



Improved and Generalized Upper Bounds on the Complexity of Policy Iteration

Bruno Scherrer

► **To cite this version:**

Bruno Scherrer. Improved and Generalized Upper Bounds on the Complexity of Policy Iteration. Neural Information Processing Systems (NIPS) 2013, Dec 2013, South Lake Tahoe, United States. 2013.

HAL Id: hal-00921261

<https://hal.inria.fr/hal-00921261>

Submitted on 20 Dec 2013

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.

Improved and Generalized Upper Bounds on the Complexity of Policy Iteration

Bruno Scherrer

Inria, Villers-lès-Nancy, F-54600, France

Université de Lorraine, LORIA, UMR 7503, Vandoeuvre-lès-Nancy, F-54506, France

bruno.scherrer@inria.fr

Abstract

Given a Markov Decision Process (MDP) with n states and m actions per state, we study the number of iterations needed by Policy Iteration (PI) algorithms to converge to the optimal γ -discounted optimal policy. We consider two variations of PI: Howard's PI that changes the actions in all states with a positive advantage, and Simplex-PI that only changes the action in the state with maximal advantage. We show that Howard's PI terminates after at most $n(m-1) \left\lceil \frac{1}{1-\gamma} \log \left(\frac{1}{1-\gamma} \right) \right\rceil = O\left(\frac{nm}{1-\gamma} \log \left(\frac{1}{1-\gamma} \right)\right)$ iterations, improving by a factor $O(\log n)$ a result by [3], while Simplex-PI terminates after at most $n^2(m-1) \left(1 + \frac{2}{1-\gamma} \log \left(\frac{1}{1-\gamma} \right)\right) = O\left(\frac{n^2m}{1-\gamma} \log \left(\frac{1}{1-\gamma} \right)\right)$ iterations, improving by a factor $O(\log n)$ a result by [11]. Under some structural assumptions of the MDP, we then consider bounds that are independent of the discount factor γ : given a measure of the maximal transient time τ_t and the maximal time τ_r to revisit states in recurrent classes under all policies, we show that Simplex-PI terminates after at most $n^2(m-1) (\lceil \tau_r \log(n\tau_r) \rceil + \lceil \tau_r \log(n\tau_t) \rceil) [(m-1)\lceil n\tau_t \log(n\tau_t) \rceil + \lceil n\tau_t \log(n^2\tau_t) \rceil] = \tilde{O}(n^3m^2\tau_t\tau_r)$ iterations. This generalizes a recent result for deterministic MDPs by [8], in which $\tau_t \leq n$ and $\tau_r \leq n$. We explain why similar results seem hard to derive for Howard's PI. Finally, under the additional (restrictive) assumption that the state space is partitioned in two sets, respectively states that are transient and recurrent for all policies, we show that Howard's PI terminates after at most $n(m-1) (\lceil \tau_t \log n\tau_t \rceil + \lceil \tau_r \log n\tau_r \rceil) = \tilde{O}(nm(\tau_t + \tau_r))$ iterations while Simplex-PI terminates after $n(m-1) (\lceil n\tau_t \log n\tau_t \rceil + \lceil \tau_r \log n\tau_r \rceil) = \tilde{O}(n^2m(\tau_t + \tau_r))$ iterations.

1 Introduction

We consider a discrete-time dynamic system whose state transition depends on a control. We assume that there is a **state space** X of finite size n . At state $i \in \{1, \dots, n\}$, the control is chosen from a **control space** A of finite size¹ m . The control $a \in A$ specifies the **transition probability** $p_{ij}(a) = \mathbb{P}(i_{t+1} = j | i_t = i, a_t = a)$ to the next state j . At each transition, the system is given a reward $r(i, a, j)$ where r is the instantaneous **reward function**. In this context, we look for a stationary deterministic policy (a function $\pi : X \rightarrow A$ that maps

¹In the works of [11, 8, 3] that we reference, the integer “ m ” denotes the total number of actions, that is nm with our notation. When we restate their result, we do it with our own notation, that is we replace their “ m ” by “ nm ”.

states into controls²) that maximizes the expected discounted sum of rewards from any state i , called the **value of policy** π at state i :

$$v_\pi(i) := \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r(i_k, a_k, i_{k+1}) \middle| i_0 = i, \forall k \geq 0, a_k = \pi(i_k), i_{k+1} \sim \mathbb{P}(\cdot | i_k, a_k) \right]$$

where $\gamma \in (0, 1)$ is a discount factor. The tuple $\langle X, A, p, r, \gamma \rangle$ is called a **Markov Decision Process (MDP)** [9, 1], and the associated problem is known as **optimal control**.

The **optimal value** starting from state i is defined as

$$v_*(i) := \max_{\pi} v_\pi(i).$$

For any policy π , we write P_π for the $n \times n$ stochastic matrix whose elements are $p_{ij}(\pi(i))$ and r_π the vector whose components are $\sum_j p_{ij}(\pi(i))r(i, \pi(i), j)$. The value functions v_π and v_* can be seen as vectors on X . It is well known that v_π is the solution of the following Bellman equation:

$$v_\pi = r_\pi + \gamma P_\pi v_\pi,$$

that is v_π is a fixed point of the affine operator $T_\pi : v \mapsto r_\pi + \gamma P_\pi v$. It is also well known that v_* satisfies the following Bellman equation:

$$v_* = \max_{\pi} (r_\pi + \gamma P_\pi v_*) = \max_{\pi} T_\pi v_*$$

where the max operator is componentwise. In other words, v_* is a fixed point of the nonlinear operator $T : v \mapsto \max_{\pi} T_\pi v$. For any value vector v , we say that a policy π is **greedy with respect to the value** v if it satisfies:

$$\pi \in \arg \max_{\pi'} T_{\pi'} v$$

or equivalently $T_\pi v = T v$. With some slight abuse of notation, we write $\mathcal{G}(v)$ for any policy that is greedy with respect to v . The notions of optimal value function and greedy policies are fundamental to optimal control because of the following property: any policy π_* that is greedy with respect to the optimal value v_* is an **optimal policy** and its value v_{π_*} is equal to v_* .

Let π be some policy. We call **advantage with respect to** π the following quantity:

$$a_\pi = \max_{\pi'} T_{\pi'} v_\pi - v_\pi = T v_\pi - v_\pi.$$

We call the **set of switchable states of** π the following set

$$S_\pi = \{i, a_\pi(i) > 0\}.$$

Assume now that π is non-optimal (this implies that S_π is a non-empty set). For any non-empty subset Y of S_π , we denote $\text{switch}(\pi, Y)$ a policy satisfying:

$$\forall i, \text{switch}(\pi, Y)(i) = \begin{cases} \mathcal{G}(v_\pi)(i) & \text{if } i \in Y \\ \pi(i) & \text{if } i \notin Y. \end{cases}$$

The following result is well known (see for instance [9]).

Lemma 1. *Let π be some non-optimal policy. If $\pi' = \text{switch}(\pi, Y)$ for some non-empty subset Y of S_π , then $v_{\pi'} \geq v_\pi$ and there exists at least one state i such that $v_{\pi'}(i) > v_\pi(i)$.*

This lemma is the foundation of the well-known iterative procedure, called Policy Iteration (PI), that generates a sequence of policies (π_k) as follows.

$$\pi_{k+1} \leftarrow \text{switch}(\pi_k, Y_k) \text{ for some set } Y_k \text{ such that } \emptyset \subsetneq Y_k \subseteq S_{\pi_k}.$$

The choice for the subsets Y_k leads to different variations of PI. In this paper we will focus on two specific variations:

²Restricting our attention to stationary deterministic policies is not a limitation. Indeed, for the optimality criterion to be defined soon, it can be shown that there exists at least one stationary deterministic policy that is optimal [9].

- When for all iterations k , $Y_k = S_{\pi_k}$, that is one switches the actions in all states with positive advantage with respect to π_k , the above algorithm is known as Howard’s PI; it can be seen then that $\pi_{k+1} \in \mathcal{G}(v_{\pi_k})$.
- When for all k , Y_k is a singleton containing a state $i_k \in \arg \max_i a_{\pi_k}(i)$, that is if we only switch one action in the state with maximal advantage with respect to π_k , we will call it Simplex-PI³.

Since it generates a sequence of policies with increasing values, any variation of PI converges to the optimal policy in a number of iterations that is smaller than the total number of policies m^n . In practice, PI converges in very few iterations. On random MDP instances, convergence often occurs in time sub-linear in n . The aim of this paper is to discuss existing and provide new upper bounds on the number of iterations required by Howard’s PI and Simplex-PI that are much sharper than m^n .

In the next sections, we describe some known results—see [11] for a recent and comprehensive review—about the number of iterations required by Howard’s PI and Simplex-PI, along with some of our original improvements and extensions.⁴

2 Bounds with respect to a Fixed Discount Factor $\gamma < 1$

A key observation for both algorithms, that will be central to the results we are about to discuss, is that the sequence they generate satisfies some contraction property⁵. For any vector $u \in \mathbb{R}^n$, let $\|u\|_\infty = \max_{1 \leq i \leq n} |u(i)|$ be the max-norm of u . Let $\mathbf{1}$ be the vector of which all components are equal to 1.

Lemma 2 (Proof in Section A). *The sequence $(\|v_* - v_{\pi_k}\|_\infty)_{k \geq 0}$ built by Howard’s PI is contracting with coefficient γ .*

Lemma 3 (Proof in Section B). *The sequence $(\mathbf{1}^T(v_* - v_{\pi_k}))_{k \geq 0}$ built by Simplex-PI is contracting with coefficient $1 - \frac{1-\gamma}{n}$.*

Though this observation is widely known for Howard’s PI, it was to our knowledge never mentioned explicitly in the literature for Simplex-PI. These contraction properties have the following immediate consequence⁶.

