

## Review

## Reversible Markov decision processes and the Gaussian free field

Venkat Anantharam

EECS Department, University of California, Berkeley, CA 94720, USA

## ARTICLE INFO

## Article history:

Received 11 July 2022

Received in revised form 11 September 2022

Accepted 12 September 2022

Available online 30 September 2022

Dedicated to the memory of Aristotle (Ari) Arapostathis

## Keywords:

Biconnectedness

Gaussian free field

Markov decision process

Markov chain Monte Carlo

Reversibility

## ABSTRACT

A Markov decision problem is called reversible if the stationary controlled Markov chain is reversible under every stationary Markovian strategy. A natural application in which such problems arise is in the control of Metropolis–Hastings type dynamics. We characterize all discrete time reversible Markov decision processes with finite state and action spaces. We show that the policy iteration algorithm for finding an optimal policy can be significantly simplified in Markov decision problems of this type. We also highlight the relation between the finite time evolution of the accrual of reward and the Gaussian free field associated to the controlled Markov chain.

© 2022 Elsevier B.V. All rights reserved.

## Contents

|                                                                               |   |
|-------------------------------------------------------------------------------|---|
| 1. Introduction.....                                                          | 1 |
| 2. Setup.....                                                                 | 2 |
| 3. Initial results.....                                                       | 2 |
| 4. A characterization of reversible Markov decision problems.....             | 3 |
| 5. Dynamic programming equations and policy iteration.....                    | 4 |
| 6. The Gaussian free field and the generalized second Ray-Knight theorem..... | 5 |
| Declaration of competing interest.....                                        | 6 |
| Data availability.....                                                        | 6 |
| Acknowledgments.....                                                          | 6 |
| Appendix A. Proof of Lemma 1.....                                             | 6 |
| Appendix B. Proof of Lemma 2.....                                             | 6 |
| Appendix C. Proof of Theorem 1.....                                           | 6 |
| Appendix D. Proof of Lemma 3.....                                             | 7 |
| Appendix E. Proof of Theorem 2.....                                           | 7 |
| Appendix F. Proof of Theorem 3.....                                           | 7 |
| Appendix G. Proof of Theorem 4.....                                           | 8 |
| References.....                                                               | 8 |

## 1. Introduction

We study Markov decision processes (MDPs) in a finite-state, finite-action framework with an average-reward criterion, when the controlled Markov chain is irreducible and reversible in stationarity under every stationary Markov control strategy. This problem was originally studied in special cases by Cogill and Peng [1], but that work does not seem to have attracted much

attention. We strengthen the main theorems in [1] by getting rid of superfluous assumptions. We characterize the class of all such problems. We also highlight the connections between such problems and the Gaussian free field of a weighted graph.

This paper is dedicated to the memory of Ari Arapostathis, a good personal friend, who was fond both of discrete-state MDPs and of the control problems arising in the Gaussian world of diffusion processes. We hope that the mix of MDPs with Gaussianity appearing in this paper – which is of a form that is unusual in the control context – would have met with his approval.

E-mail address: [ananth@berkeley.edu](mailto:ananth@berkeley.edu).

## 2. Setup

$\mathcal{X}$  and  $\mathcal{U}$  are finite sets, denoting the set of states and the set of actions respectively. To avoid dealing with corner cases, we assume that both  $\mathcal{X}$  and  $\mathcal{U}$  have cardinality at least 2. For each  $u \in \mathcal{U}$  let  $P(u) := [p_{ij}(u)]$  be a transition probability matrix (TPM) on  $\mathcal{X}$ , where  $p_{ij}(u)$  denotes the conditional probability that the next state is  $j$  when the current state is  $i$  and the action taken is  $u$ . We are also given a function  $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ , where  $r(i, u)$  denotes the reward received if the current state is  $i$  and the current action is  $u$ .

A *stationary randomized Markov strategy*  $\mu$  is defined to be a choice of conditional probability distributions  $(\mu(u|i) : u \in \mathcal{U}, i \in \mathcal{X})$  and results in the TPM  $P(\mu) := [p_{ij}(\mu)]$  on  $\mathcal{X}$ , where

$$p_{ij}(\mu) := \sum_u p_{ij}(u)\mu(u|i).$$

We write  $\mathcal{M}$  for the set of stationary randomized Markov control strategies. The interpretation of  $P(\mu)$  is as the TPM of the controlled Markov chain when the strategy  $\mu \in \mathcal{M}$  is implemented.

We make the assumption that for each  $\mu \in \mathcal{M}$  the TPM  $P(\mu)$  is irreducible and reversible. We will then say that we are dealing with a *reversible Markov decision problem* (RMDP). This assumption may seem quite restrictive, but it seems to be sufficiently interesting to merit some attention. For instance, applications to the optimal design of algorithms of the Metropolis-Hastings type to generate a target probability distribution on a large set of combinatorial configurations, i.e. the Markov Chain Monte Carlo method, are discussed in some depth in [1, Sec. 5].

By the assumption of irreducibility there is unique probability distribution  $\pi(\mu) := (\pi_i(\mu) : i \in \mathcal{X})$ , called the *stationary distribution* of  $P(\mu)$ , which can be thought of as a column vector satisfying  $\pi(\mu)^T P(\mu) = \pi(\mu)^T$ . The assumption of reversibility says that we have

$$\pi_i(\mu)p_{ij}(\mu) = \pi_j(\mu)p_{ji}(\mu), \text{ for all } i, j \in \mathcal{X}. \quad (1)$$

The conditions in (1) are often called a *detailed-balance* assumption. Note that  $(\pi_i(\mu)p_{ij}(\mu) : (i, j) \in \mathcal{X} \times \mathcal{X})$  is a probability distribution, called the *occupation measure* of  $P(\mu)$ , and the reversibility assumption for  $P(\mu)$  is equivalent to the assumption that the occupation measure is symmetric when viewed as a matrix.

We will denote the set of stationary deterministic Markov control strategies by  $\bar{\mathcal{M}}$  and write  $\bar{\mu}$  for such a strategy. Thus  $\bar{\mu} \in \bar{\mathcal{M}}$  is a function  $\bar{\mu} : \mathcal{X} \rightarrow \mathcal{U}$  and, with an abuse of notation, can also be thought of as the stationary randomized Markov control strategy  $\bar{\mu}$  where  $\bar{\mu}(u|i)$  equals 1 if  $u = \bar{\mu}(i)$  and 0 otherwise. Note that  $|\bar{\mathcal{M}}| = |\mathcal{U}|^{|\mathcal{X}|}$ , where  $|\mathcal{A}|$  denotes the cardinality of a finite set  $\mathcal{A}$ . Of course,  $P(\bar{\mu})$  need not be distinct for distinct  $\bar{\mu} \in \bar{\mathcal{M}}$ . Similarly,  $\mathcal{M}$  can be thought of as the product of  $\mathcal{X}$  copies of the probability simplex based on  $\mathcal{U}$ .

Even though irreducibility of the  $P(u)$  for  $u \in \mathcal{U}$  is not explicitly mentioned as a condition in the definition of the notion of an RMDP in [1], it seems to be implicitly assumed, since the notion of reversibility seems to be discussed there under the implicit assumption that there is a unique stationary distribution. Thus the use of the terminology “reversible Markov decision process” in this document seems to be consistent with its use in [1].

Also note that while it may be more natural to also include the reward function  $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$  in the definition of an MDP, since the notion of an RMDP depends only on a property of  $(P(u), u \in \mathcal{U})$  we choose to define an RMDP purely in terms of  $(P(u), u \in \mathcal{U})$ .

## 3. Initial results

Our first claim is the following simple observation. For completeness, a proof is provided in [Appendix A](#).

**Lemma 1.**  $P(\mu)$  is irreducible and reversible for each  $\mu \in \mathcal{M}$  iff  $P(\bar{\mu})$  is irreducible and reversible for each  $\bar{\mu} \in \bar{\mathcal{M}}$ . ■

As pointed out in [1], a natural class of examples of RMDPs arises as follows.

**Example 1.** Let  $G := (\mathcal{X}, \mathcal{E})$  be a simple connected graph with the finite vertex set  $\mathcal{X}$  and edge set  $\mathcal{E}$ . (Recall that a graph is called simple if it does not have multiple edges between any pair of vertices and does not have any self-loops.) To each edge  $(i, j) \in \mathcal{E}$  (between the vertices  $i, j \in \mathcal{X}$ ) associate the strictly positive weight  $s_{ij}$  (thus  $s_{ij} = s_{ji}$ ). Since  $G$  has no self-loops, we have  $s_{ii} = 0$  for all  $i \in \mathcal{X}$ . Write  $s_i$  for  $\sum_{j \in \mathcal{X}} s_{ij}$  and  $S$  for  $\sum_{i \in \mathcal{X}} s_i$ . Let  $P^{(0)}$  denote the transition probability matrix on  $\mathcal{X}$  with

$$p_{ij}^{(0)} = \frac{s_{ij}}{s_i}, \text{ for all } i, j \in \mathcal{X}.$$

Let  $\rho : \mathcal{X} \times \mathcal{U} \rightarrow (0, 1]$  be given. When the control action is  $u \in \mathcal{U}$ , assume that the state transitions occur according to  $P(u)$ , where

$$\begin{aligned} p_{ii}(u) &= 1 - \rho(i, u), \\ p_{ij}(u) &= \rho(i, u)p_{ij}^{(0)}, \text{ if } j \neq i. \end{aligned}$$

Finally, assume that a reward function  $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$  is given.

