


RESEARCH ARTICLE

The value of information and efficient switching in channel selection

Jiesen Wang¹ , Yoni Nazarathy² and Thomas Taimre²

¹ School of Mathematics and Statistics, The University of Melbourne, Melbourne, VIC, Australia. E-mail: jiesenwang@gmail.com

² School of Mathematics and Physics, The University of Queensland, Saint Lucia, QLD, Australia.

Keywords: Applied probability, Operations research, Stochastic modeling

Abstract

We consider a collection of statistically identical two-state continuous time Markov chains (channels). A controller continuously selects a channel with the view of maximizing infinite horizon average reward. A switching cost is paid upon channel changes. We consider two cases: full observation (all channels observed simultaneously) and partial observation (only the current channel observed). We analyze the difference in performance between these cases for various policies. For the partial observation case with two channels or an infinite number of channels, we explicitly characterize an optimal threshold for two sensible policies which we name “call-gapping” and “cool-off.” Our results present a qualitative view on the interaction of the number of channels, the available information, and the switching costs.

1. Introduction

Many scenarios in mining, finance, telecommunications, medical research, and other fields involve a situation where the reward collected from a resource fluctuates over time in a stochastic manner. For example, the yield obtained from mineral resource extraction varies as digging persists; the returns from financial investments vary based on many random market effects; the communication bit rate of communication channels varies based on physical conditions; and the findings in medical trials vary over time. Such a theme of randomly varying rewards, or a randomly modulated reward rate, reappears in multiple application areas. Hence designing, managing, and controlling such uncertain situations has been a focal point of stochastic operations research in a variety of contexts. The theory of restless bandits presents one general paradigm for dealing with such problems. See, for example, Gittins et al. [8], Whittle [21], and Weber and Weiss [20] for the general restless bandits problem. The main theme in such research is the efficient selection of channels/projects/arms over time.

In this paper, we add to the body of literature by considering problems with switching costs. Switching costs have been considered in Agrawal et al. [3], Banks and Sundaram [6], and Dusonchet and Hongler [7] in the context of multi-armed bandit processes, but to the best of our knowledge have not been studied as we do here. The general problem which we discuss is one in which a resource yields random rewards over time, with some known average reward rate and a maximal reward rate. One way to improve performance is to introduce additional independent instances of this resource and to consider a situation where at any time we use a resource of our (dynamic) choice. By doing so, we are potentially able to increase the obtained reward rate from the average reward rate towards the maximal reward rate.

As alluded to above, dynamic resource allocation problems span a variety of applications. However, one key area in which decisions often need to be made at the millisecond or second timescale is communication networks. Specifically dynamic resource allocation problems dealing with opportunistic

scheduling as in Aalto et al. [1] and further in [2] have recently attracted attention. In such a setting, wireless cellular systems utilize random channel quality variations to minimize the flow-level holding costs. Some of this research has utilized the restless bandits formulation, as in Jacko and Villar [9] (and further in [1,2]). Related is the continuous time paradigm with impulsive control as in Ayesta et al. [5]. Switching costs often need to be taken into account and this is the key contribution of our current paper since the aforementioned works did not take switching costs directly into account.

Other related work in the context of communication networks is in the domain of handover problems. See, for example, the survey Wang et al. [18] as well as Mezzavilla et al. [17] where general problems dealing with handover and their solution via Markov Decision Processes (MDPs) are formulated. A key aspect of this domain is the potential lack of information with regard to channel quality. This introduces interesting problems and has attracted much work for which a survey describing contributions up to 2015 is in Kuhn and Nazarathy [11]. In more recent years, several additional notable works in which channel information is not fully available are Kaza et al. [10], Larrnaaga et al. [12], and further in Larrañaga et al. [13]. See also recent work dealing with minimizing the age of information as Maatouk et al. [15] and Meshram and Kaza [16].

As an abstract example of a communication link, assume a link that yields an average bit rate of 4 Mbit/s and a maximal bit rate of 10 Mbit/s where the actual instantaneous bit rate varies stochastically over time and achieves the maximal bit rate for random finite durations. By introducing multiple instances of such a link and allowing the system to switch between the instances without constraint, we are able to get an effective average bit rate that exceeds 4 Mbit/s and potentially nears the maximal bit rate of 10 Mbit/s. The key is clearly to use the right instance of the channel at “the right time,” so as to on-average use channels that do better than the average bit rate. At the extreme positive case, by introducing an increasing number of instances and assuming their stochastic behavior is independent between instances, we can get arbitrarily close to the 10 Mbit/s bit rate. The other extreme is not switching between channels at all and settling for the average bit rate of 4 Mbit/s, obtained from a single channel.

While the usage of such redundant channels can be of benefit, it is clearly not without cost. First, there are the structural costs of setting up additional channels. The exact nature of these costs depends on the application and is not our focus. Then there are costs associated with setting up systems for gaining information about the instantaneous state of all channels. Finally, there are dynamic instantaneous switching costs involved when switching between channels. In this paper, we study the tradeoffs and costs associated with this problem using the simplest example model that we could consider. The elegance of our simple model is that it captures the value of real-time information and at the same time allows us to compare how the addition of more channels to the system increases rewards. Real systems are bound to be more complex than the model which we present; however, the results from our model capture the essential tradeoffs that one can expect.

After reparameterization, our model is that each of n channels yields instantaneous rewards of either 0 or 1 where switching between the reward states is according to a two-state continuous time Markov chain. The long-term average reward obtained by a channel is $\gamma \in (0, 1)$ and the instantaneous cost of switching between channels is c . We distinguish between a case of *full observation* where the state of all channels is observable, and a case of *partial observation* where only the state of the current channel is available. For each of these cases, we suggest parameterized channel selection policies, and we are able to analyze and optimize the parameters of these policies in certain situations.

An illustration of the performance of such a model for the case of $\gamma = 0.4$ is in Figure 1. There are two threshold values at $\gamma = 0.4$ and at $\gamma^2 = 0.16$. We see that in the full observation case, when $c > \gamma$, it is not worthwhile to use additional channels and otherwise the total rewards increases in an affine manner as $c \rightarrow 0$. Further, we see that in the partial observation case, usage of redundant channels is only worthwhile if $c < \gamma^2$. In this case, the total reward rate increases in a nonlinear manner as $c \rightarrow 0$. In this case, we compare two policies that are described in the sequel, namely *call-gapping* and *cool-off*. While we do not have optimality proofs for these policies within the class of all policies, we show that these policies agree in performance for $n = 2$ and $n \rightarrow \infty$. We further obtain explicit results for the case of $n = 2$ and at $n = \infty$. Note that in all cases, as the number of channels increases, the achievable

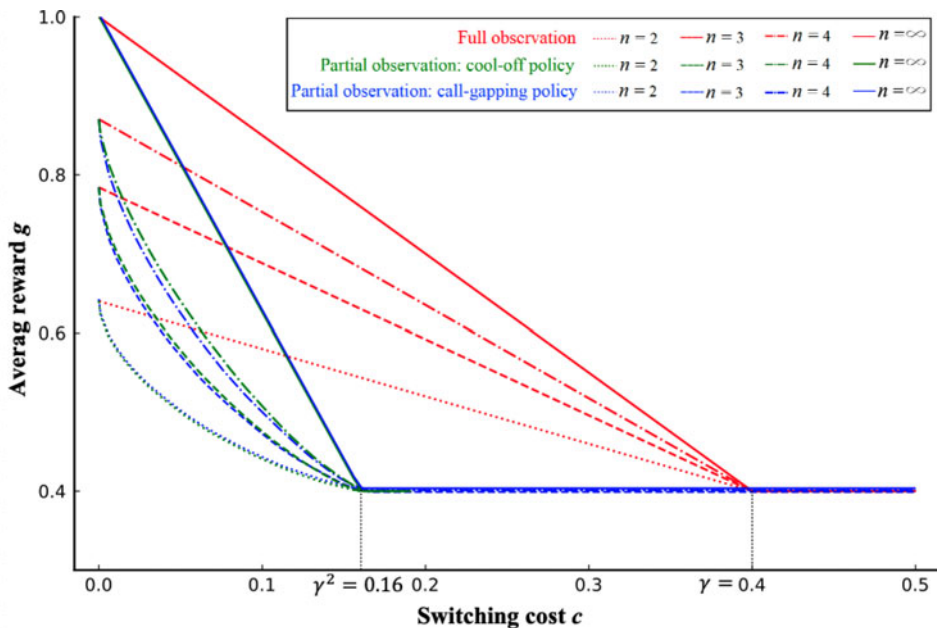


Figure 1. The optimal average reward for a $\gamma = 0.4$ system with a varying number of channels n . The red curves are for a case of full observation. The green curves are for a case of partial observation with a cool-off policy. The blue curves are for the case of partial observation with a call-gapping policy. Note that for $n = 2$ and $n = \infty$, both of the latter policies are identical.

reward rate increases towards the maximal reward for small switching costs c . Our results in this paper help understand the tradeoffs between the number of channels, switching costs, information, and policy qualitatively.

The remainder of the paper is structured as follows: In Section 2, we present the model and summarize the main results. In Section 3, we further study the full observation case by considering simple steady state results, the Whittle Index, and the optimality. In Section 4, we derive results associated with the case of partial observations. In Section 5, we present numerical results, and finally, we conclude in Section 6.

2. Model and summary of results

We consider n statistically identical channels $i = 1, 2, \dots, n$, with state $X_i(t) \in \{0, 1\}$, evolving independently according to a continuous-time, time-homogenous, two-state Markov chain model with generator

$$Q = \begin{bmatrix} -\lambda & \lambda \\ \mu & -\mu \end{bmatrix},$$

where $\lambda, \mu > 0$.