Corollary 1. *Let $V_{\max} = \frac{\max_\pi \|r_\pi\|_\infty}{1-\gamma}$ be an upper bound on $\|v_\pi\|_\infty$ for all policies π . In order to get an ϵ -optimal policy, that is a policy π_k satisfying $\|v_* - v_{\pi_k}\|_\infty \leq \epsilon$, Howard’s PI requires at most $\left\lceil \frac{\log \frac{V_{\max}}{\epsilon}}{1-\gamma} \right\rceil$ iterations, while Simplex-PI requires at most $\left\lceil \frac{n \log \frac{n V_{\max}}{\epsilon}}{1-\gamma} \right\rceil$ iterations.*

These bounds depend on the precision term ϵ , which means that Howard’s PI and Simplex-PI are *weakly polynomial* for a fixed discount factor γ . An important breakthrough was recently achieved by [11] who proved that one can remove the dependency with respect to ϵ , and thus show that Howard’s PI and Simplex-PI are *strongly polynomial* for a fixed discount factor γ .

Theorem 1 ([11]). *Simplex-PI and Howard’s PI both terminate after at most $n(m-1) \left\lceil \frac{n}{1-\gamma} \log \left(\frac{n^2}{1-\gamma} \right) \right\rceil$ iterations.*

³In this case, PI is equivalent to running the simplex algorithm with the highest-pivot rule on a linear program version of the MDP problem [11].

⁴For clarity, all proofs are deferred to the Appendix. The first proofs about bounds for the case $\gamma < 1$ are given in the Appendix of the paper. The other proofs, that are more involved, are provided in the Supplementary Material.

⁵A sequence of non-negative numbers $(x_k)_{k \geq 0}$ is contracting with coefficient α if and only if for all $k \geq 0$, $x_{k+1} \leq \alpha x_k$.

⁶For Howard’s PI, we have: $\|v_* - v_{\pi_k}\|_\infty \leq \gamma^k \|v_* - v_{\pi_0}\|_\infty \leq \gamma^k V_{\max}$. Thus, a sufficient condition for $\|v_* - v_{\pi_k}\|_\infty < \epsilon$ is $\gamma^k V_{\max} < \epsilon$, which is implied by $k \geq \frac{\log \frac{V_{\max}}{\epsilon}}{1-\gamma} > \frac{\log \frac{V_{\max}}{\epsilon}}{\log \frac{1}{\gamma}}$. For Simplex-PI, we have $\|v_* - v_{\pi_k}\|_\infty \leq \|v_* - v_{\pi_k}\|_1 \leq \left(1 - \frac{1-\gamma}{n}\right)^k \|v_* - v_{\pi_0}\|_1 \leq \left(1 - \frac{1-\gamma}{n}\right)^k n V_{\max}$, and the conclusion is similar to that for Howard’s PI.

The proof is based on the fact that PI corresponds to the simplex algorithm in a linear programming formulation of the MDP problem. Using a more direct proof, [3] recently improved the result by a factor $O(n)$ for Howard’s PI.

Theorem 2 ([3]). *Howard’s PI terminates after at most $(nm + 1) \left\lceil \frac{1}{1-\gamma} \log \left(\frac{n}{1-\gamma} \right) \right\rceil$ iterations.*

Our first two results, that are consequences of the contraction properties (Lemmas 2 and 3), are stated in the following theorems.

Theorem 3 (Proof in Section C). *Howard’s PI terminates after at most $n(m - 1) \left\lceil \frac{1}{1-\gamma} \log \left(\frac{1}{1-\gamma} \right) \right\rceil$ iterations.*

Theorem 4 (Proof in Section D). *Simplex-PI terminates after at most $n(m - 1) \left\lceil \frac{n}{1-\gamma} \log \left(\frac{n}{1-\gamma} \right) \right\rceil$ iterations.*

Our result for Howard’s PI is a factor $O(\log n)$ better than the previous best result of [3]. Our result for Simplex-PI is only very slightly better (by a factor 2) than that of [11], and uses a proof that is more direct. Using more refined argument, we managed to also improve the bound for Simplex-PI by a factor $O(\log n)$.

Theorem 5 (Proof in Section E). *Simplex-PI terminates after at most $n^2(m - 1) \left(1 + \frac{2}{1-\gamma} \log \frac{1}{1-\gamma} \right)$ iterations.*

Compared to Howard’s PI, our bound for Simplex-PI is a factor $O(n)$ larger. However, since one changes only one action per iteration, each iteration may have a complexity lower by a factor n : the update of the value can be done in time $O(n^2)$ through the Sherman-Morrison formula, though in general each iteration of Howard’s PI, which amounts to compute the value of some policy that may be arbitrarily different from the previous policy, may require $O(n^3)$ time. Overall, both algorithms seem to have a similar complexity.

It is easy to see that the linear dependency of the bound for Howard’s PI with respect to n is optimal. We conjecture that the linear dependency of both bounds with respect to m is also optimal. The dependency with respect to the term $\frac{1}{1-\gamma}$ may be improved, but removing it is impossible for Howard’s PI and very unlikely for Simplex-PI. [2] describes an MDP for which Howard’s PI requires an exponential (in n) number of iterations for $\gamma = 1$ and [5] argued that this holds also when γ is in the vicinity of 1. Though a similar result does not seem to exist for Simplex-PI in the literature, [7] consider four variations of PI that all switch one action per iteration, and show through specifically designed MDPs that they may require an exponential (in n) number of iterations when $\gamma = 1$.

3 Bounds for Simplex-PI that are independent of γ

In this section, we will describe some bounds that do not depend on γ but that will be based on some structural assumptions of the MDPs. On this topic, [8] recently showed the following result for deterministic MDPs.

Theorem 6 ([8]). *If the MDP is deterministic, then Simplex-PI terminates after at most $O(n^5 m^2 \log^2 n)$ iterations.*

Given a policy π of a deterministic MDP, states are either on cycles or on paths induced by π . The core of the proof relies on the following lemmas that altogether show that cycles are created regularly and that significant progress is made every time a new cycle appears; in other words, significant progress is made regularly.

Lemma 4. *If the MDP is deterministic, after at most $nm \lceil 2(n - 1) \log n \rceil$ iterations, either Simplex-PI finishes or a new cycle appears.*

Lemma 5. *If the MDP is deterministic, when Simplex-PI moves from π to π' where π' involves a new cycle, we have*

$$\mathbf{1}^T(v_{\pi_*} - v_{\pi'}) \leq \left(1 - \frac{1}{n}\right) \mathbf{1}^T(v_{\pi_*} - v_{\pi}).$$

Indeed, these observations suffice to prove⁷ that Simplex-PI terminates after $O(n^4 m^2 \log \frac{n}{1-\gamma}) = \tilde{O}(n^4 m^2)$. Removing completely the dependency with respect to the discount factor γ —the term in $O(\log \frac{1}{1-\gamma})$ —requires a careful extra work described in [8], which incurs an extra term of order $O(n \log(n))$.

At a more technical level, the proof of [8] critically relies on some properties of the vector $x_\pi = (I - \gamma P_\pi^T)^{-1} \mathbf{1}$ that provides a discounted measure of state visitations along the trajectories induced by a policy π starting from a uniform distribution:

$$\forall i \in X, \quad x_\pi(i) = n \sum_{t=0}^{\infty} \gamma^t \mathbb{P}(i_t = i \mid i_0 \sim U, a_t = \pi(i_t)),$$

where U denotes the uniform distribution on the state space X . For any policy π and state i , we trivially have $x_\pi(i) \in \left(1, \frac{n}{1-\gamma}\right)$. The proof exploits the fact that $x_\pi(i)$ belongs to the set $(1, n)$ when i is on a path of π , while $x_\pi(i)$ belongs to the set $\left(\frac{1}{1-\gamma}, \frac{n}{1-\gamma}\right)$ when i is on a cycle of π . As we are going to show, it is possible to extend the proof of [8] to stochastic MDPs. Given a policy π of a stochastic MDP, states are either in *recurrent classes* or *transient classes* (these two categories respectively generalize those of cycles and paths). We will consider the following structural assumption.

Assumption 1. *Let $\tau_t \geq 1$ and $\tau_r \geq 1$ be the smallest constants such that for all policies π and all states i ,*

$$\begin{aligned} (1 \leq)x_\pi(i) &\leq \tau_t && \text{if } i \text{ is transient for } \pi, \text{ and} \\ \frac{n}{(1-\gamma)\tau_r} &\leq x_\pi(i) \left(\leq \frac{n}{1-\gamma} \right) && \text{if } i \text{ is recurrent for } \pi. \end{aligned}$$

The constant τ_t (resp. τ_r) can be seen as a measure of the time needed to leave transient states (resp. the time needed to revisit states in recurrent classes). In particular, when γ tends to 1, it can be seen that τ_t is an upper bound of the expected time \mathcal{L} needed to “Leave the set of transient states”, since for any policy π ,

$$\begin{aligned} \lim_{\gamma \rightarrow 1} \tau_t &\geq \frac{1}{n} \lim_{\gamma \rightarrow 1} \sum_{i \text{ transient for } \pi} x_\pi(i) = \sum_{t=0}^{\infty} \mathbb{P}(i_t \text{ transient for } \pi \mid i_0 \sim U, a_t = \pi(i_t)) \\ &= \mathbb{E}[\mathcal{L} \mid i_0 \sim U, a_t = \pi(i_t)]. \end{aligned}$$

Similarly, when γ is in the vicinity of 1, $\frac{1}{\tau_r}$ is the minimal asymptotic frequency⁸ in recurrent states given that one starts from a random uniform state, since for any policy π and recurrent state i :

$$\begin{aligned} \lim_{\gamma \rightarrow 1} \frac{1-\gamma}{n} x_\pi(i) &= \lim_{\gamma \rightarrow 1} (1-\gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{P}(i_t = i \mid i_0 \sim U, a_t = \pi(i_t)) \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{P}(i_t = i \mid i_0 \sim U, a_t = \pi(i_t)). \end{aligned}$$

With Assumption 1 in hand, we can generalize Lemmas 4-5 as follows.