To check that this results in an RMDP we first observe that  $P^{(0)}$  is an irreducible and reversible TPM on  $\mathcal{X}$ . The irreducibility is obvious. Reversibility can be checked by observing that the stationary distribution of  $P^{(0)}$ , i.e.  $(\pi_i^{(0)} : i \in \mathcal{X})$ , is given by  $\pi_i^{(0)} = \frac{s_i}{S}$  for  $i \in \mathcal{X}$ .

Given  $\mu \in \mathcal{M}$ , write  $\rho(i, \mu)$  for  $\sum_{u \in \mathcal{U}} \rho(i, u)\mu(u|i)$ . for  $j \neq i$ . Then it can be checked that we have  $p_{ij}(\mu) = \rho(i, \mu)p_{ij}^{(0)}$  for  $j \neq i$ , while  $p_{ii}(\mu) = 1 - \rho(i, \mu)$ . Now, for each  $\mu \in \mathcal{M}$ ,  $P(\mu)$  is irreducible since, by assumption, we have  $\rho(i, u) > 0$  for all  $(i, u) \in \mathcal{X} \times \mathcal{U}$ . To check that  $P(\mu)$  is reversible it suffices to observe that its stationary distribution, i.e.  $(\pi_i(\mu) : i \in \mathcal{X})$ , is given by  $(K(\mu)\frac{\pi_i^{(0)}}{\rho(i, \mu)} : i \in \mathcal{X})$ , where  $K(\mu) := \left(\sum_{i \in \mathcal{X}} \frac{\pi_i^{(0)}}{\rho(i, \mu)}\right)^{-1}$  is the normalizing constant. ■

In the scenario of [Example 1](#), if one scales all the weights  $s_{ij}$  by the same positive constant then, with the same  $\rho : \mathcal{X} \times \mathcal{U} \rightarrow (0, 1]$  and  $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ , one gets the same RMDP, since all the  $s_i$  and  $S$  also scale by the same constant. What matters is the irreducible reversible transition probability matrix  $P^{(0)}$  with zero diagonal entries defined by the weighted graph. Conversely, one can check that any irreducible reversible TPM  $P^{(0)}$  with entries  $p_{ij}^{(0)}$ ,  $i, j \in \mathcal{X}$  and zero diagonal entries can be thought of as arising from the simple connected graph  $G := (\mathcal{X}, \mathcal{E})$  with  $(i, j) \in \mathcal{E}$  iff  $p_{ij}^{(0)} > 0$ , with weight  $s_{ij} := \pi_i^{(0)}p_{ij}^{(0)}$ , where  $(\pi_i^{(0)} : i \in \mathcal{X})$  is the stationary distribution of  $P^{(0)}$ .

As stated in the following simple lemma, whose proof is in [Appendix B](#), one can associate a simple connected graph to any RMDP. We will refer to this graph as the *canonical graph* of the RMDP.

**Lemma 2.** Consider an RMDP, defined by  $(P(u) : u \in \mathcal{U})$  as above. Then there must exist a simple connected graph  $G := (\mathcal{X}, \mathcal{E})$  such that for all  $u \in \mathcal{U}$  and distinct  $i, j \in \mathcal{X}$  we have  $p_{ij}(u) > 0$  iff  $(i, j) \in \mathcal{E}$ . ■

In fact, as stated in the following theorem, it turns out that under relatively mild conditions every RMDP must be of the form described in [Example 1](#). The proof is provided in [Appendix C](#).

**Theorem 1.** Consider an RMDP, defined by  $(P(u) : u \in \mathcal{U})$  as above. Let  $G := (\mathcal{X}, \mathcal{E})$  be the canonical graph associated to this problem, as in [Lemma 2](#). Suppose now that this graph is biconnected. (Recall that a graph is called biconnected – or 2-connected – if whenever any single vertex, together with all the edges involving that vertex, is removed the resulting graph continues to be connected.) Then there is an irreducible reversible TPM  $P^{(0)}$  on  $\mathcal{X}$  such that  $p_{ij}^{(0)} > 0$  iff  $(i, j) \in \mathcal{E}$ , and  $\rho : \mathcal{X} \times \mathcal{U} \rightarrow (0, 1]$ , such that for each  $u \in \mathcal{U}$  we have  $p_{ij}(u) = \rho(i, u)p_{ij}^{(0)}$  for  $j \neq i$ , and  $p_{ii}(u) = 1 - \rho(i, u)$ . ■

Much of the discussion in [\[1\]](#) centers around RMDPs which have a Hamilton cycle in their canonical graph. These are biconnected, and hence of the kind in [Example 1](#). However, as seen from [Example 2](#), there are RMDPs that are not of the type in [Example 1](#).

**Example 2.** Let  $\mathcal{X} = \{1, 2, 3\}$  and  $\mathcal{U} = \{1, 2\}$ . Choose  $a \neq b$  such that  $0 < a, b, < 1$ . Let

$$P(1) = \begin{bmatrix} 0 & a & 1-a \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \text{ and } P(2) = \begin{bmatrix} 0 & b & 1-b \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}.$$

It can be checked that for any  $\mu \in \mathcal{M}$  we have

$$P(\mu) = \begin{bmatrix} 0 & \lambda a + (1-\lambda)b & \lambda(1-a) + (1-\lambda)(1-b) \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix},$$

for some  $\lambda \in [0, 1]$  (depending on  $\mu$ ). Since  $P(\mu)$  is irreducible and reversible, with the stationary distribution being

$$\pi(\mu) = \left[ \frac{1}{2} \quad \frac{1}{2}(\lambda a + (1-\lambda)b) \quad \frac{1}{2}(\lambda(1-a) + (1-\lambda)(1-b)) \right],$$

this, together with some reward function  $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ , defines an RMDP. However this RMDP is not of the form in [Example 1](#) as can be seen, for instance, by noticing that  $p_{11}(\mu) = 0$  for all  $\mu$  but  $p_{12}(\mu)$  can take on distinct values for distinct  $\mu$  (since we assumed that  $a \neq b$ ).


Here the canonical graph of the RMDP has vertex set  $\mathcal{X}$  and edge set  $\mathcal{E} = \{(1, 2), (1, 3)\}$ . Note that this graph is not biconnected. ■

#### 4. A characterization of reversible Markov decision problems

Roughly speaking, a simple connected graph  $G := (\mathcal{X}, \mathcal{E})$  is as far from being biconnected as it can be if there is unique path between every pair of vertices of the graph, i.e. if the graph is a tree. This is of course not precisely true, since the graph with two vertices connected with a single edge is both biconnected and a tree and, more generally, in any tree the removal of a leaf vertex together with the edge connected to it leaves behind a connected graph. Nevertheless, this rough intuition suggests that one should pay special attention to trees. As the following simple result shows, in contrast to the case considered in [Theorem 1](#), when the canonical graph of an RMDP is a tree there are hardly any restrictions on the structure of the decision problem. The proof is in [Appendix D](#).