At any given time, the controller may only use one of the channels. The controller's choice is indicated through $U(t) \in \{1, \dots, n\}$, where $U(t) = i$ means that channel i is being used at time t . There is a fixed switching cost of $\kappa \geq 0$ for every instant at which $U(t)$ is changed; that is, every time t where $U(t^-) \neq U(t)$, an additional cost of κ is incurred.

When $X_i(t) = x \in \{0, 1\}$, the reward rate of that channel is denoted by $r(x)$, a function assumed identical for all channels. The goal of the controller is then to maximize the average reward,

$$g = \liminf_{t \rightarrow \infty} \frac{\mathbb{E} \left[\int_0^t r(X_{U(s)}(s)) ds \right] - \kappa N(t)}{t},$$

where $N(t)$ denotes the number of channel changes during the time interval $[0, t]$.

By observing that

$$\int_0^t r(X_{U(s)}(s)) ds = t r(0) + (r(1) - r(0)) \int_0^t X_{U(s)}(s) ds,$$

w.l.o.g., we are able to set $r(0) = 0$ and $r(1) = \zeta > 0$. Further, by rescaling time by a factor of $\eta = \lambda^{-1}$ and setting ζ/η to be 1, we parameterize the channel model only by $\gamma = \lambda/(\lambda + \mu)$ (the stationary probability of a channel being in state 1) and $c = \kappa/\eta$.

In what follows, we write $I(t) = X_{U(t)}(t)$, which denotes the state of the selected channel at time t . Then, the average reward g corresponding to the parametrization just outlined is expressed as:

$$g = \liminf_{t \rightarrow \infty} \frac{(\zeta/\eta) \int_0^t \tilde{I}(u) du - (\kappa/\eta) \tilde{N}(t)}{t} = \liminf_{t \rightarrow \infty} \frac{\int_0^t \tilde{I}(u) du - c \tilde{N}(t)}{t},$$

where $\tilde{I}(t) = I(\eta t)$ and $\tilde{N}(t) = N(\eta t)$.

Hence our model parameters are $\gamma \in (0, 1)$ and $c \geq 0$ with a rescaled generator

$$Q = \begin{bmatrix} -1 & 1 \\ 1/\gamma - 1 & -(1/\gamma - 1) \end{bmatrix}.$$

We distinguish two situations:

Full Observation (Case I): The controller fully observes $\{X_1(t), \dots, X_n(t)\}$.

Partial Observations (Case II): When $U(t) = i$ the controller only observes $X_i(t)$.

In considering this partial observation case, it is useful to consider belief states for $i = 1, \dots, n$, via

$$\omega_i(t) = \mathbb{P}(X_i(t) = 1 \mid \text{All information observed up to time } t),$$

and in constructing the belief states, it is useful to denote the *last time in channel* i via,

$$\mathcal{T}_i(t) = \sup\{t' \leq t : U(t') = i\}.$$

This allows the construction of the belief state of channel i via

$$\omega_i(t) = p(t - \mathcal{T}_i(t); X_i(\mathcal{T}_i(t))),$$

where $p(t; x) := \mathbb{P}(X_i(t) = 1 \mid X_i(0) = x)$ which can be represented explicitly as

$$p(t; x) = \begin{cases} \gamma(1 - e^{-t/\gamma}), & x = 0, \\ \gamma + (1 - \gamma)e^{-t/\gamma}, & x = 1. \end{cases} \tag{1}$$

Observe that if $U(t) = i$ then $\mathcal{T}_i(t) = t$ and then $\omega_i(t) = X_i(t)$ which is either 0 or 1.

Control policies: We consider several policies all of which limit channel switching to times t when $X_i(t) = 0$. This means that in the partial observation case (II), the belief state at any time t for $i \neq U(t)$ is strictly monotonically increasing in time according to the first expression in (1). In both case I and case II, one trivial policy denoted $\pi^{(s)}$ is the policy of never switching channel. Further, in case I, there

is the policy $\pi^{(a)}$ which switches to an arbitrary channel $j \neq U(t)$ if $X_j(t) = 1$ and $X_{U(t)} = 0$. That is, this policy ensures that the current channel is in state 1 if such a channel exists, and further this policy does not switch excessively if not needed.

In case II, in addition to the policy $\pi^{(s)}$, we consider a *call-gapping policy*, $\pi^{(\tau)}$, and a *cool-off policy* $\pi^{(\sigma)}$. In brief, call-gapping does not switch out of a channel prior to a call-gapping duration, parameter $\tau > 0$. Such a policy was analyzed in a different context in Lin and Ross [14] where it was termed call-gapping. Further, in brief, cool-off does not switch into a channel prior to a cooling-off period, parameter $\sigma > 0$.

Like all our policies, call-gapping does not switch out of a good channel; however, if the current channel $U(t)$ is in bad state, i.e. $X_{U(t)} = 0$ then this policy switches to the channel with the highest belief state as long as the time since the last switch is not less than τ . That is, at time t denote the *last switching time* via,

$$\mathcal{T}^\ell(t) = \sup\{t' \leq t : U(t'^-) \neq U(t')\},$$

then the call-gapping policy switches only if $X_{U(t)} = 0$ and $t - \mathcal{T}^\ell(t) \geq \tau$. Note that in our context, since the channels are homogenous and $\omega_i(t)$ is monotonically increasing, when the call-gapping policy yields a switch then it is to a channel that was not used for the longest duration. This naturally yields *round-robin* behavior as w.l.o.g., we switch from channel i to channel $i + 1$ (modulo n).

As with the call-gapping policy, the cool-off policy does not switch out of a good channel and once switching occurs we switch with a round-robin manner, i.e. to the channel that has been visited longest ago. However, in contrast to the call-gapping policy, we choose to switch into a channel only if it has not been visited for a time that is at least the cool-off parameter σ . This essentially means that the cool-off policy considers the belief states of all channels and switches into a channel i only if $U(t) = 0$ and $\omega_i(t)$ exceeds the threshold $p(\sigma ; 0)$. Note that unlike the call-gapping policy, the cool-off policy may potentially incur multiple instantaneous switches. Also note that when $n = 2$, the call-gapping and cool-off policies are identical. Both the call-gapping and cool-off policies ensure a strict upper limit on the switching costs, $c\tilde{N}(t)/t$, as with these policies, almost surely,

$$\limsup_{t \rightarrow \infty} c \frac{\tilde{N}(t)}{t} \leq \frac{c}{\tau}, \quad \text{and} \quad \limsup_{t \rightarrow \infty} c \frac{\tilde{N}(t)}{t} \leq \frac{nc}{\sigma}.$$

Our results for these policies are as follows. First for case I, the performance of the system under $\pi^{(s)}$ and $\pi^{(a)}$ is tractable, for any value of n . This also allows us to characterize, when each of these policies is preferable as a function of the switching cost c and the parameter γ .

In contrast, for case II, analysis is more difficult. When $n = 2$, we are able to characterize the optimal call-gapping parameter τ^* . For finite $n > 2$, we have not found such an analytic characterization. However, we can consider systems with large n . For this, we construct an approximate model, that we label with $n = \infty$. Such a system has only two channels, say $i = 1$ and $i = 2$. When $U(t) = 1$, then $X_2(t)$ is assumed to be in the steady state. Similarly when $U(t) = 2$ then $X_1(t)$ is assumed to be in the steady state. That is, the current channel $X_{U(t)}$ behaves normally; however, at the moment of switching to the other channel, the state of that channel is drawn from the steady state distribution $[(1 - \gamma) \ \gamma]$, and the evolution continuous. Under a round-robin-based policy such as our $\pi^{(\tau)}$ and $\pi^{(\sigma)}$, with large n , the $n = \infty$ model approximates the system because we can expect the next channel in the round-robin to be at approximate steady state.

The nature of our analytic results is in comparing the different policies (and their parameters) for Case I, and Case II, and for different values of n . That is, for Case I, we determine when $\pi^{(s)}$ is preferable to $\pi^{(a)}$ and vice versa. We also determine the value of the long-term reward g . Similarly, for Case II (under $n = 2$ and $n = \infty$), we compare $\pi^{(\tau)}$, $\pi^{(\sigma)}$, and $\pi^{(s)}$, determine the optimal parameters for τ (or σ), and obtain expressions for g .

The following two theorems are proved in Sections 3 and 4, respectively.

Theorem 2.1. *In Case I, the optimal choice between $\pi^{(a)}$ and $\pi^{(s)}$ is given by:*

$$\pi^* = \begin{cases} \pi^{(a)}, & c < \gamma, \\ \pi^{(s)}, & c \geq \gamma. \end{cases}$$

Further, under π^* , the optimal expected average reward is given by:

$$g^* = \begin{cases} 1 - (1 - \gamma)^n - c \frac{1 - \gamma - (1 - \gamma)^n}{\gamma} & c < \gamma, \\ \gamma & c \geq \gamma. \end{cases}$$

Moreover, for $n = 2$, the policy π^* is an optimal policy.

We also conjecture that in Case I, π^* is the optimal policy for any n . However, we are only able to prove this using computer algebra systems (such as Mathematica) for small n by doing symbolic computations. Additionally, as we show in Section 3, π^* also corresponds to the Whittle Index-based policy.