Lemma 6. *If the MDP satisfies Assumption 1, after at most $n[(m-1)\lceil n\tau_t \log(n\tau_t) \rceil + \lceil n\tau_t \log(n^2\tau_t) \rceil]$ iterations either Simplex-PI finishes or a new recurrent class appears.*

⁷This can be done by using arguments similar to the proof of Theorem 4 in Section D.

⁸If the MDP is aperiodic and irreducible, and thus admits a stationary distribution ν_π for any policy π , one can see that

$$\frac{1}{\tau_r} = \min_{\pi, i \text{ recurrent for } \pi} \nu_\pi(i).$$

Lemma 7. *If the MDP satisfies Assumption 1, when Simplex-PI moves from π to π' where π' involves a new recurrent class, we have*

$$\mathbf{1}^T(v_{\pi_*} - v_{\pi'}) \leq \left(1 - \frac{1}{\tau_r}\right) \mathbf{1}^T(v_{\pi_*} - v_{\pi}).$$

From these generalized observations, we can deduce the following original result.

Theorem 7 (Proof in Appendix F of the Supp. Material). *If the MDP satisfies Assumption 1, then Simplex-PI terminates after at most*

$$n^2(m-1)(\lceil \tau_r \log(n\tau_r) \rceil + \lceil \tau_r \log(n\tau_t) \rceil) [(m-1)\lceil n\tau_t \log(n\tau_t) \rceil + \lceil n\tau_t \log(n^2\tau_t) \rceil]$$

iterations.

Remark 1. *This new result is a strict generalization of the result for deterministic MDPs. Indeed, in the deterministic case, we have $\tau_t \leq n$ and $\tau_r \leq n$, and it is easy to see that Lemmas 6, 7 and Theorem 7 respectively imply Lemmas 4, 5 and Theorem 6.*

An immediate consequence of the above result is that Simplex-PI is *strongly polynomial* for sets of MDPs that are much larger than the deterministic MDPs mentioned in Theorem 6.

Corollary 2. *For any family of MDPs indexed by n and m such that τ_t and τ_r are polynomial functions of n and m , Simplex-PI terminates after a number of steps that is polynomial in n and m .*

4 Similar results for Howard's PI?

One may then wonder whether similar results can be derived for Howard's PI. Unfortunately, and as quickly mentioned by [8], the line of analysis developed for Simplex-PI does not seem to adapt easily to Howard's PI, because simultaneously switching several actions can interfere in a way that the policy improvement turns out to be small. We can be more precise on what actually breaks in the approach we have described so far. On the one hand, it is possible to write counterparts of Lemmas 4 and 6 for Howard's PI (see Appendix G of the Supp. Material).

Lemma 8. *If the MDP is deterministic, after at most n iterations, either Howard's PI finishes or a new cycle appears.*

Lemma 9. *If the MDP satisfies Assumption 1, after at most $nm\lceil \tau_t \log n\tau_t \rceil$ iterations, either Howard's PI finishes or a new recurrent class appears.*

However, on the other hand, we did not manage to adapt Lemma 5 nor Lemma 7. In fact, it is unlikely that a result similar to that of Lemma 5 will be shown to hold for Howard's PI. In a recent deterministic example due to [4] to show that Howard's PI may require at most $O(n^2)$ iterations, new cycles are created every single iteration but the sequence of values satisfies⁹ for all iterations $k < \frac{n^2}{4} + \frac{n}{4}$ and states i ,

$$v_*(i) - v_{\pi_{k+1}}(i) \geq \left[1 - \left(\frac{2}{n}\right)^k\right] (v_*(i) - v_{\pi_k}(i)).$$

Contrary to Lemma 5, as k grows, the amount of contraction gets (exponentially) smaller and smaller. With respect to Simplex-PI, this suggests that Howard's PI may suffer from subtle specific pathologies. In fact, the problem of determining the number of iterations required by Howard's PI has been challenging for almost 30 years. It was originally identified as an open problem by [10]. In the simplest—deterministic—case, the question is still open: the currently best known lower bound is the $O(n^2)$ bound by [4] we have just mentioned, while the best known upper bound is $O\left(\frac{m^n}{n}\right)$ (valid for all MDPs) due to [6].

⁹This MDP has an even number of states $n = 2p$. The goal is to minimize the long term expected cost. The optimal value function satisfies $v_*(i) = -p^N$ for all i , with $N = p^2 + p$. The policies generated by Howard's PI have values $v_{\pi_k}(i) \in (p^{N-k-1}, p^{N-k})$. We deduce that for all iterations k and states i , $\frac{v_*(i) - v_{\pi_{k+1}}(i)}{v_*(i) - v_{\pi_k}(i)} \geq \frac{1+p^{-k-2}}{1+p^{-k}} = 1 - \frac{p^{-k} - p^{-k-2}}{1+p^{-k}} \geq 1 - p^{-k}(1 - p^{-2}) \geq 1 - p^{-k}$.

On the positive side, an adaptation of the line of proof we have considered so far can be carried out under the following assumption.

Assumption 2. *The state space X can be partitioned in two sets \mathcal{T} and \mathcal{R} such that for all policies π , the states of \mathcal{T} are transient and those of \mathcal{R} are recurrent.*

Indeed, under this assumption, we can prove for Howard's PI a variation of Lemma 7 introduced for Simplex-PI.

Lemma 10. *For an MDP satisfying Assumptions 1-2, suppose Howard's PI moves from π to π' and that π' involves a new recurrent class. Then*

$$\mathbb{1}^T(v_{\pi_*} - v_{\pi'}) \leq \left(1 - \frac{1}{\tau_r}\right) \mathbb{1}^T(v_{\pi_*} - v_{\pi}).$$

And we can deduce the following original bound (that also applies to Simplex-PI).

Theorem 8 (Proof in Appendix H of the Supp. Material). *If the MDP satisfies Assumptions 1-2, then Howard's PI terminates after at most $n(m-1)(\lceil \tau_t \log n \tau_t \rceil + \lceil \tau_r \log n \tau_r \rceil)$ iterations, while Simplex-PI terminates after at most $n(m-1)(\lceil n \tau_t \log n \tau_t \rceil + \lceil \tau_r \log n \tau_r \rceil)$ iterations.*

It should however be noted that Assumption 2 is rather restrictive. It implies that the algorithms converge on the recurrent states independently of the transient states, and thus the analysis can be decomposed in two phases: 1) the convergence on recurrent states and then 2) the convergence on transient states (given that recurrent states do not change anymore). The analysis of the first phase (convergence on recurrent states) is greatly facilitated by the fact that in this case, a new recurrent class appears every single iteration (this is in contrast with Lemmas 4, 6, 8 and 9 that were designed to show under which conditions cycles and recurrent classes are created). Furthermore, the analysis of the second phase (convergence on transient states) is similar to that of the discounted case of Theorems 3 and 4. In other words, if this last result sheds some light on the practical efficiency of Howard's PI and Simplex-PI, a general analysis of Howard's PI is still largely open, and constitutes our main future work.

A Contraction property for Howard's PI (Proof of Lemma 2)

For any k , using the facts that $\{\forall \pi, T_{\pi} v_{\pi} = v_{\pi}\}$, $\{T_{\pi_*} v_{\pi_{k-1}} \leq T_{\pi_k} v_{\pi_{k-1}}\}$ and $\{\text{Lemma 1 and } P_{\pi_k} \text{ is positive definite}\}$, we have

$$\begin{aligned} v_{\pi_*} - v_{\pi_k} &= T_{\pi_*} v_{\pi_*} - T_{\pi_*} v_{\pi_{k-1}} + T_{\pi_*} v_{\pi_{k-1}} - T_{\pi_k} v_{\pi_{k-1}} + T_{\pi_k} v_{\pi_{k-1}} - T_{\pi_k} v_{\pi_k} \\ &\leq \gamma P_{\pi_*} (v_{\pi_*} - v_{\pi_{k-1}}) + \gamma P_{\pi_k} (v_{\pi_{k-1}} - v_{\pi_k}) \leq \gamma P_{\pi_*} (v_{\pi_*} - v_{\pi_{k-1}}). \end{aligned}$$

Since $v_{\pi_*} - v_{\pi_k}$ is non negative, we can take the max norm and get: $\|v_{\pi_*} - v_{\pi_k}\|_{\infty} \leq \gamma \|v_{\pi_*} - v_{\pi_{k-1}}\|_{\infty}$.

