



HAL
open science

Computing Whittle (and Gittins) Index in Subcubic Time

Nicolas Gast, Bruno Gaujal, Kimang Khun

► **To cite this version:**

Nicolas Gast, Bruno Gaujal, Kimang Khun. Computing Whittle (and Gittins) Index in Subcubic Time. 2022. hal-03602458v2

HAL Id: hal-03602458

<https://inria.hal.science/hal-03602458v2>

Preprint submitted on 29 Apr 2022 (v2), last revised 20 Jun 2023 (v5)

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.

Computing Whittle (and Gittins) Index in Subcubic Time

Nicolas Gast, Bruno Gaujal and Kimang Khun

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP*, LIG, 38000, Grenoble, France.

*Institute of Engineering Univ. Grenoble Alpes.

Contributing authors: nicolas.gast@inria.fr; bruno.gaujal@inria.fr; kimang.khun@inria.fr;

Abstract

Whittle index is a generalization of Gittins index that provides very efficient allocation rules for restless multi-armed bandits. In this work, we develop an algorithm to test the indexability and compute the Whittle indices of any finite-state Markovian bandit problem. This algorithm works in the discounted and non-discounted cases, and can compute Gittins index. Our algorithm builds on three tools: (1) a careful characterization of Whittle index that allows one to compute recursively the k th smallest index from the $(k - 1)$ th smallest, and to test indexability, (2) the use of Sherman-Morrison formula to make this recursive computation efficient, and (3) a sporadic use of the fastest matrix inversion and multiplication methods to obtain a subcubic complexity. We show that an efficient use of the Sherman-Morrison formula leads to an algorithm that computes Whittle index in $(2/3)n^3 + o(n^3)$ arithmetic operations, where n is the number of states of the arm. The careful use of fast matrix multiplication leads to the first subcubic algorithm to compute Whittle (or Gittins) index. By using the current fastest matrix multiplications, our algorithm runs in $\mathcal{O}(n^{2.5286})$. We also conduct a series of numerical experiments that demonstrate that our algorithm is very efficient and can compute indices of Markov chains with several thousands of states in at most a few seconds.

Keywords: Whittle Index, Gittins Index, Restless Bandit, Multi-armed Bandit, Sherman-Morrison, Fast Matrix Multiplication

1 Introduction

Markovian bandits form a subclass of multi-armed bandit problems in which each arm has an internal state that evolves over time in a Markovian manner, as a function of the decision maker's actions. In such a problem, at each time step, the decision maker observes the state of all arms and chooses which one to activate. When the state of an arm evolves only when this arm is chosen, one falls into the category of *rested* Markovian bandits for which an optimal policy (in the discounted case) was found by Gittins [1]. When the state of an arm can also evolve when the arm is not chosen, the problem is called a *restless bandit* problem, and computing an optimal policy is a difficult problem [2].

In his seminal paper [3], Whittle proposed a very efficient heuristic: For each arm, an index function maps each state of the arm to a real number, the Whittle index policy then consists in activating the arms having the highest index first. This heuristic generalizes Gittins index to restless bandits. Contrary to the rested case, the Whittle index policy is in general not optimal. Yet, this policy has been proven to be very efficient over the years: up to a condition called *indexability*, Whittle index has been shown to be (in the undiscounted case) asymptotically optimal as the number of arms grows to infinity [4, 5]. Moreover, the heuristic performs extremely well in practice [6–8]. Restless bandits and Whittle index have been applied to many scheduling and resource allocation problems such as wireless communication [9, 10], web crawling [11, 12], congestion control [13, 14], queueing systems [15–21], and clinical trials [22].

The above examples show that, when a problem is indexable, computing Whittle index is a very efficient way to construct an nearly-optimal heuristic. This raises important questions, that we study in this paper:

- Is testing indexability computationally hard?
- Is there an efficient algorithm to compute Whittle index?
- Is Whittle index harder to compute than Gittins index?

The computation of Gittins index has received a lot of attention in the past, see for instance [23–26] and the recent survey [27]. For a n -state Markovian bandit, the algorithm having the smallest complexity is the *fast-pivoting* implementation of [25] that performs $(2/3)n^3 + O(n^2)$ arithmetic operations. Note that the same author claims in page 4 of [28] that it is unlikely that this complexity can be improved. As we see later, we do improve upon this complexity.

The situation of Whittle index is more contrasted. To the best of our knowledge, there is no general purpose algorithm to test indexability, and most papers studying Whittle index either assume that the studied model is indexable or focus on specific classes of restless bandits for which the structure can be used to show indexability, see *e.g.* [16, 18, 29, 30]. Assuming indexability, the computation of Whittle index has been considered by a few papers.

For restless bandits with discount, $\beta \in (0, 1)$, the most efficient numerical algorithm to compute Whittle index was recently presented in [28]. This

algorithm, called *fast-pivoting* algorithm, performs $(2/3)n^3 + O(n^2)$ arithmetic operations¹ if the initialization phase is excluded from the count. This is done by using the parametric simplex method and exploiting the special structure of this linear system to reduce the complexity of simplex pivoting steps. This fast-pivoting algorithm is an efficient implementation of adaptive-greedy algorithm [31] which outputs Whittle index if and only if the restless bandit satisfies a technical condition called partial conservation law (PCL), which is more restrictive than just being indexable. So, it is not applicable for all indexable restless bandits. Based on a geometric interpretation of Whittle index, the authors in [32] propose a refinement of the adaptive-greedy algorithm of [31] to compute Whittle indices of all indexable restless bandits. The refined algorithm achieves a $O(n^3)$ complexity by using the Sherman-Morrison formula. The authors also propose a few checkable conditions to test indexability. However, those conditions are not necessary for indexability, which means that if an arm does not verify the conditions, we cannot conclude that the arm is non-indexable and an algorithm to check indexability is still needed. Also, no detailed description is given for adapting those conditions and their algorithm to restless bandit without discount. A thorough comparison between our algorithm and [28, 32] is given in Appendix D. While computing the Whittle indices of a known arm's model is still a challenge, there is interesting work in trying to learn Whittle index when only the arm's simulator is given and the arm's model is unknown. For instance, [11, 33] use Q-learning algorithm to estimate Whittle index as time evolves in finite-state restless bandit problems. Moreover, the work of [34] uses deep reinforcement learning framework to estimate the Whittle indices of the arms with large state space or convoluted transition kernel. With "strongly indexable" property introduced by the authors, deadline scheduling, recovering bandits, and wireless scheduling problems are considered and their method numerically outperforms or matches state-of-the-art control policies in each of the three problems.

For non-discounted restless bandit, the author of [28] only provides a brief description of how the fast-pivoting algorithm given in the discounted case can be adapted although no explicit algorithm is given in that paper. For continuous-time n -state restless bandits, the work of [35] proposes an algorithm to check indexability and compute Whittle index with a complexity exponential in the number of states n of each arm. According to Remark 4.1 of that paper, this complexity can be reduced to $O(n^5)$ if the restless bandit is known to be indexable and threshold-based policies are optimal. It is stated that their approach is not applicable for discounted restless bandits. For learning aspect, [36] introduces a new index function that relies on Linear Programming solution in *finite-horizon* setting and proposes an optimistic algorithm to learn their index function when the model is not given.

Contributions. In this paper, we investigate Whittle index computation in restless bandit problems and present two main contributions. Our first contribution is to propose a unified algorithm that computes the Whittle indices

¹multiplications and additions of real numbers, regardless of their values

of Markovian bandits for both discounted and non-discounted restless bandits. Our algorithm, which can be viewed as a refinement of the algorithm in [32], detects whether the input problem is indexable and computes Whittle index if the problem is. As a byproduct, our algorithm can compute Gittins index in rested bandits which are a subclass of restless bandits. This algorithm computes the indices in increasing order, and relies on an efficient use of Sherman-Morrison formula to compute Whittle index in $(2/3)n^3 + O(n^2)$ plus subcubic time [37] to solve a linear system of order n . This algorithm can detect on the fly if a computed index violates the indexability condition, which adds an extra $(1/3)n^3 + O(n^2)$ arithmetic operations. This later test is optional: the complexity of our algorithm is $n^3 + o(n^3)$ when testing indexability and $(2/3)n^3 + o(n^3)$ without the test.

As a second contribution, we show how to reduce the complexity of the above algorithm to obtain the first subcubic algorithm to compute Whittle index. This improvement is made possible by the fact that a linear system can be solved in subcubic time. By carefully reordering the computations, we show that it is possible to reduce the use of the Sherman-Morrison formula at the price of solving more linear systems. The subcubic complexity comes by striking a good balance between having too many or too few linear systems to solve. By using the current fastest matrix multiplication method, our algorithm can test indexability and compute Whittle index in $O(n^{2.5286})$. Our algorithm is also the first subcubic algorithm to compute Gittins index.

Finally, we provide an open-source implementation of our algorithm and present an empirical evaluation of the performance of our implementation. Our results show that our algorithm is very efficient in computing Whittle index and testing indexability. Moreover, our simulations indicate that the subcubic version of our algorithm does not just have an asymptotically small complexity but is also faster in practice than our original $(2/3)n^3$ algorithm. Testing the indexability and computing indices takes less than one second for $n = 1000$ states and less than 10 minutes for $n = 15000$ states. This is 15 to 20 times faster than the computation times reported in [28].

Road map. The paper is organized as follows. We introduce the problem and recall the definition of Whittle index in Section 2. In Section 3, we characterize Whittle index and provide a general idea of how to compute Whittle index. Then, we show, in Section 4, how to use the Sherman-Morrison formula to compute the indices efficiently. We then show how to reduce the complexity of the algorithm by using fast matrix multiplication algorithms in Section 5. We show how to adapt this approach to the discounted case in Section 7. We compare the numerical result of different variants of our algorithm in Section 6. Finally, we conclude in Section 8.

2 Restless Markovian Bandit and Indexability

2.1 Restless bandit arms and multi-armed bandit

In this paper, a restless bandit arm (that we denote later by RB) is a Markov decision process (MDP) with discrete state space $[n] := \{1 \dots n\}$ and binary action space $\{0, 1\}$, where 0 denotes the action “rest” and 1 denotes the action “activate”. The time is discrete and the evolution is Markovian: If the MDP is in state i and action a is chosen, the decision maker earns an instantaneous reward r_i^a and the arm transitions to a new state j with probability P_{ij}^a . We denote this MDP by the pair (\mathbf{r}, \mathbf{P}) . A policy π for a RB is a subset of $[n]$ such that the policy chooses to activate the arm in state i if $i \in \pi$. We say that π is the set of *active* states and we say that state i is *passive* if $i \notin \pi$. By abuse of notation, we note $\pi_i = 1$ if $i \in \pi$ and $\pi_i = 0$ if $i \notin \pi$.

A restless multi-armed bandit problem is a finite collection of independent RB arms. At time t , the decision maker observes the state of all arms, and chooses which arms to activate. The decision maker then earns a reward that is the sum of the rewards of all arms. This problem is called a multi-armed restless bandit. The name “restless” comes from the fact that an arm put at rest may still transition to a new state. Such a problem is notoriously difficult to solve, as its complexity grows exponentially with the number of arms. We next describe a very efficient heuristic to solve this problem, known as the Whittle index policy.

2.2 Indexability and Whittle index

Consider an arm (\mathbf{r}, \mathbf{P}) . For each $\lambda \in \mathbb{R}$, we define a λ -penalized MDP whose transition matrix are the same as in the original MDP and whose reward when taking action a in state i is $r_i^a - \lambda a$. The quantity λ can be viewed as a penalty for taking the active action. We consider that a decision maker seeks to maximize the average gain over an infinite horizon (the discounted case will be discussed in Section 7). To avoid the dependence on initial states, we assume that the arm is unichain².

For a given penalty λ , we denote by $\pi^*(\lambda)$ the minimal³ policy that is optimal for a given penalty λ . When the penalty λ is very large, activating is highly penalized, in which case the optimal policy is to put all states to rest: $\pi = \emptyset$. At the opposite, when the penalty is very low, activating is highly subsidized, in which case the optimal policy is to activate all states: $\pi = [n]$. An arm is said to be indexable when the optimal policy $\pi^*(\lambda)$ is a strictly decreasing function of the penalty λ . More precisely, we define the indexability as follows:

Definition 1 (Indexability). *An arm (\mathbf{r}, \mathbf{P}) is indexable if for all state $i \in [n]$, there exists a critical penalty λ_i such that the action “activate” is optimal in*

²a MDP is unichain if all policies induce a Markov chain with a unique stationary distribution (see [38])

³Minimal for the inclusion order. Indeed, by Lemma 6, if two policies π and π' are optimal for a given penalty λ , then the intersection $\pi \cap \pi'$ is also optimal. Hence, the minimal optimal policy is the intersection of all optimal policies.

state i for the penalty λ if and only if $\lambda \leq \lambda_i$. In this case, we call λ_i the Whittle index of state i .

For an indexable arm, the minimal optimal policy for the penalty λ is the set of states whose index is strictly higher than λ : $\pi^*(\lambda) = \{i : \lambda < \lambda_i\}$. Note that the policy $\pi^*(\lambda)$ is not necessary the unique optimal policy. In particular, for a penalty equal to the Whittle index λ_i of state i both the action rest and activate are optimal in state i .