A similar partial result is found in Case II: for any n , if $c > \gamma^2$ then the optimal policy is $\pi^{(s)}$ when compared against $\pi^{(\tau)}$ and $\pi^{(\sigma)}$. We now have,

Theorem 2.2. *For Case II, we consider two scenarios. Firstly, restrict the policy space to be the set of call-gapping policies and denote the optimal policy therein by π_C^* . Secondly, restrict the policy space to be the set of all cool-off policies and denote the optimal policy therein by π_D^* . Then we have for $n \geq 2$:*

$$\pi_C^* = \begin{cases} \pi^{(\tau^*)}, & c < \gamma^2, \\ \pi^{(s)}, & c \geq \gamma^2, \end{cases} \quad \pi_D^* = \begin{cases} \pi^{(\sigma^*)}, & c < \gamma^2, \\ \pi^{(s)}, & c \geq \gamma^2. \end{cases}$$

When $c < \gamma^2$ and $n = 2$, the optimal call-gapping time, τ^* is the unique non-negative solution of the equation

$$e^{2\tau^*/\gamma}(\gamma^2 - c)(\gamma - 2) + 2e^{\tau^*/\gamma}\gamma(\gamma - \tau^*(\gamma - 1)) - \gamma(\gamma^2 - c) = 0. \tag{2}$$

The optimal cool-off level is given by $\sigma^* = \tau^*$. When $c < \gamma^2$, $n = \infty$,

$$\sigma^* = \tau^* = 0.$$

Further with π_C^* or π_D^* , the optimal expected average reward is given by:

$$g_C^* = g_D^* = \begin{cases} \frac{A_1(\gamma, \tau^*) - cA_2(\gamma, \tau^*)}{A_3(\gamma, \tau^*)} & c < \gamma^2, n = 2, \\ 1 - c(1 - \gamma)/\gamma^2 & c < \gamma^2, n = \infty, \\ \gamma & c \geq \gamma^2, \end{cases}$$

where,

$$A_1(\gamma, \tau) = e^{2\tau/\gamma}((\tau - 1)\gamma^3 - (3\tau - 2)\gamma^2 + 2\tau\gamma) - 2e^{\tau/\gamma}\gamma^2(1 - \gamma)^2 + 2\gamma^4 + (\tau - 3)\gamma^3 - \tau\gamma^2,$$

$$A_2(\gamma, \tau) = (\gamma - 1)(e^{2\tau/\gamma}(\gamma - 2) + \gamma),$$

$$A_3(\gamma, \tau) = \gamma^3 + (\tau - 2)\gamma^2 - \tau\gamma - e^{2\tau/\gamma}(\gamma^3 - (\tau + 2)\gamma^2 + 3\tau\gamma - 2\tau).$$

Note that the expression for g_C^* or g_D^* as in Theorem 2.2 can be used to evaluate the performance for $n = 2$ for any value of τ (alt. σ), by replacing τ^* in the expression with τ (alt. σ). Also note that empirical evidence, see Figure 1, suggests that when $2 < n < \infty$ the optimal cool-off policy has a better expected average reward than the optimal call-gapping policy.

3. On the Whittle index and optimality for case I

We now prove Theorem 2.1 in two parts; first the expression for g^* is derived by using a continuous-time birth and death process. Then we prove that π^* is the optimal policy for $n = 2$, describe symbolic computations for proving it is optimal for higher order n , and conjecture it is optimal for any n . We then move on to the Whittle index associated with Case I.

Proof of correctness of g^ expression for Theorem 2.1.* We consider the process $\{M(t), t \geq 0\}$ on the state space $\{0, 1, \dots, n\}$ where,

$$M(t) = \sum_{i=1}^n X_i(t).$$

The fact that the channels are independent and satisfy the same probability law, implies that $M(t)$ is a birth and death continuous time Markov chain, with birth and death rates

$$\lambda_i = n - i, \quad \mu_i = i \left(\frac{1}{\gamma} - 1 \right), \tag{3}$$

for $i \in \{0, \dots, n - 1\}$ and $i \in \{1, \dots, n\}$, respectively. Such a birth and death process is known as an Ehrenfest model with binomial stationary distribution,

$$\varrho_i = \binom{n}{i} \gamma^i (1 - \gamma)^{n-i} \quad i = 0, \dots, n.$$

Under the policy $\pi^{(a)}$, a reward rate of 1 is accrued at times t during which $M(t) > 0$. Further, switching costs are incurred at certain transition times of $M(t)$. Specifically at times t when $M(t^-) = 0$ and $M(t) = 1$, a switching cost c is incurred with probability $(n - 1)/n$. This is the probability that the channel that changes to ‘good’ is not the currently selected channel and therefore an immediate channel change takes place. Similarly, at times t when $M(t^-) = k \geq 2$ and $M(t) = k - 1$, a cost of c is incurred with probability $1/k$. This is the probability of the current channel turning ‘bad’ and hence requiring a switch. Now considering the reward as a Markov reward process, we have that average under policy $\pi^{(a)}$ is,

$$\begin{aligned} g &= (1 - \varrho_0) - c \frac{n - 1}{n} \lambda_0 \varrho_0 - c \sum_{k=2}^n \frac{1}{k} \mu_k \varrho_k \\ &= 1 - (1 - \gamma)^n - c \frac{1 - \gamma - (1 - \gamma)^n}{\gamma}, \end{aligned}$$

where the second line follows from the structure of the birth and death rates (3). Further with $\pi^{(s)}$, we have that g is trivially γ . Hence $\pi^{(a)}$ is preferable only when $c \leq \gamma$. □

To consider the optimality of π^* , the processes $\{X_i(t)\}_{i=1}^n$ and $U(t)$ can be transformed to $\tilde{W}(t)$ and $\tilde{S}(t)$ where $\tilde{W}(t) = X_{U(t)}(t)$ is the state of the current channel and,

$$\tilde{S}(t) = \sum_{i \neq U(t)} X_i(t),$$

is the number of other channels that are in the good state. The process

$$\tilde{X}(t) = (\tilde{W}(t), \tilde{S}(t)),$$

can then be represented as part of a continuous time MDP with singular controls, similar to the framework used in Ayesta et al. [5]. The state space is $\{0, 1\} \times \{0, \dots, n - 1\}$. The action process $\{\tilde{A}(t)\}$ takes

values in the action space $\{0, 1\}$ where $\tilde{A}(t) = 0$ indicates staying on a channel at time t . Further at singular time points, t in which $\tilde{A}(t) = 1$, an instantaneous channel switch is made (and immediately at t^+ , $\tilde{A}(t^+) = 0$). In the construction of this process, we assume that if a switch is made and there is another channel available of a different state to $\tilde{W}(t)$, then we switch into such an arbitrary channel and this implies state change. Otherwise (in the border cases in which a switch is made and there is not a different channel), the state is not changed.

To construct the continuous time MDP with singular controls, we use $I^{(a)}(w, s)$ as an indicator if transitions out of state $(\tilde{W}(t), \tilde{S}(t)) = (w, s)$ are singular or not for action $a \in \{0, 1\}$. If $I^{(a)}(w, s) = 0$ then we define transition intensities, $\tilde{q}^{(a)}((w, s), \cdot)$. Alternatively if $I^{(a)}(w, s) = 1$ then there are transition probabilities, $\tilde{p}^{(a)}((w, s), \cdot)$. In our case, it is simply that $I^{(a)}(w, s) = a$ and hence for action $a = 0$ (do not switch), we have transition intensities and for action $a = 1$ (switch), we have transition probabilities. These are constructed via,

$$\begin{aligned} \tilde{q}^{(0)}((w, s), (w, s + 1)) &= \tilde{\lambda}_s \quad \text{for } s \in \{0, \dots, n - 2\}, \\ \tilde{q}^{(0)}((w, s), (w, s - 1)) &= \tilde{\mu}_s \quad \text{for } s \in \{1, \dots, n - 1\}, \\ \tilde{q}^{(0)}((0, s), (1, s)) &= 1 \quad \text{for } s \in \{0, \dots, n - 1\}, \\ \tilde{q}^{(0)}((1, s), (0, s)) &= \frac{1}{\gamma} - 1 \quad \text{for } s \in \{0, \dots, n - 1\}, \end{aligned}$$

where

$$\tilde{\lambda}_s = n - 1 - s, \quad \tilde{\mu}_s = s \left(\frac{1}{\gamma} - 1 \right),$$

and $\tilde{q}^{(0)}((w, s), \cdot) = 0$ otherwise. Further,

$$\begin{aligned} \tilde{p}^{(1)}((0, s), (1, s - 1)) &= 1 \quad \text{for } s \in \{1, \dots, n - 1\}, \\ \tilde{p}^{(1)}((1, s), (0, s + 1)) &= 1 \quad \text{for } s \in \{0, \dots, n - 2\}, \\ \tilde{p}^{(1)}((0, 0), (0, 0)) &= 1, \\ \tilde{p}^{(1)}((1, n - 1), (1, n - 1)) &= 1, \end{aligned}$$

and $\tilde{p}^{(1)}((w, s), \cdot) = 0$ otherwise.