B Contraction property for Simplex-PI (Proof of Lemma 3)

By using the fact that $\{v_{\pi} = T_{\pi} v_{\pi} \Rightarrow v_{\pi} = (I - \gamma P_{\pi})^{-1} r_{\pi}\}$, we have that for all pairs of policies π and π' .

$$\begin{aligned} v_{\pi'} - v_{\pi} &= (I - \gamma P_{\pi'})^{-1} r_{\pi'} - v_{\pi} = (I - \gamma P_{\pi'})^{-1} (r_{\pi'} + \gamma P_{\pi'} v_{\pi} - v_{\pi}) \\ &= (I - \gamma P_{\pi'})^{-1} (T_{\pi'} v_{\pi} - v_{\pi}). \end{aligned} \tag{1}$$

On the one hand, by using this lemma and the fact that $\{T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k} \geq 0\}$, we have for any k : $v_{\pi_{k+1}} - v_{\pi_k} = (I - \gamma P_{\pi_{k+1}})^{-1} (T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k}) \geq T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k}$, which implies that

$$\mathbb{1}^T (v_{\pi_{k+1}} - v_{\pi_k}) \geq \mathbb{1}^T (T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k}). \tag{2}$$

On the other hand, using Equation (1) and the facts that $\{\|(I - \gamma P_{\pi_*})^{-1}\|_{\infty} = \frac{1}{1-\gamma}\}$ and $\{(I - \gamma P_{\pi_*})^{-1} \text{ is positive definite}\}$, $\{\max_s T_{\pi_{k+1}} v_{\pi_k}(s) = \max_{s, \tilde{\pi}} T_{\tilde{\pi}} v_{\pi_k}(s)\}$ and

$\{\forall x \geq 0, \max_s x(s) \leq \mathbf{1}^T x\}$, we have:

$$\begin{aligned} v_{\pi_*} - v_{\pi_k} &= (I - \gamma P_{\pi_*})^{-1} (T_{\pi_*} v_{\pi_k} - v_{\pi_k}) \leq \frac{1}{1 - \gamma} \max_s T_{\pi_*} v_{\pi_k}(s) - v_{\pi_k}(s) \\ &\leq \frac{1}{1 - \gamma} \max_s T_{\pi_{k+1}} v_{\pi_k}(s) - v_{\pi_k}(s) \leq \frac{1}{1 - \gamma} \mathbf{1}^T (T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k}), \end{aligned}$$

which implies (using $\{\forall x, \mathbf{1}^T x \leq n \|x\|_\infty\}$) that

$$\mathbf{1}^T (T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k}) \geq (1 - \gamma) \|v_{\pi_*} - v_{\pi_k}\|_\infty \geq \frac{1 - \gamma}{n} \mathbf{1}^T (v_{\pi_*} - v_{\pi_k}). \quad (3)$$

Combining Equations (2) and (3), we get:

$$\begin{aligned} \mathbf{1}^T (v_{\pi_*} - v_{\pi_{k+1}}) &= \mathbf{1}^T (v_{\pi_*} - v_{\pi_k}) - \mathbf{1}^T (v_{\pi_{k+1}} - v_{\pi_k}) \\ &\leq \mathbf{1}^T (v_{\pi_*} - v_{\pi_k}) - \frac{1 - \gamma}{n} \mathbf{1}^T (v_{\pi_*} - v_{\pi_k}) = \left(1 - \frac{1 - \gamma}{n}\right) \mathbf{1}^T (v_{\pi_*} - v_{\pi_k}). \end{aligned}$$

C A bound for Howard's PI when $\gamma < 1$ (Proof of Theorem 3)

For any k , by using Equation (1) and the fact $\{v_* - v_{\pi_k} \geq 0$ and P_{π_k} positive definite $\}$, we have:

$$v_* - T_{\pi_k} v_* = (I - \gamma P_{\pi_k})(v_* - v_{\pi_k}) \leq v_* - v_{\pi_k}.$$

Since $v_* - T_{\pi_k} v_*$ is non negative, we can take the max norm and, using Lemma 2, Equation (1) and the fact that $\{\|(I - \gamma P_{\pi_0})^{-1}\|_\infty = \frac{1}{1 - \gamma}\}$, we get:

$$\begin{aligned} \|v_* - T_{\pi_k} v_*\|_\infty &\leq \|v_* - v_{\pi_k}\|_\infty \leq \gamma^k \|v_{\pi_*} - v_{\pi_0}\|_\infty \\ &= \gamma^k \|(I - \gamma P_{\pi_0})^{-1} (v_* - T_{\pi_0} v_*)\|_\infty \leq \frac{\gamma^k}{1 - \gamma} \|v_* - T_{\pi_0} v_*\|_\infty. \end{aligned} \quad (4)$$

By definition of the max-norm, there exists a state s_0 such that $v_*(s_0) - [T_{\pi_0} v_*](s_0) = \|v_* - T_{\pi_0} v_*\|_\infty$. From Equation (4), we deduce that for all k ,

$$v_*(s_0) - [T_{\pi_k} v_*](s_0) \leq \|v_* - T_{\pi_k} v_*\|_\infty \leq \frac{\gamma^k}{1 - \gamma} \|v_* - T_{\pi_0} v_*\|_\infty = \frac{\gamma^k}{1 - \gamma} (v_*(s_0) - [T_{\pi_0} v_*](s_0)).$$

As a consequence, the action $\pi_k(s_0)$ must be different from $\pi_0(s_0)$ when $\frac{\gamma^k}{1 - \gamma} < 1$, that is for all values of k satisfying $k \geq k^* = \left\lceil \frac{\log \frac{1}{1 - \gamma}}{\log \frac{1}{\gamma}} \right\rceil > \left\lceil \frac{\log \frac{1}{1 - \gamma}}{\log \frac{1}{\gamma}} \right\rceil$. In other words, if some policy π is not optimal, then one of its non-optimal actions will be eliminated for good after at most k^* iterations. By repeating this argument, one can eliminate all non-optimal actions (they are at most $n(m - 1)$), and the result follows.

D A bound for Simplex-PI when $\gamma < 1$ (Proof of Theorem 4)

Using $\{\forall x \geq 0, \|x\|_\infty \leq \mathbf{1}^T x\}$, Lemma 3, $\{\forall x, \mathbf{1}^T x \leq n \|x\|_\infty\}$, Equation (1) and $\{\|(I - \gamma P_{\pi_0})^{-1}\|_\infty = \frac{1}{1 - \gamma}\}$, we have for all k ,

$$\begin{aligned} \|v_{\pi_*} - T_{\pi_k} v_{\pi_*}\|_\infty &\leq \|v_{\pi_*} - v_{\pi_k}\|_\infty \leq \mathbf{1}^T (v_{\pi_*} - v_{\pi_k}) \\ &\leq \left(1 - \frac{1 - \gamma}{n}\right)^k \mathbf{1}^T (v_{\pi_*} - v_{\pi_0}) \leq n \left(1 - \frac{1 - \gamma}{n}\right)^k \|v_{\pi_*} - v_{\pi_0}\|_\infty \\ &= n \left(1 - \frac{1 - \gamma}{n}\right)^k \|(I - \gamma P_{\pi_0})^{-1} (v_* - T_{\pi_0} v_*)\|_\infty \leq \frac{n}{1 - \gamma} \left(1 - \frac{1 - \gamma}{n}\right)^k \|v_{\pi_*} - T_{\pi_0} v_{\pi_*}\|_\infty \end{aligned}$$

Similarly to the proof for Howard's PI, we deduce that a non-optimal action is eliminated after at most $k^* = \left\lceil \frac{n}{1 - \gamma} \log \frac{n}{1 - \gamma} \right\rceil \geq \left\lceil \frac{\log \frac{n}{1 - \gamma}}{\log \left(1 - \frac{1 - \gamma}{n}\right)} \right\rceil$, and the overall number of iterations is obtained by noting that there are at most $n(m - 1)$ non optimal actions to eliminate.

References

- [1] D.P. Bertsekas and J.N. Tsitsiklis. *Neurodynamic Programming*. Athena Scientific, 1996.
- [2] J. Fearnley. Exponential lower bounds for policy iteration. In *Proceedings of the 37th international colloquium conference on Automata, languages and programming: Part II*, ICALP'10, pages 551–562, Berlin, Heidelberg, 2010. Springer-Verlag.
- [3] T.D. Hansen, P.B. Miltersen, and U. Zwick. Strategy iteration is strongly polynomial for 2-player turn-based stochastic games with a constant discount factor. *J. ACM*, 60(1):1:1–1:16, February 2013.
- [4] T.D. Hansen and U. Zwick. Lower bounds for howard’s algorithm for finding minimum mean-cost cycles. In *ISAAC (1)*, pages 415–426, 2010.
- [5] R. Hollanders, J.C. Delvenne, and R. Jungers. The complexity of policy iteration is exponential for discounted markov decision processes. In *51st IEEE conference on Decision and control (CDC'12)*, 2012.
- [6] Y. Mansour and S.P. Singh. On the complexity of policy iteration. In *UAI*, pages 401–408, 1999.
- [7] M. Melekopoglou and A. Condon. On the complexity of the policy improvement algorithm for markov decision processes. *INFORMS Journal on Computing*, 6(2):188–192, 1994.
- [8] I. Post and Y. Ye. The simplex method is strongly polynomial for deterministic markov decision processes. Technical report, arXiv:1208.5083v2, 2012.
- [9] M. Puterman. *Markov Decision Processes*. Wiley, New York, 1994.
- [10] N. Schmitz. How good is howard’s policy improvement algorithm? *Zeitschrift für Operations Research*, 29(7):315–316, 1985.
- [11] Y. Ye. The simplex and policy-iteration methods are strongly polynomial for the markov decision problem with a fixed discount rate. *Math. Oper. Res.*, 36(4):593–603, 2011.

Supplementary Material

E Another bound for Simplex-PI when $\gamma < 1$ (Proof of Theorem 5)

This second bound for Simplex-PI is a factor $O(\log n)$ better, but requires a slightly more careful analysis.