**Lemma 3.** Let  $G := (\mathcal{X}, \mathcal{E})$  be a tree. Let  $(P(u) : u \in \mathcal{U})$  be any collection of TPMs on  $\mathcal{X}$  satisfying the condition that  $p_{ij}(u) > 0$  iff  $(i, j) \in \mathcal{E}$ . Then this defines an RMDP. ■

We now proceed to characterize all RMDPs. It turns out that the situations discussed in [Theorem 1](#) and [Lemma 3](#) are extreme cases and, in a sense, the general case lies between these two



**Fig. 1.** The block graph associated to a simple connected graph with nine nodes, numbered as indicated, is shown. The articulation points are the vertices 2, 3, and 6, and are depicted by thick red nodes. There are five blocks, namely  $\{1, 2\}$ ,  $\{2, 3\}$ ,  $\{3, 5, 6\}$ ,  $\{4, 6\}$ , and  $\{6, 7, 8, 9\}$ . Note that the articulation point 2 is shared by two blocks, as is the articulation point 3, while the articulation point 6 is shared by three blocks. Note that (since there is more than one block) each block has at least one articulation point, while the block  $\{3, 5, 6\}$  has two articulation points.

extremes. Underlying this is the well-known *block graph* structure of a simple connected graph  $G := (\mathcal{X}, \mathcal{E})$ . Recall that a *cutvertex* of  $G$  is a vertex such that if we remove that vertex and the edges connected to it, the resulting graph is disconnected. A *block* is defined to be a maximal connected subgraph of  $G$  that has no cutvertices. Thus a block  $B$  is biconnected; in particular, it is either a subgraph comprised of a single edge (in which case it has two vertices) or has the property that given any three distinct vertices  $i, j, k \in B$  there is path between  $j$  and  $k$  in  $B$  that does not meet  $i$ . Also, any two blocks  $B$  and  $B'$  that intersect do so at a uniquely defined vertex, called an *articulation point* of the block graph structure. An articulation point will be a cutvertex of  $G$  (but not of  $B$  or  $B'$ , since  $B$  and  $B'$  are blocks and so do not have cutvertices). If there is only one block in the block structure, then there are no articulation points. If there is more than one block then every block has at least one articulation point, but in general may have several articulation points. Every articulation point then lies in at least two blocks, but may in general lie in several blocks. An illustrative example of the block graph structure is given in [Fig. 1](#); see e.g. [\[2, Sec. 3.1\]](#) for more details (we focus on connected graphs, even though the block graph can be defined more generally).

Given the simple connected graph  $G := (\mathcal{X}, \mathcal{E})$ , we write  $\mathcal{A}$  for the set of articulation points and  $\mathcal{B}$  for the set of blocks. Note that each  $a \in \mathcal{A}$  is a vertex of  $G$ , while each  $B \in \mathcal{B}$  is a subgraph of  $G$ . Nevertheless, with an abuse of notation, we will also use  $B$  to denote the vertex set of the block  $B$ . Thus we write  $a \in B$  to indicate that the articulation point  $a$  is in the vertex set of  $B$  and similarly write  $B \ni a$  to indicate that the vertex set of  $B$  contains the articulation point  $a$ . Further, we will write  $\hat{B}$  for the subset of those vertices of the block  $B$  that are not articulation points. Note that  $\hat{B}$  can be empty.

In the following example we describe a class of RMDPs that is broader in scope than those considered in [Example 1](#) and [Lemma 3](#) (in particular [Example 2](#)), including both of these as special cases.

**Example 3.** Let  $G := (\mathcal{X}, \mathcal{E})$  be a simple connected graph. In the block graph structure of  $G$ , let  $\mathcal{A}$  denote the set of articulation points and  $\mathcal{B}$  the set of blocks. For each  $B \in \mathcal{B}$  let  $P^{(0)}(B)$  be a given irreducible reversible TPM on  $B$  with diagonal entries  $p_{ii}^{(0)}(B) = 0$  for all  $i \in B$  and with off-diagonal entries  $p_{ij}^{(0)}(B) > 0$  iff  $(i, j) \in \mathcal{E}$ . For each articulation point  $a \in \mathcal{A}$  (if any) and  $u \in \mathcal{U}$ , let  $(\nu_a(u, B) : B \ni a)$  be strictly positive numbers satisfying  $\sum_{B \ni a} \nu_a(u, B) = 1$ .

For  $i \in \hat{B}$  and  $u \in \mathcal{U}$ , define  $v_i(u, B) = 1$ , and for all  $i \in \mathcal{X}$  define  $v_i(u, B) = 0$  for all  $u \in \mathcal{U}$  if  $i \notin B$ . Let  $\rho : \mathcal{X} \times \mathcal{U} \rightarrow (0, 1]$  be given.

For each  $u \in \mathcal{U}$  define  $P(u) := [p_{ij}(u)]$ , a TPM on  $\mathcal{X}$ , by

$$\begin{aligned} p_{ij}(u) &= \rho(i, u)v_i(u, B)p_{ij}^{(0)}(B) \text{ if } i \in B, j \in B, j \neq i, B \in \mathcal{B}, \\ p_{ii}(u) &= 1 - \rho(i, u), \text{ if } i \in \mathcal{X}, \\ p_{ij}(u) &= 0, \text{ otherwise.} \end{aligned} \quad (2)$$

Then this defines an RMDP.

If there is only a single block, then there are no articulation points and we are in the scenario of [Example 1](#), where the claim has already been established. Suppose therefore that there are two or more blocks (thus every block has at least one articulation point). To verify the claim, we need to check that for each  $\mu \in \mathcal{M}$  the matrix  $P(\mu)$  on  $\mathcal{X}$  is an irreducible reversible TPM. It is straightforward to check that  $P(\mu)$  is a TPM. Noting that for  $i \neq j$  we have  $p_{ij}(\mu) > 0$  iff  $(i, j) \in \mathcal{E}$ , we see that  $P(\mu)$  is irreducible.

Let  $\tau_i(u, B) := \rho(i, u)v_i(u, B)$ , and for  $\mu \in \mathcal{M}$  let  $\tau_i(\mu, B) := \sum_u \tau_i(u, B)\mu(u|i)$ . Let  $(\psi_i(B) : i \in B)$  denote the stationary probability distribution of  $P^{(0)}(B)$ . Then, by the assumed reversibility of this matrix, we have

$$\psi_i(B)p_{ij}^{(0)}(B) = \psi_j(B)p_{ji}^{(0)}(B), \text{ for all } i, j \in B, B \in \mathcal{B}.$$

Consider the vector  $(\frac{\psi_i(B)}{\tau_i(\mu, B)} : i \in B)$ . Observe now that we have

$$\frac{\psi_i(B)}{\tau_i(\mu, B)}p_{ij}(\mu) = \frac{\psi_j(B)}{\tau_j(\mu, B)}p_{ji}(\mu) \text{ for all } i, j \in B, B \in \mathcal{B}. \quad (3)$$

We now claim that we can find positive constants  $(m(\mu, B) : B \in \mathcal{B})$  such for every  $B, B' \in \mathcal{B}$ ,  $B \neq B'$ , if they share an articulation point  $a \in \mathcal{A}$  (i.e.  $a \in B, a \in B'$ ), then we have

$$m(\mu, B)\frac{\psi_a(B)}{\tau_a(\mu, B)} = m(\mu, B')\frac{\psi_a(B')}{\tau_a(\mu, B')}. \quad (4)$$

Since this number does not depend on the choice of  $B \in \mathcal{B}$  containing  $a$ , let us denote it by  $\gamma_a(\mu)$ . Let us also write  $\gamma_i(\mu)$  for  $m(\mu, B)\frac{\psi_i(B)}{\tau_i(\mu, B)}$  for  $i \in \hat{B}$  for any  $B \in \mathcal{B}$ . With this notation in place, we further claim that we can choose the  $(m(\mu, B) : B \in \mathcal{B})$  such that

$$\sum_{i \in \mathcal{X}} \gamma_i(\mu) = 1. \quad (5)$$

It can then be checked that  $(\gamma_i(\mu) : i \in \mathcal{X})$  is then the stationary distribution of  $P(\mu)$  and, based on (3), we can conclude that  $P(\mu)$  is reversible.

To find the scaling factors  $(m(\mu, B) : B \in \mathcal{B})$  with the claimed properties, pick any block and call it the root. Because there are at least two blocks, this block has at least one articulation point, and each such articulation point is associated with a unique block other than the root. Call these blocks the ones at depth 1. If any such block has additional articulation points (other than the one it shares with the root), each of these will be associated with a new block, and we will call the blocks identified in this way (from all the blocks at depth 1) the blocks at depth 2, and so on. We start with a scaling factor 1 for the root, and see that we can set the scaling factor uniquely for each of the blocks at depth 1 in order to get the matching condition in (4) to hold at all the articulation points that are shared between the root and the blocks at depth 1. We can next set the scaling factor for each of the blocks at depth 2 uniquely in order to get the matching condition in (4) to hold at all the articulation points that are shared between a block at depth 1 and a block at depth 2 and so on. At the end of this process we have scaling factors such that the condition in (4) holds at all articulation points and then we can finally scale all the scaling factors jointly by the same positive constant to get the condition in (5). ■

The class of RMDPs arising as in [Example 3](#) also includes those arising as in [Lemma 3](#). This corresponds to the case where every block is a single edge, which is equivalent to the case where the given graph  $G := (\mathcal{X}, \mathcal{E})$  is a tree. Each  $P^{(0)}(B)$  is then of the form  $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ . If  $|\mathcal{X}| = 2$  then there are no articulation points and the scenario is covered in [Theorem 1](#) (and also in [Lemma 3](#)). If  $|\mathcal{X}| \geq 3$  the articulation points are precisely the non-leaf vertices of the tree.