Remark: By Definition 1, it is possible that in some indexable arms, there are states whose index value equals to positive or negative infinite. This is not possible for the arms that are *unichain* because when a λ -penalized MDP is *unichain* there exists the smallest and biggest penalties, $\mu > -\infty$ and $\mu' < +\infty$, such that the only optimal policy for $\lambda \leq \mu$ is $\pi^*(\lambda) = [n]$ and for $\lambda \geq \mu'$ is $\pi^*(\lambda) = \emptyset$. We should mention that *unichain* is the minimal assumption in the sense that there exists a 3-state MDP that are is not unichain and for which $\lim_{\lambda \rightarrow -\infty} \pi^*(\lambda) \neq [n]$ and $\lim_{\lambda \rightarrow +\infty} \pi^*(\lambda) \neq \emptyset$.

2.3 Characterization of indexability

For a given policy π and a given penalty λ , we denote by $g^\pi(\lambda)$ the average gain⁴ of the policy π for the λ -penalized MDP. We denote by $g^*(\lambda) = \max_{\pi \subset [n]} g^\pi(\lambda)$ the optimal average gain for the penalty λ .

Lemma 1. *An arm is indexable if and only if there is a non decreasing sequence of values $\mu^0 := -\infty < \mu_{\min}^1 \leq \mu_{\min}^2 \leq \dots \leq \mu_{\min}^n < \mu^{n+1} := +\infty$ and a sequence of policies $\pi^1 = [n] \supseteq \pi^2 \supseteq \dots \supseteq \pi^{n+1} = \emptyset$ such that π^k is an optimal policy for all penalty $\lambda \in [\mu_{\min}^{k-1}, \mu_{\min}^k]$ and the minimal optimal policy for all $\lambda \in [\mu_{\min}^{k-1}, \mu_{\min}^k]$ if $\mu_{\min}^{k-1} \neq \mu_{\min}^k$.*

Moreover, if the arm is indexable, then the index of state i is μ_{\min}^k , where $k \geq 1$ is such that $i \in \pi^k \setminus \pi^{k+1}$.

Proof The lemma is a direct consequence of the definition of indexability. Indeed, assume first that there exist sequences μ_{\min} and π that satisfy the condition of the lemma. For a state i , let $\lambda_i := \mu_{\min}^k$, where k is such that $i \in \pi^k \setminus \pi^{k+1}$. Then, λ_i satisfies all conditions of the critical penalty of Definition 1. This implies that the arm is indexable. Assume now that the arm is indexable, and let σ^k be the state with the k th smallest index (where ties are broken arbitrarily). Let $\mu_{\min}^k := \lambda_{\sigma^k}$ be the index of the state σ^k and define the policy $\pi^{k+1} = [n] \setminus \{\sigma^1, \dots, \sigma^k\}$. Then the sequences μ_{\min}^k and π^k satisfy the conditions of the lemma. \square

To illustrate Lemma 1, we show in Figure 1 two arms: one indexable (on the left) and one non-indexable (on the right). In both figures, we plot in dashed-dot black line the average gain $g^\pi(\lambda)$ as a function of λ for all policies $\pi \subset [n]$ (there are $2^4 = 16$ of them). We also plot in solid red line the average gain

⁴Note that as we assume that the arm is unichain, the average gain exists and does not depend on the state [38].




Fig. 1: The average gain as a function of penalty for two examples, one which is indexable (Figure 1a) and the other one which is not (Figure 1b).

of the optimal policy. In Figure 1a), we observe that the sequence of optimal policies is decreasing: $\{1, 2, 3, 4\} \supset \{1, 2, 3\} \supset \{1, 2\} \supset \{2\} \supset \emptyset$, which means that this arm is indexable. The vertical red lines represent the indices of all states: for a given λ , the optimal policy (written in red) is to activate the arm at all states whose λ_i is greater than the penalty λ . For Figure 1b), the sequence of optimal policies is $\{1, 2, 3\} \supset \{1, 2\} \supset \{1\} \subset \{1, 3\} \supset \{3\} \supset \emptyset$: this example is not indexable because of what happens around $\lambda \approx 0.7$: it is optimal to put 3 at rest when $\lambda \in [0.41, 0.66] \cup [0.71, +\infty]$ but not when $\lambda \in [0.66, 0.71]$.

3 Condition for Indexability and Basic Algorithm

This section aims at providing a basic algorithm to detect whether an arm is indexable or not and if it is the case, compute the Whittle index of any state i of the arm.

3.1 Overview of the algorithm

Our algorithm computes Whittle index in increasing order by navigating through the optimal average gain and using the characterization provided by Lemma 1. It follows the graphical construction given in Figure 1a. It also uses the following facts:

- For a given policy π , the average gain $g^\pi(\lambda)$ is an affine function of λ . As a result, it is possible to compute the penalty λ at which two curves $g^\pi(\lambda)$ and $g^{\pi'}(\lambda)$ intersect.

- For an indexable arm, computing the index can be done by a greedy algorithm that constructs a sequence of penalties $\mu_{\min}^1 \leq \mu_{\min}^2 \leq \dots \leq \mu_{\min}^n$ and a sequence of policies $\pi^1 \supseteq \dots \supseteq \pi^n$ by looking at where $g^{\pi^k}(\lambda)$ and $g^{\pi^k \setminus \{i\}}(\lambda)$ intersect for all $i \in \pi^k$.
- The arm is indexable if and only if for all k , π^k is an optimal policy for penalty μ_{\min}^k .

Note that the last fact above is derived from Lemma 1 and Lemma 2. Effectively, from Lemma 1, the arm is indexable if and only if for all k , π^k is optimal for penalty $\lambda \in [\mu_{\min}^{k-1}, \mu_{\min}^k]$. By Lemma 2, π^k is optimal for penalty $\lambda \in [\mu_{\min}^{k-1}, \mu_{\min}^k]$ if and only if it is optimal for penalty $\lambda \in \{\mu_{\min}^{k-1}, \mu_{\min}^k\}$. Finally, by construction, for all k , π^k is optimal for penalty $\lambda = \mu_{\min}^{k-1}$.

This leads to Algorithm 1, that we write in pseudo-code. This algorithm relies on two subroutines: on Line 4, to compute the next index and on Line 10 to test if a policy is optimal. We will describe later in the paper how to implement these functions in an efficient manner. Note that in all the paper, we use the superscript k (e.g., π^k, μ^k, σ^k) to refer to the quantities computed at iteration k . We use the subscripts i or j (e.g., $\pi_i, \lambda_i, \mu_i, \pi_j$) to refer to the quantities related to states i or j .


Algorithm 1 Test indexability and compute Whittle index if the arm is indexable.

```

1: Set  $\pi^1 := [n]$ ,  $k := 1$ , and  $\mu_{\min}^0 := -\infty$ 
2: for  $k = 1$  to  $n$  do
3:   for  $i \in \pi^k$  do
4:     Let  $\mu_i^k = \inf\{\lambda : g^{\pi^k}(\lambda) = g^{\pi^k \setminus \{i\}}(\lambda)\}$ 
5:   end for
6:   if  $\arg \min_{i \in \pi^k : \mu_i^k \geq \mu_{\min}^{k-1}} \mu_i^k = \emptyset$  then
7:     return the arm is not indexable
8:   end if
9:   Let  $\sigma^k := \arg \min_{i \in \pi^k : \mu_i^k \geq \mu_{\min}^{k-1}} \mu_i^k$  and  $\lambda_{\sigma^k} := \mu_{\min}^k := \mu_{\sigma^k}^k$ 
10:  if  $\pi^k$  is not an optimal policy for the penalty  $\mu_{\min}^k$  then
11:    return the arm is not indexable
12:  end if
13:  Set  $\pi^{k+1} := \pi^k \setminus \{\sigma^k\}$ 
14: end for
15: return the arm is indexable and the indices are  $\lambda$ .
```

To illustrate how the algorithm works, we plot the first three iterations in Figure 2 for the arm represented in Figure 1a. The optimal policy for small λ is $\pi = [n]$ so the first iteration starts with $\pi^1 = \{1, 2, 3, 4\}$ and computes the equation of the affine function $g^{\pi^1}(\cdot)$. We then compute the equations of the affine functions $g^{\pi^1 \setminus \{i\}}(\cdot)$ for $i \in \pi^1$. This is represented in Figure 2a. We observe that $g^{\pi^1 \setminus \{4\}}(\cdot)$ is the first affine function to cross $g^{\pi^1}(\cdot)$ at $\mu_{\min}^1 = \lambda_4 \approx$

0.9. After that, we set $\pi^2 = \{1, 2, 3\}$ and compute the equation of the affine functions $g^{\pi^2 \setminus \{i\}}(\cdot)$ for $i \in \pi^2$ and use them to compute $\mu_{min}^2 = \lambda_3 \approx 0.21$ (Figure 2b). We then continue with the affine functions $g^{\{2\}}(\cdot)$ and $g^{\{1\}}(\cdot)$ (note that $g^{\{1\}}(\cdot)$ is not visible on the figure because $g^{\{1\}}(\lambda) < 0.35$ for $\lambda \in [0, 0.3]$, see Figure 1a). We find $\mu_{min}^3 = \lambda_1 \approx 0.28$ (Fig. 2c). We are left with the last iteration (not shown here) to compute the intersection between $g^{\{2\}}(\lambda)$ and $g^{\emptyset}(\lambda)$ which occurs at $\mu_{min}^4 = \lambda_2 \approx 0.57$ (see Figure 1a).



(a) After the 1st iteration. (b) After the 2nd iteration. (c) After the 3rd iteration.

Fig. 2: Illustration of the behavior of our algorithm for an arm with 4 states. We plot the average gain as a function of penalty λ of the various policies and of the policy computed by the algorithm (the model is the same as the one of Figure 1a.)

3.2 Correctness of Algorithm 1

The following result shows that Algorithm 1 is correct when testing indexability or computing Whittle index.

Theorem 1. *For an unichain arm, if the arm is indexable, then the vector λ returned by Algorithm 1 corresponds to the Whittle indices. If the arm is not indexable, then the algorithm returns that the arm is not indexable.*

Proof Recall that $g^*(\lambda) = \max_{\pi} g^{\pi}(\lambda)$ is the optimal average gain for penalty λ (i.e., the solid red curve of Figure 1). We define the function $g^{algo}(\cdot)$ as the average gain of the policies computed by Algorithm 1 (i.e., the blue curve on Figure 2) as follows: For all $k \in \{1, \dots, n\}$ and $\lambda \in [\mu_{min}^{k-1}, \mu_{min}^k]$: $g^{algo}(\lambda) := g^{\pi^k}(\lambda)$. Our proof relies on the following lemma.

Lemma 2. *The three following facts are equivalent:*

1. For all $k \in [n]$: $g^*(\mu_{\min}^k) = g^{algo}(\mu_{\min}^k)$
2. For all $\lambda \in \mathbb{R}$: $g^*(\lambda) = g^{algo}(\lambda)$.
3. The arm is indexable.

Indeed, assume that Lemma 2 holds. Then, if the arm is not indexable Line 6 or 10 will return “non-indexable” because of 1. If the arm is indexable, then the sequence of policies π^k satisfy the assumptions of Lemma 1 and the values λ_{σ^k} are therefore the Whittle indices.

Proof of Lemma 2 1 \Leftrightarrow 2: $g^{algo}(\cdot)$ is piecewise affine and convex. As such, it is the largest affine function that goes through the points $g^{algo}(\mu_{\min}^k)$. Since g^* is the optimal average gain, we have $g^{algo}(\lambda) \leq g^*(\lambda)$. Since g^* is also convex and goes through $g^{algo}(\mu_{\min}^k)$, both functions are equal.


2 \Rightarrow 3: If 2 holds, then the sequence of policies π^k satisfies the assumptions of Lemma 1 and the values λ_{σ^k} are therefore the Whittle indices.

3 \Rightarrow 2: Assume that the arm is indexable. We prove by induction on k that

$$g^*(\lambda) = g^{algo}(\lambda) \quad \text{for all } \lambda \leq \mu_{\min}^k. \quad (1)$$

This is clearly true for $k = 0$ since $\mu_{\min}^0 = -\infty$. Because μ_{\min}^1 is the smallest penalty at which one of the curves $g^{\pi^1 \setminus \{i\}}(\cdot)$ crosses $g^{\pi^1}(\cdot)$, (1) also holds for $k = 1$. Since the arm is unichain, we have $\mu_{\min}^1 > \mu_{\min}^0$.

Assume that (1) holds for some $k \geq 0$. If $\mu_{\min}^{k+1} = \mu_{\min}^k$, then (1) holds for $k + 1$. Assume now that $\mu_{\min}^{k+1} > \mu_{\min}^k$. We argue that π^{k+1} is the minimal optimal policy for the penalty μ_{\min}^k . Indeed, if there exists another policy $\pi' \subsetneq \pi^{k+1}$ that is also optimal, then by Lemma 6, there would exist $i \in \pi^{k+1} \setminus \pi'$ such that the policy $\pi^{k+1} \setminus \{i\}$ is optimal for μ_{\min}^k which would imply that $\mu_{\min}^{k+1} = \mu_{\min}^k$. Hence, we can assume that π^{k+1} is the minimal optimal policy for the penalty μ_{\min}^k . Now, let $\nu = \inf\{\lambda > \mu_{\min}^k \text{ such that } \pi^{k+1} \text{ is not optimal}\}$. By continuity of the functions $g^{\pi}(\lambda)$, there are two distinct optimal policies at ν (π^{k+1} and another policy π'). By definition of indexability, $\pi^{k+1} \supseteq \pi'$. By Lemma 6, this implies that there exists $\pi^{k+1} \setminus \{i\}$ that is optimal at ν , which implies that $\mu_{\min}^{k+1} = \nu$.



