



HAL
open science

Optimal bounds for bit-sizes of stationary distributions in finite Markov chains

Mateusz Skomra

► **To cite this version:**

Mateusz Skomra. Optimal bounds for bit-sizes of stationary distributions in finite Markov chains. 2022. hal-03765666

HAL Id: hal-03765666

<https://laas.hal.science/hal-03765666v1>

Preprint submitted on 31 Aug 2022

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.

OPTIMAL BOUNDS FOR BIT-SIZES OF STATIONARY DISTRIBUTIONS IN FINITE MARKOV CHAINS

MATEUSZ SKOMRA

ABSTRACT. An irreducible stochastic matrix with rational entries has a stationary distribution given by a vector of rational numbers. We give an upper bound on the lowest common denominator of the entries of this vector. Bounds of this kind are used to study the complexity of algorithms for solving stochastic mean payoff games. They are usually derived using the Hadamard inequality, but this leads to suboptimal results. We replace the Hadamard inequality with the Markov chain tree formula in order to obtain optimal bounds. We also adapt our approach to obtain bounds on the absorption probabilities of finite Markov chains and on the gains and bias vectors of Markov chains with rewards.

1. INTRODUCTION

In this note, we study the following problem. Suppose that $P \in [0, 1]^{n \times n}$ is an irreducible stochastic matrix whose entries are rational numbers with a common denominator $M \in \mathbb{N}$. Then, the stationary distribution $\pi \in]0, 1]^n$ of P is a vector with rational entries. Our aim is to obtain an optimal upper bound on the lowest common denominator of the numbers $(\pi_i)_{i=1}^n$, which bounds the number of bits needed to encode π .

1.1. Context and motivation. Our main motivation to study the problem stated above comes from the area of stochastic mean payoff games, which form a generalization of finite Markov decision processes. A stochastic mean payoff game is a zero-sum game played by two players (Min and Max) who move a token along the edges of a finite directed graph $\vec{\mathcal{G}} = (V, E)$. Some vertices of the graph are controlled by player Min, some are controlled by player Max, and some are controlled by nature, which moves the token according to some fixed probability distribution. Furthermore, each vertex v of the graph is equipped with an integer payoff $r_v \in \mathbb{Z}$. The players are supposed to play according to positional strategies, i.e., their decisions depend only on the current position of the token. In particular, if the token lands twice on the same vertex controlled by one of the two players, then this player makes the same decision on both occasions. As a consequence, once the strategies of the players are fixed, the movement of the token is described by a Markov chain (X_0, X_1, \dots) on V , where X_0 is the starting position of the token, and the randomness of this process comes only from the random decisions made by the nature. The payoff of player Max is given by the average reward criterion

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}(r_{X_0} + \dots + r_{X_N}).$$

Player Max aims to maximize this quantity, while player Min wants to minimize it. It is known that stochastic mean payoff games always have optimal strategies [27]. In other words, there exists a vector $\chi \in \mathbb{R}^V$, known as the *value* of the game, such that player Max has a strategy that guarantees that the payoff is not smaller than χ_v for all initial states v . Likewise, player Min has a strategy that guarantees that the payoff is not greater than χ_v . In particular, if both players play optimally, then the final payoff is equal to χ_v . Stochastic mean payoff games attracted a significant interest in the computer science literature thanks to their uncertain complexity status. Even though optimal strategies exist, finding them algorithmically is a nontrivial task. In particular, it is not known if these strategies can be found in polynomial time and this problem has been open for 30 years, even in some restrictive cases (deterministic mean payoff games and parity games) [15, 22, 18]. We refer the reader to [6, 34, 14, 24, 10, 23, 11, 16, 19] for more information about mean payoff games and the related algorithmic issues. We also note that the one-player variant of these games is equivalent to Markov decision processes with average reward criterion, studied for instance in [29].

Numerous algorithms for solving stochastic mean payoff games that are proposed in the literature, such as the value iteration algorithms or the pumping algorithm, approximate the value χ without knowing the optimal strategies of the game. When the value is approximated to a sufficient precision, a rounding procedure is used to find χ exactly. We refer to [15, 24, 8, 10] for examples of such algorithms. In order to use a rounding procedure, one needs to have a bound on the precision needed to recover χ . This is done by bounding the denominators of χ . Such a bound can be obtained using the Hadamard inequality, and this approach is used in [15, 8, 3, 10], but it leads to suboptimal results in many cases of interest. In this note, we propose to use a more combinatorial approach, based on the Markov chain tree formula [12, Lemma 3.2], to obtain optimized bounds. As noted above, χ is the payoff of player Max obtained when both players play optimally. Let (X_0, X_1, \dots) be the Markov chain obtained under the optimal strategies and let P be its transition matrix. Our basic case of interest arises when P is irreducible. Then, [29, Appendix A.4] shows that χ does not depend on the initial state, $\chi = \eta(1, 1, \dots, 1)$, and $\eta = r^T \pi$, where π is the stationary distribution of P . Thus, the denominator of η is not greater than the lowest common denominator of $(\pi_v)_{v \in V}$, which leads to the problem stated in the first paragraph of this note.

1.2. Main results. Throughout this note, we use the following notation. We denote $[m] = \{1, \dots, m\}$ for $m \in \mathbb{N}$. Furthermore, let $P \in [0, 1]^n$ be a stochastic matrix with rational entries and let $\mathcal{X} := (X_0, X_1, \dots)$ be a Markov chain on the state space $V := [n]$ with transition matrix P . In general, we do not suppose that P is irreducible, since most of our results do not require this assumption. For every $i \in [n]$, let M_i be the lowest common denominator of the entries in the i th row of P , and let M be the lowest common denominator of all the entries of P . We also put

$$D := M_1 M_2 \dots M_n.$$

Moreover, we denote by $C_1, \dots, C_p \subset [n]$ the recurrent classes of \mathcal{X} and, for all $\ell \in [p]$, we denote by $\pi^{(\ell)} \in]0, 1]^{C_\ell}$ the stationary distribution on C_ℓ . Furthermore, let $r \in \mathbb{Z}^n$ be a vector of integer numbers and let $\chi \in \mathbb{R}^n$ be the *gain* vector defined as

$$\forall i, \chi_i := \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}(r_{X_0} + \dots + r_{X_N} \mid X_0 = i).$$

We note that χ is well defined and given by

$$\forall i, \chi_i = \sum_{\ell=1}^p \psi(i, C_\ell) \eta^{(\ell)}, \quad (1)$$

where $\psi(i, C_\ell)$ denotes the probability that the Markov chain starting at i reaches C_ℓ , and $\eta^{(\ell)} := \sum_{j \in C_\ell} r_j \pi_j^{(\ell)}$ for all $\ell \in [p]$. The formula (1) follows from the ergodic theorem of finite Markov chains, see [29, Appendix A.4] and [13, Part I, §6–§9] for detailed information. Our main result is the following theorem for irreducible matrices and its corollary, which holds even if P is not irreducible.

