(n)} E_n$ for some $E_n \in \text{conv}(\mathcal{E})$. Initially $x^{(0)} = (1, \dots, 1)$, and clearly $|\mathbf{F}_n| = x^{(n)} \cdot (1, \dots, 1)^\top$, hence

$$|\mathbf{F}_{2n}| = (1, \dots, 1) A_0 E_0 A_1 E_1 \cdots A_{n-1} E_{n-1} (1, \dots, 1)^\top = \|A_0 E_0 A_1 E_1 \cdots A_{n-1} E_{n-1}\|.$$

Taking in the MMG over $(\text{conv}(\mathcal{A}), \text{conv}(\mathcal{E}))$ the strategies ς and θ , which choose matrices $A_0, E_0, A_1, E_1, \dots$ we obtain the required:

$$P_{\text{EG}}(\sigma, \tau) = \limsup_{n \rightarrow \infty} |\mathbf{F}_{2n}|^{1/n} = \limsup_{n \rightarrow \infty} \|A_0 E_0 A_1 E_1 \cdots A_{n-1} E_{n-1}\|^{1/n} = P_{\text{MMG}}(\varsigma, \theta).$$

It is easy to see that for positional σ our construction gives $A_n = A_\sigma$ for all n . \square

Proof of Lemma 10. As stated in [13, Corollary 8.1.29], for any nonnegative matrix A and $u > 0$

$$\alpha u \leq Au \leq \beta u \Rightarrow \alpha \leq \rho(A) \leq \beta, \quad (19)$$

our statement (i) is now immediate. Let us prove the three remaining assertions.

- (ii) Let $Au \leq \rho u$ for $u > 0$ with $A > 0$ and $Au \neq \rho u$. Then at least one coordinate of the vector $Au - \rho u \leq 0$ is strictly negative. Therefore the condition $A > 0$ implies strict negativity of all coordinates of the vector $A(Au - \rho u)$. Then there exists $\varepsilon > 0$ such that $A(Au - \rho u) \leq -\varepsilon u$ and therefore $A^2 u = A(Au - \rho u) + \rho Au \leq (\rho^2 - \varepsilon)u$. Then, by (19), we get $\rho(A^2) \leq \rho^2 - \varepsilon$, and thus $\rho(A) \leq \sqrt{\rho^2 - \varepsilon} < \rho$, q.e.d.
- (iii) The condition $Au \geq \rho u$ with non-zero $u \geq 0$ implies $A^n u \geq \rho^n u$ for any $n \geq 1$. Then $\|A^n\| \cdot \|u\| \geq \|A^n u\| \geq \rho^n \|u\|$. Therefore $\|A^n\| \geq \rho^n$, and by Gelfand's formula $\rho(A) \geq \rho$, q.e.d.
- (iv) Now let $A > 0$ and $Au \neq \rho u$. Then at least one coordinate of the vector $Au - \rho u \geq 0$ is strictly positive. Therefore the condition $A > 0$ implies strict positivity of all the coordinates of the vector $A(Au - \rho u)$. Then there exists $\varepsilon > 0$ such that $A(Au - \rho u) \geq \varepsilon u$ and therefore $A^2 u = A(Au - \rho u) + \rho Au \geq (\rho^2 + \varepsilon)u$. This, by (iii) applied to the matrix A^2 , implies $\rho(A^2) \geq \rho^2 + \varepsilon$, and thus $\rho(A) \geq \sqrt{\rho^2 + \varepsilon} > \rho$, q.e.d. \square

Proof of Lemma 11. To prove (i), we represent the vectors u and v in coordinate form:

$$u = (u_1, u_2, \dots, u_M)^\top, \quad v = (v_1, v_2, \dots, v_N)^\top.$$

Suppose that for some matrix $A = (a_{ij}) \in \mathcal{A}$ the inequality $Au \geq v$ fails. Then

$$a_{i1}u_1 + a_{i2}u_2 + \cdots + a_{iM}u_M < v_i$$

for some $i \in \{1, 2, \dots, N\}$; we may assume $i = 1$ without loss of generality. In this case, the matrix