(a) Overview of the visited μ_{\min}^k .

(b) Zoom on $\lambda \in [0.65, 0.73]$

Fig. 3: Algorithm v.s. optimal average gains for a non-indexable arm. We plot the average gain as a function of penalty for the non-indexable arm with 3 states of Figure 1b.

□

We should note that

- if $\mu_i^k < \mu_{\min}^{k-1}$ for all $i \in \pi^k$, then π^k remains optimal for all $\lambda > \mu_{\min}^{k-1}$. This is not possible for indexable unichain arms because for an infinite penalty the only optimal strategy is to put all states at rest. That is why Line 6 returns that the arm is not indexable
- by Line 9, it is possible to have $\mu_{\min}^k = \mu_{\min}^{k-1}$. This happens when several states have the same value of Whittle index. This is not problematic because we are sure that $\sigma^k \neq \sigma^{k-1}$ by Line 13.

□

3.3 Naive implementation of Algorithm 1 (in $O(n^4)$)

For a given penalty λ and a policy π , let $g^\pi(\lambda)$ be a scalar (called the average gain), and let $\mathbf{h}^\pi(\lambda)$ be a vector (called the bias) such that $h_1^\pi(\lambda) = 0$, and for all state $i \in [n]$:

$$g^\pi(\lambda) + h_i^\pi(\lambda) = r_i^{\pi_i} - \lambda\pi_i + \sum_{j=1}^n P_{ij}^{\pi_j} h_j^\pi(\lambda). \quad (2)$$

The above system is a system of $n+1$ linear equations with $n+1$ variables (the additional equation begins $h_1^\pi(\lambda) = 0$). As the MDP is unichain, according to [38], the average gain and bias are uniquely defined by the system of linear equations (2), together with the condition that $h_1^\pi(\lambda) = 0$. Note that in (2) the sum is for $j = 1$ to n . Since $h_1^\pi(\lambda) = 0$, it can be transformed into a sum from $j = 2$ to n .

Let us define the vector $\mathbf{v}^\pi(\lambda) := [g^\pi(\lambda), h_2^\pi(\lambda), \dots, h_n^\pi(\lambda)]^T$ which is similar to the vector $\mathbf{h}(\lambda)$ in which we replaced $h_1^\pi(\lambda)$ by $g^\pi(\lambda)$. We can write Equation (2) under a matrix form as:

$$\mathbf{A}^\pi \mathbf{v}^\pi(\lambda) = \mathbf{r}^\pi - \lambda\boldsymbol{\pi}, \quad (3)$$

where \mathbf{r}^π is the reward vector of policy π : $\mathbf{r}^\pi = [r_1^{\pi_1}, \dots, r_n^{\pi_n}]^T$, $\boldsymbol{\pi} = [\pi_1, \dots, \pi_n]^T$, and \mathbf{A}^π is the following square matrix:

$$\mathbf{A}^\pi := \begin{bmatrix} 1 & & & \\ 1 & 1 & & \\ & & \ddots & \\ 1 & & & 1 \end{bmatrix} - \begin{bmatrix} 0 & P_{12}^{\pi_1} & \dots & P_{1n}^{\pi_1} \\ 0 & P_{22}^{\pi_2} & \dots & P_{2n}^{\pi_2} \\ \vdots & & \ddots & \vdots \\ 0 & P_{n2}^{\pi_n} & \dots & P_{nn}^{\pi_n} \end{bmatrix} = \begin{bmatrix} 1 & -P_{12}^{\pi_1} & \dots & -P_{1n}^{\pi_1} \\ 1 & 1 - P_{22}^{\pi_2} & \dots & -P_{2n}^{\pi_2} \\ \vdots & & \ddots & \vdots \\ 1 & -P_{n2}^{\pi_n} & \dots & 1 - P_{nn}^{\pi_n} \end{bmatrix} \quad (4)$$

As we show in Lemma 7, in unichain MDP, the matrix \mathbf{A}^π is invertible for any policy π . In consequence, \mathbf{v}^π is an affine function of λ :

$$\mathbf{v}^\pi(\lambda) = (\mathbf{A}^\pi)^{-1}(\mathbf{r}^\pi - \lambda\boldsymbol{\pi}) = (\mathbf{A}^\pi)^{-1}\mathbf{r}^\pi - \lambda(\mathbf{A}^\pi)^{-1}\boldsymbol{\pi} \quad (5)$$

The following lemma gives a characterization of the indifference between activating or resting a state under a given penalty. Its proof is in Appendix A.3.

Lemma 3. *Let π be a policy, $i \in \pi$ be any active state and λ be a penalty. Then the two policies π and $\pi \setminus \{i\}$ have the same average gain for the penalty λ (i.e., $g^\pi(\lambda) = g^{\pi \setminus \{i\}}(\lambda)$) if and only if*

$$\delta_i - \lambda + \Delta_i v^\pi(\lambda) = 0, \quad (6)$$

where $\delta_i = r_i^1 - r_i^0$ and Δ_i is the vector such that $\Delta_{i1} = 0$ and $\Delta_{ij} = P_{ij}^1 - P_{ij}^0$ for $j \in \{2, \dots, n\}$.

Using this result, computing a value μ_i^k at Line 4 of Algorithm 1 can be constructed by finding a penalty that satisfies (6) for the policy π^k , which is a (one-dimensional) linear equation. Indeed, let $\mathbf{d}^{\pi^k} = -(A^{\pi^k})^{-1} \boldsymbol{\pi}^k$. By (5), $\mathbf{v}^{\pi^k}(\lambda)$ is a linear function of λ whose derivative is \mathbf{d}^{π^k} . In particular, one has: $\mathbf{v}^{\pi^k}(\lambda) = \mathbf{v}^{\pi^k}(\mu_{\min}^{k-1}) + (\lambda - \mu_{\min}^{k-1}) \mathbf{d}^{\pi^k}$. Using the lemma above for μ_{\min}^{k-1} , we have $\mathbf{v}^{\pi^k}(\mu_{\min}^{k-1}) = \mathbf{v}^{\pi^{k-1}}(\mu_{\min}^{k-1})$. Hence, Equation (6) can be written as $\delta_i - \lambda + \Delta_i (\mathbf{v}^{\pi^{k-1}}(\mu_{\min}^{k-1}) + (\lambda - \mu_{\min}^{k-1}) \mathbf{d}_i^{\pi^k}) = 0$ which is satisfied if and only if $\lambda = \mu_i^k$, where

$$\mu_i^k := \frac{\delta_i + \Delta_i \mathbf{v}^{\pi^{k-1}}(\mu_{\min}^{k-1}) - \mu_{\min}^{k-1} \Delta_i \mathbf{d}_i^{\pi^k}}{1 - \Delta_i \mathbf{d}_i^{\pi^k}}. \quad (7)$$

This shows that, for a given k , computing all μ_i^k of Line 4 can be done in $O(n^3)$: A $O(n^3)$ to compute the inverse of matrix A^π and compute \mathbf{d}^π , plus some smaller order terms to compute the solutions of (7). Similarly, the test in Line 10 of Algorithm 1 can also be implemented in $O(n^3)$ by using Lemma 4 below. This leads to an overall complexity of $O(n^4)$ for Algorithm 1 that contains n loops each having a $O(n^3)$ complexity.

Lemma 4. *Let π be a policy, λ be a penalty and $\mathbf{v}^\pi(\lambda)$ as in (5). Then the policy π is optimal for the penalty λ if and only if for all state i :*

$$\delta_i - \lambda + \Delta_i \mathbf{v}^\pi(\lambda) \leq 0 \text{ if } i \notin \pi \quad \text{and} \quad \delta_i - \lambda + \Delta_i \mathbf{v}^\pi(\lambda) \geq 0 \text{ if } i \in \pi \quad (8)$$

where δ_i and Δ_i are as in Lemma 3.

Proof The proof is a direct consequence of *Policy Improvement* as described in Theorem 8.4.4 of [38]: a policy that satisfies (8) cannot be improved and is therefore optimal. \square

4 The $(2/3)n^3 + o(n^3)$ Algorithm

This section describes a way to implement Algorithm 1 efficiently using $O(n^3)$ operations. The main idea is to use the Sherman-Morrison formula to compute in $O(n^2)$ the bias vector associated to $g^{\pi \setminus \{i\}}(\lambda)$ from the one associated to $g^\pi(\lambda)$. This leads to a $O(n^3)$ algorithm. Once this main idea is in place, we show how to avoid unnecessary computations to obtain an algorithm that performs $(2/3)n^3 + o(n^3)$ arithmetic operations.

4.1 Additional notations

In order to obtain a more efficient and compact algorithm, for an iteration k and a state i , we define $y_i^k := \Delta_i \mathbf{d}^{\pi^k}$ and $z_i^k := \Delta_i \mathbf{v}^{\pi^k}(\mu_{\min}^k)$, where \mathbf{d}^{π^k} , Δ_i and \mathbf{v}^{π^k} are as in (7). Equation (7) can be rewritten as

$$\mu_i^k = \frac{\delta_i + z_i^{k-1} - \mu_{\min}^{k-1} y_i^k}{1 - y_i^k}. \quad (9)$$

The above equation can be used to compute μ_i^k and μ_{\min}^k easily from y_i^k and z_i^{k-1} .

For a given policy π , \mathbf{v}^π is a linear function of λ whose derivative is \mathbf{d}^π . Hence, the value z_i^k can be computed from y_i^k and z_i^{k-1} as:

$$\begin{aligned} z_i^k &= \Delta_i \mathbf{v}^{\pi^k}(\mu_{\min}^k) \\ &= \Delta_i \mathbf{v}^{\pi^k}(\mu_{\min}^{k-1}) + \Delta_i \mathbf{d}^{\pi^k}(\mu_{\min}^k - \mu_{\min}^{k-1}) \\ &= \Delta_i \mathbf{v}^{\pi^{k-1}}(\mu_{\min}^{k-1}) + (\mu_{\min}^k - \mu_{\min}^{k-1}) y_i^k \\ &= z_i^{k-1} + (\mu_{\min}^k - \mu_{\min}^{k-1}) y_i^k. \end{aligned} \quad (10)$$

This shows that the critical values to compute are the variables y_i^k . In the remainder of this section, we show that the quantity y_i^k can be computed efficiently by a recursive formula.

4.2 Application of Sherman-Morrison Formula

To compute $y_i^{k+1} = \Delta_i \mathbf{d}^{\pi^{k+1}}$, we need to compute the quantities $\mathbf{d}^{\pi^{k+1}} = -(\mathbf{A}^{\pi^{k+1}})^{-1} \boldsymbol{\pi}^{k+1}$. This requires the inverse of $\mathbf{A}^{\pi^{k+1}}$. From the definition of π^k , two policies π^k and π^{k+1} differs by exactly one state: $\boldsymbol{\pi}^{k+1} = \boldsymbol{\pi}^k - \mathbf{e}_{\sigma^k}$ where \mathbf{e}_j denotes the column vector with a 1 in j th coordinate and 0's elsewhere. Also by definition of \mathbf{A}^π in (4), the two matrices \mathbf{A}^{π^k} and $\mathbf{A}^{\pi^{k+1}}$ differs only at the line σ^k :

$$\mathbf{A}^{\pi^{k+1}} = \mathbf{A}^{\pi^k} + \mathbf{e}_{\sigma^k} \Delta_{\sigma^k} \quad (11)$$

where Δ_{σ^k} is a row vector defined as in the previous section.