At each iteration k , let s_k be the state in which an action is switched. We have (by definition of the algorithm):

$$T_{\pi_{k+1}} v_{\pi_k}(s_k) - v_{\pi_k}(s_k) = \max_{\pi, s} T_{\pi} v_{\pi_k}(s) - v_{\pi_k}(s).$$

Starting with arguments similar to those for the contraction property of Simplex-PI, we have:

$$\begin{aligned} v_{\pi_*} - v_{\pi_k} &= (I - \gamma P_{\pi_*})^{-1} (T_{\pi_*} v_{\pi_k} - v_{\pi_k}) && \{\text{Lemma 1}\} \\ &\leq \frac{1}{1 - \gamma} \max_s T_{\pi_*} v_{\pi_k}(s) - v_{\pi_k}(s) && \{\|(I - \gamma P_{\pi_*})^{-1}\|_{\infty} = \frac{1}{1 - \gamma} \text{ and } (I - \gamma P_{\pi_*})^{-1} \succ 0\} \\ &\leq \frac{1}{1 - \gamma} (T_{\pi_{k+1}} v_{\pi_k}(s_k) - v_{\pi_k}(s_k)), && \{\text{By definition of } s_k\} \end{aligned}$$

which implies that

$$\|v_{\pi_*} - v_{\pi_k}\|_{\infty} \leq \frac{1}{1 - \gamma} (T_{\pi_{k+1}} v_{\pi_k}(s_k) - v_{\pi_k}(s_k)). \quad (5)$$

On the other hand, we have:

$$\begin{aligned} v_{\pi_{k+1}} - v_{\pi_k} &= (I - \gamma P_{\pi_{k+1}})^{-1} (T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k}) && \{\text{Lemma 1}\} \\ &\geq T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k}, && \{(I - \gamma P_{\pi_{k+1}})^{-1} \succ 0 \text{ and } T_{\pi_{k+1}} v_{\pi_k} - v_{\pi_k} \geq 0\} \end{aligned}$$

which implies that

$$v_{\pi_{k+1}}(s_k) - v_{\pi_k}(s_k) \geq T_{\pi_{k+1}} v_{\pi_k}(s_k) - v_{\pi_k}(s_k). \quad (6)$$

Write $\Delta_k = v_{\pi_*} - v_{\pi_k}$. From Equations (5) and (6), we deduce that:

$$\begin{aligned} \Delta_{k+1}(s_k) &\leq \Delta_k(s_k) - (1 - \gamma) \|\Delta_k\|_{\infty} \\ &= \left(1 - (1 - \gamma) \frac{\|\Delta_k\|_{\infty}}{\Delta_k(s_k)}\right) \Delta_k(s_k). \end{aligned}$$

This implies in particular that

$$\Delta_{k+1}(s_k) \leq \gamma \Delta_k(s_k),$$

but also—since $\Delta_k(s_k)$ and $\Delta_{k+1}(s_k)$ are non-negative—that

$$\|\Delta_k\|_{\infty} \leq \frac{1}{1 - \gamma} \Delta_k(s_k).$$

Now, write n_k the vector on the state space such that $n_k(s)$ is the number of times state s has been switched until iteration k (including k). Since by Lemma 1 the sequence $(\Delta_k)_{k \geq 0}$ is non-increasing, we have

$$\begin{aligned} \|\Delta_k\|_{\infty} &\leq \frac{1}{1 - \gamma} \Delta_k(s_k) \\ &\leq \frac{\gamma^{n_{k-1}(s_k)}}{1 - \gamma} \Delta_0(s_k) \\ &\leq \frac{\gamma^{n_{k-1}(s_k)}}{1 - \gamma} \|\Delta_0\|_{\infty}. \end{aligned}$$

At any iteration k , let $s_k^* = \arg \max_s n_{k-1}(s)$. Since at each iteration k , one of the n component is increased by 1, we necessarily have

$$n_{k-1}(s_k^*) \geq \left\lfloor \frac{k-1}{n} \right\rfloor.$$

Write $k^* \leq k-1$ the last iteration when the state s_k^* was updated, such that we have

$$n_{k-1}(s_k^*) = n_{k^*-1}(s_k^*).$$

Since $(\|\Delta_k\|_\infty)_{k \geq 0}$ is non-increasing (using again Lemma 1), we have

$$\begin{aligned} \|\Delta_k\|_\infty &\leq \|\Delta_{k^*}\|_\infty \\ &\leq \frac{\gamma^{n_{k^*-1}(s_k^*)}}{1-\gamma} \|\Delta_0\|_\infty \\ &= \frac{\gamma^{n_{k-1}(s_k^*)}}{1-\gamma} \|\Delta_0\|_\infty \\ &\leq \frac{\gamma^{\lfloor \frac{k-1}{n} \rfloor}}{1-\gamma} \|\Delta_0\|_\infty. \end{aligned}$$

We are now ready to finish the proof. We have

$$\begin{aligned} \|v_{\pi_*} - T_{\pi_k} v_{\pi_*}\|_\infty &\leq \|\Delta_k\|_\infty \\ &\leq \frac{\gamma^{\lfloor \frac{k-1}{n} \rfloor}}{1-\gamma} \|\Delta_0\|_\infty \\ &\leq \frac{\gamma^{\lfloor \frac{k-1}{n} \rfloor}}{(1-\gamma)^2} \|v_{\pi_*} - T_{\pi_0} v_{\pi_*}\|_\infty. \end{aligned}$$

Using the relation $n \lfloor \frac{k-1}{n} \rfloor \geq k-n$ and arguments similar to the previous proofs, we deduce that a non-optimal action is eliminated after at most $n \left(1 + \frac{2}{1-\gamma} \log \frac{1}{1-\gamma}\right)$ iterations, and the result follows from the fact that there are at most $n(m-1)$ non-optimal actions.

F A general bound for Simplex-PI (Proof of Theorem 7)

The proof we give here is strongly inspired by that for the deterministic case of [8]: the steps (a series of lemmas) are similar. There are essentially two differences. First our arguments are somewhat more direct in that we do not refer to linear programming. Second, it is more general: for any policy π , we need to consider the set of transient states (respectively recurrent classes) instead of the set of path states (respectively cycles).

For any policy π , write $\mathcal{R}(\pi)$ for the set of states that are recurrent for π . Recall that $x_\pi = (I - \gamma P_\pi^T)^{-1} \mathbf{1}$. A useful corollary of Lemma 1 is that for any pair of policies π and π' ,

$$\mathbf{1}^T (v_{\pi'} - v_\pi) = x_{\pi'}^T (T_{\pi'} v_\pi - v_\pi). \quad (7)$$

With some slight abuse of notation, we will write that $s \in \mathcal{R}(\pi)$ if there exists a recurrent class $R \in \mathcal{R}(\pi)$ that contains s . We will repeatedly exploit Assumption 1, that we restate here for clarity:

$$\forall s \in \mathcal{R}(\pi), \frac{n}{(1-\gamma)\tau_r} \leq x_\pi(s) \leq \frac{n}{1-\gamma}, \quad (8)$$

$$\forall s \notin \mathcal{R}(\pi), x_\pi(s) \leq \tau_t. \quad (9)$$

As mentioned, before, the proof is structured in two steps: first, we will show that recurrent classes are created often; then we will show that significant progress is made every time a new recurrent class appears.

F.1 Part 1: Recurrent classes are created often

Lemma 11. *Suppose one moves from policy π to policy π' without creating any recurrent class. Let π_\dagger be the final policy before either a new recurrent class appears or the algorithm terminates. Then*

$$\mathbf{1}^T (v_{\pi_\dagger} - v_{\pi'}) \leq \left(1 - \frac{1}{n\tau_t}\right) \mathbf{1}^T (v_{\pi_\dagger} - v_\pi).$$

Proof. The arguments are similar to those for the proof of Theorem 4. On the one hand, we have:

$$\mathbf{1}^T(v_{\pi'} - v_{\pi}) \geq \mathbf{1}^T(T_{\pi'}v_{\pi} - v_{\pi}). \quad (10)$$

On the other hand, we have

$$\begin{aligned} \mathbf{1}^T(v_{\pi_{\dagger}} - v_{\pi}) &= \mathbf{1}^T(I - \gamma P_{\pi_{\dagger}})^{-1}(T_{\pi_{\dagger}}v_{\pi} - v_{\pi}) \\ &= x_{\pi_{\dagger}}^T(T_{\pi_{\dagger}}v_{\pi} - v_{\pi}) \\ &= \sum_{s \notin \mathcal{R}(\pi_{\dagger})} x_{\pi_{\dagger}}(s)(T_{\pi_{\dagger}}v_{\pi}(s) - v_{\pi}(s)) + \sum_{s \in \mathcal{R}(\pi_{\dagger})} x_{\pi_{\dagger}}(s)(T_{\pi_{\dagger}}v_{\pi}(s) - v_{\pi}(s)) \\ &\leq n\tau_t \max_{s \notin \mathcal{R}(\pi_{\dagger})} T_{\pi_{\dagger}}v_{\pi}(s) - v_{\pi}(s) + \frac{n^2}{1-\gamma} \max_{s \in \mathcal{R}(\pi_{\dagger})} T_{\pi_{\dagger}}v_{\pi}(s) - v_{\pi}(s). \end{aligned} \quad \{\text{Equations (8)-(9)}\}$$

Since by assumption cycles of π_{\dagger} are also cycles of π , we deduce that for all $s \in \mathcal{R}(\pi_{\dagger})$, $\pi_{\dagger}(s) = \pi(s)$, so that $\max_{s \in \mathcal{R}(\pi_{\dagger})} T_{\pi_{\dagger}}v_{\pi}(s) - v_{\pi}(s) = 0$. Thus, the second term of the above r.h.s. is null and

$$\begin{aligned} \mathbf{1}^T(v_{\pi_{\dagger}} - v_{\pi}) &\leq n\tau_t \max_s T_{\pi_{\dagger}}v_{\pi}(s) - v_{\pi}(s) \\ &\leq n\tau_t \max_s T_{\pi'}v_{\pi}(s) - v_{\pi}(s) \quad \{\max_s T_{\pi'}v_{\pi}(s) = \max_{s, \bar{\pi}} T_{\bar{\pi}}v_{\pi}(s)\} \\ &= n\tau_t \mathbf{1}^T(T_{\pi'}v_{\pi} - v_{\pi}). \end{aligned} \quad (11)$$

Combining Equations (10) and (11), we get:

$$\begin{aligned} \mathbf{1}^T(v_{\pi_{\dagger}} - v_{\pi'}) &= \mathbf{1}^T(v_{\pi_{\dagger}} - v_{\pi}) - \mathbf{1}^T(v_{\pi'} - v_{\pi}) \\ &\leq \left(1 - \frac{1}{n\tau_t}\right) \mathbf{1}^T(v_{\pi_{\dagger}} - v_{\pi}). \end{aligned}$$