This continuous time MDP with singular controls can be converted to a continuous time MDP without singular controls using the same approach as in Ayesta et al. [5]. Further, we can uniformize with a uniformization rate of n/γ to obtain a standard discrete time MDP with the same state space and action space. In this MDP, we denote the state process via

$$X(k) = (W(k), S(k)), \quad k = 0, 1, 2, \dots$$

and the action sequence denoted $\{A(n)\}$ takes values in $\{0, 1\}$. The transition probabilities for this MDP are,

$$\begin{aligned} p^{(0)}((w, s), (w, s + 1)) &= \lambda_s \quad \text{for } s \in \{0, \dots, n - 2\}, \\ p^{(0)}((w, s), (w, s - 1)) &= \mu_s \quad \text{for } s \in \{1, \dots, n - 1\}, \\ p^{(0)}((0, s), (1, s)) &= \frac{\gamma}{n} \quad \text{for } s \in \{0, \dots, n - 1\}, \\ p^{(0)}((1, s), (0, s)) &= \frac{1 - \gamma}{n} \quad \text{for } s \in \{0, \dots, n - 1\}, \end{aligned}$$

$$\begin{aligned}
 p^{(0)}((w, s), (w, s)) &= \frac{(w + s)\gamma + (n - w - s)(1 - \gamma)}{n} \quad \text{for } s \in \{0, \dots, n - 1\}, \\
 p^{(1)}((0, s), (1, s)) &= \lambda_{s-1} \quad \text{for } s \in \{1, \dots, n - 1\}, \\
 p^{(1)}((0, s), (1, s - 2)) &= \mu_{s-1} \quad \text{for } s \in \{2, \dots, n - 1\}, \\
 p^{(1)}((0, s), (0, s - 1)) &= \frac{1 - \gamma}{n} \quad \text{for } s \in \{1, \dots, n - 1\}, \\
 p^{(1)}((0, s), (1, s - 1)) &= \frac{s\gamma + (n - s)(1 - \gamma)}{n} \quad \text{for } s \in \{1, \dots, n - 1\}, \\
 p^{(1)}((1, s), (0, s + 2)) &= \lambda_{s+1} \quad \text{for } s \in \{0, \dots, n - 3\}, \quad (*) \\
 p^{(1)}((1, s), (0, s)) &= \mu_{s+1} \quad \text{for } s \in \{0, \dots, n - 2\}, \\
 p^{(1)}((1, s), (1, s + 1)) &= \frac{\gamma}{n} \quad \text{for } s \in \{0, \dots, n - 2\}, \\
 p^{(1)}((1, s), (0, s + 1)) &= \frac{(s + 1)\gamma + (n - s - 1)(1 - \gamma)}{n} \quad \text{for } s \in \{0, \dots, n - 2\},
 \end{aligned}$$

where,

$$\lambda_s = \gamma \frac{n - 1 - s}{n}, \quad \mu_s = (1 - \gamma) \frac{s}{n},$$

and $p^{(a)}((w, s), \cdot) = 0$ otherwise. Note that the transition above marked via (*) only occurs for systems with $n \geq 3$.

The rewards of this uniformized MDP are,

$$\begin{aligned}
 r^{(0)}(w, s) &= \frac{\gamma}{n} w, \\
 r^{(1)}(w, s) &= \begin{cases} \frac{\gamma}{n} - c & \text{if } (w = 0 \text{ and } s > 0) \text{ or } (w = 1 \text{ and } s = n - 1), \\ -c & \text{otherwise.} \end{cases}
 \end{aligned}$$

We are now able to prove that the optimal policy for this MDP is π^* in the case of $n = 2$.

Proof of optimality of π^ for $n = 2$ in Theorem 2.1.* We carry out policy evaluation for this MDP under π^* and this yields a relative value function, $V^* : \{0, 1\} \times \{0, 1\} \rightarrow \mathbb{R}$, interpreted as,

$$V^*(w, s) = \lim_{\ell \rightarrow \infty} \mathbb{E} \left[\left(\sum_{k=0}^{\ell-1} r^{(A(k))}(W(k), S(k)) \right) - \ell \bar{g}_{(w,s)}^* \mid W(0) = w, S(0) = s \right],$$

where,

$$\bar{g}_{(w,s)}^* = \lim_{\ell \rightarrow \infty} \mathbb{E} \left[\frac{1}{\ell} \sum_{k=0}^{\ell-1} r^{(A(k))}(W(k), S(k)) \mid W(0) = w, S(0) = s \right].$$

Note that this MDP is a uni-chain and thus $\bar{g}_{(w,s)}^*$ does not depend on the initial state and can be denoted as \bar{g}^* . Now we carry out explicit policy evaluation (under the policy π^*) where we constrain $V^*(0, 0) = 0$. This is done separately for $c < \gamma$ (where $\pi^* = \pi^{(a)}$) and $c \geq \gamma$ (where $\pi^* = \pi^{(s)}$). The

resulting expressions for the relative value function and gains are,

$$\begin{aligned} V^*(0, 1) &= \frac{(c - \gamma)(\gamma - 2)}{2}, \\ V^*(1, 0) &= \frac{\gamma(c - \gamma + 2)}{2}, \\ V^*(1, 1) &= \frac{3\gamma - c}{2} + c\gamma - \gamma^2, \\ \bar{g}^* &= \frac{\gamma(1 - (1 - \gamma)^2 - c(1 - \gamma))}{2}, \end{aligned} \quad \text{for } c < \gamma,$$

and

$$\begin{aligned} V^*(0, 1) &= 0, \\ V^*(1, 0) &= \gamma, \\ V^*(1, 1) &= \gamma, \quad \text{for } c < \gamma. \\ \bar{g}^* &= \frac{\gamma^2}{2}, \end{aligned}$$

Note that as expected, taking $n = 2$, we have that $\bar{g}^* = (\gamma/n)g^*$, where g^* is as in Theorem 2.1.

We now consider the Bellman equation for the MDP,

$$V + g \mathbf{1} = \max \left\{ \underbrace{r^{(0)} + P^{(0)}V}_{R_0}, \underbrace{r^{(1)} + P^{(1)}V}_{R_1} \right\}, \tag{4}$$

where $\mathbf{1}$ is a vector of ones, V and $r^{(\cdot)}$ are the appropriate vectors and $P^{(\cdot)}$ is the appropriate transition probability matrix (using the lexicographic ordering of the state space). Using $V = V^*$ from above, expressions for the vector difference $R_1 - R_0$ are as follows:

$$\begin{aligned} c & \text{ for } (w, s) = (0, 0), \\ \frac{c - \gamma}{2} & \text{ for } (w, s) = (0, 1), \\ \frac{3c + \gamma}{2} & \text{ for } (w, s) = (1, 0), \\ c & \text{ for } (w, s) = (1, 1), \end{aligned} \quad \text{for } c < \gamma,$$

and

$$\begin{aligned} c & \text{ for } (w, s) = (0, 0), \\ c - \gamma & \text{ for } (w, s) = (0, 1), \\ c + \gamma & \text{ for } (w, s) = (1, 0), \\ c & \text{ for } (w, s) = (1, 1), \end{aligned} \quad \text{for } c \geq \gamma.$$

The respective signs of these entries are $(+, -, +, +)$ for $c < \gamma$ and $(+, +, +, +)$ for $c \geq \gamma$. These signs agree with π^* meaning that the only case for switching is in state $(0, 1)$ when $c < \gamma$. Hence, V^* is indeed an optimal solution of (4) because it was computed for π^* which agrees with these signs. Hence π^* is an optimal policy. □

We note that using computer algebra systems (Mathematica in our case), we are able to explicitly carry out policy evaluation and obtain expressions for V^* for $n = 2, 3, 4, \dots, 8$ in both the $c < \gamma$ and $c \geq \gamma$ cases.¹ In all these cases, we can also compute $R_1 - R_0$ and see that the signs of the entries agree with π^* and hence π^* is optimal. Note that these computations also involve finding an expression for the gain, \bar{g}^* . In all cases, we obtain $\bar{g}^* = (\gamma/n)g^*$ (g^* is as in Theorem 2.1), similarly to the case $n = 2$ in the proof above .

¹Higher values of n can potentially be handled as well with greater computation cost (we did not carry this out).

Hence for $n \in \{2, \dots, 8\}$, we essentially have a proof of optimality (in addition to the explicit proof written above for $n = 2$). However, we were not able to distill explicit expressions for V^* using a simple formula and hence we do not have a proof for arbitrary n . Nevertheless, we conjecture that π^* is optimal in Case I for any n . See the paper's associated GitHub repository Wang et al. [19] for the associated Mathematica code.

3.1. The Whittle index

It is also interesting to consider the system as if controlled by the Whittle index (see [8] for an overview). As this is a continuous time system with impulsive controls, we use a setup similar to Ayesta et al. [5].

In general, an index-based policy operates as follows. There is an index function $I : \mathcal{X} \rightarrow \mathbb{R}$, where \mathcal{X} is the state space of the channel. The index-based policy continuously monitors the index of each channel, say I_i for $i = 1, \dots, n$ by evaluating $I(\cdot)$ on the channel state. It then ensures that the chosen channel is the channel i with the highest index I_i . In our continuous time setup, we assume that the policy only switches channels if necessary. That is, if there is a tie, there is not a switch.

The Whittle index is one such index function. A key component for the evaluation of the Whittle index is the one-arm subsidy problem, parameterized via $\nu \in \mathbb{R}$. In this problem, instead of considering the full system of n channels we focus on a single channel and assume the controller has a choice between two actions: $a = 0$ which means not using the channel and receiving a subsidy at rate ν , or $a = 1$ which means using the channel.

In the one-arm subsidy problem for our case, the state space is

$$\mathcal{X} = \{(x, u) : x \in \{0, 1\}, u \in \{0, 1\}\},$$

where x and u are indicators of whether the channel is in a good state and whether the channel is in use, respectively.

Since we are dealing with impulsive controls, similarly to the n channel continuous time impulsive MDP above, we use $I^{(a)}(x, u)$ as indicators for impulsive control. Here, in the one-arm subsidy problem, $I^{(a)}(x, u) = \mathbb{1}_{\{a \neq u\}}$. With these, we have transition rates (for state-action pairs without impulsive controls),

$$\begin{aligned} \tilde{q}^{(0)}((0, 0), (1, 0)) &= 1, & \tilde{q}^{(0)}((1, 0), (0, 0)) &= \frac{1}{\gamma} - 1, \\ \tilde{q}^{(1)}((0, 1), (1, 1)) &= 1, & \tilde{q}^{(1)}((1, 1), (0, 1)) &= \frac{1}{\gamma} - 1, \end{aligned}$$

and $\tilde{q}^{(\cdot)}((x, u), \cdot) = 0$ otherwise for all other state-action pairs without impulsive control. Further, the transition probabilities (for state-action pairs with impulsive control) are,

$$\begin{aligned} \tilde{p}^{(0)}((0, 1), (0, 0)) &= 1, & \tilde{p}^{(0)}((1, 1), (1, 0)) &= 1, \\ \tilde{p}^{(1)}((0, 0), (0, 1)) &= 1, & \tilde{p}^{(1)}((1, 0), (1, 1)) &= 1, \end{aligned}$$

and $\tilde{p}^{(\cdot)}((x, u), \cdot) = 0$ otherwise for all other state-action pairs with impulsive control.