One can compute efficiently the inverse of matrix $\mathbf{A}^{\pi^{k+1}}$ from the one of \mathbf{A}^{π^k} by using Sherman-Morrison formula, which says that if $\mathbf{A} \in \mathbb{R}^{n \times n}$ is an invertible square matrix and $\mathbf{p}, \mathbf{q} \in \mathbb{R}^n$ are two column vectors such that $\mathbf{A} + \mathbf{p}\mathbf{q}^\top$ is invertible, then:

$$(\mathbf{A} + \mathbf{p}\mathbf{q}^\top)^{-1} = \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1}\mathbf{p}\mathbf{q}^\top\mathbf{A}^{-1}}{1 + \mathbf{q}^\top\mathbf{A}^{-1}\mathbf{p}}.$$

Let $X_{ij}^k = \Delta_i(\mathbf{A}^{\pi^k})^{-1}\mathbf{e}_j$. Following (11), we can apply Sherman-Morrison formula with matrix \mathbf{A}^{π^k} ; and vectors $\mathbf{p} = \mathbf{e}_{\sigma^k}$, and $\mathbf{q}^\top = \Delta_{\sigma^k}$. After some simplification, we get:

$$\begin{aligned} X_{ij}^{k+1} &= \Delta_i(\mathbf{A}^{\pi^{k+1}})^{-1}\mathbf{e}_j = \Delta_i(\mathbf{A}^{\pi^k} + \mathbf{e}_{\sigma^k}\Delta_{\sigma^k})^{-1}\mathbf{e}_j \\ &= X_{ij}^k - \frac{X_{i\sigma^k}^k X_{\sigma^k j}^k}{1 + X_{\sigma^k \sigma^k}^k}. \end{aligned} \quad (12)$$

For $y_i^{k+1} = \Delta_i \mathbf{d}^{\pi^{k+1}} = -\Delta_i(\mathbf{A}^{\pi^{k+1}})^{-1}\boldsymbol{\pi}^{k+1}$, we use the fact that $\boldsymbol{\pi}^{k+1} = \boldsymbol{\pi}^k - \mathbf{e}_{\sigma^k}$ and apply Sherman-Morrison formula to get:

$$\begin{aligned} y_i^{k+1} &= -\Delta_i(\mathbf{A}^{\pi^k} + \mathbf{e}_{\sigma^k}\Delta_{\sigma^k})^{-1}(\boldsymbol{\pi}^k - \mathbf{e}_{\sigma^k}) \\ &= -\Delta_i(\mathbf{A}^{\pi^k})^{-1}(\boldsymbol{\pi}^k - \mathbf{e}_{\sigma^k}) + \frac{\Delta_i(\mathbf{A}^{\pi^k})^{-1}\mathbf{e}_{\sigma^k}\Delta_{\sigma^k}(\mathbf{A}^{\pi^k})^{-1}}{1 + \Delta_{\sigma^k}(\mathbf{A}^{\pi^k})^{-1}\mathbf{e}_{\sigma^k}}(\boldsymbol{\pi}^k - \mathbf{e}_{\sigma^k}) \\ &= y_i^k + X_{i\sigma^k}^k + \frac{X_{i\sigma^k}^k(-y_{\sigma^k}^k - X_{\sigma^k \sigma^k}^k)}{1 + X_{\sigma^k \sigma^k}^k} \\ &= y_i^k + \frac{X_{i\sigma^k}^k(1 - y_{\sigma^k}^k)}{1 + X_{\sigma^k \sigma^k}^k} = y_i^k + (1 - y_{\sigma^k}^k)X_{i\sigma^k}^{k+1} \end{aligned} \quad (13)$$

The above formula indicate how to compute \mathbf{y}^{k+1} from \mathbf{y}^k . To complete this analysis, let us show that $\mathbf{y}^1 = \mathbf{0}$. For a given policy π , the vector \mathbf{d}^π satisfies the same equation as Equation (3) but replacing $r_i^{\pi_i} - \lambda\pi_i$ by $-\pi_i$. This implies that for $\pi = \pi^1 = [1, \dots, 1]^T$, one has $\mathbf{d}^{\pi^1} = [-1, 0, \dots, 0]^T$ as $d_1^{\pi^1}$ is the average gain of a MDP whose reward is one in all states and $d_2^{\pi^1}, \dots, d_n^{\pi^1}$ is the bias of this MDP. This shows that for all i , we have $y_i^1 = \Delta_i \mathbf{d}^{\pi^1} = -\Delta_{i1} = 0$.

4.3 Detailed algorithm

Equation (9) shows how to compute μ_i^k from the values of y_i^k and z_i^{k-1} while (13), (12) and (10) show how to compute the values of y , z and \mathbf{X} recursively in k . In order to compute μ_{\min}^k and σ^k , one only needs to compute the values μ_i^k for $i \in \pi^k$. Once μ_{\min}^k is determined as in Line 9, we need to verify the optimality of π^k for penalty μ_{\min}^k . That is, we check whether for any passive state $i \notin \pi^k$, the ‘‘rest’’ action is worst than the ‘‘activate’’ action in state i for

penalty μ_{\min}^k . If it is the case, then the arm is non-indexable. By Lemma 4, it is sufficient to check that for $i \in [n] \setminus \pi^k$:

$$\mu_{\min}^k < \delta_i + \Delta_i \mathbf{v}^{\pi^k}(\mu_{\min}^k) = \delta_i + z_i^k. \quad (14)$$

Algorithm 2 Test indexability and compute Whittle indices (if indexable).

- 1: Set $\pi^1 = [n]$, $k_0 = 1$, $\delta_i = r_i^1 - r_i^0$, $\Delta_{i1} = 0$ and $\Delta_{ij} = P_{ij}^1 - P_{ij}^0, \forall i \in [n], j \in \{2 \dots n\}$.
 - 2: Set $\mathbf{X}^1 = \Delta(\mathbf{A}^{\pi^1})^{-1}$, where \mathbf{A}^{π^1} is defined by (4).
 - 3: Set $\mu_{\min}^0 = -\infty$, $\mathbf{y}^1 = \mathbf{0}$, and $\mathbf{z}^0 = \mathbf{X}^1 \mathbf{r}^{\pi^1}$
 - 4: **for** $k = 1$ **to** n **do**
 - 5: **for** $i \in \pi^k$ **do**
 - 6: Set $\mu_i^k = \frac{\delta_i + z_i^{k-1} - \mu_{\min}^{k-1} y_i^k}{1 - y_i^k}$ \triangleright For $k = 1$: $\mu_{\min}^0 y_i^1 = -\infty \times 0 = 0$.
 - 7: **end for**
 - 8: **if** $\arg \min_{i \in \pi^k: \mu_i^k \geq \mu_{\min}^{k-1}} \mu_i^k = \emptyset$ **then**
 - 9: **return** the arm is not indexable
 - 10: **end if**
 - 11: Set $\sigma^k := \arg \min_{i \in \pi^k: \mu_i^k \geq \mu_{\min}^{k-1}} \mu_i^k$ and $\lambda_{\sigma^k} = \mu_{\min}^k = \mu_{\sigma^k}^k$
 - 12: Set $\mathbf{z}^k = \mathbf{z}^{k-1} + (\mu_{\min}^k - \mu_{\min}^{k-1}) \mathbf{y}^k$ \triangleright For $k = 1$, one has $\mathbf{z}^1 = \mathbf{z}^0$.
 - 13: **for** $i \in [n] \setminus \pi^k$ **do**
 - 14: **if** $\mu_{\min}^k < \delta_i + z_i^k$ **then**
 - 15: **return** the arm is not indexable
 - 16: **end if**
 - 17: **end for**
 - 18: **if** $k < n$ **then**
 - 19: Set $\pi^{k+1} = \pi^k \setminus \{\sigma^k\}$
 - 20: Update $\mathbf{X}(k)$ \triangleright Here we call Subroutine 3 or Subroutine 4
 - 21: Set $\mathbf{y}^{k+1} = \mathbf{y}^k + (1 - y_{\sigma^k}^k) \mathbf{X}_{:\sigma^k}^{k+1}$
 - 22: **end if**
 - 23: **end for**
 - 24: **return** The arm is indexable and the indices are λ .
-

This leads to Algorithm 2 that can be decomposed as follows:

1. In Line 1 to 3, we initialize the various variables. The main complexity of this part is to compute the matrix \mathbf{X}^1 , which is equivalent to solving the linear system $\mathbf{X} \mathbf{A}^{\pi^1} = \Delta$. It can be done by inverting the matrix \mathbf{A}^{π^1} and multiplying this by the matrix Δ . This can be done in a subcubic complexity by using for instance Strassen's algorithm [37].
2. We then enter the main loop:

Subroutine 3 Update_X(k)

```

1: for  $\ell = 1$  to  $k$  do
2:   for  $i \in [n]$  do  $\triangleright$  or  $i \in \pi^{\ell+1}$  if we do not test indexability.
3:      $X_{i\sigma^k}^{\ell+1} := X_{i\sigma^k}^\ell - \frac{X_{i\sigma^\ell}^\ell}{1 + X_{\sigma^\ell\sigma^\ell}^\ell} X_{\sigma^\ell\sigma^k}^\ell$ 
4:   end for
5: end for

```

- We update the vectors μ, \mathbf{z} by using Equations (9) and (10), and test indexability. This costs $O(n)$ operations per iteration, thus $O(n^2)$ in total.
- We update the vector \mathbf{X} according to (12). The “naive” way to do so is to use Subroutine 3. At iteration k this costs $2kn$ arithmetic operations if we want to test indexability, and $2 \sum_{l=1}^k (n-l)$ if we do not want to test indexability. The total complexity of computing \mathbf{X} is $n^3 + O(n^2)$ arithmetic operations if we test indexability and $(2/3)n^3 + O(n^2)$ if we do not. A detailed study of the arithmetic complexity is provided in Appendix B, where we also provide details on how to efficiently implement the algorithm, including how to optimize memory costs.
- In Line 21, we update the vector \mathbf{y} by using Equation (13), which costs $O(n)$ per iteration.

Hence, the total complexity of this algorithm is $n^3 + o(n^3)$ if we test indexability and $(2/3)n^3 + o(n^3)$ if we do not test indexability. Without testing the indexability, our algorithm has the same main complexity as [28]. However, the algorithm of [28] computes Whittle index only for an arm that is PCL-indexable. Hence we can claim that our algorithm is the first algorithm that computes Whittle index with cubic complexity for all indexable restless bandits. It is also the first algorithm that detects non-indexability of restless bandits.

Remark:

Equivalently, instead of calling Subroutine 3 at Line 20, one could do the following update for all $j \in \pi^k$ and $i \in [n]$ (or $i \in \pi^k$ if we do not test indexability):

$$X_{ij}^{k+1} := X_{ij}^k - \frac{X_{i\sigma^k}^k}{1 + X_{\sigma^k\sigma^k}^k} X_{\sigma^k j}^k. \quad (15)$$

This iterative update is very close to the one used in [28, 32] for discounted restless bandit. This results in an algorithm that has the same total complexity as Subroutine 3 (of $n^3 + O(n^2)$ or $(2/3)n^3 + O(n^2)$ with or without the indexability test) because both algorithms will have computed the same values of X_{ij}^k . The reason to use Subroutine 3 is that, as we will see in the next section, not all values of X_{ij}^k are needed: in particular, for $\ell < k$, the computation of the value $X_{i\sigma^k}^\ell$ has no interest per say and is only useful because it allows to

recursively compute $X_{i\sigma^k}^{k+1}$. In the section below, we show how to reduce the cost by avoiding the computation of $X_{i\sigma^k}^\ell$ when ℓ is much smaller than k . We comment more on the differences with [28, 32] in Appendix D and in particular we explain why our approach can be tuned into a subcubic algorithm while (15) cannot.

5 The Subcubic Algorithm

5.1 Main idea: recomputing \mathbf{X}^k from \mathbf{A}^{π^k} periodically

The main computational burden of Algorithm 2 is concentrated on two lines: on Line 2 where we compute \mathbf{X}^1 by solving a linear system, and on Line 20 where we compute $\mathbf{X}_{:\sigma^k}^{k+1}$ from $\mathbf{X}_{:\sigma^k}^1$. The remainder of the code runs in $O(n^2)$ operations and is therefore negligible for large matrices. In fact, the computation of \mathbf{X}^1 is done by solving a linear system, which can be computed by using a subcubic algorithm (like Strassen [37]). In this section, we show how to optimize our algorithm by reducing the complexity of the update_X() function, at the price of recomputing the full matrix \mathbf{X}^{k_0} from $\mathbf{A}^{\pi^{k_0}}$ periodically.

At iteration k of Algorithm 2, the quantities $X_{i\sigma^k}^{k+1}$ are used at Line 21 to obtain the values \mathbf{y}^{k+1} . The matrix \mathbf{X}^k is defined as $\mathbf{X}^k = \Delta(\mathbf{A}^{\pi^k})^{-1}$. It also satisfies Equation (12), that is, for an iteration ℓ and states i and j , we have:

$$X_{ij}^{\ell+1} = X_{ij}^\ell - \frac{X_{i\sigma^\ell}^\ell X_{\sigma^\ell j}^\ell}{1 + X_{\sigma^\ell \sigma^\ell}^\ell}. \quad (16)$$

The way Subroutine 3 is implemented is to initialize $\mathbf{X}^1 = \Delta(\mathbf{A}^{\pi^1})^{-1}$ and then use (16) recursively to compute the vector $\mathbf{X}_{:\sigma^k}^{k+1}$ from $\mathbf{X}_{:\sigma^k}^1$ at each iteration.

Here, we propose an alternative formulation which consists in recomputing the whole matrix $\mathbf{X}^k = \Delta(\mathbf{A}^{\pi^k})^{-1}$ every K iterations. In the meantime, we use (16) to compute the values $X_{i\sigma^{k-1}}^\ell$ for $\ell \in \{k_0 + 1, \dots, k\}$ where k_0 is the iteration at which we recomputed the whole matrix $\mathbf{X}^{k_0} = \Delta(\mathbf{A}^{\pi^{k_0}})^{-1}$. This can be implemented by changing the call to Subroutine 3 on Line 20 by a call to Subroutine 4.