It turns out that the scenarios covered in [Example 3](#) completely characterize all the ways in which an RMDP can arise. This is stated in the following theorem, whose proof is in [Appendix E](#).

**Theorem 2.** Consider an RMDP, defined by the TPMs  $(P(u) : u \in \mathcal{U})$ . Let  $G := (\mathcal{X}, \mathcal{E})$  be the canonical graph associated to this problem, as in [Lemma 2](#). In the block graph structure of  $G$ , let  $\mathcal{A}$  denote the set of articulation points and  $\mathcal{B}$  the set of blocks. Then for each  $B \in \mathcal{B}$  there will be an irreducible reversible TPM  $P^{(0)}(B)$  on  $B$  with diagonal entries  $p_{ii}^{(0)}(B) = 0$  for all  $i \in B$  and off-diagonal entries  $p_{ij}^{(0)}(B) > 0$  iff  $(i, j) \in \mathcal{E}$ ; for each articulation point  $a \in \mathcal{A}$  (if any) and  $u \in \mathcal{U}$  there will be strictly positive numbers  $(v_a(u, B) : B \ni a)$ , satisfying  $\sum_{B \ni a} v_a(u, B) = 1$ ; and there will be  $\rho : \mathcal{X} \times \mathcal{U} \rightarrow (0, 1]$  such that for each  $u \in \mathcal{U}$  the entries of the matrix  $P(u)$  are given by (2). ■

## 5. Dynamic programming equations and policy iteration

Consider an MDP defined by a family  $(P(u) : u \in \mathcal{U})$  where each  $P(u)$  is an irreducible TPM on  $\mathcal{X}$ , and a reward function  $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ . Here  $\mathcal{X}$  and  $\mathcal{U}$  are finite sets each assumed to be of cardinality at least 2. Given  $\mu \in \mathcal{M}$ , let  $\beta(\mu)$  denote the long term average reward associated to the stationary randomized Markov strategy  $\mu$ . Then we have  $\beta(\mu) = \sum_i \pi_i(\mu)r(i, \mu)$ , where  $\pi(\mu) = (\pi_i(\mu) : i \in \mathcal{X})$  denotes the stationary distribution of  $P(\mu)$  and  $r(i, \mu) := \sum_u r(i, u)\mu(u|i)$ . Further, there is function  $h(\mu) : \mathcal{X} \rightarrow \mathbb{R}$  such that for all  $i \in \mathcal{X}$  we have

$$\beta(\mu) = r(i, \mu(i)) + \sum_j p_{ij}(\mu)(h_j(\mu) - h_i(\mu)). \quad (6)$$

The family of Eqs. (6), one for each  $i \in \mathcal{X}$ , is often viewed as needing to be solved for  $\beta(\mu)$  and  $h(\mu) := (h_i(\mu) : i \in \mathcal{X})$ , in which case it is called *Poisson's equation* associated to the TPM  $P(\mu)$ . Note that the number of variables is one more than the number of equations and, indeed, one can add the same fixed constant to each  $h_i(\mu)$  in any solution to find another solution.

A natural choice for  $h(\mu)$ , thought of as a column vector, is given by the Cesàro limit of the sequence  $(\sum_{k=0}^{K-1} (P(\mu)^k - \mathbb{1}\pi(\mu)^T)r(\mu), K \geq 1)$ , where  $r(\mu)$  is thought of as the column vector with  $r(\mu) = (r(i, \mu(i)) : i \in \mathcal{X})$ , and  $\mathbb{1}$  denotes the all-ones column vector. This Cesàro limit exists because the sequence  $(\frac{1}{K} \sum_{k=0}^{K-1} P(\mu)^k, K \geq 1)$  converges geometrically fast to  $\mathbb{1}\pi(\mu)^T$  as  $K \rightarrow \infty$ . Taking the Cesàro limit is needed to deal with the phenomenon of periodicity.

The *average cost dynamic programming equation* characterizes an optimal stationary randomized Markov strategy  $\mu$  as one having the property that for each  $i \in \mathcal{X}$  if  $\mu(u|i) > 0$  then we must have

$$\begin{aligned} r(i, u) + \sum_j p_{ij}(u)(h_j(\mu) - h_i(\mu)) \\ = \max_v \left( r(i, v) + \sum_j p_{ij}(v)(h_j(\mu) - h_i(\mu)) \right), \end{aligned} \quad (7)$$

which implies the form in which it is usually written, namely

$$\beta(\mu) = \max_v \left( r(i, v) + \sum_j p_{ij}(v)(h_j(\mu) - h_i(\mu)) \right). \quad (8)$$



The characterization of optimal stationary randomized Markov strategies in Eq. (7) leads to the *policy iteration algorithm* to find an optimal stationary deterministic strategy. Namely, starting with  $\bar{\mu}^{(0)} \in \bar{\mathcal{M}}$ , consider the sequence  $(\bar{\mu}^{(k)} \in \bar{\mathcal{M}}, k \geq 0)$  where to get  $\bar{\mu}^{(k+1)}$  from  $\bar{\mu}^{(k)}$  we pick some state  $i$  (if possible) for which  $\arg\max_v (r(i, v) + \sum_j p_{ij}(v)(h_j(\bar{\mu}^{(k)}) - h_i(\bar{\mu}^{(k)})))$  does not equal  $\bar{\mu}^{(k)}(i)$ , and replace  $\bar{\mu}^{(k)}(i)$  by an action achieving the argmax. It is well-known that we will then have  $\beta_{\bar{\mu}^{(k+1)}} > \beta_{\bar{\mu}^{(k)}}$  (a proof is given in [1], for instance) and that this iteration will terminate in a finite number of steps to a stationary deterministic optimal strategy.

We turn now to the case where the MDP is an RMDP, i.e. when  $P(\mu)$  is reversible for all  $\mu \in \mathcal{M}$ . Consider first the case where the canonical graph  $G := (\mathcal{X}, \mathcal{E})$  associated to the RMDP is biconnected. Then, according to Theorem 1, we have  $\rho : \mathcal{X} \times \mathcal{U} \rightarrow (0, 1]$  and an irreducible reversible TPM  $P^{(0)}$  on  $\mathcal{X}$  such that  $p_{ij}(u) = \rho(i, u)p_{ji}^{(0)}$  for all distinct  $i, j \in \mathcal{X}$  and  $p_{ii}(u) = 1 - \rho(i, u)$  for all  $i \in \mathcal{X}$ . As stated in the following theorem, the policy iteration algorithm can be dramatically simplified in this case. The proof is in Appendix F.

**Theorem 3.** Consider an RMDP whose associated canonical graph is biconnected. Let  $P^{(0)}$  and  $\rho : \mathcal{X} \times \mathcal{U} \rightarrow (0, 1]$  be as in Theorem 1. Let  $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$  be the reward function. Then any sequence  $(\bar{\mu}^{(k)} \in \bar{\mathcal{M}}, k \geq 0)$  of stationary deterministic Markov strategies, starting from some  $\bar{\mu}^{(0)} \in \bar{\mathcal{M}}$ , where  $\bar{\mu}^{(k+1)}$  is got from  $\bar{\mu}^{(k)}$  by picking some state  $i$  (if possible) for which

$$\frac{r(i, \bar{\mu}^{(k)}) - \beta(\bar{\mu}^{(k)})}{\rho(i, \bar{\mu}^{(k)})} < \arg\max_v \left( \frac{r(i, v) - \beta(\bar{\mu}^{(k)})}{\rho(i, v)} \right) \quad (9)$$

and replacing  $\bar{\mu}^{(k)}(i)$  by some action achieving the argmax, has the property that  $\beta(\bar{\mu}^{(k+1)}) > \beta(\bar{\mu}^{(k)})$ , and this iteration will terminate in a finite number of steps to a stationary deterministic optimal strategy. ■

A weaker version of Theorem 3 is proved in [1, Thm.4.2] under the assumption that there is a Hamilton cycle in the canonical graph associated to the RMDP (which implies, but is a strictly stronger requirement than biconnectedness) and that at each step of the policy iteration the actions at all states are updated simultaneously in a specific way related to this Hamilton cycle, see [1, Sec. 4.1].