$$\bar{A} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1M} \\ \tilde{a}_{21} & \tilde{a}_{22} & \cdots & \tilde{a}_{2M} \\ \cdots & \cdots & \cdots & \cdots \\ \tilde{a}_{N1} & \tilde{a}_{N2} & \cdots & \tilde{a}_{NM} \end{bmatrix},$$

obtained from the matrix $\tilde{A} = (\tilde{a}_{ij})$ replacing the first row by $a_1 = [a_{11}, a_{12}, \dots, a_{1M}]$, yields the inequalities $a_{11}u_1 + a_{12}u_2 + \cdots + a_{1M}u_M < v_1$; and $\tilde{a}_{i1}u_1 + \tilde{a}_{i2}u_2 + \cdots + \tilde{a}_{iM}u_M = v_i$ for $i = 2, 3, \dots, N$. Consequently, $\bar{A}u \leq v$ and $\bar{A}u \neq v$, which completes the proof of the first statement of the lemma. The proof of statement (ii) is similar. \square

Proof of Corollary 12. We denote $V = \text{mm}(\mathcal{A}, \mathcal{B})$ and $V' = \text{mm}(\text{conv}(\mathcal{A}), \text{conv}(\mathcal{B}))$. Then

$$V' \stackrel{1}{=} \min_{A \in \text{conv}(\mathcal{A})} \max_{B \in \text{conv}(\mathcal{B})} \rho(BA) \stackrel{2}{\leq} \min_{A \in \mathcal{A}} \max_{B \in \text{conv}(\mathcal{B})} \rho(BA) \stackrel{3}{=} \min_{A \in \mathcal{A}} \max_{B \in \mathcal{B}} \rho(BA) = V,$$

where 1 follows from the equality $\rho(AB) = \rho(BA)$, 2 from the inclusion $\mathcal{A} \subseteq \text{conv}(\mathcal{A})$, 3 from Lemma 1 and equalities (3). Symmetrically,

$$V' = \max_{B \in \text{conv}(\mathcal{B})} \min_{A \in \text{conv}(\mathcal{A})} \rho(AB) \geq \max_{B \in \mathcal{B}} \min_{A \in \text{conv}(\mathcal{A})} \rho(AB) = \max_{B \in \mathcal{B}} \min_{A \in \mathcal{A}} \rho(AB) = V. \quad \square$$

Proof of Lemma 15. Let σ' and τ' be arbitrary strategies in the EG, then by Lemma 7 for the strategy pair (σ', τ) there is a corresponding pair (ζ', E_τ) with some strategy ζ' having the same value in the MMG. Symmetrically for the pair (σ, τ') there is a corresponding pair (A_σ, θ') . We have:

$$P(\sigma', \tau) = P(\zeta', E_\tau) \leq P(A_\sigma, E_\tau) = P(\sigma, \tau) = P(A_\sigma, E_\tau) \leq P(A_\sigma, \theta') = P(\sigma, \tau'),$$

where the equalities come from Lemma 7 and the inequalities from the optimality of E_τ and A_σ , respectively. Thus σ and τ are optimal. \square

We will need the following result in order to prove Lemma 17.

Lemma A.3. *Let \mathcal{A} be a compact IRU-set of positive $(N \times N)$ -matrices.*

- (i) *If $\tilde{A} \in \mathcal{A}$ is a matrix satisfying $\rho(\tilde{A}) = \check{\rho}(\mathcal{A})$ and \tilde{v} is its positive eigenvector corresponding to the eigenvalue $\rho(\tilde{A})$, then $A\tilde{v} \geq \check{\rho}(\mathcal{A})\tilde{v}$ for all $A \in \mathcal{A}$.*
- (ii) *If $\tilde{A} \in \mathcal{A}$ is a matrix satisfying $\rho(\tilde{A}) = \hat{\rho}(\mathcal{A})$ and \tilde{v} is its positive eigenvector corresponding to the eigenvalue $\rho(\tilde{A})$, then $A\tilde{v} \leq \hat{\rho}(\mathcal{A})\tilde{v}$ for all $A \in \mathcal{A}$.*