□

Lemma 12. *While Simplex-PI does not create any recurrent class nor finishes:*

- *an action is eliminated from policies after at most $\lceil n\tau_t \log(n\tau_t) \rceil$ iterations;*
- *a recurrent class is broken after at most $\lceil n\tau_t \log(n^2\tau_t) \rceil$ iterations.*

Proof. Let π be the policy in some iteration, π_{\dagger} be the last policy before a new recurrent class appears, and π' any policy between π and π_{\dagger} . Since

$$\begin{aligned} 0 &\leq \mathbf{1}^T(v_{\pi_{\dagger}} - v_{\pi}) && \{v_{\pi_{\dagger}} \geq v_{\pi}\} \\ &= x_{\pi}^T(v_{\pi_{\dagger}} - T_{\pi}v_{\pi_{\dagger}}) && \{\text{Equation (7)}\} \\ &= \sum_{s \notin \mathcal{R}(\pi)} x_{\pi}(s)(v_{\pi_{\dagger}}(s) - T_{\pi}v_{\pi_{\dagger}}(s)) + \sum_{C \in \mathcal{R}(\pi)} \sum_{s \in C} x_{\pi}(s)(v_{\pi_{\dagger}}(s) - T_{\pi}v_{\pi_{\dagger}}(s)), \end{aligned}$$

there must exist either a state $s_0 \notin \mathcal{R}(\pi)$ such that

$$x_{\pi}(s_0)(v_{\pi_{\dagger}}(s_0) - T_{\pi}v_{\pi_{\dagger}}(s_0)) \geq \frac{1}{n} x_{\pi}^T(v_{\pi_{\dagger}} - T_{\pi}v_{\pi_{\dagger}}) \geq 0. \quad (12)$$

or a recurrent class R_0 such that

$$\sum_{s \in R_0} x_{\pi}(s)(v_{\pi_{\dagger}}(s) - T_{\pi}v_{\pi_{\dagger}}(s)) \geq \frac{1}{n} x_{\pi}^T(v_{\pi_{\dagger}} - T_{\pi}v_{\pi_{\dagger}}) \geq 0. \quad (13)$$

We consider these two cases separately below.

- **case 1:** Equation (12) holds for some $s_0 \notin \mathcal{R}(\pi)$. If $\pi'(s_0) = \pi(s_0)$, then

$$\begin{aligned} \mathbf{1}^T(v_{\pi_{\dagger}} - v_{\pi'}) &\geq v_{\pi_{\dagger}}(s_0) - v_{\pi'}(s_0) && \{v_{\pi_{\dagger}} \geq v_{\pi'}\} \\ &= v_{\pi_{\dagger}}(s_0) - T_{\pi'}v_{\pi'}(s_0) && \{v_{\pi'} = T_{\pi'}v_{\pi'}\} \\ &\geq v_{\pi_{\dagger}}(s_0) - T_{\pi'}v_{\pi_{\dagger}}(s_0) && \{v_{\pi_{\dagger}} \geq v_{\pi'}\} \\ &= v_{\pi_{\dagger}}(s_0) - T_{\pi}v_{\pi_{\dagger}}(s_0) && \{\pi(s_0) = \pi'(s_0)\} \\ &\geq \frac{1}{\tau_t} x_{\pi}(s_0)(v_{\pi_{\dagger}}(s_0) - T_{\pi}v_{\pi_{\dagger}}(s_0)) && \{\text{Equation (9)}\} \\ &\geq \frac{1}{n\tau_t} x_{\pi}^T(v_{\pi_{\dagger}} - T_{\pi}v_{\pi_{\dagger}}) && \{\text{Equation (12)}\} \\ &= \frac{1}{n\tau_t} \mathbf{1}^T(v_{\pi_{\dagger}} - v_{\pi}). && \{\text{Equation (7)}\} \end{aligned}$$

If there is no recurrent class creation, the contraction property given in Lemma 11 implies that after $k = \lceil n\tau_t \log(n\tau_t) \rceil > \frac{\log(n\tau_t)}{\log \frac{1}{1-\frac{1}{n\tau_t}}}$ iterations we have

$$\mathbb{1}^T(v_{\pi_{\dagger}} - v_{\pi'}) < \frac{1}{n\tau_t} \mathbb{1}^T(v_{\pi_{\dagger}} - v_{\pi}),$$

and thus $\pi'(s_0) \neq \pi(s_0)$.

- **case 2:** Equation (13) holds for some $R_0 \in \mathcal{R}(\pi)$. Write \mathcal{T} be the set of states that are transient for π (formally, $\mathcal{T} = X \setminus \mathcal{R}(\pi)$). For any subset Y of the state space X , write P_{π}^Y for the stochastic matrix of which the i^{th} row is equal to that of P_{π} if $i \in Y$, and is 0 otherwise, and write $\mathbb{1}_Y$ the vectors of which the i^{th} component is equal to 1 if $i \in Y$ and 0 otherwise. Using the fact that $P_{\pi}^{R_0} P_{\pi}^T = 0$, one can first observe that

$$(I - \gamma P_{\pi}^{R_0})(I - \gamma P_{\pi}^T) = I - \gamma(P_{\pi}^{R_0} + P_{\pi}^T),$$

from which we can deduce that

$$\begin{aligned} \mathbb{1}_{\mathcal{T} \cup R_0}^T (I - \gamma P)^{-1} &= \mathbb{1}_{\mathcal{T} \cup R_0}^T (I - \gamma(P_{\pi}^{R_0} + P_{\pi}^T))^{-1} \\ &= \mathbb{1}_{\mathcal{T} \cup R_0}^T (I - \gamma P_{\pi}^T)^{-1} (I - \gamma P_{\pi}^{R_0})^{-1}. \end{aligned} \quad (14)$$

Also, writing $h_{\mathcal{T}} = (I - \gamma P_{\pi}^{T^T})^{-1} \mathbb{1}_{\mathcal{T}}$, that satisfies

$$h_{\mathcal{T}} = \mathbb{1}_{\mathcal{T}} + \gamma P_{\pi}^{T^T} h_{\mathcal{T}},$$

we can see that:

$$\forall s \in R_0, h_{\mathcal{T}}(s) = \gamma \sum_{s' \in \mathcal{T}} p_{s's}(\pi(s')) h_{\mathcal{T}}(s'), \quad \{s \in R_0 \Rightarrow \mathbb{1}_{\mathcal{T}}(s) = 0\} \quad (15)$$

and thus:

$$\begin{aligned} (I - \gamma P_{\pi}^{T^T})^{-1} \mathbb{1}_{\mathcal{T} \cup R_0}(s) &= (I - \gamma P_{\pi}^{T^T})^{-1} \mathbb{1}_{\mathcal{T}}(s) + 1 && \{P_{\pi}^{T^T} \mathbb{1}_{R_0} = 0\} \\ &= h_{\mathcal{T}}(s) + 1 \\ &\leq \gamma \sum_{s' \in \mathcal{T}} p_{s's}(\pi(s')) h_{\mathcal{T}}(s') + 1 && \{\text{Equation (15)}\} \\ &\leq \sum_{s' \in \mathcal{T}} h_{\mathcal{T}}(s') + 1 \\ &= \sum_{s' \in \mathcal{T}} x_{\pi}(s') + 1 && \{\forall s' \in \mathcal{T}, h_{\mathcal{T}}(s') = x_{\pi}(s')\} \\ &\leq (n-1)\tau_t + 1 && \{|\mathcal{T}| \leq (n-1) \text{ and Equation (9)}\} \\ &\leq n\tau_t. && \{\tau_t \geq 1\} \end{aligned} \quad (16)$$

Writing δ the vector that equals $v_{\pi_{\dagger}} - T_{\pi}v_{\pi_{\dagger}}$ on R_0 and that is null everywhere else, we have

$$\begin{aligned}
& \sum_{s \in R_0} x_{\pi}(s)(v_{\pi_{\dagger}}(s) - T_{\pi}v_{\pi_{\dagger}}(s)) \\
&= \sum_{s \in R_0} [(I - \gamma P_{\pi}^T)^{-1} \mathbf{1}](s) \delta(s) \\
&= \sum_{s \in R_0} [(I - \gamma P_{\pi}^T)^{-1} \mathbf{1}_{\mathcal{T} \cup R_0}](s) \delta(s) && \{\forall s \in R_0, [(I - \gamma P_{\pi}^T)^{-1} \mathbf{1}_{X \setminus (\mathcal{T} \cup R_0)}](s) = 0\} \\
&= \sum_s [(I - \gamma P_{\pi}^T)^{-1} \mathbf{1}_{\mathcal{T} \cup R_0}](s) \delta(s) && \{\forall s \notin R_0, \delta(s) = 0\} \\
&= \mathbf{1}_{\mathcal{T} \cup R_0}^T (I - \gamma P_{\pi})^{-1} \delta \\
&= \mathbf{1}_{\mathcal{T} \cup R_0}^T (I - \gamma P_{\pi}^T)^{-1} (I - \gamma P_{\pi}^{R_0})^{-1} \delta && \{\text{Equation (14)}\} \\
&= \sum_s [(I - \gamma P_{\pi}^{T^T})^{-1} \mathbf{1}_{\mathcal{T} \cup R_0}](s) [(I - \gamma P_{\pi}^{R_0})^{-1} \delta](s) \\
&= \sum_{s \in R_0} [(I - \gamma P_{\pi}^{T^T})^{-1} \mathbf{1}_{\mathcal{T} \cup R_0}](s) [(I - \gamma P_{\pi}^{R_0})^{-1} \delta](s) && \{\forall s \notin R_0, \delta(s) = 0\} \\
&= \sum_{s \in R_0} [(I - \gamma P_{\pi}^{T^T})^{-1} \mathbf{1}_{\mathcal{T} \cup R_0}](s) (v_{\pi_{\dagger}}(s) - v_{\pi}(s)) && \{\text{Lemma 1}\} \\
&\leq n\tau_t \mathbf{1}_{R_0} (v_{\pi_{\dagger}} - v_{\pi}). && \{\text{Equation (16)}\} \\
& && (17)
\end{aligned}$$