To see why this can be more efficient, we illustrate in Figure 4 the pairs (ℓ, σ^k) for which we compute the vector $\mathbf{X}_{:\sigma^k}^\ell$, either for Subroutine 3 or Subroutine 4. In each case, a vertical blue line indicates that we recompute the whole matrix \mathbf{X}^k by solving a linear system. The gray zone corresponds to the values (ℓ, σ^k) for which we compute $\mathbf{X}_{:\sigma^k}^\ell$ using Equation (16) and the red squares represent the vector $\mathbf{X}_{:\sigma^k}^{k+1}$ used at Line 21 of Algorithm 2. For Subroutine 3, we do one matrix inversion at the beginning and then compute for all (ℓ, σ^k) with $\ell \leq k$ because the red square at value $(k + 1, \sigma^k)$ is computed starting from the vertical blue line at value $(1, \sigma^k)$. For Subroutine 4, we do


Subroutine 4 Update_X_FMM(k)

```

1: if  $k$  is an iteration at which we recompute the whole  $\mathbf{X}^{k+1}$  then
2:   Set  $k_0 := k + 1$ 
3:   Set  $\mathbf{X}^{k_0} = \Delta(\mathbf{A}^{\pi^{k_0}})^{-1}$ 
4: end if
5: for  $\ell = k_0$  to  $k$  do
6:   for  $i \in [n]$  do ▷ or  $i \in \pi^{\ell+1}$  if we do not test indexability.
7:      $X_{i\sigma^k}^{\ell+1} := X_{i\sigma^k}^\ell - \frac{X_{i\sigma^\ell}^\ell}{1 + X_{\sigma^\ell\sigma^\ell}^\ell} X_{\sigma^\ell\sigma^k}^\ell$ 
8:   end for
9: end for

```

n/K (here $n/K = 4$) full recomputation of \mathbf{X}^k , which correspond to the vertical blue lines. We gain in terms of operations because the surface of the gray zone to compute is divided by n/K .



(a) Implementation of Subroutine 3. (b) Implementation of Subroutine 4

Fig. 4: Illustration of the improvement proposed by replacing Subroutine 3 by Subroutine 4. For Subroutine 4, we solve more linear systems (each vertical blue line corresponds to solving a linear system) but we reduce the gray zone to compute. A linear arrow corresponds to the internal loop of Line 5 of Subroutine 4.

Note that the y -axis of Figure 4 is ordered by increasing value of σ^k (and not by increasing value of k). The value of σ^k is computed at iteration k but is unknown before. This explains why in Subroutine 4, when we recompute the matrix (X_{ij}^{k+1}) at an iteration k (vertical blue lines in Figure D2(b)), we recompute it for all i, j and not just $i, j \in \{\sigma^{k+1}, \dots, \sigma^K\}$ (which are the only values that we will use): Indeed, $\sigma^{k+1}, \dots, \sigma^K$ are not known at iteration k .

5.2 A subcubic algorithm for Whittle index

We now assume to have access to a subcubic matrix multiplication algorithm that satisfies the following property:

1. There exists an algorithm to multiply a matrix of size $n \times n$ by a matrix of size $n \times n^\alpha$ that runs in $O(n^{\omega(\alpha)})$, where $\omega : [0, 1] \rightarrow [2, 3]$ is a non-decreasing function.

Going back to Algorithm 2 where Line 20 is Subroutine 4, we now assume that we recompute the whole matrix \mathbf{X}^k every $O(n^\alpha)$ iterations. The new algorithm has a subcubic complexity:

Theorem 2. *Algorithm 2 with Subroutine 4 checks indexability and computes Whittle (and Gittins) index in time at least $\Omega(n^{2.5})$ and at most $O(n^{2.5286})$ when choosing $\alpha = 0.5286$.*

We believe that Theorem 2 is the first result that shows that Whittle index can be computed in subcubic time. As we show in Section 7, this algorithm can be directly extended to discounted index. As a byproduct, we also obtain the first subcubic algorithm to compute Gittins index.

Proof The algorithm starts by computing \mathbf{X}^1 which can be done in $O(n^{\omega(1)})$. Then, there are $n^{1-\alpha}$ times that we do:

1. We fill the "gray" mini matrices by using (16). This amounts to three for loops of size n (for i), n^α (for k) and n^α (for ℓ). Hence, each small gray matrix costs $O(n^{1+2\alpha})$.
2. At the end of a cycle, we recompute the full inverse by updating $(\mathbf{A}^{\pi^k})^{-1}$ from $(\mathbf{A}^{\pi^{k-n^\alpha}})^{-1}$. As we show in Lemma 5 (stated below), this can be done in $O(n^{\omega(\alpha)})$.

This implies that the algorithm has a complexity:

$$O(n^{\omega(1)}) + n^{1-\alpha} \left(O(n^{1+2\alpha}) + O(n^{\omega(\alpha)}) \right) = O(n^{\max\{2+\alpha, 1-\alpha+\omega(\alpha)\}}).$$

To compute the optimal value of α minimizing this expression requires the knowledge of the function $\omega(\alpha)$ which is not known. The current state of the art only gives a lower bound ($\omega(\alpha) \geq 2$) and an upper bound described in [39].

It is shown in [39] that $\alpha = 0.5286$ is the smallest currently known value of α for which $\omega(\alpha) < 1 + 2\alpha$. This implies that the complexity is at most $O(n^{2.5286})$.

As for the lower bound, $\omega(\alpha) \geq 2$ implies that the complexity of the algorithm is at least $\Omega(n^{2.5})$. \square

In the next lemma, \mathbf{B} plays the role of \mathbf{A}^{π^k} and \mathbf{A} the role of $\mathbf{A}^{\pi^{k-n^\alpha}}$. Note that as required in the lemma, exactly n^α rows and columns are changed between the two.

Lemma 5. *Assume 1. Let \mathbf{A} be a square matrix whose inverse \mathbf{A}^{-1} has already been computed, and let \mathbf{B} be an invertible square matrix such that $\mathbf{A} - \mathbf{B}$ is*

of rank smaller than n^α . Then, it is possible to compute the inverse of \mathbf{B} in $O(n^{\omega(\alpha)})$.

Proof The matrix \mathbf{B} can be written as $\mathbf{B} = \mathbf{A} + \mathbf{UCV}$ where \mathbf{U} is a $n \times n^\alpha$ matrix, \mathbf{C} is $n^\alpha \times n^\alpha$ and \mathbf{V} is $n^\alpha \times n$. The Sherman–Morrison–Woodbury formula [40] states that

$$\mathbf{B}^{-1} = (\mathbf{A} + \mathbf{UCV})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{U} \left(\mathbf{C}^{-1} + \mathbf{VA}^{-1}\mathbf{U} \right)^{-1} \mathbf{VA}^{-1}.$$

This shows that \mathbf{B}^{-1} can be computed by:

- Computing $\mathbf{D} := \mathbf{A}^{-1}\mathbf{U}$ and $\mathbf{E} := \mathbf{VA}^{-1}$: this takes $O(n^{\omega(\alpha)})$.
- Computing $\mathbf{F} := (\mathbf{C}^{-1} + \mathbf{VA}^{-1}\mathbf{U})^{-1}$: as this is the inversion of a $n^\alpha \times n^\alpha$ matrix, it can be done in $O(n^{\alpha\omega(1)})$ where $\alpha\omega(1) \leq \omega(\alpha)$.
- Computing $\mathbf{G} := \mathbf{DF}$ and then \mathbf{GE} : this again takes $O(n^{\omega(\alpha)})$.

Hence, computing \mathbf{B}^{-1} can be done in $O(n^{\omega(\alpha)})$ operations for the inversion and all multiplications plus an additional $O(n^2)$ term for the subtraction and the addition. As $\omega(\alpha) \geq 2$, this concludes the proof of the lemma. \square

5.3 The subcubic algorithm in practice

The complexity of $O(n^{2.5286})$ given in Theorem 2 is mainly of theoretical interest. The value $\alpha = 0.5286$ is obtained by using the best upper bound on $\omega(\alpha)$ known today which is based on the Coppersmith–Winograd algorithm and its variants. The Coppersmith–Winograd algorithm (or its variants) are, however, known as a *galactic algorithm*: the hidden constant in the $O()$ is so large that their runtime is prohibitive for any reasonable value of n . Hence, the existence of these algorithms is of theoretical interest but has limited applicability.

This does not discard the practical improvement provided by Subroutine 4 which is based on the mere fact that multiplying two matrices (or inverting a matrix) is faster than three nested loops even for matrices of moderate size. To verify this, we launched a detailed profiling of the code of Algorithm 2 with the non-optimized Subroutine 3. It shows that for a problem of dimensions 5000, the update of Line 20 takes more than 90% of the computation time, the initialization of \mathbf{X}^1 on Line 2 takes about 5% of the time and the rest of the code takes less than 1% of the running time.

Now, if inverting the full matrix takes about 5% of the execution time, and updating the gray zone takes 95%, then by doing 5 updates, one can hope to obtain an algorithm whose running time is roughly $5 \times 5 + 95/5 \approx 43\%$ the one of the original implementation. As we observe in Section 6, this is close to the gain that we obtain in practice. A general way to choose the best number of updates is used in the numerical section. It is based on the following reasoning. For large matrices (say $n \geq 10^3$), the fastest implementations of matrix multiplication and inversion are based on Strassen’s algorithm [41, 42]. As we report in Appendix C, the time to solve a linear system of size n by using the default installation of `scipy` seems to run in $O(n^{2.8})$. By replacing

the function $\omega(\alpha)$ used in Theorem 2 by a more practical bound ($\omega(\alpha) = 2.8$), the best value for α becomes $\alpha = 0.9$. This indicates that our algorithm can be implemented in $O(n^{2.9})$ by doing $O(n^{0.1})$ recomputation of \mathbf{X}^k from \mathbf{A}^{π^k} . Note that even for very large values of n (like $n = 15000$), $n^{0.1}$ remains quite small, *e.g.*, $15000^{0.1} \approx 2.6$. In practice, we observe that updating $\text{int}(2n^{0.1})$ times (the notation $\text{int}(x)$ indicates that it is rounded to the closest integer) gives the best performance among all algorithms, as reported in the next section.

6 Numerical Experiments

In complement to our theoretical analysis, we developed a python package that implements Algorithm 2 and gives the choice of using the variant of Subroutine 3 or of Subroutine 4 to do the "update_X()" function. This package relies on three python libraries: `scipy` and `numpy` for matrix operations, and `numba` to compile the python code. To facilitate its use, this package can be installed by using `pip install markovianbandit-pkg`.

All experiments were run on a laptop (Macbook Pro 2020) with an Intel Core i9 CPU at 2.3 GHz with 16GB of Memory using Python 3.6.9 :: Anaconda custom (64-bit) under macOS Big Sur version 11.6.2. The version of the packages are `scipy` version 1.5.4, `numpy` version 1.19.5 and `numba` version 0.53.1. The code of all experiments is available at <https://gitlab.inria.fr/markovianbandit/efficient-whittle-index-computation>.

6.1 Time to compute Whittle indices

To test the implementation of our algorithm, we randomly generate restless bandit instances with n states where $n \in \{100, 1000, \dots, 15000\}$. In each case, both transition matrices are uniform probabilistic matrices: for each row of each matrix, we generate n i.i.d. entries following the exponential distribution and divide the row by its sum. This means that all matrices are dense. We use dense matrix since it is the worst case for computational interest. By running our algorithms on sparse matrix, we would expect to have faster running time. Note that all tested matrices are indexable. This is coherent with [31, 43] that report that for uniform matrices, the probability of finding a non-indexable example decreases very rapidly with the dimension n . Finally, reward vectors were generated from random $\text{Uniform}[0,1]$ entries.

We record the runtime of the different variants of our algorithm and report the results in Table 1. Note that these results present the whole execution time of the algorithm, including the initialization phase in which \mathbf{X}^1 is computed. For each value of n , we run Algorithm 2 with four variants:

- The first two columns correspond to the $O(n^3)$ algorithm (that uses Subroutine 3 for Line 20), either (a) with the indexability test, or (b) without the indexability test.

Table 1: Running time (in seconds) of the variants of Algorithm 2

n	$O(n^3)$ algorithm (Subroutine 3)		Subcubic algorithm (Subroutine 4)	
	(a) With index. test	(b) Without test	(c) With index. test	(d) Without test
100	0.006	0.004	0.007	0.005
1000	0.2	0.2	0.2	0.2
2000	1.9	1.2	1.1	1.0
3000	7.2	5.0	3.2	2.6
4000	17	12	8	6
5000	34	25	16	12
6000	60	43	27	20
7000	95	68	42	33
8000	142	99	63	49
9000	199	141	92	70
10000	275	191	122	95
11000	361	257	164	133
12000	471	339	225	190
13000	620	428	286	243
14000	790	553	408	317
15000	965	685	501	403


- The next two columns correspond to the subcubic algorithm (that uses Subroutine 4 for Line 20 with $\text{int}(2n^{0.1})$ updates), either (c) with the indexability test, or (d) without the indexability test.

Our numbers show that our algorithm can compute the Whittle index in less than one second for $n = 1000$ states and slightly less than 7 minutes for $n = 15000$ states with variant (d). As expected, not doing the indexability test does improve the performance compared to doing the indexability test (here by a factor approximately $1/3$ for Subroutine 3 and $1/5$ for Subroutine 4). More importantly, this table shows that the time when using the subcubic variant, Subroutine 4, diminishes the computation time by about 40% to 50% compared to when using Subroutine 3. Note that for $n = 100$, using the Subroutine 3 is slightly faster than using the Subroutine 4 (note that both takes only a few milliseconds). This indicates that the subcubic algorithm becomes interesting when n is large enough (say $n \geq 2000$).