Proof of Lemma A.3. To prove (i) let us note that $\tilde{A}\tilde{v} = \check{\rho}(\mathcal{A})\tilde{v}$. Then by Lemma 11(i) either $A\tilde{v} \geq \check{\rho}(\mathcal{A})\tilde{v}$ for all $A \in \mathcal{A}$ or there exists a matrix $\bar{A} \in \mathcal{A}$ such that $\bar{A}\tilde{v} \leq \check{\rho}(\mathcal{A})\tilde{v}$ and $\bar{A}\tilde{v} \neq \check{\rho}(\mathcal{A})\tilde{v}$. In the latter case, by Lemma 10 the inequality $\rho(\bar{A}) < \check{\rho}(\mathcal{A})$ would hold, which contradicts to the definition of $\check{\rho}(\mathcal{A})$. Hence, the inequality $A\tilde{v} \geq \check{\rho}(\mathcal{A})\tilde{v}$ holds for all $A \in \mathcal{A}$, q.e.d. Assertion (ii) is proved similarly. \square

Proof of Lemma 17. For positive matrices, implications \Leftarrow follow from Lemma 10. As for \Rightarrow , it suffices to take v the eigenvector (corresponding to the spectral radius) of the matrix $\tilde{A} \in \mathcal{A}$ with the largest (smallest) spectral radius, and to apply Lemma A.3.

As for non-negative matrices, we have four implications to prove:

(12), \Rightarrow Denote, for any $\varepsilon > 0$, $\mathcal{A}_\varepsilon = \{A + \varepsilon\mathbf{1} \mid A \in \mathcal{A}\}$. If $\hat{\rho}(\mathcal{A}) < \alpha$ then due to compactness of the set \mathcal{A} there exists $\varepsilon > 0$ such that $\hat{\rho}(\mathcal{A}_\varepsilon) = \hat{\rho}(\mathcal{A} + \varepsilon\mathbf{1}) < \alpha$. Then by (12) (already proved for positive matrices), there exists $v > 0$ such that $(A + \varepsilon\mathbf{1})v < \alpha v$ for all $A \in \mathcal{A}$. Since $Av \leq (A + \varepsilon\mathbf{1})v$, then $Av < \alpha v$ for all $A \in \mathcal{A}$, q.e.d.

(12), \Leftarrow Suppose there exists $v > 0$ such that $Av < \alpha v$ for all $A \in \mathcal{A}$. Then due to compactness of the set \mathcal{A} there exists $\varepsilon > 0$ such that $(A + \varepsilon\mathbf{1})v < \alpha v$ for all $A \in \mathcal{A}$. Therefore by (12) (for positive matrices) $\hat{\rho}(\mathcal{A} + \varepsilon\mathbf{1}) < \alpha$, and hence by the monotonicity of the spectral radius we obtain $\hat{\rho}(\mathcal{A}) < \alpha$, q.e.d.

(16), \Rightarrow Let $\check{\rho}(\mathcal{A}) \geq \alpha$, then by the monotonicity of the spectral radius it holds that $\check{\rho}(\mathcal{A} + \varepsilon\mathbf{1}) \geq \alpha$ for any $\varepsilon > 0$. Then by (15) (for positive matrices) for any $\varepsilon > 0$ exists a vector $v_\varepsilon > 0$ such that $\|v_\varepsilon\| = 1$ and

$$(A + \varepsilon\mathbf{1})v_\varepsilon \geq \alpha v_\varepsilon \tag{20}$$

for all $A \in \mathcal{A}$. Choose a sequence $\varepsilon_n \rightarrow 0$ for which the corresponding vectors v_{ε_n} converge to some vector $v \geq 0$ (let us point out that $\|v\| = 1$ and so it is non-zero). Then passing to the limit in (20) we obtain $Av \geq \alpha v$ for all $A \in \mathcal{A}$, q.e.d.

(16), \Leftarrow Suppose there exists a non-zero vector $v \geq 0$ such that $Av \geq \alpha v$ for all $A \in \mathcal{A}$. Then by Lemma 10, $\rho(A) \geq \alpha$ for all $A \in \mathcal{A}$ and hence $\check{\rho}(\mathcal{A}) \geq \alpha$, q.e.d. \square