Theorem 1.1. *Suppose that P is irreducible and let $\pi \in]0, 1]^n$ denote its stationary distribution. Then, π is a vector of rational numbers whose lowest common denominator is not greater than $\min\{nD, nM^{n-1}\}$.*

Corollary 1.2. *Suppose that $\chi = \eta(1, 1, \dots, 1)$ for some $\eta \in \mathbb{R}$. Then, η is a rational number with denominator not greater than $\min\{nD, nM^{n-1}\}$.*

Before discussing these results, observe that $M \leq D \leq M^k$, where k is the number of rows of P that have at least two nonzero entries, $k := \{i \in [n] : \exists j, 0 < P_{ij} < 1\}$. In particular, we have the inequality

$$\min\{nD, nM^{n-1}\} \leq nM^{\min\{k, n-1\}}. \quad (2)$$

As noted above, our proof of Theorem 1.1 relies on a combinatorial formula for stationary distributions, known as the Markov chain tree formula [12, Lemma 3.2]. By comparison, [3, Lemma 4.10] uses the Hadamard inequality to obtain a bound $n^{n/2} M^n$ for the same problem. A more precise application of the Hadamard inequality is used in [10, Lemma 6] to obtain a bound of the form $kn(2M)^{k+1}$. The inequality (2) shows that our estimate is better than both of these bounds. Even more, in Proposition 3.2 we show that our bound is essentially optimal, in the sense that it cannot be improved even by a multiplicative constant. The proof of Proposition 3.2 also shows that this bound remains optimal even if we only want to bound the denominators of π_i separately. We also note that the interest of having bounds that depend on k is that these types of bounds may be used for stochastic mean payoff games with bounded number of states controlled by nature, see [21, 24, 8, 10] for more discussion. Furthermore, we point out that (1) shows that the assumption of Corollary 1.2 is satisfied if \mathcal{X} is irreducible or has only one recurrent class, but it may also be satisfied even if \mathcal{X} has multiple recurrent classes. We refer to [2, 10] for conditions that ensure that a stochastic mean payoff game has a value that does not depend on the initial state.

Our next result is especially useful in the situation in which \mathcal{X} has an absorbing state, $P_{jj} = 1$, and we put $r_j := 1$ and $r_i := 0$ for all $i \neq j$.

In this case, (1) shows that the gain χ_i is equal to the probability that the Markov chain starting at i reaches j . This situation arises in simple stochastic games [15] and some of its generalizations [21, 8, 7]. In particular, in order to bound the denominator of χ_i , we want to bound the denominators of absorption probabilities (which are rational numbers). To do so, we use an adaptation of the Markov chain tree formula to absorption probabilities, following the approach presented in [12]. In order to state our estimate, let $T \subset [n]$ denote the set of transient states of the Markov chain \mathcal{X} and put $D_T := \prod_{i \in T} M_i$. Our approach gives the following result.

Theorem 1.3. *The numbers $(\psi(i, C_\ell))_{i \in [n], \ell \in [p]}$ are rational and their lowest common denominator is not greater than $\min\{D_T, M^{n-2}\}$.*

As previously, we have the inequality $D_T \leq M^{k_T}$, where $k_T \leq k$ denotes the number of rows of P that have at least two nonzero entries and represent transient states of the Markov chain, i.e.,

$$k_T = \{i \in [n] : i \text{ is transient and there exists } j \in [n] \text{ such that } 0 < P_{ij} < 1\}.$$

In particular, we have

$$\min\{D_T, M^{n-2}\} \leq M^{\min\{k_T, n-2\}} \quad (3)$$

and this bound improves the bounds obtained in the literature using the Hadamard inequality [15, 8]. Furthermore, this bound is tight as shown in Example 3.5. Combining Corollary 1.2, Theorem 1.3, and (1), we obtain the following estimate on the gain vector in general chains.

Corollary 1.4. *The numbers $(\chi_i)_{i \in [n]}$ are rational and their lowest common denominator is not greater than $3^{s/2}D$, where s denotes the number of recurrent states of a Markov chain with transition matrix P .*

The main difference between the bounds of Corollary 1.2 and Corollary 1.4 is that the latter bound is exponential in s . The example presented in [9] shows that this is unavoidable in general chains even if $k = 1$.

To state our final result, we recall the notion of a bias vector. Given P and r , [29, Theorem 8.2.6] shows that the gain χ can be found by solving the system of equalities

$$\begin{cases} P\chi &= \chi \\ Pu &= \chi + u - r \end{cases} \quad (4)$$

in variables $(\chi, u) \in \mathbb{R}^{2n}$. More precisely, (4) has a solution and any such solution (χ', u) satisfies $\chi' = \chi$. If (χ, u) is a solution of (4), then we say that u is a *bias* vector. In general, a bias vector is not unique, even up to an additive constant. Bias vectors play an important role in the policy iteration algorithms for Markov decision processes [29, Chapter 9] and for stochastic mean payoff games [1]. Moreover, in a recent work, Allamigeon, Gaubert, Katz, and Skomra proposed a condition number for stochastic mean payoff games that governs the complexity of the value iteration algorithm [4, 5]. This condition number depends on the quantity $\inf_u \|u\|_H$, where the infimum goes over all bias vectors of the Shapley operator associated with a stochastic mean payoff game, and $\|\cdot\|_H$ denotes the *Hilbert seminorm*,

$\|u\|_H := \max_i u_i - \min_i u_i$. Since $\|u\|_H \leq 2\|u\|_\infty$, one can use the supremum norm to bound $\inf_u \|u\|_H$. The complexity estimates on value iteration obtained in [5] rely on Theorem 1.1 and on the following result.