For a general RMDP it turns out that a simplification of policy iteration similar to that in Theorem 3 is possible at vertices that are not articulation points. This is stated in the following theorem, whose proof is in Appendix G

**Theorem 4.** For a general RMDP, let  $\mathcal{A}$  denote the set of articulation points and  $\mathcal{B}$  the set of blocks in the block graph structure of the canonical graph  $G := (\mathcal{X}, \mathcal{E})$  associated to it. Let  $(P^{(0)}(B) : B \in \mathcal{B})$ ,  $(v_a(u, B) : a \in \mathcal{A}, B \ni a, u \in \mathcal{U})$ , and  $\rho : \mathcal{X} \times \mathcal{U} \rightarrow (0, 1]$  be as in Theorem 2. Let  $r : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$  be the reward function. Let  $\bar{\mu} \in \bar{\mathcal{M}}$  and suppose that for some  $B \in \mathcal{B}$  and  $k \in \bar{B}$  we have

$$\frac{r(k, \bar{\mu}) - \beta(\bar{\mu})}{\rho(k, \bar{\mu})} < \arg\max_v \left( \frac{r(k, v) - \beta(\bar{\mu})}{\rho(k, v)} \right). \quad (10)$$

Let  $\bar{\mu}' \in \bar{\mathcal{M}}$  be defined by setting  $\bar{\mu}'(k)$  equal to some  $v$  achieving the argmax, and let  $\bar{\mu}'(j) = \bar{\mu}(j)$  for all  $j \neq k$ . Then we have  $\beta(\bar{\mu}') > \beta(\bar{\mu})$ . ■

## 6. The Gaussian free field and the generalized second Ray-Knight theorem

For every  $\mu \in \mathcal{M}$  the Cesàro limit of the sequence  $(\sum_{k=0}^{K-1} (P(\mu)^k - \mathbb{1}\pi(\mu)^T), K \geq 1)$  exists and is called the *fundamental matrix* associated to  $P(\mu)$  [3, Sec. 2.2.2]. Denote this matrix

by  $Z(\mu)$ , with entries  $z_{ij}(\mu)$ . From the discussion in Section 5, note that  $h(\mu) := Z(\mu)r(\mu)$ , together with  $\beta(\mu) = \pi(\mu)^T r(\mu)$ , solves Poisson's equation for  $P(\mu)$ , i.e. Eq. (6). Thus, understanding the fundamental matrix  $Z(\mu)$  is central to understanding the dynamics of the RMDP under  $\mu \in \mathcal{M}$ .

$Z(\mu)$  is best understood by moving to continuous time, replacing the TPM  $P(\mu)$  by the rate matrix  $P(\mu) - I$ , where  $I$  denotes the identity matrix on  $\mathcal{X}$ . Let  $(X_t(\mu), t \geq 0)$  denote the corresponding continuous time Markov chain. Then one can check that

$$z_{ij}(\mu) = \lim_{T \rightarrow \infty} \left( E_i \left[ \int_0^T 1(X_t(\mu) = j) dt \right] - \pi_j(\mu)T \right). \quad (11)$$

Since  $P(\mu)$  is reversible, it is straightforward to show that the matrix on  $\mathcal{X}$  with entries  $\frac{z_{ij}(\mu)}{\pi_j(\mu)}$  is symmetric [3, Sec. 3.1]. Based on (11), we may now write, for the choice of  $h(\mu)$  above, for each  $i \in \mathcal{X}$ , the formula

$$\begin{aligned} h_i(\mu) &= \sum_j \frac{z_{ij}(\mu)}{\pi_j(\mu)} \pi_j(\mu) r_j(\mu) \\ &= \sum_j \lim_{T \rightarrow \infty} \left( E_i \left[ \frac{1}{\pi_j(\mu)} \int_0^T 1(X_t(\mu) = j) dt \right] - T \right) \pi_j(\mu) r_j(\mu). \end{aligned} \quad (12)$$

While this may seem a peculiar thing to do, one natural aspect of the formula on the RHS of (12) is that  $\pi_j(\mu)r_j(\mu)$  has the interpretation, in continuous time, of the rate at which reward is generated in stationarity while in state  $j$ . Another natural aspect is that the centering of the integral is the actual time and not a state-dependent scaled version of it. However, the real value of this way of writing the formula comes from the observation that the matrix with entries  $\frac{z_{ij}(\mu)}{\pi_j(\mu)}$  is a positive semidefinite matrix [3, Eqn. (3.42)]. This means that we can find a multivariate mean zero Gaussian random variable, call it  $(V_i(\mu) : i \in \mathcal{X})$ , with this covariance matrix. This points to an intriguing and unusual connection between Gaussianity and Markov decision theory in the case of RMDP. As we will see shortly, while the  $h_i(\mu)$  are expressed as asymptotic limits in (12), the introduction of Gaussian methods gives, in a sense, much more detailed information about the behavior of the functions  $T \rightarrow \frac{1}{\pi_j(\mu)} \int_0^T 1(X_t(\mu) = j)$  and thus a much more detailed picture of the role of the initial condition in causing deviations from the stationary rate of generation of reward in an RMDP.

Notice that we have  $\sum_i \sum_j \pi_i(\mu) \frac{z_{ij}(\mu)}{\pi_j(\mu)} \pi_j(\mu) = 0$ , and so  $\sum_i \pi_i(\mu) V_i(\mu) = 0$  as a random variable. Thus, to work with  $(V_i(\mu) : i \in \mathcal{X})$  involves, in a sense, a choice of coordinates to capture the underlying multivariate Gaussian structure. Other natural choices of coordinates are possible. For instance, for each  $k \in \mathcal{X}$  we may define the multivariate Gaussian  $(V_i^{[k]}(\mu) : i \in \mathcal{X})$  via  $V_i^{[k]}(\mu) := V_i(\mu) - V_k(\mu)$  (so the choice of coordinates in this case makes  $V_k^{[k]}(\mu) = 0$ ).

Instead of making a choice of coordinates, the Gaussian object of interest can be constructed in an intrinsic way. One starts with independent mean zero Gaussian random variables on the edges of the canonical graph of the RMDP, with the variance of the Gaussian on edge  $(i, j)$  being  $(\pi_i(\mu)p_{ij}(\mu))^{-1}$ . To each edge one associates a direction in an arbitrary way, with the understanding that traversing the edge along its direction corresponds to adding this Gaussian, while traversing it in the opposite direction corresponds to subtracting this Gaussian. One then conditions on being in the subspace of  $\mathbb{R}^{\mathcal{E}}$  where the total sum of the Gaussians over every loop in the canonical graph equals zero. This will allow us to construct a multivariate Gaussian on the vertices of the canonical graph with the property that the Gaussian on each edge is the difference between those at its endpoints. This

multivariate Gaussian on the vertices is defined only up to one degree of freedom and this is what corresponds to the freedom in the choice of coordinates discussed above. See [4, Sec. 9.4] for more details. This Gaussian object is called the *Gaussian free field* associated to  $P(\mu)$ . It is discussed in many sources, e.g. [5, Chap. 5], [6, Sec. 2.8], [4, Sec. 9.4], [7, Sec. 2.8].

For each  $k \in \mathcal{X}$  the representation of the Gaussian free field via the multivariate Gaussian  $(V_i^{[k]}(\mu) : i \in \mathcal{X})$  also has a natural probabilistic interpretation. Consider the transient continuous time Markov chain on  $\mathcal{X}$ , with absorption in state  $k$ , with the rate  $\pi_i(\mu)p_{ij}(\mu)$  of jumping from state  $i$  to state  $j$  for all  $i \neq k$ . Let  $g_{ij}^{[k]}(\mu)$ , for  $i, j \neq k$ , denote the mean time spent in state  $j$  before absorption. Then it can be checked that the matrix on  $\mathcal{X} \setminus \{k\}$  with entries  $g_{ij}^{[k]}(\mu)$  is a symmetric positive definite matrix. It is, indeed, the covariance matrix of  $(V_i^{[k]}(\mu) : i \neq k)$ . See [5,7] for more details.

Let us also observe that the recurrent continuous time Markov chain  $(\tilde{X}_t(\mu), t \geq 0)$  on  $\mathcal{X}$  with the rate of jumping from state  $i$  to state  $j$  being  $\pi_i(\mu)p_{ij}(\mu)$  for  $j \neq i$  satisfies

$$\frac{z_{ij}(\mu)}{\pi_j(\mu)} = \lim_{T \rightarrow \infty} \left( E_i \left[ \int_0^T 1(\tilde{X}_t(\mu) = j) dt \right] - T \right). \quad (13)$$

This is basically a consequence of (11) but is somewhat more subtle that it might seem.  $(X_t(\mu), t \geq 0)$  can be coupled to  $(\tilde{X}_t(\mu), t \geq 0)$  by creating the latter from the former by stretching out each duration of time spent in state  $i$  by the factor  $\pi_i(\mu)^{-1}$ , for each  $i \in \mathcal{X}$ . But then the integral to a fixed time  $T$  in (13) makes the corresponding integral in (11) be to a random time. Nevertheless, since we take the asymptotic limit in  $T$ , (13) follows from (11).