To give a visual idea of how the various variants of the algorithms compare, we plot in Figure 5a the runtime of the four variants along with the runtime reported [28], that we believe was the fastest algorithm to compute Whittle index so far. We plot in Figure 5b the runtime of each variant divided by the runtime of the algorithm of [28]. For large n , our best implementation is about 10 to 20 times faster than the one of [28]. For instance, for $n = 15000$, our implementation takes about 7 minutes to compute the index (or 9 minutes when checking indexability) whereas [28] reports 2 hours. Recall that the algorithm of [28] cannot check indexability on the fly. In our implementation, not testing indexability reduces the computation time of 15 – 20% for Subroutine 4 or 25 – 30% for Subroutine 3 compared to the version that tests indexability.

It should be clear that the comparison between our implementation and the numbers reported in [28] has its limits. First, there is no public implementation of the algorithm of [28] and we did not re-implement it. Hence,

the computation was not done on the same machine, and not implemented in the same programming language. The characteristic of the machine described in [28] seems comparable if not better to ours, and our implementation is in python/numpy whereas the one of [28] is in matlab. Second, it is written in [28] that the reported numbers do not “count the initialization stage of computing the initial tableau”. There is no mention of the cost of the initialization stage in [28]. On the contrary, the numbers that we report in this paper include all initialization stages, which puts our running time at disadvantage.



(a) Runtime of our implementations as a function of the state size n .

(b) Speedup of each implementation over the data from [28] as a function of the state size n .

Fig. 5: Numerical result over 7 simulations taking around 21 hours in total: in each simulation, we run the algorithm over randomly generated RBs with the state size ranging over $\{1000, \dots, 15000\}$. We plot the average runtime over 7 simulations. The solid lines represent the result of Subroutine 3 and the dashed-dot lines represent the one of Subroutine 4. The marker “+” indicates that algorithms test indexability and the triangles indicates that algorithms do not test indexability.

6.2 Statistics of indexable problems

To the best of our knowledge, our algorithm provides the first indexability test that scales well with the dimension n . We used this to answer a very natural question: given a randomly generated arm, how likely is it to be indexable? This question was partially answered in [31] that shows that when generating dense arms, the probability of generating a non-indexable arm is close to 10^{-n} for $n \in \{3 \dots 7\}$. This suggests that most arms are indexable. Below, we answer two questions: what happens for larger values of n , and more importantly, what happens when the arms are not dense?

To answer these questions, we consider randomly generated arms where the matrices \mathbf{P}^0 and \mathbf{P}^1 are b -diagonal matrices with b non-null diagonals. In particular, $b = 3$ corresponds to tridiagonal matrices, $b = 5$ corresponds to pentadiagonal matrices and $b = 7$ corresponds to septadiagonal matrices. We also compare with the classical case of dense matrices (which corresponds to

$b = 2n - 1$). For each model, the entries are generated from the exponential distribution for each row and we divide the row by the sum of generated entries for this row. We vary n from 3 to 50 and for each case, we generate 100000 arms. We report in Table 2 the number of indexable arms for each. Note that pentadiagonal matrices are dense matrices for $n = 3$ and septadiagonal matrices are dense matrices for $n = 4$ and do not make sense for $n = 3$, which is why no numbers are reported.

Table 2: Number of indexable problems among 100 000 randomly generated problems.

n	Tridiagonal	5-diagonal	7-diagonal	Dense
3	98 731	–	–	99 883
4	95 067	99 655	–	99 931
5	89 198	99 309	99 902	99 969
10	54 129	90 377	98 914	100 000
30	7 094	29 699	66 143	100 000
50	1 823	9 332	32 069	100 000

Based on these results, we can assert that dense models are essentially always indexable which conforms with the data reported in [31]. The situation is, however, radically different for sparse models: the number of indexable problems decreases quickly with the number of states. For instance, there are only 1 823 indexable 50-state problems among 100 000 generated tridiagonal models (*i.e.* around 1.8% are indexable). Note that a tridiagonal model is a birth-death Markov chain which is frequently used for queueing systems. Hence, it is very important to check the indexability of the problem because it is not a prevalent property for sparse models. This also calls for new efficient policies in multi-arm bandits problems that are not based on Whittle indices.

7 Extension to the Discounted Case

The model described in Section 2 corresponds to the definition of Whittle index for a time-average criterion, for which Whittle index is known to be asymptotically optimal [5] for restless bandits. Yet, Whittle index can also be defined for the discounted case [28, 32]. Notably, the discounted Whittle index simplifies into Gittins index when the bandit is rested (*i.e.*, when $\mathbf{P}^0 = \mathbf{I}$ and $\mathbf{r}^0 = \mathbf{0}$.) In this section, we show how to adapt our algorithm to the discounted case. As a by product, we obtain the first subcubic algorithm to compute Gittins index in the discounted case.

7.1 Discounted Whittle index

We now consider a λ -penalized MDP in which the instantaneous reward received at time $t \geq 0$ is discounted by a factor β^t , where $\beta \in (0, 1)$ is called the discount factor: when executing action a in state i at time $t \geq 0$, the decision maker earns a reward $\beta^t(r_i^a - \lambda a)$. For a given policy π , we denote by $u_i^\pi(\lambda)$

the expected discounted reward earned by the decision maker when the MDP starts in state i at time 0. The vector $\mathbf{u}^\pi(\lambda) = [u_1^\pi(\lambda), \dots, u_n^\pi(\lambda)]^T$ is called the value function of the policy π . From [38], it satisfies Bellman's equation, that is, for all state i we have:

$$u_i^\pi(\lambda) = r_i^{\pi_i} - \lambda\pi_i + \beta \sum_{j=1}^n P_{ij}^{\pi_i} u_j^\pi(\lambda). \quad (17)$$

The above equation is a linear equation, whose solution is unique because $\beta < 1$. It is given by:

$$\mathbf{u}^\pi(\lambda) = (\mathbf{I} - \beta\mathbf{P}^\pi)^{-1}(\mathbf{r}^\pi - \lambda\boldsymbol{\pi}). \quad (18)$$

For a given penalty λ and a state i , we denote by $u_i^*(\lambda) = \max_\pi u_i^\pi(\lambda)$ be the optimal value of state i . A policy π is optimal for the penalty λ if for all state $i \in [n]$, $u_i^*(\lambda) = u_i^\pi(\lambda)$. By [38] such a policy exists. Similarly to the time-average criterion studied before, a β -discounted RB is called indexable if for all state $i \in [n]$, there exists a critical penalty λ_i such that the action "activate" is optimal in state i for penalty λ if and only if $\lambda \leq \lambda_i$.

7.2 Analogy between the time-average and the discounted versions

Let π be a policy and $i \in \pi$ be an active state. By Equation (17), the two actions "rest" and "activate" are equivalent in state i if and only if

$$r_i^1 - \lambda + \beta \sum_{j=1}^n P_{ij}^1 u_j^*(\lambda) = r_i^0 + \beta \sum_{j=1}^n P_{ij}^0 u_j^*(\lambda),$$

which is equivalent to

$$\lambda = \delta_i + \tilde{\Delta}_i \mathbf{u}^\pi(\lambda), \quad (19)$$

where δ_i is defined as for the time-average case and $\tilde{\Delta}$ is such that for all⁵ states $i, j \in [n]$: $\tilde{\Delta}_{ij} := \beta(P_{ij}^1 - P_{ij}^0)$. This Equation (19) is the same as its time-average analogue (6), but replacing Δ by $\tilde{\Delta}$.

To finish the derivation of the algorithm, one should note that the value function $\mathbf{u}(\lambda)$ plays the same role as the vector $\mathbf{v}(\lambda)$ defined for the time-average case. In particular, the definition of \mathbf{u} in Equation (18) is the analogue of the definition of \mathbf{v} in (5) up to the replacement of the matrix \mathbf{A}^π in (5) by the matrix $\mathbf{I} - \beta\mathbf{P}^\pi$. This means that similarly to \mathbf{v} , the value function $\mathbf{u}(\lambda)$ is affine in λ .

⁵Note that the definition of $\tilde{\Delta}$ is identical to Δ except when $j = 1$, for which $\Delta_{i1} := 0$ but $\tilde{\Delta}_{i1} = \beta(P_{i1}^1 - P_{i1}^0)$.

Hence, following the same development in Section 4, we can modify Algorithm 2 to compute the discounted Whittle index by modifying only the initialization phase:

$$\tilde{\Delta} = \beta(\mathbf{P}^1 - \mathbf{P}^0) \text{ and } \mathbf{X}^1 = \tilde{\Delta}(\mathbf{I} - \beta\mathbf{P}^{\pi^1})^{-1}.$$

Note that we still have $\mathbf{y}^1 = \mathbf{0}$ because $(\mathbf{I} - \beta\mathbf{P}^{\pi^1})^{-1}\boldsymbol{\pi}^1 = \frac{1}{1-\beta}\mathbf{1}$ (value function in a β -discounted Markov chain with reward equals to 1 in all states) and $\tilde{\Delta}\mathbf{1} = \mathbf{0}$. Also, if Subroutine 4 is used, Line 3 should be changed to $\mathbf{X}^{k_0} = \tilde{\Delta}(\mathbf{I} - \beta\mathbf{P}^{\pi^{k_0}})^{-1}$. Last but not least, in the discounted case, the unichain property is no longer needed because the matrix $(\mathbf{I} - \beta\mathbf{P}^{\pi})$ is invertible for any policy π as long as $\beta < 1$ (from Perron-Frobenius' Theorem).

7.3 Gittins Index

The notion of “restless” bandit comes from the fact that even when the action “rest” is taken, the Markov chain can still change state and generate rewards. When this is not the case (*i.e.*, when $\mathbf{P}^0 = \mathbf{I}$ and $\mathbf{r}^0 = \mathbf{0}$), an arm is no longer restless and is simply called a Markovian bandit (or a *rested* Markovian bandit if one wants to emphasize that it is not restless).

In a discounted rested bandit, the notion of Whittle index coincides with the notion of Gittins index (In fact, Whittle index was first introduced as a generalization of Gittins index to restless bandit in [3]). In such a case, there is no notion of indexability: a discounted rested bandit is always indexable. Its index can be computed by Algorithm 2 without testing indexability. The best known algorithm to compute Gittins index runs in $(2/3)n^3 + O(n^2)$ [27]. When using fast multiplication, our algorithm computes Gittins index in $O(n^{2.5286})$ which makes it the first algorithm to compute Gittins index in subcubic time.

Note that when $\mathbf{P}^1 = \mathbf{I}$, it is possible to compute $\mathbf{X}^1 = \tilde{\Delta}/(1-\beta)$ without having to solve a linear system which means that, when $\mathbf{P}^1 = \mathbf{I}$, our algorithm with the variant Subroutine 3 has complexity $(2/3)n^3 + O(n^2)$. This shows that our cubic algorithm can compute the Gittins index in $(2/3)n^3 + O(n^2)$ instead of $(2/3)n^3 + o(n^3)$.

8 Conclusion

In this paper, we present an algorithm that is efficient for detecting the non-indexability and computing Whittle index of all indexable finite-state restless bandits. Our algorithm is based on the efficient application of Sherman-Morrison formula. This is a unified algorithm that works for both discounted and non-discounted restless bandits, and can be used for Gittins index computation. We present a first version of our algorithm that runs in $n^3 + o(n^3)$ arithmetic operations (or in $(2/3)n^3 + o(n^3)$ if we do not test the indexability). So, we conclude that Whittle index is not harder to compute than Gittins index. The second version of our algorithm uses the fastest matrix

multiplication method and has a complexity of $O(n^{2.5286})$. This makes it the first subcubic algorithm to compute Whittle index or Gittins index. We provide numerical simulations that show that our algorithm is very efficient in practice: it can test indexability and compute the index of a restless bandit process in less than one second for $n = 1000$ states, and in a few minutes for $n = 15000$. These numbers are provided for dense matrices. One might expect to have more efficient algorithms if the arm has a sparse structure. We leave this question for future work.

Acknowledgments. This work is supported by the ANR project REFINO (ANR-19-CE23-0015).

Appendix A Technical Lemmas

A.1 Minimal optimal policy

Lemma 6. *Let λ be a fixed penalty and consider that there exist two distinct optimal policies $\pi \neq \pi'$. Then:*

1. *The policy $\pi \cap \pi'$ is also optimal for this penalty.*
2. *If $\pi \supset \pi'$, then for $i \in \pi \setminus \pi'$, $\pi \setminus \{i\}$ is also optimal for this penalty.*

Proof By Section 8.4 of [38], and since the arm is unichain, there exists a bias vector $h^*(\lambda)$ such that π is an optimal policy if and only if for all state i :

$$r_i^1 - \lambda + \sum_{j=2}^n P_{ij}^1 h_j^*(\lambda) > r_i^0 + \sum_{j=2}^n P_{ij}^0 h_j^*(\lambda) \quad \text{implies } i \in \pi; \quad (\text{A1})$$

$$r_i^1 - \lambda + \sum_{j=2}^n P_{ij}^1 h_j^*(\lambda) < r_i^0 + \sum_{j=2}^n P_{ij}^0 h_j^*(\lambda) \quad \text{implies } i \notin \pi. \quad (\text{A2})$$

It should be clear that if both π and π' satisfy (A1)–(A2), then so does $\pi \cap \pi'$ which implies that it is also an optimal policy.