Now, one can deduce from this that if R_0 is also a recurrent class of π' , which implies $\mathbf{1}_{R_0}^T v_{\pi} = \mathbf{1}_{R_0}^T v_{\pi'}$, then

$$\begin{aligned}
\mathbf{1}^T (v_{\pi_{\dagger}} - v_{\pi'}) &\geq \mathbf{1}_{R_0}^T (v_{\pi_{\dagger}} - v_{\pi'}) && \{v_{\pi_{\dagger}} \geq v_{\pi'}\} \\
&= \mathbf{1}_{R_0}^T (v_{\pi_{\dagger}} - v_{\pi}) && \{\mathbf{1}_{R_0}^T v_{\pi} = \mathbf{1}_{R_0}^T v_{\pi'}\} \\
&\geq \frac{1}{n\tau_t} \sum_{s \in R_0} x_{\pi}(s) (v_{\pi_{\dagger}}(s) - T_{\pi}v_{\pi_{\dagger}}(s)) && \{\text{Equation (17)}\} \\
&\geq \frac{1}{n^2\tau_t} x_{\pi}^T (v_{\pi_{\dagger}} - T_{\pi}v_{\pi_{\dagger}}) && \{\text{Equation (13)}\} \\
&= \frac{1}{n^2\tau_t} \mathbf{1}^T (v_{\pi_{\dagger}} - v_{\pi}). && \{\text{Equation (7)}\}
\end{aligned}$$

If there is no recurrent class creation, the contraction property given in Lemma 11 implies that after $k = \lceil n\tau_t \log(n^2\tau_t) \rceil > \frac{\log(n^2\tau_t)}{\log \frac{1}{1 - \frac{1}{n\tau_t}}}$ iterations we have

$$\mathbf{1}^T (v_{\pi_{\dagger}} - v_{\pi'}) < \frac{1}{n^2\tau_t} \mathbf{1}^T (v_{\pi_{\dagger}} - v_{\pi}),$$

and thus R_0 cannot be a recurrent class of π' . □

A direct consequence of the above result is Lemma 6 that we originally stated page 5, and that we restate for clarity.

Lemma 6. *After at most $n \lceil (m-1) \lceil n\tau_t \log(n\tau_t) \rceil + \lceil n\tau_t \log(n^2\tau_t) \rceil$ iterations, either Simplex-PI finishes or a new recurrent class appears.*

Proof. Before a recurrent class is created, at most n recurrent classes need to be broken and $n(m-1)$ actions to be eliminated, and the time required by these events is bounded thanks to the previous lemma. □

F.2 Part 2: A new recurrent class implies a significant step towards the optimal value

We now proceed to the second part of the proof, and begin by proving Lemma 7 (originally stated page 6).

Lemma 7. *Suppose Simplex-PI moves from π to π' and that π' involves a new recurrent class. Then*

$$\mathbf{1}^T(v_{\pi_*} - v_{\pi'}) \leq \left(1 - \frac{1}{\tau_r}\right) \mathbf{1}^T(v_{\pi_*} - v_{\pi}).$$

Proof. Let s_0 be the state such that $\pi'(s) \neq \pi(s)$. On the one hand, since π' contains a new recurrent class R (necessarily containing s_0), we have

$$\begin{aligned} \mathbf{1}^T(v_{\pi'} - v_{\pi}) &= x_{\pi'}^T(T_{\pi'}v_{\pi} - v_{\pi}) && \{\text{Equation (7)}\} \\ &= x_{\pi'}(s_0)(T_{\pi'}v_{\pi}(s_0) - v_{\pi}(s_0)) && \{\text{Simplex-PI switches 1 action}\} \\ &\geq \frac{n}{(1-\gamma)\tau_r}(T_{\pi'}v_{\pi}(s_0) - v_{\pi}(s_0)). && \{\text{Equation 8 with } s_0 \in R \subset \mathcal{R}(\pi')\} \end{aligned} \quad (18)$$

On the other hand,

$$\begin{aligned} v_{\pi_*} - v_{\pi} &= (I - \gamma P_{\pi_*})^{-1}(T_{\pi_*}v_{\pi} - v_{\pi}) && \{\text{Lemma 1}\} \\ &\leq \frac{1}{1-\gamma} \max_s T_{\pi_*}v_{\pi}(s) - v_{\pi}(s) \quad \{\|(I - \gamma P_{\pi_*})^{-1}\|_{\infty} \leq \frac{1}{1-\gamma} \text{ and } (I - \gamma P_{\pi_*})^{-1} \succ 0\} \\ &\leq \frac{1}{1-\gamma} \max_s T_{\pi'}v_{\pi}(s) - v_{\pi}(s) && \{\max_s T_{\pi'}v_{\pi}(s) = \max_{s, \bar{\pi}} T_{\bar{\pi}}v_{\pi}(s)\} \\ &= \frac{1}{1-\gamma}(T_{\pi'}v_{\pi}(s_0) - v_{\pi}(s_0)). && \{\text{Simplex-PI switches 1 action}\} \end{aligned} \quad (19)$$

Combining these two observations, we obtain:

$$\begin{aligned} \mathbf{1}^T(v_{\pi_*} - v_{\pi'}) &= \mathbf{1}^T(v_{\pi_*} - v_{\pi'}) - \mathbf{1}^T(v_{\pi'} - v_{\pi}) \\ &\leq \mathbf{1}^T(v_{\pi_*} - v_{\pi'}) - \frac{n}{(1-\gamma)\tau_r}(T_{\pi'}v_{\pi}(s_0) - v_{\pi}(s_0)) && \{\text{Equation (18)}\} \\ &\leq \mathbf{1}^T(v_{\pi_*} - v_{\pi'}) - \frac{n}{\tau_r} \max_s v_{\pi_*}(s) - v_{\pi'}(s) && \{\text{Equation (19)}\} \\ &\leq \left(1 - \frac{1}{\tau_r}\right) \mathbf{1}^T(v_{\pi_*} - v_{\pi'}). && \{\forall x, \mathbf{1}^T x \leq n \max_s x(s)\} \end{aligned}$$

□

Lemma 14. *While the algorithm does not terminate,*

- *some non-optimal action is eliminated from recurrent states after at most $\lceil \tau_r \log(n\tau_r) \rceil$ recurrent class creations;*
- *some non-optimal action is eliminated from policies after at most $\lceil \tau_r \log(n\tau_t) \rceil$ recurrent class creations.*

Proof. Let π be the policy in some iteration and π' be any policy between π and π_* . Let $s_0 = \arg \max_s x_{\pi}(s)(v_{\pi_*}(s) - T_{\pi}v_{\pi_*}(s))$. We have

$$\begin{aligned} x_{\pi}(s_0)(v_{\pi_*}(s_0) - T_{\pi}v_{\pi_*}(s_0)) &\geq \frac{1}{n} x_{\pi}^T(v_{\pi_*} - T_{\pi}v_{\pi_*}) && \{\forall x, \mathbf{1}^T x \leq n \max_s x(s)\} \\ &= \mathbf{1}^T(v_{\pi_*} - v_{\pi}). && \{\text{Equation (7)}\} \end{aligned} \quad (20)$$

We now consider two cases.

- **case 1:** $s_0 \notin \mathcal{R}(\pi)$. If $\pi'(s_0) = \pi(s_0)$, then

$$\begin{aligned} \mathbf{1}^T(v_{\pi_*} - v_{\pi'}) &= x_{\pi'}^T(v_{\pi_*} - T_{\pi'}v_{\pi_*}) && \{\text{Equation (7)}\} \\ &\geq x_{\pi'}(s_0)(v_{\pi_*}(s_0) - T_{\pi'}v_{\pi_*}(s_0)) && \{v_{\pi_*} \geq T_{\pi'}v_{\pi_*}\} \\ &\geq v_{\pi_*}(s_0) - T_{\pi'}v_{\pi_*}(s_0) && \{x_{\pi'}(s_0) \geq 1\} \\ &= v_{\pi_*}(s_0) - T_{\pi}v_{\pi_*}(s_0) && \{\pi(s_0) = \pi'(s_0)\} \\ &\geq \frac{1}{\tau_t} x_{\pi}(s_0)(v_{\pi_*}(s_0) - T_{\pi}v_{\pi_*}(s_0)) && \{\text{Equation (9)}\} \\ &\geq \frac{1}{n\tau_t} \mathbf{1}^T(v_{\pi_*} - v_{\pi}). && \{\text{Equation (20)}\} \end{aligned}$$

After $k = \lceil \tau_r \log n \tau_t \rceil > \frac{\log n \tau_t}{\log \frac{1}{1-\frac{1}{\tau_r}}}$ new recurrent classes are created, we have by the contraction property of Lemma 7 that

$$\mathbf{1}^T(v_{\pi_*} - v_{\pi'}) < \frac{1}{n\tau_t} \mathbf{1}^T(v_{\pi_*} - v_{\pi}).$$

This implies that $\pi'(s_0) \neq \pi(s_0)$.