Now, the generalized second Ray-Knight theorem [7, Thm. 2.17] gives us the promised insight into the transient rates at which rewards are generated in the individual states. For  $i \in \mathcal{X}$  and  $t \geq 0$ , let  $L_{i,t}(\mu) := \int_0^t 1(\tilde{X}_s(\mu) = i) ds$ . For  $k \in \mathcal{X}$  and  $s \geq 0$  define

$$\Gamma_{k,s}(\mu) := \inf\{t \geq 0 : L_{k,t}(\mu) \geq s\},$$

which is the first time at which the time spent in state  $k$  by the process  $(\tilde{X}_t(\mu), t \geq 0)$  is at least  $s$ . We then have

$$\left( L_{i,\Gamma_{k,s}(\mu)}(\mu) + \frac{1}{2}(V_i^{[k]}(\mu))^2 : i \in \mathcal{X} \right) \stackrel{d}{=} \left( \frac{1}{2}(V_i^{[k]}(\mu) + \sqrt{2s})^2 : i \in \mathcal{X} \right) \quad (14)$$

for all  $s \geq 0$ , where  $\stackrel{d}{=}$  denotes equality in distribution of the vector random variables on each side. Here  $(V_i^{[k]}(\mu) : i \in \mathcal{X})$  is the Gaussian free field, as described earlier, and is assumed to be independent of  $(L_{i,\Gamma_{k,s}(\mu)}(\mu) : i \in \mathcal{X})$ , whose law is taken assuming that the process  $(\tilde{X}_t(\mu), t \geq 0)$  starts at  $k \in \mathcal{X}$ .

In (14) note that  $L_{i,\Gamma_{k,s}(\mu)}(\mu)$  denotes the time spent in state  $i$  by  $(\tilde{X}_t(\mu), t \geq 0)$  (which is assumed to start in state  $k$ ) up to the first time at which the time spent in state  $k$  by this process is at least  $s$  (in particular, we have  $L_{k,\Gamma_{k,s}(\mu)}(\mu) = s$ ).  $(X_t(\mu), t \geq 0)$  can be coupled to  $(\tilde{X}_t(\mu), t \geq 0)$ , as described above, and  $(X_t(\mu), t \geq 0)$  is just the continuous time Markov chain corresponding to the discrete sequence  $(X_n, n \geq 0)$  (started at state  $k$  and under the stationary Markov control strategy  $\mu \in \mathcal{M}$ ). It is in this sense that the distributional identity in (14) gives us detailed information into the transient rates at which rewards are generated in the individual states under the stationary Markov strategy  $\mu$ , for each initial state.

This unusual way in which Gaussians plays a role in the context of RMDP to give insight into the transient behavior of the generation of reward is quite striking. Our purpose in this paper has only been to highlight this connection. We leave the exploration of its implications to future research.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Data availability

No data was used for the research described in the article.

## Acknowledgments

This research was supported by National Science Foundation, USA grants CCF-1618145, CCF-1901004, CIF-2007965, and the NSF Science & Technology Center grant CCF-0939370 (Science of Information). The author would like to thank Devon Ding for several discussions centered around the monographs [5] and [7], and also for reading the completed paper for a sanity check.

## Appendix A. Proof of Lemma 1

It is well known that the set of occupation measures as  $\mu$  ranges over  $\mathcal{M}$  is a closed convex set and every extreme point of this convex set corresponds to the occupation measure of some  $\bar{\mu} \in \bar{\mathcal{M}}$  (see e.g. [8, Lemma 5.2]). Now suppose that  $P(\bar{\mu})$  is irreducible and reversible for each  $\bar{\mu} \in \bar{\mathcal{M}}$ . Every  $\mu \in \mathcal{M}$  can be expressed as a finite convex combination of  $(\bar{\mu} : \bar{\mu} \in \bar{\mathcal{M}})$ , and thus  $P(\mu)$  is expressed as the corresponding convex combination of  $(P(\bar{\mu}) : \bar{\mu} \in \bar{\mathcal{M}})$ . It follows that  $P(\mu)$  is irreducible.

But the occupation measure of  $P(\mu)$  is in the convex hull of the occupation measures  $(P(\bar{\mu}) : \bar{\mu} \in \bar{\mathcal{M}})$ , and these are symmetric by assumption, so the occupation measure of  $P(\mu)$  is symmetric and hence  $P(\mu)$  is reversible. (Note that the convex combination expressing the occupation measure of  $P(\mu)$  in terms of the occupation measures of stationary deterministic Markov control strategies may be different from that used above to express  $P(\mu)$  in terms of  $(P(\bar{\mu}) : \bar{\mu} \in \bar{\mathcal{M}})$ .) This proves that  $P(\mu)$  irreducible and reversible for each  $\mu \in \mathcal{M}$ . The claim in the opposite direction is obvious. This completes the proof of the lemma. ■

## Appendix B. Proof of Lemma 2

We claim that for every  $i \in \mathcal{X}$  the set of neighbors of  $i$  under  $(p_{ij}(u) : j \in \mathcal{X})$ , namely  $\{j \neq i : p_{ij}(u) > 0\}$  is the same for all  $u \in \mathcal{U}$ . Suppose, to the contrary, that for some distinct  $u, v \in \mathcal{U}$  and some  $i \in \mathcal{X}$  we have  $p_{ij}(u) > 0$  but  $p_{ij}(v) = 0$ , for some  $j \neq i$ . Pick some  $\tilde{u} \in \mathcal{U}$  (which could be either  $u$  or  $v$  if desired) and consider the two stationary deterministic Markov strategies  $\bar{\mu}^{(a)}$  and  $\bar{\mu}^{(b)}$  given by

$$\bar{\mu}^{(a)}(i) = u, \bar{\mu}^{(b)}(i) = v, \text{ and } \bar{\mu}^{(a)}(l) = \bar{\mu}^{(b)}(l) = \tilde{u} \text{ for } l \neq i. \quad (15)$$

Since  $P(\bar{\mu}^{(a)})$  is reversible and  $p_{ij}(u) > 0$  it follows that  $p_{ji}(\tilde{u}) > 0$ . But then, since  $P(\bar{\mu}^{(b)})$  is reversible, it would follow that  $p_{ij}(v) > 0$ , a contradiction. This establishes the claim. This also establishes the existence of a simple connected graph  $G := (\mathcal{X}, \mathcal{E})$  such that for all  $u \in \mathcal{U}$  and distinct  $i, j \in \mathcal{X}$  we have  $p_{ij}(u) > 0$  iff  $(i, j) \in \mathcal{E}$ , as claimed. ■

## Appendix C. Proof of Theorem 1

Suppose first that all the  $P(u)$  for  $u \in \mathcal{U}$  are the same, and let  $P = [p_{ij}]$  denote this common TPM over  $\mathcal{X}$ . Thus  $P$  is irreducible and reversible. Let  $(\pi_i : i \in \mathcal{X})$  denote the stationary distribution of  $P$ . We have  $\pi_i p_{ij} = \pi_j p_{ji}$  for all  $i, j \in \mathcal{X}$ .

For each  $i \in \mathcal{X}$  we can choose  $\rho(i, u) \in (0, 1]$  to be  $1 - p_{ii}$ . (Note that we have  $p_{ii} < 1$  since  $P$  is irreducible and  $|\mathcal{X}| \geq 2$ .)

We then let  $p_{ij}^{(0)} := \frac{p_{ij}}{1-p_{ii}}$  for  $i \neq j$ , with  $p_{ii}^{(0)} := 0$  for all  $i \in \mathcal{X}$ . It can be checked that this defines an irreducible TPM  $P^{(0)}$  on  $\mathcal{X}$  with  $p_{ij}^{(0)} > 0$  iff  $(i, j) \in \mathcal{E}$ , where  $G := (\mathcal{X}, \mathcal{E})$  denotes the canonical graph associated to this RMDP. It can be checked that the stationary distribution  $(\pi_i^{(0)} : i \in \mathcal{X})$  of  $P^{(0)}$  is given by  $\pi_i^{(0)} = K\pi_i(1-p_{ii})$ , where  $K$  is the normalizing constant. Further, we have  $\pi_i^{(0)} p_{ij}^{(0)} = \pi_j^{(0)} p_{ji}^{(0)}$  for all  $i, j \in \mathcal{X}$ , which establishes that  $P^{(0)}$  is reversible. This completes the proof in this case.

We may thus turn to the case when not all the  $P(u)$  are the same.

Suppose first that  $|\mathcal{X}| = 2$ , and write  $\mathcal{X} = \{1, 2\}$ . Then, for any  $u \in \mathcal{U}$ ,  $P(u)$  is irreducible and reversible iff we have both  $p_{12}(u) > 0$  and  $p_{21}(u) > 0$  (the corresponding stationary distribution is  $\left[ \frac{p_{21}(u)}{p_{12}(u)+p_{21}(u)}, \frac{p_{12}(u)}{p_{12}(u)+p_{21}(u)} \right]$ ). It can be checked that any collection  $\{P(u) : u \in \mathcal{U}\}$ , where each  $P(u)$  is irreducible and reversible, defines an RMDP (because  $P(\mu)$  will then be irreducible and reversible for each  $\mu \in \mathcal{M}$ ). We can then define  $P^{(0)} := \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ , with  $\rho(1, u) = p_{12}(u)$ , and  $\rho(2, u) = p_{21}(u)$ , thus establishing the main claim of the theorem in this case. Note that the graph associated to this RMDP is biconnected.