Theorem 1.5. *Suppose that $u \in \mathbb{R}^n$ is a bias vector of (r, P) . We have the following estimates:*

- i) *if for every $\ell \in [p]$ there exists $i_\ell \in C_\ell$ such that $u_{i_\ell} = 0$, then $\|u\|_\infty \leq 2\|r\|_\infty n \min\{D, M^{n-1}\}$;*
- ii) *if $\sum_{i \in C_\ell} u_i \pi_i^{(\ell)} = 0$ for all $\ell \in [p]$, then $\|u\|_\infty \leq 4\|r\|_\infty n \min\{D, M^{n-1}\}$.*

The discussion in [29, Section 8.2.3] implies that bias vectors of both kinds exist for any pair (r, P) . Furthermore, these types of bias vectors are particularly useful in the policy iteration algorithms for Markov decision processes [29, Section 9.2].

A preliminary version of the results presented in this note appeared in the PhD thesis of the author [32, Chapter 8]. We note that Theorem 1.3, which was only briefly mentioned in [32, Remark 8.46], has since been obtained independently by Auger, Badin de Montjoye, and Strozecski [7, Theorem 23], using a similar technique (the proof in [7] is based on the matrix tree theorem).

1.3. Organization of the paper. The rest of the paper is organized as follows. In Section 2.1 we present the necessary notions on directed trees and forests, which are used in the Markov chain tree formula. In Section 2.2 we present this formula and its adaptation to absorption probabilities. Section 3.1 contains the proofs of Theorems 1.1 and 1.3 and their corollaries. Finally, we present the proof of Theorem 1.5 in Section 3.2.

2. PRELIMINARIES

2.1. Rooted trees and forests. Let $\vec{\mathcal{G}} = (V, E)$ be a directed graph. In this paper, we allow a directed graph to have loops, but not multiple edges, i.e., we suppose that E is a subset of $\{(u, v) : u, v \in V\}$. If $E' \subset E$ is a subset of edges, then we denote by $\vec{\mathcal{G}}(E') = (V, E')$ the subgraph that consists of all the vertices of $\vec{\mathcal{G}}$, but the edges taken only from E' .

Definition 2.1. Let $E' \subset E$. We say that the graph $\vec{\mathcal{G}}(E')$ is a *rooted forest* if it does not have any directed cycles and every vertex of $\vec{\mathcal{G}}(E')$ has at most one outgoing edge. We say that a vertex $v \in V$ is a *root* of $\vec{\mathcal{G}}(E')$ if it has no outgoing edges. We say that $\vec{\mathcal{G}}(E')$ is a *rooted tree* if it is a rooted forest and has exactly one root.

(In Definition 2.1 we use the convention that a loop is a directed cycle, so that a rooted forest does not have any loops.) Figure 1 depicts a rooted forest. Before presenting the relationship between the rooted forests and Markov chains, let us give a few comments about Definition 2.1. First, we point out that a rooted forest is, indeed, a forest, i.e., it does not contain any undirected cycle. This follows from the fact that any undirected cycle that does not come from a directed cycle contains a vertex with two outgoing edges. Second, we note that the number of connected components of a rooted forest is equal to the number of its roots. This follows from the

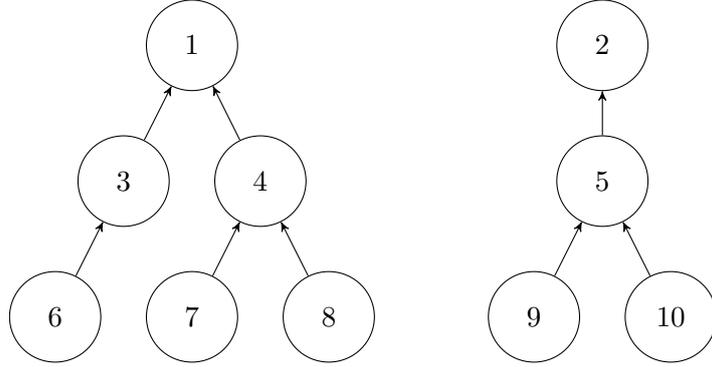


FIGURE 1. A rooted forest with two roots.

fact that from every vertex there is a unique directed path leading to a root. In particular, a rooted tree has one connected component, i.e., it is a tree. The following remark explains the vocabulary and conventions used in Definition 2.1.

Remark 2.2. Since we suppose that $\vec{\mathcal{G}}(E')$ contains all the vertices of $\vec{\mathcal{G}}$, the objects that we are considering are, in fact, *spanning* rooted forests and trees. Since we never consider forests that are not spanning, we drop the word “spanning” from the definition. We also point out that rooted forests are also called *branchings* and rooted trees are called *arborescences*. Furthermore, we note that in our definition, rooted forests are oriented in such a way that every edge points “towards the root.” However, the opposite convention is commonly used in the literature, cf. [17] or [30, Section 3.2]. The choice of orientation that we made in Definition 2.1 is justified by the fact that the orientation “towards the root” corresponds to the direction of transition between states in the Markov chains that we discuss in Section 2.2.

2.2. Combinatorial formulas for Markov chains. As in the introduction, we denote by $\mathcal{X} := (X_0, X_1, \dots)$ a Markov chain on the space $V := [n]$ with transition matrix P . We start by recalling the definition of an open set, see [25, § 3.5].

Definition 2.3. We say that a nonempty subset of sets $W \subset V$ is *open* if the chain starting at any state $v \in W$ can leave W with nonzero probability,

$$\forall v \in W, \mathbb{P}(\exists \ell, X_\ell \notin W \mid X_0 = v) > 0.$$

We note that the definition above does not imply that the chain cannot come back to W after leaving it. In particular, W can contain some recurrent states, but not a whole recurrent class. This also implies that the condition of Definition 2.3 can be replaced by a stronger one—if W is open and $v \in W$, then the chain starting at v will leave W almost surely.

If $W \subset V$ is open, we denote by $\tau_W := \inf\{\ell \geq 1 : X_\ell \notin W\}$ the moment when the chain leaves W for the first time. Furthermore, for every $v, w \in W$ we let

$$\zeta_{vw}^W := \mathbb{E}\left(\sum_{\ell=0}^{\tau_W-1} \mathbf{1}_{\{X_\ell=w\}} \mid X_0 = v\right)$$

be the expected number of visits in w before leaving W , provided that the chain starts at v . The following result gives a formula for computing ζ_{vw}^W .