In such a one-arm subsidy problem, the continuous time reward is taken as

$$\tilde{r}^{(a)}(x, u) = xu \quad x \in \{0, 1\}, u \in \{0, 1\}. \tag{5}$$

Further in adapting the one-arm subsidy problem to our general control problem, we set instantaneous rewards at transitions where $I^{(a)}(x, u) = 1$ at a value of $-c/2$. Here, the factor of $1/2$ captures the fact that a one-arm subsidy problem is not a real problem being used but rather an abstraction of a single channel vs. all channels together. In the real system, every channel switch would result in a cost of c

(reward of $-c$) and then the one-arm subsidy problem of the current channel “becomes” the problem of the next channel.

The continuous time MDP with singular controls of the one-arm subsidy problem can now be converted to a continuous time MDP without singular controls using the same approach as Ayesta et al. [5] and our analysis above. Further, similar to the MDP above, we can uniformize, now using a uniformization rate of $1/\gamma$. This yields a standard discrete time MDP with the same state space and action space. The transition probabilities for this uniformized MDP have for every $(x, u) \in \mathcal{X}$,

$$\begin{aligned}
 p^{(0)}((x, u), (0, 0)) &= 1 - \gamma, & \text{and} & & p^{(1)}((x, u), (0, 1)) &= 1 - \gamma, \\
 p^{(0)}((x, u), (1, 0)) &= \gamma, & & & p^{(1)}((x, u), (1, 1)) &= \gamma,
 \end{aligned}
 \tag{6}$$

and $p^{(\cdot)}((x, u), \cdot) = 0$ otherwise. Further, the rewards are now,

$$r^{(a)}(x, u) = \gamma x - \mathbf{1}_{\{a \neq u\}} \frac{c}{2}.$$

Observe that the first term captures the uniformized continuous time reward and the second term captures the impulsive costs.

Now for any fixed $\nu \in \mathbb{R}$, consider the Bellman equation for the relative value function $V^\nu(x, u)$,

$$V^\nu + g^\nu \mathbf{1} = \max \left\{ \underbrace{r^{(0)} + \nu \mathbf{1} + P^{(0)} V^\nu}_{R_0^\nu}, \underbrace{r^{(1)} + P^{(1)} V^\nu}_{R_1^\nu} \right\},
 \tag{7}$$

where as previously, $\mathbf{1}$ is a vector of ones, V^ν and $r^{(\cdot)}$ are the appropriate 4-vectors and $P^{(\cdot)}$ is the appropriate transition probability matrix (using the lexicographic ordering of the state space) using (6), and g^ν is the (scalar) long-term average reward for the problem parameterized by ν .

It can now be verified that the Whittle index for $c < \gamma$ is:

$$\nu(0, 0) = 0, \quad \nu(0, 1) = c\gamma, \quad \nu(1, 0) = c\gamma + (\gamma - c), \quad \nu(1, 1) = \gamma.
 \tag{8}$$

And further, the Whittle index for $c \geq \gamma$ is:

$$\nu(0, 0) = 0, \quad \nu(0, 1) = \gamma^2, \quad \nu(1, 0) = \gamma^2, \quad \nu(1, 1) = \gamma.
 \tag{9}$$

That is, for each state $(x, u) \in \mathcal{X}$, $\nu(x, u)$ in (8) or (9) defines a Bellman equation (7) which has a solution in which $R_0^\nu(x, u) = R_1^\nu(x, u)$ (again using the lexicographic ordering for the vectors R_0^ν and R_1^ν). This can be verified, also by noting that the associated value function solutions (V^ν, g^ν) of (7) are as follows for $c < \gamma$:

$$\begin{aligned}
 &\left(\left[0 \quad \frac{c}{2} \quad \gamma \quad \gamma + \frac{c}{2} \right]^\top, \gamma^2 \right) && \text{for } \nu(0, 0) = 0, \\
 &\left(\left[0 \quad -\frac{c}{2} \quad \gamma - c \quad \gamma - \frac{c}{2} \right]^\top, \gamma^2 \right) && \text{for } \nu(0, 1) = c\gamma, \\
 &\left(\left[0 \quad -\frac{c}{2} \quad 0 \quad \frac{c}{2} \right]^\top, \gamma - c + c\gamma \right) && \text{for } \nu(1, 0) = c\gamma + (\gamma - c), \\
 &\left(\left[0 \quad -\frac{c}{2} \quad 0 \quad -\frac{c}{2} \right]^\top, \gamma \right) && \text{for } \nu(1, 1) = \gamma.
 \end{aligned}$$

And when $c \geq \gamma$ these are,

$$\begin{aligned} & \left(\left[0 \quad \frac{c}{2} \quad \gamma \quad \gamma + \frac{c}{2} \right]^T, \gamma^2 \right) && \text{for } \nu(0, 0) = 0, \\ & \left(\left[0 \quad -\frac{c}{2} \quad 0 \quad \gamma - \frac{c}{2} \right]^T, \gamma^2 \right) && \text{for } \nu(0, 1) = \gamma^2, \\ & \left(\left[0 \quad \frac{c}{2} - \gamma \quad 0 \quad \frac{c}{2} \right]^T, \gamma^2 \right) && \text{for } \nu(1, 0) = \gamma^2, \\ & \left(\left[0 \quad -\frac{c}{2} \quad 0 \quad -\frac{c}{2} \right]^T, \gamma \right) && \text{for } \nu(1, 1) = \gamma. \end{aligned}$$

With these Whittle index expressions at hand, we now observe that a Whittle index-based policy agrees with π^* of Theorem 2.1.

Proposition 1. *Consider Case I, and assume the system is controlled via a Whittle index policy, which ensures that at any time t the chosen channel is $U(t)$ with the highest index as given by (8) for $c < \gamma$ or (9) for $c \geq \gamma$, and performs switching only when necessary. Then the system operating under this policy is equivalent to a system operating under π^* of Theorem 2.1.*

Proof. Observe that both in the $c < \gamma$ and $c \geq \gamma$ cases,

$$\nu(1, 0) < \nu(1, 1) \quad \text{and} \quad \nu(0, 0) < \nu(1, 1).$$

Hence the system controlled by a Whittle index policy never switches out of a good channel ($X_{U(t)}(t) = 1$).

Further, at time t if $X_{U(t^-)}(t^-) = 1$ and then $X_{U(t)}(t) = 0$ (the state of the current channel turns from the good state to the bad state), then if for all the other channels $j \neq U(t)$, $X_j(t) = 0$, then the index-based policy chooses to stay with $U(t)$, since both in the $c < \gamma$ and $c \geq \gamma$ cases, $\nu(0, 0) < \nu(0, 1)$.

Conversely in such a situation, if there is at least one good channel in the system, say $i \neq U(t)$ with $X_i(t) = 1$, then the index-based policy chooses to switch to channel i (or any other good channel) only if $c < \gamma$, and otherwise stays. This is because, $\nu(0, 1) < \nu(1, 0)$ only if $c < \gamma$ and $\nu(0, 1) = \nu(1, 0)$ when $c \geq \gamma$. □

4. Renewal reward analysis for case II

In this section, we prove Theorem 2.2 via a regenerative analysis for $n = 2$ and later for $n = \infty$. We consider the regenerative structure based on regeneration points where switching has just occurred into a bad channel ($= 0$). That is, take the time-axis and consider time points at which the controller switched (from a bad channel) into a bad channel. These are regeneration points which we denote by $T_0 = 0, T_1, T_2, \dots$. We can assume that at time $t = 0$ the system starts at such a point because we are looking at the infinite horizon average cost. Note that between T_{k-1} and T_k it is possible that there were other switches—namely switches into a channel that is in a good state ($= 1$).

The net reward obtained during the interval $(T_{k-1}, T_k]$, is given by

$$V_k = \int_{T_{k-1}}^{T_k} \tilde{I}(u) \, du - c(\tilde{N}(T_k) - \tilde{N}(T_{k-1})).$$

That is the total time during the interval when the channel was in a good state less the switching costs.

From the regenerative property of the process, it turns out that the sequence $(T_k - T_{k-1}, V_k)$ is an i.i.d. sequence. We denote W as a generic random variable distributed as $T_k - T_{k-1}$ and V as a generic

random variable distributed as V_k . It then follows from the ‘‘Renewal Reward Theorem,’’ see for example Theorem 1.2, Chap VI [4], that the average reward is given by

$$g = \frac{\mathbb{E}[V]}{\mathbb{E}[W]}. \tag{10}$$

Note that by construction W is non-lattice and both V and W have a finite mean. Hence our goal is now to compute the expectation of V and W under a call gapping policy with parameter τ . For $n = 2$ or $n = \infty$ this is equivalent to the cool-off policy with $\sigma = \tau$.

We now construct two generic random variables via their probability distributions that come to aid. We denote these as W_0 and W_1 and their CDFs by $F_{W_i}(t)$ for $i = 0, 1$. We have,

$$F_{W_i}(t) = \mathbb{P}(W_i \leq t) = \begin{cases} 0, & t < \tau, \\ (1 - p(\tau; i)) + p(\tau; i)(1 - e^{-(1-\gamma)/\gamma(t-\tau)}), & \tau \leq t. \end{cases} \tag{11}$$

These random variables are mixtures of a mass at τ and a shifted exponential. The random variable W_i denotes the time of switching under a policy with call-gapping parameter τ when at time 0 a switch occurred into a channel in state i . It is constructed by observing that if at time $t = \tau$ the state is 0 we switch with probability $1 - p(\tau; i)$; otherwise we wait an exponentially distributed duration until switching. Note that,

$$\mathbb{E}[W_i] = \tau + p(\tau; i) \frac{\gamma}{1 - \gamma}. \tag{12}$$

With the generic random variables W_0 and W_1 at hand, we can analyze a complete regenerative cycle of duration W . For this denote,

$$W = \tilde{W}_0 + \tilde{W}_1 + \tilde{W}_2 + \dots + \tilde{W}_M,$$

with \tilde{W}_i denoting the inter-switching times within the cycle. Here M is a random variable with support $0, 1, 2, \dots$, denoting the number of times we switch into a good channel within such a cycle. It then holds by construction that,

$$\tilde{W}_0 =^d W_0, \quad \tilde{W}_i =^d W_1, \quad i = 1, 2, \dots$$

This follows since our regeneration points are the time points at which the controller switched from a bad channel into a bad channel. Then each cycle starts with a duration distributed as W_0 . It is then followed by M additional durations which are switches from bad states to good states, each distributed as W_1 .