Having dealt with the case  $|\mathcal{X}| = 2$ , we may henceforth assume that  $|\mathcal{X}| \geq 3$ . Fix  $i \in \mathcal{X}$ . By the assumption that  $G$  is biconnected there must exist distinct  $j, k \in \mathcal{X}$  such that  $(i, j) \in \mathcal{E}$  and  $(i, k) \in \mathcal{E}$ . This means that for all  $u \in \mathcal{U}$  we have  $p_{ij}(u) > 0$  and  $p_{ik}(u) > 0$ . We claim that the ratio  $\frac{p_{ij}(u)}{p_{ik}(u)}$  does not depend on  $u$ . To see this, let  $u, v \in \mathcal{U}$  be distinct and pick some  $\tilde{u} \in \mathcal{U}$  (which could be either  $u$  or  $v$  if desired) and consider the two stationary deterministic Markov strategies  $\tilde{\mu}^{(a)}$  and  $\tilde{\mu}^{(b)}$  given as in (15). Write  $p_{lm}^{(a)}$  for  $p_{lm}(\tilde{\mu}^{(a)})$  for  $l, m \in \mathcal{X}$ , and  $\pi^{(a)}$  for the stationary distribution of  $P(\tilde{\mu}^{(a)})$ ; similarly for  $\tilde{\mu}^{(b)}$ . By the assumption that  $G$  is biconnected, there is a path in  $G$  from  $j$  to  $k$  that does not touch  $i$ , i.e. one can find a sequence  $(l_0, l_1, \dots, l_R)$  of vertices of  $G$ , where  $R \geq 1$ , with  $l_0 = j$ ,  $l_R = k$ ,  $l_r \neq i$  for  $0 \leq r \leq R$ , and such that  $(l_r, l_{r+1}) \in \mathcal{E}$  for  $0 \leq r \leq R-1$ . Then we have the equations

$$\pi_{l_r}^{(a)} p_{l_r l_{r+1}}^{(a)} = \pi_{l_{r+1}}^{(a)} p_{l_{r+1} l_r}^{(a)} \quad \text{and} \quad \pi_{l_r}^{(b)} p_{l_r l_{r+1}}^{(b)} = \pi_{l_{r+1}}^{(b)} p_{l_{r+1} l_r}^{(b)}, \quad (16)$$

for all  $0 \leq r \leq R-1$  (these follow from the reversibility of  $P(\tilde{\mu}^{(a)})$  and  $P(\tilde{\mu}^{(b)})$  respectively). Since for all  $l, m \in \mathcal{X}$  with  $l \neq i$  and  $m \neq i$  we have  $p_{lm}^{(a)} = p_{lm}^{(b)} = p_{lm}(\tilde{u})$ , we can conclude from the equations in (16) that

$$\frac{\pi_j^{(a)}}{\pi_k^{(a)}} = \frac{\pi_j^{(b)}}{\pi_k^{(b)}}. \quad (17)$$

But the reversibility of  $P(\tilde{\mu}^{(a)})$  and  $P(\tilde{\mu}^{(b)})$  also gives us the equations

$$\begin{aligned} \pi_i^{(a)} p_{ij}(u) &= \pi_j^{(a)} p_{ji}(\tilde{u}), \\ \pi_i^{(a)} p_{ik}(u) &= \pi_k^{(a)} p_{ki}(\tilde{u}), \\ \pi_i^{(b)} p_{ij}(v) &= \pi_j^{(b)} p_{ji}(\tilde{u}), \\ \pi_i^{(b)} p_{ik}(v) &= \pi_k^{(b)} p_{ki}(\tilde{u}). \end{aligned}$$

Dividing the first of these by the second (on each side) and the third of these by the fourth and comparing the resulting equations, using (17) we conclude that  $\frac{p_{ij}(u)}{p_{ik}(u)}$  equals  $\frac{p_{ij}(v)}{p_{ik}(v)}$ . Since  $u, v \in \mathcal{U}$ ,  $u \neq v$ , were arbitrarily chosen, we conclude that  $\frac{p_{ij}(u)}{p_{ik}(u)}$  does not depend on  $u$ , as claimed.

Now, for each  $i \in \mathcal{X}$ , pick an arbitrary  $u_i \in \mathcal{U}$  (all of these could be the same action, if one wishes). Having made such a choice, define  $\tilde{\mu} \in \tilde{\mathcal{M}}$  by  $\tilde{\mu}(i) = u_i$  for all  $i \in \mathcal{X}$ . Since  $P(\tilde{\mu})$  is irreducible and reversible, we have the equations  $\pi_i(\tilde{\mu}) p_{ij}(u_i) = \pi_j(\tilde{\mu}) p_{ji}(u_j)$  for all distinct  $i, j \in \mathcal{X}$ , where  $(\pi_i(\tilde{\mu}) : i \in \mathcal{X})$

denotes the stationary distribution of  $P(\tilde{\mu})$  as usual. For  $i \neq j$ , define  $p_{ij}^{(0)} := \frac{p_{ij}(u_i)}{1-p_{ii}(u_i)}$ , and let  $p_{ii}^{(0)} = 0$  for all  $i \in \mathcal{X}$ . Note that  $p_{ij}^{(0)} > 0$  iff  $(i, j) \in \mathcal{E}$ , where  $G = (\mathcal{X}, \mathcal{E})$  is the canonical graph associated to this RMDP. The resulting matrix  $P^{(0)}$  based on  $\mathcal{X}$  is an irreducible TPM with zero diagonal entries, and it is reversible with stationary distribution  $(K\pi_i(\tilde{\mu})(1-p_{ii}(u_i)) : i \in \mathcal{X})$ , where  $K$  is the proportionality constant. We can now set  $\rho(i, u) = 1-p_{ii}(u)$  for all  $(i, u) \in \mathcal{X} \times \mathcal{U}$ . Indeed, we have already proved that the  $(p_{ij}(u) : j \neq i)$  for  $u \in \mathcal{U}$  are proportional (for fixed  $i \in \mathcal{X}$ ), and so we will have  $p_{ij}(u) = p_{ij}(u_i) \frac{1-p_{ii}(u)}{1-p_{ii}(u_i)}$  for all  $(i, j) \in \mathcal{E}$ , which gives  $\rho(i, u) p_{ij}^{(0)} = p_{ij}(u)$  for all  $u \in \mathcal{U}$  and all distinct  $i, j \in \mathcal{X}$ . Note that we have  $\rho(i, u) \in (0, 1]$  for all  $(i, u)$ , as required.

This concludes the proof of the theorem. ■

#### Appendix D. Proof of Lemma 3

For all  $\mu \in \mathcal{M}$  we have  $p_{ij}(\mu) > 0$  iff  $(i, j) \in \mathcal{E}$ , and so  $P(\mu)$  is an irreducible TPM on  $\mathcal{X}$ . For  $i \in \mathcal{X}$  and  $k \neq i$  define  $p_{k \rightarrow i}(\mu)$  to be  $p_{kj}(\mu)$ , where  $j \in \mathcal{X}$  is defined as the vertex adjacent to  $k$  on the unique path from  $k$  to  $i$  in the tree. It can be checked that the stationary distribution of  $P(\mu)$  is proportional to  $(\prod_{k \neq i} p_{k \rightarrow i}(\mu) : i \in \mathcal{X})$  and so  $P(\mu)$  is reversible. This concludes the proof. ■

#### Appendix E. Proof of Theorem 2

If there is only one block then we are in biconnected case covered in Theorem 1, where we have already proved that the structure of the RMDP must be consistent with the type described in Example 3. We may therefore assume that there are at least two blocks, and so every block has at least one articulation point. For each block  $B \in \mathcal{B}$  an argument similar to that in Theorem 1 shows that for each  $i \in B$  the  $(p_{ij}(u) : j \in B)$  as  $u$  ranges over  $\mathcal{U}$  are all proportional. We can therefore find a TPM  $P^{(0)}(B) = \left[ p_{ij}^{(0)}(B) \right]$  on  $B$ , with zero diagonal entries, such that  $p_{ij}(u) = \left( \sum_{k \in B} p_{ik}(u) \right) p_{ij}^{(0)}(B)$  for all distinct  $i, j \in B$ . Since  $p_{ij}^{(0)}(B) > 0$  iff  $(i, j)$  is an edge in  $B$  (viewed as a subgraph), and since  $B$  is connected, we see that  $P^{(0)}(B)$  is irreducible. Define  $\rho(i, u)$  to be  $\sum_{j \in B} p_{ij}(u)$  for  $i \in B$  (if any) and, for each articulation point  $a \in B$ , define  $\rho(a, u)$  to be  $\sum_{j \in \mathcal{X}, j \neq a} p_{aj}(u)$  (this quantity does not depend on which  $B$  containing  $a$  is being considered), and define  $v_a(u, B)$  to be  $\frac{\sum_{j \in B} p_{ij}(u)}{\rho(a, u)}$ . Note that the  $v_a(u, B)$  are strictly positive and  $\sum_{B \ni a} v_a(u, B) = 1$ , as required. Also note that  $\rho(i, u) \in (0, 1]$  for all  $(i, u) \in \mathcal{X} \times \mathcal{U}$ .