Lemma 2.4 ([25, Theorem 3.5.4(1)]). *Let $\hat{P}^{W \times W}$ denote the submatrix of P formed by the rows and columns from W . Then, the matrix $(I - \hat{P})$ is invertible and $(I - \hat{P})_{vw}^{-1} = \zeta_{vw}^W$ for all $v, w \in W$.*

An alternative combinatorial formula for ζ_{vw}^W is given by Catoni in [12], where it is used to derive the Markov chain tree formula. In order to introduce it, we recall that a finite Markov chain is naturally associated with a directed graph representing its transitions. More precisely, we define the graph $\vec{\mathcal{G}} = (V, E)$ by the condition $(v, w) \in E \iff P_{vw} > 0$. We can consider that this graph is weighted, with weights $p(e)$ that correspond to the probabilities of transitions, i.e., if $e = (v, w) \in E$, then $p(e) := P_{vw}$. We also extend this definition to subsets of edges by taking the product. More precisely, if $E' \subset E$, then we define the *weight* of E' as the product of the weights of its elements,

$$p(E') := \prod_{e \in E'} p(e) = \prod_{(v,w) \in E'} P_{vw}.$$

The formulas that we present below consider only subsets of edges that give rise to rooted forests and trees. To this end, we denote

$$\mathcal{F}(R) := \{E' \subset E : \vec{\mathcal{G}}(E') \text{ is a rooted forest} \\ \text{and its set of roots is equal to } R\}.$$

It is also useful to consider the forests that contain a directed path between two fixed vertices. We denote

$$\mathcal{F}_{vw}(R) := \{E' \in \mathcal{F}(R) : \vec{\mathcal{G}}(E') \text{ contains a directed path from } v \text{ to } w\}.$$

We use the convention that if $v = w$, then $\mathcal{F}_{vw}(R) := \mathcal{F}(R)$. If $R = \{w'\}$ is a singleton, then we use the notation $\mathcal{F}(w')$ and $\mathcal{F}_{vw}(w')$ instead of $\mathcal{F}(\{w'\})$, $\mathcal{F}_{vw}(\{w'\})$. In this way, we can think of $\mathcal{F}(w')$ as the set of rooted trees whose root is w' .

Lemma 2.5 ([12, Lemma 3.1]). *Suppose that $W \subset V$ is open and denote $S := V \setminus W$. Then, for every $v, w \in W$ we have*

$$\zeta_{vw}^W = \left(\sum_{E' \in \mathcal{F}_{vw}(S \cup \{w\})} p(E') \right) \left(\sum_{E' \in \mathcal{F}(S)} p(E') \right)^{-1}.$$

Remark 2.6. We note that Lemma 2.5 is stated in [12] only for nontrivial subsets of irreducible chains, but the proof presented in [12] applies to arbitrary open sets of finite Markov chains.

Lemma 2.5 leads to the following two corollaries. The first one characterizes the stationary distributions of irreducible Markov chains. This corollary is known as the Markov chain tree formula, and was discovered by numerous authors [20, 26, 31, 33], see also [28] for more information. The second one characterizes the probabilities of absorption in different recurrent classes. We give the proof of the second corollary, since it is not stated in [12].

Corollary 2.7 (Markov chain tree formula, [12, Lemma 3.2]). *Suppose that the Markov chain is irreducible. Then, its stationary distribution $\pi \in]0, 1]^V$ is given by the formula*

$$\forall w \in V, \pi_w = \left(\sum_{E' \in \mathcal{F}(w)} p(E') \right) \left(\sum_{v \in V} \sum_{E' \in \mathcal{F}(v)} p(E') \right)^{-1}.$$

Corollary 2.8. *Let $S \subset V$ denote the set of all recurrent states of the Markov chain and let $C \subset S$ be a recurrent class. Suppose that $v \in V$ is a transient state and let $\psi(v, C)$ be the probability that the chain starting at v reaches C . Then, we have the equality*

$$\psi(v, C) = \left(\sum_{w \in C} \sum_{E' \in \mathcal{F}_{vw}(S)} p(E') \right) \left(\sum_{E' \in \mathcal{F}(S)} p(E') \right)^{-1}.$$

Proof. Let $W := V \setminus S$ denote the set of transient states and fix $w \in C$. Then, [25, Theorem 3.5.4] implies that the probability that the chain starting at $v \in W$ goes to w when it leaves W is equal to

$$\sum_{w' \in W} P_{w'w} \zeta_{vw'}^W.$$

Moreover, by Lemma 2.5 we have

$$\begin{aligned} \sum_{w' \in W} P_{w'w} \zeta_{vw'}^W &= \sum_{w' \in W} P_{w'w} \left(\sum_{E' \in \mathcal{F}_{vw'}(S \cup \{w'\})} p(E') \right) \left(\sum_{E' \in \mathcal{F}(S)} p(E') \right)^{-1} \\ &= \left(\sum_{w' \in W} \sum_{E' \in \mathcal{F}_{vw'}(S \cup \{w'\})} P_{w'w} p(E') \right) \left(\sum_{E' \in \mathcal{F}(S)} p(E') \right)^{-1} \\ &= \left(\sum_{E' \in \mathcal{F}_{vw}(S)} p(E') \right) \left(\sum_{E' \in \mathcal{F}(S)} p(E') \right)^{-1}. \end{aligned}$$

We obtain the claimed result by summing over $w \in C$. \square

3. PROOFS OF THE MAIN THEOREMS

In this section, we give the proofs of our main theorems. Section 3.1 contains the proofs of Theorems 1.1 and 1.3 and Section 3.2 contains the proof of Theorem 1.5.

3.1. Stationary distributions and absorption probabilities. The proof of our main theorem for irreducible chains relies on the following observation.