Note that in particular, if $\pi \supseteq \pi'$ are two optimal policies, then let $i \in \pi \setminus \pi'$. This state must be such that $r_i^1 - \lambda + \sum_{j=2}^n P_{ij}^1 h_j^*(\lambda) = r_i^0 + \sum_{j=2}^n P_{ij}^0 h_j^*(\lambda)$ which implies that $\pi \setminus \{i\}$ satisfies (A1)–(A2) and is therefore also an optimal policy. \square

A.2 Unichain property

Lemma 7. *Given a MDP $\langle [n], \{0, 1\}, \mathbf{r}, \mathbf{P} \rangle$, let \mathbf{P}^π be the transition matrix under policy π . If the MDP is unichain, the matrix*

$$\mathbf{A}^\pi = \begin{bmatrix} 1 & -P_{12}^\pi & \dots & -P_{1n}^\pi \\ 1 & 1 - P_{22}^\pi & \dots & -P_{2n}^\pi \\ \vdots & & & \\ 1 & -P_{n2}^\pi & \dots & 1 - P_{nn}^\pi \end{bmatrix}$$

is invertible for any policy π .

Proof \mathbf{A}^π is not invertible, if there exists a vector $\mathbf{u} \neq \mathbf{0}$ such that $\mathbf{u}^T \mathbf{A}^\pi = \mathbf{0}$. We prove that such \mathbf{u} does not exist when the MDP is unichain. Let $\mathbf{u} \in \mathbb{R}^n$ be an arbitrary vector such that $\mathbf{u}^T \mathbf{A}^\pi = \mathbf{0}$. Then, we have

$$\begin{cases} \sum_{i=1}^n u_i & = 0 \\ u_i - \sum_{j=1}^n u_j P_{ji}^\pi & = 0, \text{ for } 2 \leq i \leq n \end{cases}$$

Then,

$$\begin{aligned} \sum_{i=1}^n u_i &= u_1 + \sum_{i=2}^n u_i \\ 0 &= u_1 + \sum_{i=2}^n \sum_{j=1}^n u_j P_{ji}^\pi = u_1 + \sum_{j=1}^n u_j \sum_{i=2}^n P_{ji} = u_1 + \sum_{j=1}^n u_j (1 - P_{j1}^\pi) \\ 0 &= u_1 - \sum_{j=1}^n u_j P_{j1}^\pi \end{aligned}$$

Then, we get

$$\begin{cases} \sum_{i=1}^n u_i & = 0 \\ \mathbf{u}^T \mathbf{P}^\pi & = \mathbf{u}^T \end{cases}$$

In unichain MDP, for any policy π , the only vectors \mathbf{u} verifying $\mathbf{u}^T \mathbf{P}^\pi = \mathbf{u}^T$ are multiples of the stationary distribution under policy π [38]. Therefore $\sum_{i=1}^n u_i = 0$ implies $\mathbf{u} = \mathbf{0}$. \square

A.3 Proof of Lemma 3

Proof Let $\mathbf{v}^\pi(\lambda)$ be as in Equation (5). By definition, the matrices \mathbf{A}^π and $\mathbf{A}^{\pi \setminus \{i\}}$ differ only in row i , that is, for $i \in \pi$

$$\begin{aligned} (\mathbf{A}^\pi)_i &= [1 \quad -P_{i2}^1 \quad \dots \quad -P_{ii-1}^1 \quad 1 - P_{ii}^1 \quad -P_{ii+1}^1 \quad \dots \quad -P_{in}^1] \\ \text{and } (\mathbf{A}^{\pi \setminus \{i\}})_i &= [1 \quad -P_{i2}^0 \quad \dots \quad -P_{ii-1}^0 \quad 1 - P_{ii}^0 \quad -P_{ii+1}^0 \quad \dots \quad -P_{in}^0]. \end{aligned}$$

This implies that (3) with $\pi \setminus \{i\}$ holds for all $j \neq i$: $(\mathbf{A}^\pi \mathbf{v}^\pi(\lambda))_j = (\mathbf{A}^{\pi \setminus \{i\}} \mathbf{v}^\pi(\lambda))_j = r_j^{\pi_j} - \lambda \pi_j$. Moreover, if λ is such that (6) holds, then

$$r_i^1 - \lambda + \sum_{j=2}^n P_{ij}^1 v_j^\pi(\lambda) = r_i^0 + \sum_{j=2}^n P_{ij}^0 v_j^\pi(\lambda).$$

So, for any $i \in \pi$,

$$\begin{aligned} g^\pi(\lambda) + v_i^\pi(\lambda) &= r_i^1 - \lambda + \sum_{j=2}^n P_{ij}^1 v_j^\pi(\lambda) = r_i^0 + \sum_{j=2}^n P_{ij}^0 v_j^\pi(\lambda) \\ g^\pi(\lambda) + v_i^\pi(\lambda) - \sum_{j=2}^n P_{ij}^0 v_j^\pi(\lambda) &= r_i^0 \text{ or } (\mathbf{A}^{\pi \setminus \{i\}} \mathbf{v}^\pi(\lambda))_i = r_i^0. \end{aligned}$$

This implies that $\mathbf{v}^\pi(\lambda)$ is a solution of $\mathbf{A}^{\pi \setminus \{i\}} \mathbf{v}^\pi(\lambda) = \mathbf{r}^{\pi \setminus \{i\}} - \lambda(\boldsymbol{\pi} - \mathbf{e}_i)$, which implies that $\mathbf{v}^\pi(\lambda) = \mathbf{v}^{\pi \setminus \{i\}}(\lambda)$. Note that \mathbf{e}_i is the column vector with a 1 in i th coordinate and 0's elsewhere. \square

Appendix B Implementations

B.1 Arithmetic complexity of Subroutine 3 and memory usage

Recall that in Subroutine 3, we compute the values X_{ij}^ℓ by doing the update (for all iteration k , for all $\ell = 1$ to k and for all $i \in [n]$ or all $i \in \pi^{\ell+1}$ if we do not test indexability):

$$X_{i\sigma^k}^{\ell+1} := X_{i\sigma^k}^\ell - \frac{X_{i\sigma^\ell}^\ell}{1 + X_{\sigma^\ell\sigma^k}^\ell} X_{\sigma^\ell\sigma^k}^\ell \quad (\text{B3})$$

If we test indexability, there are $\sum_{k=1}^n kn = n^3/2 + O(n^2)$ such updates. If we do not test indexability, there are $\sum_{k=1}^n \sum_{\ell=1}^k (n - \ell) = n^3/3 + O(n^2)$ such updates. Below, we show each update of Equation (B3) can be done in two arithmetic operations (one addition and one multiplication), which leads to the complexity of $n^3 + O(n^2)$ (or $(2/3)n^3 + O(n^2)$) arithmetic operations for the computation of all the needed X_{ij}^k . We also show how to reduce the memory size to $O(n^2)$.

Let $W_{ik} = X_{i\sigma^k}^k / (1 + X_{\sigma^k\sigma^k}^k)$ and $V_{i\ell}^k = X_{i\sigma^\ell}^\ell$. Using this, Equation (B3) can be rewritten as:

$$V_{i\ell+1}^k = V_{i\ell}^k - W_{i\ell}^k V_{\sigma^\ell\ell}^k, \quad (\text{B4})$$

with $W_{ik} = V_{ik}^k / (1 + V_{\sigma^k k}^k)$.

This results in the following loop at iteration k :

- Initialize V_{i1}^k from $X_{i\sigma^k}^1$.
- For all $\ell \in \{1 \dots k-1\}$, and all $i \in [n]$ (or $i \in \pi^{\ell+1}$), apply (B4).
- Compute $W_{ik} = V_{ik}^k / (1 + V_{\sigma^k k}^k)$
- For $\ell = k$, and all $i \in [n]$ (or $i \in \pi^{\ell+1}$), apply (B4).

Note that the value of V^{k-1} is not necessary for iteration k (only the values of $W_{i\ell}$ are needed). This shows that the algorithm can be implemented with a memory $O(n^2)$. Moreover, each update of (B3) can be implemented by 2 arithmetic operations (one multiplication and one subtraction).

B.2 Speedup when not checking the indexability: First found go last

When the indexability is not tested, the update (B4) is computed for all $i \in \pi^{\ell+1}$. This creates inefficiencies (due to inefficient cache usage) because the elements $V_{i\ell+1}^k$ are not accessed sequentially.

To speedup the memory accesses, our solution is to sort the items during the execution of the algorithm. At iteration k , the algorithm computes σ^k . When this is done, our implementation switches all quantities in positions σ^k and $n - k + 1$. These quantities are δ, y, z, \mathbf{W} and \mathbf{X} . For instance, once σ^1 is

found, we know that the state at position n is state n and we do the following switches:

$$\begin{aligned} \delta_{\sigma^1}, \delta_n &\rightarrow \delta_n, \delta_{\sigma^1} \\ y_{\sigma^1}^1, y_n^1 &\rightarrow y_n^1, y_{\sigma^1}^1 \\ z_{\sigma^1}^1, z_n^1 &\rightarrow z_n^1, z_{\sigma^1}^1 \\ \mathbf{W}_{\sigma^1:}, \mathbf{W}_n &\rightarrow \mathbf{W}_n, \mathbf{W}_{\sigma^1:} \\ \text{and } \mathbf{X}_{\sigma^1:}, \mathbf{X}_n &\rightarrow \mathbf{X}_n, \mathbf{X}_{\sigma^1:}. \end{aligned}$$

To do so, we need an array to store all states such that at iteration k , the first $n - k$ states of the array are the active states. We will need to track the position of each state in such array.

Appendix C Analysis of the experimental time to solve a linear system

In this section, we report in Figure C1 the time taken by the default implementation to solve a linear system of the form $\mathbf{AX} = \mathbf{B}$ where \mathbf{A} and \mathbf{B} are two square matrices. To obtain this figure, we generated random (full) matrices where each entry is between 0 and 1 and use the function `scipy.linalg.solve` from the library `scipy`. The reported numbers suggest that the complexity of the solver is closer to $O(n^{2.8})$ than to $O(n^3)$, although we agree that the difference between the $O(n^{2.8})$ and the $O(n^3)$ curves is small. Note that this is in accordance with the papers [41, 42] that claim that the fastest implementations of matrix multiplication and inversion are based on Strassen's algorithm and should therefore be in $O(n^{2.8})$.




Fig. C1: Time taken of the default implementations `scipy.linalg.solve` of `scipy` to solve a linear system $\mathbf{AX} = \mathbf{B}$ where \mathbf{A} and \mathbf{B} are two square $n \times n$ matrices.

Appendix D Detailed comparison with [32] and [28]

In this section, we compare our algorithm with two main related works for finite-state restless bandits problem.

D.1 Comparison with [32]

The paper presents an algorithm that computes Whittle indices in $O(n^3)$ (no explicit constant before n^3 is given) for all indexable problems. Despite following a different approach, our algorithm for computing Whittle index can be viewed as a refinement of this work. Let us recall once again that our approach also allows one to check the indexability of general restless bandits.

In the following, we show how we can refine the work of [32] to obtain an algorithm that is exactly the same as ours. Let D^π and N^π be two vectors defined as in [32],

$$D^\pi = (1 - \beta)(\mathbf{I} - \beta\mathbf{P}^\pi)^{-1}\mathbf{r}^\pi, \quad \text{and} \quad N^\pi = (1 - \beta)(\mathbf{I} - \beta\mathbf{P}^\pi)^{-1}\boldsymbol{\pi}.$$

Then, we have $D^\pi - \lambda N^\pi = (1 - \beta)\mathbf{u}^\pi(\lambda)$ where $\mathbf{u}^\pi(\lambda)$ is defined as in (18). In our proposition, at each iteration k , we compute μ_i^k by Line 6. Instead, it is defined in [32] by two steps:

1. for all state $j \in [n]$ such that $N_j^{\pi^k \setminus \{i\}} \neq N_j^{\pi^k}$, compute $\mu_{ij}^k = \frac{D_j^{\pi^k \setminus \{i\}} - D_j^{\pi^k}}{N_j^{\pi^k \setminus \{i\}} - N_j^{\pi^k}}$
2. compute $\mu_i^k = \arg \min_{j \in [n]: N_j^{\pi^k \setminus \{i\}} \neq N_j^{\pi^k}} \mu_{ij}^k$.

From [32, Theorem 2], in an indexable problem, for state σ^k , there exists a state $j \in [n]$ such that $N_j^{\pi^k \setminus \{\sigma^k\}} \neq N_j^{\pi^k}$. Now, suppose that for any active state $i \in \pi^k$, there exists $j \in [n]$ such that $N_j^{\pi^k \setminus \{i\}} \neq N_j^{\pi^k}$. Using Sherman-Morrison formula, we have⁶

$$D^{\pi^k \setminus \{i\}} - D^{\pi^k} = -\frac{(1 - \beta)\delta_i + \tilde{\Delta}_i D^{\pi^k}}{1 + \tilde{\Delta}_i [(\mathbf{I} - \beta\mathbf{P}^{\pi^k})^{-1}]_{:i}} [(\mathbf{I} - \beta\mathbf{P}^{\pi^k})^{-1}]_{:i},$$

$$\text{and} \quad N^{\pi^k \setminus \{i\}} - N^{\pi^k} = -\frac{(1 - \beta) + \tilde{\Delta}_i N^{\pi^k}}{1 + \tilde{\Delta}_i [(\mathbf{I} - \beta\mathbf{P}^{\pi^k})^{-1}]_{:i}} [(\mathbf{I} - \beta\mathbf{P}^{\pi^k})^{-1}]_{:i}.$$

Then, for any $j \in [n]$ such that $N_j^{\pi^k \setminus \{i\}} \neq N_j^{\pi^k}$, $\mu_{ij}^k = \frac{(1 - \beta)\delta_i + \tilde{\Delta}_i D^{\pi^k}}{(1 - \beta) + \tilde{\Delta}_i N^{\pi^k}}$ which does not depend on j . Then, we simply have $\mu_i^k = \frac{(1 - \beta)\delta_i + \tilde{\Delta}_i D^{\pi^k}}{(1 - \beta) + \tilde{\Delta}_i N^{\pi^k}}$.