It remains to show that each  $P^{(0)}(B)$  is reversible. Pick any  $u \in \mathcal{U}$ . Let  $(\pi_i(u) : i \in \mathcal{X})$  denote the stationary distribution of  $P(u)$ . Fix  $B \in \mathcal{B}$ . By the reversibility of  $P(u)$  we have  $\pi_i(u) p_{ij}(u) = \pi_j(u) p_{ji}(u)$  for all  $i, j \in B$ . It follows that

$$\left( \pi_i(u) \sum_{k \in B} p_{ik}(u) \right) p_{ij}^{(0)} = \left( \pi_j(u) \sum_{k \in B} p_{jk}(u) \right) p_{ji}^{(0)}$$

This means that if  $(\psi_i(B) : i \in B)$  denotes the stationary distribution of  $P^{(0)}(B)$  then it is proportional to  $(\pi_i(u) \sum_{k \in B} p_{ik}(u) : i \in B)$  and thus that  $\psi_i(B) p_{ij}^{(0)}(B) = \psi_j(B) p_{ji}^{(0)}(B)$  for all  $i, j \in B$ , which establishes that  $P^{(0)}(B)$  is reversible. This concludes the proof. ■

#### Appendix F. Proof of Theorem 3

Let  $(\pi_i^{(0)} : i \in \mathcal{X})$  denote the stationary distribution of  $P^{(0)}$ , and recall that for any  $\mu \in \mathcal{M}$  the stationary distribution of  $P(\mu)$  is

given by  $(K(\mu) \frac{\pi_i^{(0)}}{\rho(i, \mu(i))} : i \in \mathcal{X})$ , where  $K(\mu) := \left( \sum_i \frac{\pi_i^{(0)}}{\rho(i, \mu(i))} \right)^{-1}$  is the normalizing constant. Since

$$\beta(\mu) = \sum_i r(i, \mu(i)) \pi_i(\mu) = K(\mu) \sum_i r(i, \mu(i)) \frac{\pi_i^{(0)}}{\rho(i, \mu(i))},$$

we get

$$\sum_i \frac{r(i, \mu(i)) - \beta(\mu)}{\rho(i, \mu(i))} \pi_i^{(0)} = 0.$$

Thus  $\beta(\mu)$  can be characterized as

$$\beta(\mu) = \sup\{\beta \in \mathbb{R} : \sum_i \frac{r(i, \mu(i)) - \beta}{\rho(i, \mu(i))} \pi_i^{(0)} \geq 0\}.$$

If we can find  $i \in \mathcal{X}$  for which Eq. (9) holds, then pick  $u \in \mathcal{U}$  achieving the argmax on the RHS of Eq. (9) and let  $\bar{\mu}^{(k+1)}(i) = u$  and  $\bar{\mu}^{(k+1)}(j) = \bar{\mu}^{(k)}(j)$  for all  $j \neq i$ , as in the simplified policy iteration algorithm. We then have

$$\begin{aligned} \sum_j \frac{r(j, \bar{\mu}^{(k+1)}(j)) - \beta(\bar{\mu}^{(k)})}{\rho(j, \bar{\mu}^{(k+1)}(j))} \pi_j^{(0)} &= \sum_{j \neq i} \frac{r(j, \bar{\mu}^{(k)}(j)) - \beta(\bar{\mu}^{(k)})}{\rho(j, \bar{\mu}^{(k)}(j))} \pi_j^{(0)} \\ &\quad + \frac{r(i, u) - \beta(\bar{\mu}^{(k)})}{\rho(i, u)} \pi_i^{(0)} \\ &> \sum_j \frac{r(j, \bar{\mu}^{(k)}(j)) - \beta(\bar{\mu}^{(k)})}{\rho(j, \bar{\mu}^{(k)}(j))} \pi_j^{(0)} \\ &= 0. \end{aligned}$$

It follows that  $\beta(\bar{\mu}^{(k+1)}) > \beta(\bar{\mu}^{(k)})$ , which concludes the proof.

## Appendix G. Proof of Theorem 4

Define  $p_{ij}^{(0)}(\bar{\mu}|_{\mathcal{A}})$  to be  $v_i(\bar{\mu}(i), B) p_{ij}^{(0)}(B)$  for  $i, j \in B$ ,  $i \neq j$ , for each  $B \in \mathcal{B}$ . Here we recall that we defined  $v_i(u, B) = 1$  for

all  $i \in \hat{B}$  and  $u \in \mathcal{U}$ , and so we realize that  $p_{ij}^{(0)}(\bar{\mu}|_{\mathcal{A}})$  depends only on the restriction of  $\bar{\mu}$  to the articulation nodes, which is indicated by the notation  $\bar{\mu}|_{\mathcal{A}}$ . It is straightforward to check that the  $p_{ij}^{(0)}(\bar{\mu}|_{\mathcal{A}})$  define a TPM on  $\mathcal{X}$ . Let  $(\pi_i^{(0)}(\bar{\mu}|_{\mathcal{A}}) : i \in \mathcal{X})$  denote the stationary distribution associated to this TPM. It is straightforward to check that the stationary distribution of  $P(\bar{\mu})$  is proportional to  $(\frac{\pi_i^{(0)}(\bar{\mu}|_{\mathcal{A}})}{\rho(i, \bar{\mu}(i))} : i \in \mathcal{X})$ . Further, we can check that for all  $\bar{\eta} \in \bar{\mathcal{M}}$  such that  $\bar{\eta}|_{\mathcal{A}} = \bar{\mu}|_{\mathcal{A}}$  the stationary distribution of  $P(\bar{\mu})$  is proportional to  $(\frac{\pi_i^{(0)}(\bar{\mu}|_{\mathcal{A}})}{\rho(i, \bar{\eta}(i))} : i \in \mathcal{X})$ .

From this, as in the proof of Theorem 3, we can check that for all  $\bar{\eta} \in \bar{\mathcal{M}}$  such that  $\bar{\eta}|_{\mathcal{A}} = \bar{\mu}|_{\mathcal{A}}$  we have the characterization

$$\beta(\bar{\eta}) = \sup\{\beta \in \mathbb{R} : \sum_i \frac{r(i, \bar{\eta}(i)) - \beta}{\rho(i, \bar{\eta}(i))} \pi_i^{(0)}(\bar{\mu}|_{\mathcal{A}}) \geq 0\}.$$

The rest of the proof then follows as in the proof of Theorem 3, allowing us to conclude the desired strict inequality. ■

## References

- [1] Randy Cogill, Cheng Peng, Reversible Markov decision processes with an average-reward criterion, *SIAM J. Control Optim.* 51 (1) (2013) 402–418.
- [2] Reinhard Diestel, *Graph Theory*, fifth ed., Springer, 2017.
- [3] David Aldous, James Allen Fill, Reversible Markov chains and random walks on graphs. Unfinished monograph (2002). Recompiled version, 2014, <https://www.stat.berkeley.edu/users/aldous/RWG/book.pdf>.
- [4] Svante Janson, Gaussian Hilbert Spaces, in: *Cambridge Tracts in Mathematics*, vol. 129, Cambridge University Press, 1997.
- [5] Yves Le Jan, Markov Path, Loops and Fields. *École d'été de Probabilités de Saint-Flour, XXXVIII*, Springer, 2008, 2011.
- [6] Russell Lyons, Yuval Peres, *Probability on Trees and Networks*, Cambridge University Press, 2016.
- [7] Alain-Sol Sznitman, *Topics in Occupation Times and Gaussian Free Fields*. Notes of the Course Special Topics in Probability at ETH Zurich, Spring, 2011.
- [8] Aristotle Arapostathis, Vivek S. Borkar, Emmanuel Fernández-Gaucherand, Mrinal K. Ghosh, Steven I. Marcus, Discrete-time controlled Markov processes with average cost criterion: A survey, *SIAM J. Control Optim.* 31 (2) (1993) 282–344.