Lemma 3.1. *Let $Q \in \{D, M^{n-1}\}$ and suppose that $E' \subset E$ is such that $\vec{\mathcal{G}}(E')$ is a rooted tree. Then, $p(E')Q$ is a natural number.*

Proof. Since $\vec{\mathcal{G}}(E')$ has exactly $n - 1$ edges, $p(E')$ is a product of $n - 1$ rational numbers with common denominator M . Hence, $p(E')M^{n-1}$ is a natural number. Moreover, since every vertex of $\vec{\mathcal{G}}(E')$ has at most one outgoing edge, the product in $p(E')$ involves at most one number taken from the first row of P , at most one number taken from the second row of P , at most one number taken from the third row of P and so on. Therefore, $p(E')D$ is also a natural number. \square

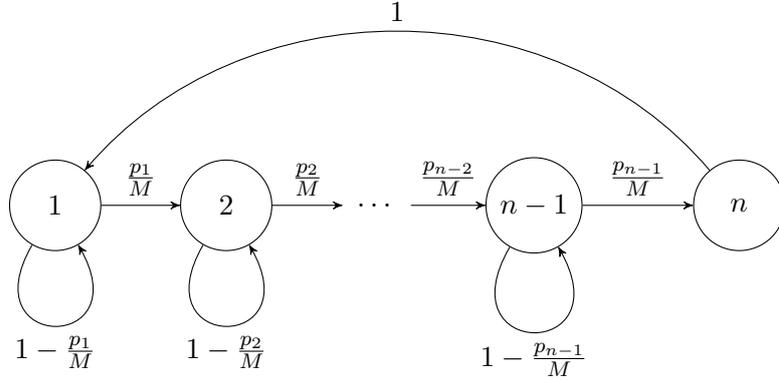


FIGURE 2. An irreducible Markov chain from Proposition 3.2.

Proof of Theorem 1.1. Let $Q \in \{D, M^{n-1}\}$. By Corollary 2.7, for every $w \in V$ we have

$$\pi_w = \frac{Q \sum_{E' \in \mathcal{F}(w)} p(E')}{Q \sum_{v \in V} \sum_{E' \in \mathcal{F}(v)} p(E')}.$$

By Lemma 3.1, the numerator and the denominator of the above fraction are natural numbers. In particular,

$$Q \sum_{v \in V} \sum_{E' \in \mathcal{F}(v)} p(E')$$

is a common denominator $(\pi_v)_v$. Furthermore, note that for every $v \in V$ we have

$$\sum_{E' \in \mathcal{F}(v)} p(E') \leq \prod_{w \neq v} (\sum_{w'} P_{ww'}) = 1. \quad (5)$$

Indeed, $\prod_{w \neq v} (\sum_{w'} P_{ww'})$ is the total weight of all graphs $\vec{\mathcal{G}}(E')$ in which v has no outgoing edges and every other vertex has exactly one outgoing edge. Since every tree rooted at v has these properties, we get (5). Therefore, the lowest common denominator of $(\pi_v)_v$ is not greater than

$$Q \sum_{v \in V} \sum_{E' \in \mathcal{F}(v)} p(E') \leq nQ. \quad \square$$

The next proposition shows that the bound of Theorem 1.1 is optimal for any fixed n . More precisely, let $g(n, M)$ denote the optimal bound that could be obtained in Theorem 1.1 for any fixed n, M , under the additional assumption that $D = M^{n-1}$. Then, we have the following result.

Proposition 3.2. *For every $n \geq 2$ we have*

$$\limsup_{M \rightarrow +\infty} \frac{g(n, M)}{nM^{n-1}} = 1.$$

Proof. We have $g(n, M) \leq nM^{n-1}$ by Theorem 1.1. To prove the opposite inequality, fix $n \geq 2$ and $\varepsilon > 0$. The following construction shows that we can find an arbitrarily large $M \geq 1$ and an irreducible Markov chain on n states such that the lowest common denominator of $(\pi_v)_v$ is not smaller than $(1 - \varepsilon)nM^{n-1}$. To do this, let $(m_1, m_2, \dots) = (2, 3, 5, \dots)$ be the

sequence of prime numbers and let $\delta > 0$ be such that $(\frac{1}{1+\delta})^{n-1} \geq 1 - \varepsilon$. For sufficiently large $q \geq n$ we have $m_n \leq \delta(m_q!)$. Take any such q and let $M = m_q! + m_n$, so that $M \leq (1 + \delta)m_q!$. Furthermore, let $p_1 = m_q! + m_1$, $p_2 = m_q! + m_2, \dots, p_n = m_q! + m_n = M$. We note that the numbers (p_1, \dots, p_n) are pairwise coprime. Indeed, if a prime number m divides $m_q! + m_i$ and $m_q! + m_j$, then it also divides $|m_i - m_j|$. Since $|m_i - m_j|$ is smaller than $m_n \leq m_q$, the number m also divides $m_q!$ and so it divides both m_i and m_j , which is a contradiction. Consider the irreducible Markov chain shown in Figure 2 and note that this chain satisfies $D = M^{n-1}$. Moreover, in this chain every state $v \in [n]$ is a root of exactly one rooted tree and the weight of this tree is equal to $\prod_{w \neq v} \frac{p_w}{M}$. Hence, if we denote $Q_v = p_1 \dots p_n / p_v$ for all v , then Corollary 2.7 shows that the stationary distribution of this chain is given by

$$\forall v, \pi_v = \frac{Q_v}{\sum_{w \in V} Q_w}. \quad (6)$$

We note that the numbers Q_v are natural. Even more, the fraction in (6) is simple. Indeed, if a prime number m divides Q_v , then it divides exactly one of p_1, \dots, p_n because these numbers are pairwise coprime. Suppose that m divides p_i . Then, m divides Q_w for all $w \neq i$ and it does not divide Q_i . In particular, m does not divide $\sum_{w \in V} Q_w$, showing that the fraction in (6) cannot be simplified. Therefore, the lowest common denominator of $(\pi_v)_v$ is equal to $\sum_{w \in V} Q_w$. Furthermore, we have

$$\sum_{w \in V} Q_w \geq np_1 \dots p_{n-1} \geq n(m_q!)^{n-1} \geq n\left(\frac{1}{1+\delta}\right)^{n-1} M^{n-1} \geq (1 - \varepsilon)nM^{n-1}. \quad \square$$

Remark 3.3. The fact that the fractions in (6) are simple implies that the bound of Theorem 1.1 remains optimal even if one only wants to bound the denominators of π_v separately.

The theorem for absorption probabilities follows by a similar argument. The following lemma is analogous to Lemma 3.1.