- **case 2:** $s_0 \in \mathcal{R}(\pi)$. If $\pi'(s_0) = \pi(s_0)$ and $s_0 \in \mathcal{R}(\pi')$, then

$$\begin{aligned} \mathbf{1}^T(v_{\pi_*} - v_{\pi'}) &= x_{\pi'}^T(v_{\pi_*} - T_{\pi'}v_{\pi_*}) && \{\text{Equation (7)}\} \\ &= \sum_s x_{\pi'}(s)(v_{\pi_*}(s) - T_{\pi'}v_{\pi_*}(s)) \\ &\geq \sum_{s \in R_0} x_{\pi'}(s)(v_{\pi_*}(s) - T_{\pi'}v_{\pi_*}(s)) && \{v_{\pi_*} \geq T_{\pi'}v_{\pi_*}\} \\ &\geq \frac{n}{(1-\gamma)\tau_r} \sum_{s \in R_0} v_{\pi_*}(s) - T_{\pi'}v_{\pi_*}(s) && \{\text{Equation 8}\} \\ &\geq \frac{n}{(1-\gamma)\tau_r} v_{\pi_*}(s_0) - T_{\pi'}v_{\pi_*}(s_0) && \{v_{\pi_*} \geq T_{\pi'}v_{\pi_*}\} \\ &= \frac{n}{(1-\gamma)\tau_r} v_{\pi_*}(s_0) - T_{\pi}v_{\pi_*}(s_0) && \{\pi(s_0) = \pi'(s_0)\} \\ &= \frac{1}{\tau_r} x_{\pi}(s_0)(v_{\pi_*}(s_0) - T_{\pi}v_{\pi_*}(s_0)) && \{x_{\pi}(s_0) \leq \frac{n}{1-\gamma}\} \\ &\geq \frac{1}{n\tau_r} \mathbf{1}^T(v_{\pi_*} - v_{\pi}). && \{\text{Equation (20)}\} \end{aligned}$$

After $k = \lceil \tau_r \log n \tau_r \rceil > \frac{\log n \tau_r}{\log \frac{1}{1-\frac{1}{\tau_r}}}$ new recurrent classes are created, we have by the contraction property of Lemma 7 that

$$\mathbf{1}^T(v_{\pi_*} - v_{\pi'}) < \frac{1}{n\tau_r} \mathbf{1}^T(v_{\pi_*} - v_{\pi}).$$

This implies that $\pi'(s_0) \neq \pi(s_0)$ if s_0 is recurrent for π' .

□

We are ready to conclude: At most, the $n(m-1)$ non-optimal actions may need to be eliminated from recurrent and transient states, requiring at most a total of $n(m-1)(\lceil \tau_r \log(n\tau_r) \rceil + \lceil \tau_r \log(n\tau_t) \rceil)$ recurrent classes creations. The result follows from the fact that each class creation requires at most $n \lceil (m-1) \lceil n\tau_t \log(n\tau_t) \rceil + \lceil n\tau_t \log(n^2\tau_t) \rceil \rceil$ iterations.

G Cycle and recurrent classes creations for Howard's PI (Proofs of Lemmas 8 and 9)

Lemma 8. *If the MDP is deterministic, after at most n iterations, either Howard's PI finishes or a new cycle appears.*

Proof. Consider a sequence of l generated policies π_1, \dots, π_l from an initial policy π_0 such that no new cycle appears. By induction, we have

$$\begin{aligned} v_{\pi_l} - v_{\pi_k} &= T_{\pi_l}v_{\pi_l} - T_{\pi_l}v_{\pi_{k-1}} + T_{\pi_l}v_{\pi_{k-1}} - T_{\pi_k}v_{\pi_{k-1}} + T_{\pi_k}v_{\pi_{k-1}} - T_{\pi_k}v_{\pi_k} && \{\forall \pi, T_{\pi}v_{\pi} = v_{\pi}\} \\ &\leq \gamma P_{\pi_l}(v_{\pi_l} - v_{\pi_{k-1}}) + \gamma P_{\pi_k}(v_{\pi_{k-1}} - v_{\pi_k}) && \{T_{\pi_l}v_{\pi_{k-1}} \leq T_{\pi_k}v_{\pi_{k-1}}\} \\ &\leq \gamma P_{\pi_l}(v_{\pi_l} - v_{\pi_{k-1}}). && \{\text{Lemma 1 and } P_{\pi_k} > 0\} \\ &\leq (\gamma P_{\pi_l})^k (v_{\pi_l} - v_{\pi_0}). && \{\text{by induction on } k\} \end{aligned} \tag{21}$$

Since the MDP is deterministic and has n states, $(P_{\pi_l})^n$ will only have non-zero values on columns that correspond to $\mathcal{R}(\pi_l)$. Furthermore, since no cycle is created, $\mathcal{R}(\pi_l) \subset \mathcal{R}(\pi_0)$, which implies that $v_{\pi_l}(s) - v_{\pi_0}(s) = 0$ for all $s \in \mathcal{R}(\pi_l)$. As a consequence, we have $(P_{\pi_l})^n(v_{\pi_l} - v_{\pi_0}) = 0$. By Equation (21), this implies that $v_{\pi_l} = v_{\pi_n}$. If $l > n$, then the algorithm must have terminated. \square

Lemma 9. *If the MDP satisfies Assumption 1, after at most $nm\lceil\tau_l \log n\tau_l\rceil$ iterations, either Howard's PI finishes or a new recurrent class appears.*

Proof. It can be seen that the proof of Lemma 6 also applies to Howard's PI. \square

H A bound for Howard's PI and Simplex-PI under Assumption 2 (proof of Theorem 8)

We here consider that the state space is decomposed into 2 sets: \mathcal{T} is the set of states that are transient under all policies, and \mathcal{R} is the set of states that are recurrent under all policies. From this assumption, it can be seen that when running Howard's PI or Simplex-PI, the values and actions chosen on \mathcal{T} have no influence on the evolution of the values and policies on \mathcal{R} . So we will study the convergence of both algorithms in two steps: We will first bound the number of iterations to converge on \mathcal{R} . We will then add the number of iterations for converging on \mathcal{T} given that convergence has occurred on \mathcal{R} .

Convergence on the set \mathcal{R} of recurrent states Without loss of generality, we consider that the state space is only made of the set of recurrent states.

First consider Simplex-PI. If all states are recurrent, new recurrent classes are created at every iteration, and Lemma 6 holds. Then, in a way similar to the proof of Lemma 14, it can be shown that every $\lceil\tau_r \log n\tau_r\rceil$ iterations, a non-optimal action can be eliminated. As there are at most $n(m-1)$ non-optimal actions, we deduce that Simplex-PI converges in at most $n(m-1)\lceil\tau_r \log n\tau_r\rceil$ iterations on \mathcal{R} .

Now consider Howard's PI. We can prove Lemma 10, that we restate for clarity.

Lemma 10. *For an MDP satisfying Assumptions 1-2, suppose Howard's PI moves from π to π' and that π' involves a new recurrent class. Then*

$$\mathbf{1}^T(v_{\pi_*} - v_{\pi'}) \leq \left(1 - \frac{1}{\tau_r}\right) \mathbf{1}^T(v_{\pi_*} - v_{\pi}).$$

Proof. In the case we focus on the convergence on the set \mathcal{R} of recurrent states, new recurrent classes are created at every iteration. So we will prove that the inequality holds for every k . On the one hand, we have for all iterations k ,

$$\begin{aligned} \mathbf{1}^T(v_{\pi_{k+1}} - v_{\pi_k}) &= x_{\pi_{k+1}}^T(T_{\pi_{k+1}}v_{\pi_k} - v_{\pi_k}) && \{\text{Equation (7)}\} \\ &\geq \frac{n}{(1-\gamma)\tau_r} \mathbf{1}^T(T_{\pi_{k+1}}v_{\pi_k} - v_{\pi_k}) && \{\text{Equation (8)}\} \\ &\geq \frac{n}{(1-\gamma)\tau_r} \|T_{\pi_{k+1}}v_{\pi_k} - v_{\pi_k}\|_{\infty}. && \{\forall x \geq 0, \mathbf{1}^T x \geq \|x\|_{\infty}\} \end{aligned} \quad (22)$$

On the other hand,

$$\begin{aligned} \mathbf{1}^T(v_{\pi_*} - v_{\pi_k}) &= x_{\pi_*}^T(T_{\pi_*}v_{\pi_k} - v_{\pi_k}) && \{\text{Equation (7)}\} \\ &\leq \frac{n}{1-\gamma} \|T_{\pi_*}v_{\pi_k} - v_{\pi_k}\|_{\infty} && \left\{\sum_i x_{\pi_*}(i) \leq \frac{n}{1-\gamma}\right\} \\ &\leq \frac{n}{1-\gamma} \|T_{\pi_{k+1}}v_{\pi_k} - v_{\pi_k}\|_{\infty}. \end{aligned} \quad (23)$$

By combining Equations (22) and (23), we obtain:

$$\begin{aligned} \mathbf{1}^T(v_{\pi_*} - v_{\pi_{k+1}}) &= \mathbf{1}^T(v_{\pi_*} - v_{\pi_k}) - \mathbf{1}^T(v_{\pi_{k+1}} - v_{\pi_k}) \\ &\leq \left(1 - \frac{1}{\tau_r}\right) \mathbf{1}^T(v_{\pi_*} - v_{\pi_k}). \end{aligned}$$

\square

Then, similarly to Simplex-PI, we can prove that after every $\lceil\tau_r \log n\tau_r\rceil$ iterations a non-optimal action must be eliminated. And as there are at most $n(m-1)$ non-optimal actions, we deduce that Howard's PI converges in at most $n(m-1)\lceil\tau_r \log n\tau_r\rceil$ iterations on \mathcal{R} .

Convergence on the set \mathcal{T} of transient states Consider now that convergence has occurred on the recurrent states \mathcal{R} . A simple variation of the proof of Lemma 6/Lemma 9 (where we use the fact that we don't need to consider the events where recurrent classes are broken since recurrent classes do not evolve anymore) allows to show that the extra number of iterations to converge on the transient states is $n(m-1)\lceil\tau_t \log n\tau_t\rceil$ for Howard's PI and $n^2(m-1)\lceil\tau_t \log n\tau_t\rceil$ for Simplex-PI, and the result follows.