The following two lemmas yield explicit expressions for the denominator and numerator of (10).

Lemma 4.1. *The denominator of (10) can be represented as,*

$$\begin{aligned} \mathbb{E}[W] &= \mathbb{E}[W_0] + \frac{\mathbb{E}[p(W_0; 0)]}{1 - \mathbb{E}[p(W_1; 0)]} \mathbb{E}[W_1] \\ &= \frac{e^{2\tau/\gamma}(\gamma^3 - (\tau + 2)\gamma^2 + 3\tau\gamma - 2\tau) - \gamma^3 - (\tau - 2)\gamma^2 + \tau\gamma}{(\gamma - 1)^2(e^{2\tau/\gamma}(\gamma - 2) - 2e^{\tau/\gamma}\gamma + \gamma)}. \end{aligned}$$

Proof. Observe that,

$$W =^d \tilde{W}_0 + \tilde{I}_0 \tilde{W}_1 + \tilde{I}_0 \tilde{I}_1 \tilde{W}_2 + \dots + \tilde{I}_0 \tilde{I}_1 \dots \tilde{I}_{k-1} \tilde{W}_k + \dots,$$

where the indicator, \tilde{I}_i , for $i = 0, 1, 2, \dots$ is equal to 1 if the jump at the end of the interval associated with \tilde{W}_i was to a good channel. Note that the random variables \tilde{I}_i for $i = 1, 2, \dots$ are identically distributed

and \tilde{I}_0 follows a different distribution. We can now denote generic random variables, I_0 and I_1 satisfying

$$\tilde{I}_0 =^d I_0, \quad \tilde{I}_i =^d I_1, \quad i = 1, 2, \dots$$

For these two random variables, by conditioning on W_i (for $i = 0, 1$), it holds that

$$\begin{aligned} \mathbb{E}[I_i] &= \mathbb{E}[\mathbb{E}[I_i | W_i]] \\ &= \mathbb{E}[p(W_i; 0)] \\ &= (1 - p(\tau; i))p(\tau; 0) + p(\tau; i) \int_{\tau}^{\infty} p(t; 0) \frac{1 - \gamma}{\gamma} e^{-((1-\gamma)/\gamma)(t-\tau)} dt \\ &= \gamma - \frac{e^{-\frac{\tau}{\gamma}} \gamma (\gamma + p(\tau; i) - 2)}{\gamma - 2}, \end{aligned} \tag{13}$$

where the second step follows because every switch is out of a bad channel and the third step follows from (11).

Observe that the sequence $\{\tilde{I}_i\}$ is a mutually independent sequence of random variables and for each i , \tilde{I}_i is independent of W_j for all $j \neq i$. Based on the fact that $\mathbb{E}[\tilde{W}_i]$ is same as $\mathbb{E}[W_1]$ for $i = 1, 2, \dots$, we have

$$\begin{aligned} \mathbb{E}[W] &= \mathbb{E}[W_0] + \mathbb{E}[I_0]\mathbb{E}[W_1] + \mathbb{E}[I_0]\mathbb{E}[I_1]\mathbb{E}[W_1] + \dots + \mathbb{E}[I_0](\mathbb{E}[I_1])^{k-1}\mathbb{E}[W_1] + \dots \\ &= \mathbb{E}[W_0] + \mathbb{E}[I_0] \frac{1}{1 - \mathbb{E}[I_1]} \mathbb{E}[W_1]. \end{aligned} \tag{14}$$

After manipulation using (1), (12), and (13), the result follows. □

Lemma 4.2. *The numerator, $\mathbb{E}[V]$ of (10) can be represented as*

$$\mathbb{E}[V] = \mathbb{E}[R] - c(\mathbb{E}[M] + 1),$$

where,

$$\begin{aligned} \mathbb{E}[R] &= \mathbb{E}[R_0] + \frac{\mathbb{E}[p(W_0; 0)]}{1 - \mathbb{E}[p(W_1; 0)]} \mathbb{E}[R_1] \\ &= \frac{e^{2\tau/\gamma}((1 - \tau)\gamma^3 + (3\tau - 2)\gamma^2 - 2\tau\gamma) + 2e^{\tau/\gamma}\gamma^2(\gamma - 1)^2 - 2\gamma^4 + (3 - \tau)\gamma^3 + \tau\gamma^2}{(\gamma - 1)^2(e^{2\tau/\gamma}(\gamma - 2) - 2e^{\tau/\gamma}\gamma + \gamma)}, \end{aligned}$$

and,

$$\begin{aligned} \mathbb{E}[M] &= \frac{\mathbb{E}[p(W_0; 0)]}{1 - \mathbb{E}[p(W_1; 0)]}, \\ &= \frac{e^{2\tau/\gamma}(2 - \gamma) - \gamma}{(e^{2\tau/\gamma}(\gamma - 2) - 2e^{\tau/\gamma}\gamma + \gamma)(\gamma - 1)} - 1. \end{aligned}$$

Proof. The proof follows similar lines to the previous proof. We have that $V = R - c, (M + 1)$ with M as defined above and the reward R being the net time during $[0, W]$ in which a good channel is selected. As with the analysis of the cycle duration, W , we denote,

$$R = \tilde{R}_0 + \tilde{R}_1 + \tilde{R}_2 + \dots + \tilde{R}_M,$$

where \tilde{R}_i is the reward accumulated over the period corresponding to \tilde{W}_i . Similar to the previous analysis, we construct generic random variables R_0 and R_1 with

$$\tilde{R}_0 =^d R_0, \quad \tilde{R}_i =^d R_1, \quad i = 1, 2, \dots$$

The expectations of these can be computed to be

$$\begin{aligned} \mathbb{E}[R_0] &= \int_0^\tau p(t; 0) dt + p(\tau; 0) \frac{1-\gamma}{\gamma} = \gamma(\tau - p(\tau; 0)) + \frac{p(\tau; 0)(1-\gamma)}{\gamma}, \\ \mathbb{E}[R_1] &= \int_0^\tau p(t; 1) dt + p(\tau; 1) \frac{1-\gamma}{\gamma} = \gamma(\tau - p(\tau; 1) + 1) + \frac{p(\tau; 1)(1-\gamma)}{\gamma}. \end{aligned}$$

Using \tilde{I}_i as in the previous proof, we observe,

$$R = {}^d \tilde{R}_0 + \tilde{I}_0 \tilde{R}_1 + \tilde{I}_0 \tilde{I}_1 \tilde{R}_2 + \dots + \tilde{I}_0 \tilde{I}_1 \dots \tilde{I}_{k-1} \tilde{R}_k + \dots,$$

where again for any i , \tilde{I}_i is independent of \tilde{R}_j for $j \neq i$. Now a similar geometric series to (14) is applied. The expression for M is computed in a similar manner by observing

$$M = {}^d \tilde{I}_0 + \tilde{I}_0 \tilde{I}_1 + \tilde{I}_0 \tilde{I}_1 \tilde{I}_2 + \dots$$

□

We can now prove Theorem 2.2 for $n = 2$.

Proof. Combining and manipulating the explicit expressions from the lemmas above we obtain that under a policy $\pi^{(\tau^*)}$ (alt. $\pi^{(\sigma^*)}$), the reward as in (10) as a function of τ (alt. σ) is,

$$g(\tau) = \frac{A_1(\gamma, \tau) - c A_2(\gamma, \tau)}{A_3(\gamma, \tau)} \tag{15}$$

where $A_i(\cdot, \cdot)$, $i = 1, 2, 3$ are as defined in Theorem 2.2. We consider first the case $c \geq \gamma^2$. Observe that at $c = \gamma^2$,

$$g(\tau)|_{c=\gamma^2} = \gamma - \frac{2e^{\tau/\gamma}(\gamma - 1)^2\gamma^2}{A_3(\gamma, \tau)} < \gamma$$

and hence the $\pi^{(s)}$ policy is preferable to the $\pi^{(\tau)}$ policy for any finite τ (note that as $\tau \rightarrow \infty$ the inequality above becomes an equality). Further $g(\tau)$ is monotonically decreasing in c and hence for $c > \gamma^2$ it remains optimal to use $\pi^{(s)}$.