Lemma 3.4. *Let $Q \in \{D_T, M^{n-2}\}$ and let $S \subset V$ denote the set of recurrent states of \mathcal{X} . Suppose that $|S| \geq 2$ and that $E' \subset E$ is such that $\vec{\mathcal{G}}(E')$ is a forrest rooted at S . Then, $p(E')Q$ is a natural number.*

Proof. Since $|S| \geq 2$, the graph $\vec{\mathcal{G}}(E')$ has at least two roots, and so $p(E')$ is a product of at most $n - 2$ rational numbers with common denominator M . Thus, $p(E')M^{n-2}$ is a natural number. Furthermore, $p(E')$ is a product obtained by taking one number from each row of P that corresponds to a transient state of \mathcal{X} . Therefore, $p(E')D_T$ is also a natural number. \square

Proof of Theorem 1.3. Let $S \subset V$ denote the set of recurrent states of \mathcal{X} . If \mathcal{X} has exactly one recurrent state $w \in V$, then $\psi(v, \{w\}) = 1$ for all $v \neq w$ and the claim is trivial. From now on we suppose that $|S| \geq 2$. Let $Q \in \{D_T, M^{n-2}\}$. Then, Corollary 2.8 shows that for every $v \in V$ we have

$$\psi(v, C) = \frac{Q \sum_{w \in C} \sum_{E' \in \mathcal{F}_{vw}(S)} p(E')}{Q \sum_{E' \in \mathcal{F}(S)} p(E')}.$$

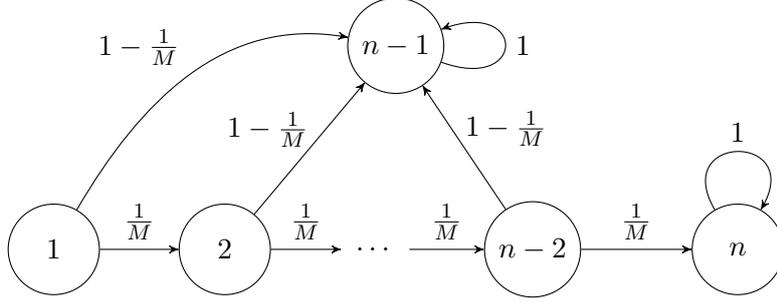


FIGURE 3. A Markov chain from Example 3.5.

By Lemma 3.4, both the numerator and the denominator of the fraction above are natural numbers. Furthermore, we have

$$\sum_{E' \in \mathcal{F}(S)} p(E') \leq \prod_{w \notin S} \left(\sum_{w'} P_{ww'} \right) = 1.$$

Hence, the lowest common denominator of $(\psi(v, C))_{v \in V, C \in \mathcal{C}}$ is not greater than

$$Q \sum_{E' \in \mathcal{F}(S)} p(E') \leq Q. \quad \square$$

The bound of Theorem 1.3 is attained for every value of n, M as shown by the next example.

Example 3.5. Consider the Markov chain depicted in Figure 3. It is clear that we have $D_T = M^{n-2}$ and $\psi(1, n) = 1/M^{n-2}$, showing that the bound of Theorem 1.3 is attained.

We now present the proofs of Corollaries 1.2 and 1.4.

Proof of Corollary 1.2. Let $v \in V$ be a recurrent state of \mathcal{X} belonging to some recurrent class $C \subset V$. Then, (1) shows that $\eta = \sum_{w \in C} r_w \pi_w$, where $\pi \in \mathbb{R}^C$ is the stationary distribution of C . Hence, by applying Theorem 1.1 to C we get that η is a rational number and that its denominator is not greater than $\min\{nD, nM^{n-1}\}$. \square

Proof of Corollary 1.4. Let C_1, \dots, C_p denote the recurrent classes of \mathcal{X} . Furthermore, for every $\ell \in [p]$ let $\pi^{(\ell)} \in \mathbb{R}^{C_\ell}$ be the stationary distribution on C_ℓ and let $D_\ell \in \mathbb{N}$ be defined as $D_\ell := \prod_{v \in C_\ell} M_v$. By combining Theorems 1.1 and 1.3 we get that the lowest common denominator of the numbers $\left((\psi(v, C_\ell))_{v \in V, \ell \in [p]}, (\pi_v^{(\ell)})_{\ell \in [p], v \in C_\ell} \right)$ is not greater than

$$|C_1| \dots |C_p| D_1 \dots D_p D_T = |C_1| \dots |C_p| D.$$

Therefore, (1) implies that the numbers $(\chi_v)_{v \in V}$ are rational and that their lowest common denominator is not greater than $|C_1| \dots |C_p| D$. By the inequality of arithmetic and geometric means we get $|C_1| \dots |C_p| \leq (s/p)^p$. Furthermore, the function $f: \mathbb{R}_{>0} \rightarrow \mathbb{R}$ defined as $f(p) := p \ln(s) - p \ln(p)$ achieves its maximum when $\ln(p) = \ln(s) - 1$, i.e., $p = s/e$. Therefore, we get $(s/p)^p \leq e^{s/e}$. Since $e^{1/e} < 1.5 < \sqrt{3}$, we obtain $|C_1| \dots |C_p| \leq 3^{s/2}$, which finishes the proof. \square

3.2. Estimating a bias vector. We now give our estimates concerning bias vectors. Our proof of Theorem 1.5 is based on the following lemma.

Lemma 3.6. *Suppose that $W \subset V$ is open and let $\hat{P}^{W \times W}$ denote the submatrix of P formed by the rows and columns from W . Then, we have $\|(I - \hat{P})^{-1}\|_\infty \leq \prod_{w \in W} M_w \leq \min\{D, M^{n-1}\}$.*

Proof. Let $D_W := \prod_{w \in W} M_w$. We have $D_W \leq \min\{D, M^{n-1}\}$ because W has at most $n - 1$ states. By Lemma 2.4, we have $(I - \hat{P})_{vw}^{-1} = \zeta_{vw}^W$ for all $v, w \in W$. Fix v, w and let $\tau_W = \inf\{\ell \geq 1: X_\ell \notin W\}$ denote the moment when the Markov chain leaves W for the first time. Moreover, let $Z := \sum_{\ell=0}^{\tau_W-1} \mathbf{1}_{\{X_\ell=w\}}$ denote the number of times the Markov chain visits w before leaving W . Under this notation, we have $\zeta_{vw}^W = \mathbb{E}(Z|X_0 = v)$. Furthermore, let $q \in [0, 1]$ denote the probability that the Markov chain starting at w goes back to w before leaving W , i.e., $q := \mathbb{P}(Z \geq 2|X_0 = w)$. Since W is open, there exists a simple path in $\vec{\mathcal{G}}$ that starts in w and ends in some state that is outside W . The probability that the Markov chain starting from w follows this path is not smaller than $1/D_W$. Therefore, $q \leq 1 - D_W^{-1}$. Furthermore, note that for all $t \geq 1$ we have $\mathbb{P}(Z \geq t|X_0 = v) \leq q^{t-1}$, because in order to achieve $Z \geq t$ the chain starting from v has to reach w and subsequently go back to w at least $t - 1$ times. Thus,

$$\zeta_{vw}^W = \mathbb{E}(Z|X_0 = v) = \sum_{t=1}^{\infty} \mathbb{P}(Z \geq t|X_0 = v) \leq \sum_{t=1}^{\infty} q^{t-1} = \frac{1}{1-q} \leq D_W. \quad \square$$

To prove Theorem 1.5, we start with the irreducible case and then move to the general case.