⁶the expression of $D^{\pi^k \setminus \{i\}}$ and $N^{\pi^k \setminus \{i\}}$ given by Equation (18) in [32] are erroneous.

Also, we have

$$\begin{aligned}\tilde{\Delta}_i N^{\pi^k} &= (1 - \beta)\tilde{\Delta}_i(\mathbf{I} - \beta\mathbf{P}^{\pi^k})^{-1}\boldsymbol{\pi}^k = -(1 - \beta)y_i^k \quad \text{and} \\ \tilde{\Delta}_i D^{\pi^k} &= (1 - \beta)\tilde{\Delta}_i \mathbf{u}^{\pi^k} (\mu_{\min}^{k-1}) + \mu_{\min}^{k-1} \tilde{\Delta}_i N^{\pi^k} = (1 - \beta)z_i^{k-1} - (1 - \beta)\mu_{\min}^{k-1} y_i^k.\end{aligned}$$

So, replacing these terms in μ_i^k , we get the formula given at Line 6.

Note that the algorithm of [32] was only developed for the discounted case. Our approach for the time-average reward case is different because we use the average gain and bias instead of directly the expected discounted total reward D^{π^k} and total number of activations N^{π^k} under policy π^k . Moreover, our code is also optimized to avoid unnecessary computation and to reduce memory usage. Finally, the way we do the update of our matrix \mathbf{X} makes it possible to obtain a subcubic algorithm whereas their approach does not (see also below).

D.2 Comparison with algorithm of [28]

The algorithm [28] had the best complexity up to date, there is a square matrix \mathbf{A} that plays a similar role as the square matrix \mathbf{X} in our proposed algorithm. The most costly operations in the algorithm of [28] is to update their matrix \mathbf{A} at each iteration and it is done by Equation (15) that we recall here (using the same notation \mathbf{A} as in their paper):

$$\text{for } i, j \in \pi^k, \mathbf{A}_{ij}^{k+1} = \mathbf{A}_{ij}^k - \frac{\mathbf{A}_{i\sigma^k}^k}{\mathbf{A}_{\sigma^k\sigma^k}^k} \mathbf{A}_{\sigma^k j}^k. \quad (\text{D5})$$

This incurs a total complexity of $(2/3)n^3 + O(n^2)$ arithmetic operations. As mentioned in Section 4.3, if we updated \mathbf{X}^{k+1} as given by (15), our algorithm would also have a $(2/3)n^3 + O(n^2)$ complexity but this version of the update cannot be optimized by using fast matrix multiplication.

D.3 The approach of [28, 32] cannot be directly transformed into a subcubic algorithm

In addition to all the previously cited differences, one of the major contribution of our algorithm with respect to [28, 32] is that the most advanced version of our algorithm runs in a subcubic time. The approach⁷ used in [28, 32] is to update the full matrix \mathbf{X}^{k+1} at iteration k , by using (D5). This idea is represented in Figure D2(a): for a given ℓ , their algorithm compute $\mathbf{X}_{\cdot\sigma^k}^\ell$ for all k (i.e., the full vertical lines represented by arrows). Our first Subroutine 3 uses an horizontal approach based on (16), which we recall here:

$$X_{i\sigma^k}^{\ell+1} = X_{i\sigma^k}^\ell - \frac{X_{i\sigma^\ell}^\ell X_{\sigma^\ell\sigma^k}^\ell}{1 + X_{\sigma^\ell\sigma^\ell}^\ell}.$$

⁷Equation (18) of [32], which is central to their algorithm is the same as the above equation (D5).

At iteration k , we use $\mathbf{X}_{:\sigma^k}^1$ to compute all values of $\mathbf{X}_{:\sigma^k}^\ell$ up to $\ell = k + 1$. This is represented in Figure D2(b). Our approach can be used to obtain the subcubic algorithm illustrated in Figure D2(c) by using subcubic algorithms for multiplication.




Fig. D2: Comparison of the implementation of (D5) used in [28, 44] with Subroutine 3 and Subroutine 4.

This leads to the next fundamental question: why should the computation of Whittle index be harder than matrix inversion (or multiplication)? To us, the main difference is that when computing Whittle indices, the permutation σ is not known a priori but discovered as the algorithm progresses: σ^k is only known at iteration k . Hence, while all terms of the matrices X_{ij}^k are not needed, it is difficult to know a priori which ones are needed and which ones are not. Hence, a simple divide and conquer algorithm cannot be used. This is why when recomputing \mathbf{X}^{k+1} in Subroutine 4, we recompute the whole matrix (the vertical blue line) and not just the part that will be used to compute the gray zone: we do not know *a priori* what part of X_{ij}^{k+1} will be useful or not.

References

- [1] Gittins, J.C.: Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society: Series B (Methodological)* **41**(2), 148–164 (1979)
- [2] Papadimitriou, C.H., Tsitsiklis, J.N.: The complexity of optimal queueing network control. In: *Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory*, pp. 318–322 (1994). IEEE
- [3] Whittle, P.: Restless bandits: Activity allocation in a changing world. *Journal of applied probability* **25**(A), 287–298 (1988)
- [4] Lott, C., Teneketzis, D.: On the optimality of an index rule in multichannel allocation for single-hop mobile networks with multiple service classes.

- Probability in the Engineering and Informational Sciences **14**(3), 259–297 (2000)
- [5] Weber, R.R., Weiss, G.: On an index policy for restless bandits. *Journal of applied probability*, 637–648 (1990)
- [6] Glazebrook, K.D., Ruiz-Hernandez, D., Kirkbride, C.: Some indexable families of restless bandit problems. *Advances in Applied Probability* **38**(3), 643–672 (2006)
- [7] Ansell, P., Glazebrook, K.D., Nino-Mora, J., O’Keeffe, M.: Whittle’s index policy for a multi-class queueing system with convex holding costs. *Mathematical Methods of Operations Research* **57**(1), 21–39 (2003)
- [8] Glazebrook, K., Mitchell, H.: An index policy for a stochastic scheduling model with improving/deteriorating jobs. *Naval Research Logistics (NRL)* **49**(7), 706–721 (2002)
- [9] Aalto, S., Lassila, P., Taboada, I.: Whittle index approach to opportunistic scheduling with partial channel information. *Performance Evaluation* **136**, 102052 (2019)
- [10] Liu, K., Zhao, Q.: Indexability of restless bandit problems and optimality of whittle index for dynamic multichannel access. *IEEE Transactions on Information Theory* **56**(11), 5547–5567 (2010)
- [11] Avrachenkov, K.E., Borkar, V.S.: Whittle index based q-learning for restless bandits with average reward. *Automatica* **139**, 110186 (2022)
- [12] Niño-Mora, J.: A dynamic page-refresh index policy for web crawlers. In: *International Conference on Analytical and Stochastic Modeling Techniques and Applications*, pp. 46–60 (2014). Springer
- [13] Avrachenkov, K., Ayesta, U., Doncel, J., Jacko, P.: Congestion control of tcp flows in internet routers by means of index policy. *Computer Networks* **57**(17), 3463–3478 (2013)
- [14] Avrachenkov, K., Piunovskiy, A., Zhang, Y.: Impulsive control for g-aimd dynamics with relaxed and hard constraints. In: *2018 IEEE Conference on Decision and Control (CDC)*, pp. 880–887 (2018). IEEE
- [15] Scully, Z., Harchol-Balter, M., Scheller-Wolf, A.: Soap: One clean analysis of all age-based scheduling policies. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* **2**(1), 1–30 (2018)
- [16] Aalto, S., Ayesta, U., Richter, R.: Properties of the gittins index with application to optimal scheduling. *Probability in the Engineering and*

- Informational Sciences **25**(3), 269–288 (2011)
- [17] Aalto, S., Ayesta, U., Righter, R.: On the gittins index in the m/g/1 queue. *Queueing Systems* **63**(1-4), 437 (2009)
- [18] Borkar, V.S., Pattathil, S.: Whittle indexability in egalitarian processor sharing systems. *Annals of Operations Research*, 1–21 (2017)
- [19] Larrañaga, M., Ayesta, U., Verloop, I.M.: Asymptotically optimal index policies for an abandonment queue with convex holding cost. *Queueing systems* **81**(2), 99–169 (2015)
- [20] Archibald, T.W., Black, D., Glazebrook, K.D.: Indexability and index heuristics for a simple class of inventory routing problems. *Operations research* **57**(2), 314–326 (2009)
- [21] Glazebrook, K.D., Kirkbride, C., Ouenniche, J.: Index policies for the admission control and routing of impatient customers to heterogeneous service stations. *Operations Research* **57**(4), 975–989 (2009)
- [22] Villar, S.S., Bowden, J., Wason, J.: Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Statistical science: a review journal of the Institute of Mathematical Statistics* **30**(2), 199 (2015)
- [23] Chen, Y.R., Katehakis, M.N.: Linear programming for finite state multi-armed bandit problems. *Mathematics of Operations Research* **11**(1), 180–183 (1986)
- [24] Katehakis, M.N., Veinott Jr, A.F.: The multi-armed bandit problem: decomposition and computation. *Mathematics of Operations Research* **12**(2), 262–268 (1987)
- [25] Niño-Mora, J.: A $(2/3) n^3$ fast-pivoting algorithm for the gittins index and optimal stopping of a markov chain. *INFORMS Journal on Computing* **19**(4), 596–606 (2007)
- [26] Sonin, I.M.: A generalized gittins index for a markov chain and its recursive calculation. *Statistics & Probability Letters* **78**(12), 1526–1533 (2008)
- [27] Chakravorty, J., Mahajan, A.: Multi-armed bandits, gittins index, and its calculation. *Methods and applications of statistics in clinical trials: Planning, analysis, and inferential methods* **2**(416-435), 455 (2014)
- [28] Niño-Mora, J.: A fast-pivoting algorithm for whittle’s restless bandit index. *Mathematics* **8**(12), 2226 (2020)

- [29] Akbarzadeh, N., Mahajan, A.: Restless bandits with controlled restarts: Indexability and computation of whittle index. In: 2019 IEEE 58th Conference on Decision and Control (CDC), pp. 7294–7300 (2019). IEEE
- [30] Akbarzadeh, N., Mahajan, A.: Maintenance of a collection of machines under partial observability: Indexability and computation of whittle index. arXiv preprint arXiv:2104.05151 (2021)
- [31] Niño-Mora, J.: Dynamic priority allocation via restless bandit marginal productivity indices. *Top* **15**(2), 161–198 (2007)
- [32] Akbarzadeh, N., Mahajan, A.: Conditions for indexability of restless bandits and an $O(K^3)$ algorithm to compute Whittle index. arXiv (2020)
- [33] Fu, J., Nazarathy, Y., Moka, S., Taylor, P.G.: Towards q-learning the whittle index for restless bandits. In: 2019 Australian & New Zealand Control Conference (ANZCC), pp. 249–254 (2019). IEEE
- [34] Nakhleh, K., Ganji, S., Hsieh, P.-C., Hou, I., Shakkottai, S., et al.: Neurwin: Neural whittle index network for restless bandits via deep rl. *Advances in Neural Information Processing Systems* **34** (2021)
- [35] Ayesta, U., Gupta, M.K., Verloop, I.M.: On the computation of whittle’s index for markovian restless bandits. *Mathematical Methods of Operations Research* **93**(1), 179–208 (2021)
- [36] Xiong, G., Li, J., Singh, R.: Reinforcement Learning for Finite-Horizon Restless Multi-Armed Multi-Action Bandits. arXiv (2021)
- [37] Strassen, V.: Gaussian elimination is not optimal. *Numerische mathematik* **13**(4), 354–356 (1969)
- [38] Puterman, M.L.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 1st edn. John Wiley & Sons, Inc., USA (1994)
- [39] Gall, F.L., Urrutia, F.: Improved rectangular matrix multiplication using powers of the coppersmith-winograd tensor. In: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1029–1046 (2018). SIAM
- [40] Woodbury, M.A.: *Inverting modified matrices*. Statistical Research Group (1950)
- [41] Huang, J., Smith, T.M., Henry, G.M., Van De Geijn, R.A.: Strassen’s algorithm reloaded. In: *SC’16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 690–701 (2016). IEEE

- [42] Huang, J., et al.: Practical fast matrix multiplication algorithms. PhD thesis (2018)
- [43] Gast, N., Gaujal, B., Yan, C.: Exponential convergence rate for the asymptotic optimality of whittle index policy. arXiv preprint arXiv:2012.09064 (2020)
- [44] Akbarzadeh, N., Mahajan, A.: Restless bandits: indexability and computation of whittle index. arXiv e-prints, 2008 (2020)