Moving to the $c < \gamma^2$ case, we optimize the (continuous and differentiable) function $g(\tau)$ over $(0, \infty)$ to obtain Eq. (2) from Theorem 2.2. We do this by representing the derivative as $g'(\tau) = h(\tau)f(\tau)$ where,

$$h(\tau) = \frac{(\gamma - 1)^2(e^{2\tau/\gamma}(2 - \gamma) + \gamma)}{(\gamma((\gamma - 2)\gamma + \tau(1 - \gamma)) + e^{2\tau/\gamma}(2 - \gamma)(\gamma^2 + \tau - \tau\gamma))^2},$$

and,

$$f(\tau) = e^{2\tau/\gamma}(\gamma^2 - c)(\gamma - 2) + 2e^{\tau/\gamma}\gamma(\gamma - \tau(\gamma - 1)) - \gamma(\gamma^2 - c).$$

It can be shown that for $c < \gamma^2$ the derivative $f'(\tau) < 0$. Now since $h(\cdot) > 0$ we have that, $g'(\tau) = 0$ if and only if $f(\tau) = 0$ and Eq. (2) is $f(\tau^*) = 0$. Since $f(0) > 0$ and $\lim_{\tau \rightarrow \infty} f(\tau) < 0$, it is evident that for $c < \gamma^2$ there is a single root to $f(\tau) = 0$ and further at the solution τ^* the second order conditions hold,

$$g''(\tau^*) = 2h(\tau^*)((1 - e^{\tau^*/\gamma})\gamma^2 - c - e^{\tau^*/\gamma}\tau^*(1 - \gamma)) < 0.$$

□

Note that at the extremes of τ , the reward $g(\tau)$ as in (15) yields the expected results. With zero switching costs,

$$\lim_{\tau \rightarrow 0} g(\tau) \Big|_{c=0} = \lim_{\tau \rightarrow 0} \frac{A_1(\gamma, \tau)}{A_3(\gamma, \tau)} = 1 - (1 - \gamma)^2,$$

and further without switching,

$$\lim_{\tau \rightarrow \infty} \frac{A_1(\gamma, \tau) - cA_2(\gamma, \tau)}{A_3(\gamma, \tau)} = \lim_{\tau \rightarrow \infty} \frac{A_1(\gamma, \tau)}{A_3(\gamma, \tau)} - c0 = \gamma.$$

In the case of $n = \infty$, a similar (yet simpler) analysis to all of the above sections pursues. As described in Section 2, the system may be viewed as a two-channel system where every transition is always to a steady state channel. The usage of the transient probabilities $p(\tau ; i)$ as in (11) is then replaced by γ . This results in correspondently simpler expressions in all the cases where transient probabilities are used ($n = 2$) and can now be replaced by γ . The resulting reward expression is then,

$$g(\tau) = \frac{(\gamma - 1)c + \gamma(\gamma - \gamma\tau + \tau)}{\gamma^2 - \gamma\tau + \tau},$$

with the derivative,

$$g'(\tau) = \frac{(\gamma - 1)^2(c - \gamma^2)}{(\gamma^2 - \gamma\tau + \tau)^2}.$$

This shows that $g(\cdot)$ is monotonic decreasing in τ if $c < \gamma^2$ and monotonic increasing in τ if $c > \gamma^2$. Hence for $c < \gamma^2$ the optimal switching parameter is $\tau^* = 0$ and yields reward as in Theorem 2.2. Further, since $\lim_{\tau \rightarrow \infty} g(\tau) = \gamma$, for $c > \gamma^2$ it is not optimal to use $\pi^{(\tau)}$ for any finite τ , and $\pi^{(s)}$ is preferable.

5. Numerical results for case II

We now carry out numerical experiments for our system under various policies, focusing on Case II. We first describe the numerical computations and simulation experiment used for Figure 1 and then expand with further numerical results.

The Case I curves in the figure are obtained directly from Theorem 2.1. The Case II curves for $n = \infty$ are created directly using Theorem 2.2 and the $n = 2$ curves are created by numerically solving Eq. (2) of Theorem 2.2 (using the bisection method). The remaining curves for $n = 3$ and $n = 4$ using both the call-gapping and the cool-off policies are obtained via Monte–Carlo simulation and optimization to find the optimal τ and σ parameters for these policies.

This Monte–Carlo simulation (as well as other code) can be found in the GitHub repository Wang et al. [19]. We simulated the system over a grid of c in the range $[0, \gamma^2]$ where $\gamma^2 = 0.16$, with a spacing of 0.005. Then for each value of c , we simulated both the call-gapping and cool-off policies over respective grids of τ and σ in the range $[0, 1.55]$, each with a spacing of 0.001. In each of these individual simulation runs we kept a constant seed to reduce variability in the curves using common random numbers, and simulated for a time horizon of $T = 10^5$, with an arbitrary fixed initial condition. We then obtained the average reward g using a crude Monte–Carlo estimator over the continuous time $[0, T]$ and then plotted the reward for the optimal τ and σ in each case.

In addition to the simulations leading to Figure 1, we also carried out extensive additional numerical experiments for more parameter values. Key results from these experiments are summarized in Table 1. Our purpose is to compare the cool-off and call-gapping policies for finite $n > 2$. As we observe from Figure 1, the cool-off policy is able to achieve a slightly higher average reward for $\gamma = 0.4$. We now considered the system for several additional γ values and several values of n .

In each case, the cool-off policy is indeed slightly better than call-gapping. The numerical results in Table 1 report the maximal gap, as well as error estimates, over a range of c values between the rewards

Table 1. Evaluating the difference between the cool-off and call-gapping policy.

System	Number of channels	Maximal gap	Worst cost	
$\gamma = 0.2$	$n = 3$	0.00385 ± 0.00029	$c = 0.0048$	
	$n = 4$	0.00823 ± 0.00033	$c = 0.0048$	
	$n = 5$	0.01295 ± 0.00045	$c = 0.0056$	
	$n = 6$	0.01767 ± 0.00045	$c = 0.0056$	
	$n = 7$	0.02188 ± 0.00044	$c = 0.0056$	
	$n = 8$	0.02619 ± 0.00047	$c = 0.0064$	
	$n = 9$	0.02879 ± 0.00051	$c = 0.0072$	
	$n = 10$	0.03136 ± 0.00057	$c = 0.0072$	
	$n = 11$	0.03277 ± 0.00047	$c = 0.0072$	
	$n = 12$	0.03414 ± 0.0006	$c = 0.008$	
	$n = 13$	0.03554 ± 0.00048	$c = 0.008$	
	$n = 14$	0.03609 ± 0.0006	$c = 0.0088$	
	$n = 15$	0.03611 ± 0.00072	$c = 0.0096$	
	$n = 16$	0.03611 ± 0.00061	$c = 0.0104$	
	$n = 17$	0.03532 ± 0.00064	$c = 0.0112$	
	$n = 18$	0.0339 ± 0.00066	$c = 0.0112$	
	$n = 19$	0.0332 ± 0.00064	$c = 0.0112$	
	$\gamma = 0.3$	$n = 20$	0.03208 ± 0.0007	$c = 0.0128$
		$n = 3$	0.00679 ± 0.00035	$c = 0.009$
$n = 4$		0.01318 ± 0.00044	$c = 0.009$	
$n = 5$		0.01927 ± 0.00063	$c = 0.018$	
$n = 6$		0.02351 ± 0.00055	$c = 0.0144$	
$n = 7$		0.0268 ± 0.00074	$c = 0.018$	
$n = 8$		0.02829 ± 0.00073	$c = 0.0198$	
$n = 9$		0.03012 ± 0.00072	$c = 0.0216$	
$n = 10$		0.02924 ± 0.00056	$c = 0.0216$	
$n = 11$		0.0293 ± 0.00084	$c = 0.0252$	
$n = 12$		0.02798 ± 0.00096	$c = 0.0288$	
$n = 13$		0.026 ± 0.00078	$c = 0.0288$	
$n = 14$		0.02524 ± 0.00078	$c = 0.0306$	
$n = 15$	0.02216 ± 0.00074	$c = 0.027$		
$\gamma = 0.4$	$n = 3$	0.00899 ± 0.0005	$c = 0.0224$	
	$n = 4$	0.01606 ± 0.00066	$c = 0.032$	
	$n = 5$	0.02114 ± 0.00064	$c = 0.0256$	
	$n = 6$	0.02289 ± 0.00072	$c = 0.032$	
	$n = 7$	0.02426 ± 0.00071	$c = 0.0416$	
	$n = 8$	0.02294 ± 0.0007	$c = 0.0352$	
	$n = 9$	0.0213 ± 0.00074	$c = 0.0448$	
	$n = 10$	0.01849 ± 0.00076	$c = 0.0544$	
$\gamma = 0.6$	$n = 3$	0.01001 ± 0.00053	$c = 0.0648$	
	$n = 4$	0.01274 ± 0.00055	$c = 0.0648$	
	$n = 5$	0.01267 ± 0.00064	$c = 0.1008$	
$\gamma = 0.8$	$n = 3$	0.00452 ± 0.00042	$c = 0.1792$	
	$n = 4$	0.00389 ± 0.0005	$c = 0.2432$	
	$n = 5$	0.00236 ± 0.00035	$c = 0.2048$	

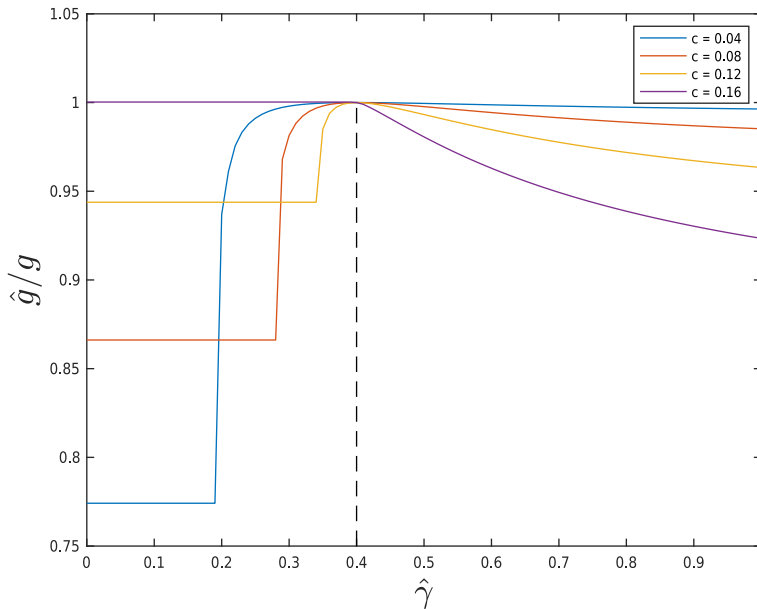


Figure 2. Robustness analysis when $\gamma = 0.4$ for Case II and $n = 2$ under call-gapping. The system is controlled with a perceived value of $\hat{\gamma}$ and the ratio of the reward and the optimal reward is plotted for several values of c .

of the policies. We also present an estimate of the c value in which the gap is maximal and highlight in bold the values of n for which the gap is largest.² As an overarching summary, we see that while the cool-off policy is slightly better, the difference between the rewards is never high and hence in practice using the simpler call-gapping policy is probably preferable. As expected from Figure 1, for each value of γ there is some finite n in which the maximal gap is highest (for $n = 2$ and $n \rightarrow \infty$ there is no gap as the policies are identical). It is also evident that for low γ values, the worst case n is higher than that of higher γ values. This observation informed our selection of the range of n for which to run the experiments.