Lemma 3.7. *Theorem 1.5 is true when P is irreducible.*

Proof. We start by proving the case *i*). Let $u \in \mathbb{R}^V$ be a bias vector such that $u_n = 0$ (the proof is analogous if $u_v = 0$ for some other $v \in V$). Denote $W := V \setminus \{n\}$ and let $\hat{P} \in \mathbb{R}^{W \times W}$ be the matrix obtained from P by deleting the last row and column. Likewise, let $\hat{u}, \hat{r} \in \mathbb{R}^W$ be the vectors obtained from u, r by deleting their last coordinates. Let $\pi \in]0, 1]^V$ be the stationary distribution of P and denote $\eta = \sum_{v \in V} \pi_v r_v \in \mathbb{R}$, so that $\chi = \eta(1, 1, \dots, 1)$. We note that $|\eta| \leq \|r\|_\infty$. The definition of the bias vector and the fact that $u_n = 0$ imply the equality $\hat{P}\hat{u} = \eta + \hat{u} - \hat{r}$. Since W is an open set, Lemma 2.4 gives $\hat{u} = (I - \hat{P})^{-1}(-\eta + \hat{r})$. Hence, Lemma 3.6 shows that

$$\|\hat{u}\|_\infty \leq 2\|r\|_\infty(n-1) \min\{D, M^{n-1}\} \leq 2\|r\|_\infty n \min\{D, M^{n-1}\}.$$

To prove the second case, note that the kernel of the matrix $(I - P)$ is equal to $\{\lambda(1, 1, \dots, 1): \lambda \in \mathbb{R}\}$. Therefore, the set of bias vectors of (r, P) is given by $\{\lambda + \bar{u}: \lambda \in \mathbb{R}\}$, where \bar{u} is a bias such that $\bar{u}_n = 0$. Thus, all bias vectors have the same Hilbert seminorm and the previous case gives

$$\|u\|_H = \|\bar{u}\|_H \leq 2\|\bar{u}\|_\infty \leq 4\|r\|_\infty n \min\{D, M^{n-1}\}.$$

Since $\pi^T u = 0$, for all $v \in V$ we get

$$|u_v| = \left| \sum_{w \in V} \pi_w (u_v - u_w) \right| \leq \max_{w \in V} |u_v - u_w| \leq \|u\|_H \leq 4\|r\|_\infty n \min\{D, M^{n-1}\}.$$

□

Proof of Theorem 1.5. The proof of both cases follows from Lemma 3.7 using the same argument, so we focus only on i). Let $S := \cup_{\ell=1}^p C_\ell$ denote the set of recurrent states of \mathcal{X} and $W := V \setminus S$ denote the set of transient states. For every ℓ , let $P^{(\ell)} \in \mathbb{R}^{C_\ell \times C_\ell}$ denote the submatrix of P formed by the rows and columns with indices in C_ℓ , and let $u^{(\ell)}, r^{(\ell)} \in \mathbb{R}^{C_\ell}$ denote the restrictions of u, r to the indices from C_ℓ . The equations (1) and (4) imply that $u^{(\ell)}$ is a bias of $(r^{(\ell)}, P^{(\ell)})$. Therefore, Lemma 3.7 implies that for all $\ell \in [p]$ we have $\|u^{(\ell)}\|_\infty \leq 2\|r\|_\infty |S| \min\{D, M^{n-1}\}$. Let $\bar{P} \in \mathbb{R}^{W \times S}$ denote the submatrix of P formed by the rows from W and columns from S and $\hat{P} \in \mathbb{R}^{W \times W}$ denote the submatrix of P formed by the rows and columns from W . Define a vector $\bar{u} \in \mathbb{R}^S$ as $\forall v \in S, \bar{u}_v := u_v^{(\ell)}$, where ℓ is such that $v \in C_\ell$. Furthermore, let $\hat{\chi}, \hat{u}, \hat{r} \in \mathbb{R}^W$ denote the vectors χ, u, r restricted to the coordinates from W . By the definition of the bias vector we have $\hat{P}\hat{u} + \bar{P}\bar{u} = \hat{\chi} + \hat{u} - \hat{r}$. Since W is an open set, Lemma 2.4 gives $\hat{u} = (I - \hat{P})^{-1}(-\hat{\chi} + \hat{r} + \bar{P}\bar{u})$. Moreover, (1) implies that $\|\chi\|_\infty \leq \|r\|_\infty$. Hence, by Lemma 3.6 we get

$$\|(I - \hat{P})^{-1}(-\hat{\chi} + \hat{r})\|_\infty \leq 2\|r\|_\infty |W| \min\{D, M^{n-1}\}.$$

Furthermore, let $R := (I - \hat{P})^{-1}\bar{P} \in \mathbb{R}^{W \times S}$. By [25, Theorem 3.5.4], for every $w \in W, v \in S$, R_{wv} is the probability that the Markov chain starting at w goes to v when it leaves W . Hence $R_{wv} \geq 0$ and $\sum_{v \in S} R_{wv} = 1$ for all $w \in W$. Therefore, we get

$$\|(I - \hat{P})^{-1}\bar{P}\bar{u}\|_\infty \leq \|\bar{u}\|_\infty \leq 2\|r\|_\infty |S| \min\{D, M^{n-1}\}$$

and so $\|\hat{u}\|_\infty \leq 2\|r\|_\infty n \min\{D, M^{n-1}\}$. \square

REFERENCES

- [1] M. Akian, J. Cochet-Terrasson, S. Detournay, and S. Gaubert. Solving multichain stochastic games with mean payoff by policy iteration. In *52nd IEEE Annual Conference on Decision and Control (CDC)*, pages 1834–1841. IEEE, 2013.
- [2] M. Akian, S. Gaubert, and A. Hochart. Ergodicity conditions for zero-sum games. *Discrete Contin. Dyn. Syst.*, 35(9):3901–3931, 2015.
- [3] M. Akian, S. Gaubert, and A. Hochart. Generic uniqueness of the bias vector of finite zero-sum stochastic games with perfect information. *J. Math. Anal. Appl.*, 457:1038–1064, 2018.
- [4] X. Allamigeon, S. Gaubert, R. D. Katz, and M. Skomra. Condition numbers of stochastic mean payoff games and what they say about nonarchimedean semidefinite programming. In *Proceedings of the 23rd International Symposium on Mathematical Theory of Networks and Systems (MTNS)*, pages 160–167, 2018.
- [5] X. Allamigeon, S. Gaubert, R. D. Katz, and M. Skomra. Universal complexity bounds based on value iteration and application to entropy games. In *Proceedings of the 49th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 229 of *LIPICs. Leibniz Int. Proc. Inform.*, pages 110:1–110:20, Wadern, 2022. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- [6] D. Andersson and P. B. Miltersen. The complexity of solving stochastic games on graphs. In *Proceedings of the 20th International Symposium on Algorithms and Computation (ISAAC)*, volume 5878 of *Lecture Notes in Comput. Sci.*, pages 112–121. Springer, 2009.
- [7] D. Auger, X. Badin de Montjoye, and Y. Strozecki. A generic strategy iteration method for simple stochastic games. In *Proceedings of the 46th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 202

- of *LIPICs. Leibniz Int. Proc. Inform.*, pages 12:1–12:22, Wadern, 2021. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- [8] D. Auger, P. Coucheney, and Y. Strozecki. Finding optimal strategies of almost acyclic simple stochastic games. In *Proceedings of the 11th Annual Conference on Theory and Applications of Models of Computation (TAMC)*, volume 8402 of *Lecture Notes in Comput. Sci.*, pages 67–85. Springer, 2014.
 - [9] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. On discounted approximations of undiscounted stochastic games and Markov decision processes with limited randomness. *Oper. Res. Lett.*, 41(4):357–362, 2013.
 - [10] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. A pseudo-polynomial algorithm for mean payoff stochastic games with perfect information and few random positions. *Inform. and Comput.*, 267:74–95, 2019.
 - [11] C. S. Calude, S. Jain, B. Khousainov, W. Li, and F. Stephan. Deciding parity games in quasipolynomial time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 252–263. ACM, 2017.
 - [12] O. Catoni. Simulated annealing algorithms and Markov chains with rare transitions. In J. Azéma, M. Émery, M. Ledoux, and M. Yor, editors, *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Math.*, pages 69–119. Springer, Berlin, 1999.
 - [13] K. L. Chung. *Markov Chains With Stationary Transition Probabilities*, volume 104 of *Grundlehren Math. Wiss.* Springer, Heidelberg, 1967.
 - [14] C. Comin and R. Rizzi. Improved pseudo-polynomial bound for the value problems and optimal strategy synthesis in mean payoff games. *Algorithmica*, 77(4):995–1021, 2017.
 - [15] A. Condon. The complexity of stochastic games. *Inform. and Comput.*, 96(2):203–224, 1992.
 - [16] W. Czerwiński, L. Daviaud, N. Fijalkow, M. Jurdziński, R. Lazić, and P. Parys. Universal trees grow inside separating automata: Quasi-polynomial lower bounds for parity games. In *Proceedings of the 2019 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2333–2349. SIAM, 2019.
 - [17] J. Edmonds. Optimum branchings. *J. Res. Bur. Stand.*, 71B(4):233–240, 1967.
 - [18] E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy. In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 368–377. IEEE, 1991.
 - [19] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, New York, 2007.
 - [20] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems*, volume 260 of *Grundlehren Math. Wiss.* Springer, Heidelberg, 2012.
 - [21] H. Gimbert and F. Horn. Simple stochastic games with few random vertices are easy to solve. In *Proceedings of the 11th International Conference on Foundations of Software Science and Computational Structures (FoSSaCS)*, volume 4962 of *Lecture Notes in Comput. Sci.*, pages 5–19. Springer, 2008.
 - [22] V. A. Gurvich, A. V. Karzanov, and L. G. Khachiyan. Cyclic games and finding minimax mean cycles in digraphs. *Zh. Vychisl. Mat. Mat. Fiz.*, 28(9):1406–1417, 1988.
 - [23] 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.
 - [24] R. Ibsen-Jensen and P. B. Miltersen. Solving simple stochastic games with few coin toss positions. In *Proceedings of the 20th Annual European Symposium on Algorithms (ESA)*, volume 7501 of *Lecture Notes in Comput. Sci.*, pages 636–647. Springer, 2012.
 - [25] J. G. Kemeny and J. L. Snell. *Finite Markov Chains*. Undergrad. Texts Math. Springer, New York, 1976.
 - [26] H.-H. Kohler and E. Vollmerhaus. The frequency of cyclic processes in biological multistate systems. *J. Math. Biol.*, 9:275–290, 1980.
 - [27] T. M. Liggett and S. A. Lippman. Stochastic games with perfect information and time average payoff. *SIAM Rev.*, 11(4):604–607, 1969.
 - [28] J. Pitman and W. Tang. Tree formulas, mean first passage times and Kemeny’s constant of a Markov chain. *Bernoulli*, 24(3):1942–1972, 2018.

- [29] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Ser. Probab. Stat. Wiley, Hoboken, NJ, 2005.
- [30] A. Schrijver. *Combinatorial Optimization. Polyhedra and Efficiency*, volume 24 of *Algorithms Combin.* Springer, Berlin, 2003.
- [31] B. O. Shubert. A flow-graph formula for the stationary distribution of a Markov chain. *IEEE Trans. Syst. Man Cybern. Syst.*, SMC-5(5):565–566, 1975.
- [32] M. Skomra. *Tropical spectrahedra: Application to semidefinite programming and mean payoff games*. PhD thesis, Université Paris-Saclay, 2018.
- [33] J. J. Solberg. A graph theoretic formula for the steady state distribution of finite Markov processes. *Manag. Sci.*, 21(9):1040–1048, 1975.
- [34] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoret. Comput. Sci.*, 158(1–2):343–359, 1996.

LAAS-CNRS, UNIVERSITÉ DE TOULOUSE, CNRS, TOULOUSE, FRANCE
Email address: `mateusz.skomra@laas.fr`