The computational experiments required non-negligible compute time because for any system setting (γ, n, c) optimization over the best τ (resp. σ) parameter was required. We thus relied on an empirical observation associated with the nature of the system. We observed that for any fixed γ and c , as n increases, the optimal parameter (τ^* or σ^*) decreases with n (with the optimal value tending to 0 as $n \rightarrow \infty$). This allowed us to use Theorem 2.2 to first compute the optimal parameter for $n = 2$ and use it as an upper bound for the (stochastic) search for the optimal parameter for $n = 3$. Further for any $n = k \geq 4$, we used the estimated optimal parameter of $n = k - 1$ to determine an upper bound. This allowed to use a fixed size search grid of size 20 for the optimal τ (resp. σ) of each level $n \geq 3$. Other elements of the simulation involved the search grid over c . For this, after initial experimentation indicating the location of the maximal gap, we always considered $c \in [0, \gamma^2/2]$ and searched over 25 points within this grid. Each sample path involved a time horizon of 10^3 time units, and for each parameter setting (γ, n, c) , we simulated 100 repetitions, allowing us to obtain 95% error bounds using a standard normal approximation. The simulation duration was in the order of 24 hours on a standard laptop using the (compiled) Julia language.

In addition to the simulations, we also carried out a robustness analysis for Case II with $n = 2$. In this case, the optimal call-gapping (and cool-off) parameter τ is easily obtained via Eq. (2). However, we also wish to see how misspecification of the system parameter γ , leading to a misspecification of the

²For $\gamma = 0.2$, the maximal gap is at either $n = 14$, $n = 15$, or $n = 16$ and in considering the error estimates, the exact value is not determined based on the simulations we ran.

optimal τ , affects performance. For this, we again consider the case where $\gamma = 0.4$; however, we now assume that the system is perceived to operate at a potentially different system parameter, namely $\hat{\gamma}$. The controller then optimizes τ based on (2) and the perceived value, $\hat{\gamma}$. We then compare the ratio of the actual reward, \hat{g} and the ideal reward, g obtained using $\gamma = 0.4$. Both are obtained via the expressions in Theorem 2.2. The ratio is plotted in Figure 2 where we considered cost values $c = 0.04, 0.08, 0.12,$ and 0.16 . As we see in Figure 2, as we might expect, for larger deviations of $\hat{\gamma}$ from $\gamma = 0.4$ we get larger relative losses of the reward. However, the relative loss at a few percentage points in most cases. The plateaus observed on each of the curves are due to values of $\hat{\gamma}$ during which $c > \hat{\gamma}^2$ in which case the $\pi^{(s)}$ policy is employed.

6. Conclusion and extensions

Our aim with this study was to obtain a qualitative view on the interaction of the number of channels, the available information, and the switching costs in a channel selection context. In such a setting, system designers can consider the value of information as well as the effect of switching costs (the value of efficient switching) in lieu of various control policies. Through a simple model, we obtained a qualitative relationship between the various system factors and in certain cases, we were able to explicitly analyze the system.

The case of full observation is straightforward to analyze; however, for the case of partial observation, to the best of our knowledge, explicit analysis is only attainable with $n = 2$ and $n = \infty$ as we have done here. This is for the call-gapping and cool-off policies that we introduced. Further numerical analysis was carried out for other small finite n illustrating that in practice there are only minor efficiency gains to be had by using cool-off instead of call-gapping. This is despite the fact that call-gapping policy only requires local information of the current state of the channel while the cool-off policy requires information for all channels. We complement our analysis with a numerical robustness experiment hinting that inexact knowledge of system parameters would not hinder performance greatly when using a call-gapping policy.

Our work did not focus on optimality of the policies discussed within the class of all possible policies. We conjecture that for $n = 2$ call-gapping (equivalent to cool-off) is optimal in the case of partial observation. However, proving this remains a challenge for a future study. Further, real systems will typically exhibit a more complicated structure than two-state Markov chains. For such systems, adaptations of the call-gapping and cool-off policies may still be employed; however, the analysis is more complicated. Nevertheless, we believe that our qualitative and quantitative results based on two-state Markov chains may help serve as a guide for the tradeoffs between information, switching costs, policy complexity, and the number of channels in a system.

Acknowledgments. Jiesen Wang would like to thank the University of Melbourne for supporting her work through the Melbourne Research Scholarship and the Albert Shimmins Fund. Thomas Taimre is supported by the Australian Research Council (ARC) Centre of Excellence for the Mathematical and Statistical Frontiers (ACEMS). Yoni Nazarathy and Thomas Taimre are supported by ARC grant DP180101602.

References

- [1] Aalto, S., Lassila, P., & Taboada, I. (2019). Indexability of an opportunistic scheduling problem with partial channel information. In *Proceedings of the 12th EAI International Conference on Performance Evaluation Methodologies and Tools*. Palma, Spain: Association for Computing Machinery, pp. 95–102.
- [2] Aalto, S., Lassila, P., & Taboada, I. (2019). Whittle index approach to opportunistic scheduling with partial channel information. *Performance Evaluation* 136: 102052.
- [3] Agrawal, R., Hegde, M., & Teneketzis, D. (1990). Multi-armed bandit problems with multiple plays and switching cost. *Stochastics and Stochastic Reports* 29(4): 437–459.
- [4] Asmussen, S. (2008). *Applied probability and queues*, vol. 51. New York, NY, US: Springer Science & Business Media.
- [5] Ayesta, U., Gupta, M.K., & Verloop, I.M. (2021). On the computation of Whittle's index for Markovian restless bandits. *Mathematical Methods of Operations Research* 93(1): 179–208.

- [6] Banks, J.S. & Sundaram, R.K. (1994). Switching costs and the Gittins index. *Econometrica: Journal of the Econometric Society* 62(3): 687–694.
- [7] Dusonchet, F. & Hongler, M.O. (2003). Optimal hysteresis for a class of deterministic deteriorating two-armed bandit problem with switching costs. *Automatica* 39(11): 1947–1955.
- [8] Gittins, J., Glazebrook, K., & Weber, R. (2011). *Multi-armed bandit allocation indices*. New York, NY, US: John Wiley & Sons.
- [9] Jacko, P. & Villar, S.S. (2012). Opportunistic schedulers for optimal scheduling of flows in wireless systems with ARQ feedback. In *24th International Teletraffic Congress (ITC 24)*, Krakow, Poland: International Teletraffic Congress, pp. 1–8.
- [10] Kaza, K., Meshram, R., Mehta, V., & Merchant, S.N. (2019). Sequential decision making with limited observation capability: Application to wireless networks. *IEEE Transactions on Cognitive Communications and Networking* 5(2): 237–251.
- [11] Kuhn, J. & Nazarathy, Y. (2015). Wireless channel selection with reward-observing restless multi-armed bandits. In R. Boucherie & N. van Dijk (eds), *Markov decision processes in practice*. Cham: Springer.
- [12] Larrañaga, M., Assaad, M., Destounis, A., & Paschos, G.S. (2017). Asymptotically optimal pilot allocation over Markovian fading channels. *IEEE Transactions on Information Theory* 64(7): 5395–5418.
- [13] Larrañaga, M., Ayesta, U., & Verloop, I.M. (2016). Dynamic control of birth-and-death restless bandits: Application to resource-allocation problems. *IEEE/ACM Transactions on Networking* 24(6): 3812–3825.
- [14] Lin, K.Y. & Ross, S.M. (2004). Optimal admission control for a single-server loss queue. *Journal of Applied Probability* 41(2): 535–546.
- [15] Maatouk, A., Kriouile, S., Assad, M., & Ephremides, A. (2020). On the optimality of the Whittle’s index policy for minimizing the age of information. *IEEE Transactions on Wireless Communications* 20(2): 1263–1277.
- [16] Meshram, R. & Kaza, K. (2021). Indexability and rollout policy for multi-state partially observable restless bandits. In *60th IEEE Conference on Decision and Control (CDC)*, Athens, Greece: Institute of Electrical and Electronics Engineers (IEEE), pp. 2342–2347.
- [17] Mezzavilla, M., Goyal, S., Panwar, S., Rangan, S., & Zorzi, M. (2016). An MDP model for optimal handover decisions in mmWave cellular networks. In *European Conference on Networks and Communications (EuCNC)*. Austin, Texas, USA: Institute of Electrical and Electronics Engineers (IEEE), pp. 100–105.
- [18] Wang, M., Chen, J., Aryafar, E., & Chiang, M. (2016). A survey of client-controlled HetNets for 5G. *IEEE Access* 5: 2842–2854.
- [19] Wang, J., Nazarathy, Y., & Taimre, T. (2020). The value of information and efficient switching in channel selection – Github. <https://github.com/yoninazarathy/ValueOfInformationAndEfficientSwitching>
- [20] Weber, R.R. & Weiss, G. (1990). On an index policy for restless bandits. *Journal of Applied Probability* 27(3): 637–648.
- [21] Whittle, P. (1988). Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability* 25A: 287–298.