An Opportunistic and Non-Anticipating Size-Aware Scheduling Proposal for Mean Holding Cost Minimization in Time-Varying Channels

Ianire Taboada[a,*], Fidel Liberal[a], Peter Jacko[b]

[a] University of the Basque Country, ETSI Bilbao, Alameda Urquijo s/n, 48013 Bilbao, Spain
[b] Lancaster University Management School, Bailrigg, Lancaster, Lancashire LA1 4YX, UK

Abstract

In this paper we study how to design a scheduling strategy aimed at minimizing the average holding cost for flows with general size distribution when the feasible transmission rate of each user varies randomly over time. We employ a Whittle-index-based approach in order to achieve an opportunistic and non-anticipating size-aware scheduling index rule proposal. When the flow size distribution belongs to the Decreasing Hazard Rate class, we propose the so-called Attained Service Potential Improvement index rule, which consists in giving priority to the flows with the highest ratio between the current attained-service-dependent completion probability and the expected potential improvement of this completion probability. We further analyze the performance of the proposed scheduler, concluding that it outperforms well-known opportunistic disciplines.

Keywords: opportunistic scheduling, non-anticipating size-aware scheduling, mean holding cost minimization, Whittle index, Markov Decision Process

1. Introduction

Undoubtedly, due to the massive use of mobile Internet applications, one of the fundamental challenges that networks providers nowadays face is the management for sharing radio resources among users' traffic flows. Thus, motivated by the necessity of obtaining an implementable scheduler in channels with randomly time-varying capacity such as wireless links, in this paper we aim at characterizing in closed-form a novel channel-aware or opportunistic scheduler for the problem of minimizing the expected holding cost in a scenario where flows arrive and depart upon service completion.

Although in time-varying transmission conditions taking advantage of the channel opportunistic gains seems good, short-term disciplines that serve the user with the best instantaneous rate, such as Max Rate, perform very poorly in this setting (see for example [6]). Moreover, due to the complexity of the present problem, flow-level opportunistic scheduling in time-varying systems has been analyzed by approximate techniques ([13, 3, 2]) to design simple schedulers, and in the asymptotic regimes to study optimality and maximal stability ([1, 5]). Nevertheless, all these works deal with unrealistic assumptions in reference to traffic flow sizes. On the one hand, exponential flow size distributions are considered for traffic modeling, which, even though they simplify the resolution of those problems, are far from reality. On the other hand, it is assumed that flow sizes are known by the scheduler, while in current network systems they are not.

In this paper we take a step forward towards removing the assumption of exponential sizes. Furthermore, we incorporate non-anticipating size-awareness by taking into account the bits that have been transferred of a flow: the attained service. In the context of non-anticipating strategies the work [8] of

*Corresponding author. Telf: +34 94 6017352; fax: +34 94 6014259.
Email addresses: `ianire.taboada@ehu.es` (Ianire Taboada), `fidel.liberal@ehu.es` (Fidel Liberal), `peter.jacko@gmail.com` (Peter Jacko)

Gittins is relevant, which based on the attained service of jobs proposed an index rule that minimizes the mean holding cost when channel capacity is constant. [4, 15] propose some heuristics using Gittins approach for the case of time-varying capacity, the first work for deterministic channels. However, to the best of our knowledge, there is no strong analytically founded and well-performing scheduling proposal for randomly time-varying channels that combines this kind of size-awareness with channel-awareness using Gittins approach.

Therefore, in this paper we aim at developing in closed form a simple opportunistic and non-anticipating size-aware scheduler for the problem of minimizing the expected holding cost in random time-varying channels for flows with general size distribution. In order to achieve our goal, the work presented in [3] has been relevant. [3] considers a finite number of channel conditions and exponentially distributed flow sizes, and its associated optimal scheduling problem is formulated as a Markov Decision Process (MDP). Due to the impossibility of solving the general model for being PSPACE-hard [12], the authors of [3] propose a simple Whittle-index-based [19, 11] heuristic scheduler, which they show to perform well in several simulation scenarios. Moreover, as shown in [5, 10], this Whittle-index-based proposal is maximal stable and fluid-optimal, as well as asymptotically optimal under some assumptions as the number of flows and servers grows to infinity [18]. Hence, so as to achieve our aim, we design a simple Whittle-index-based scheduler extending the framework presented in [3] to the case of general size distribution.

The rest of the paper is structured as follows. In Section 2 we present the problem description. We formulate the problem as a MDP model in Section 3. We design the Whittle-index-based scheduler in Section 4, and we evaluate its performance in Section 5. Finally, Section 6 gathers the main conclusions of the paper. For the sake of readability, some of the proofs are postponed to the appendix.

2. Problem Description

We analyze a discrete-time job scheduling problem aimed at minimizing the expected holding cost, in which the feasible transmission rate of each user varies randomly over time. Scheduling decisions are taken at the beginning of time slots $t \in \mathcal{T} := \{0, 1, \dots\}$, and are applied during a slot duration.

We consider a system without arrivals with K jobs waiting for service, which incur a holding cost $c_k > 0$ per slot while the flow transmission is not completed. We will use terms job/flow/user interchangeably throughout the paper. The job size in bits x_k follows a general distribution with $\mathbb{E}[x_k] < \infty$, characterized by its probability density function $f_k(x)$.

The channel of a user k can take N_k conditions from a finite set $\mathcal{N}_k := \{1, 2, \dots, N_k\}$. These channel conditions are associated to different transmission rates $r_{k,n}$ (in bits), where $r_{k,1} \leq r_{k,2} \leq \dots \leq r_{k,N_k}$. The channel condition of a user k evolves randomly and independently of other users. We denote the probability of being in state n by $q_{k,n}$, having $\sum\limits_{n \in \mathcal{N}} q_{k,n} = 1$.

The server makes use of the instantaneous channel information ($r_{k,n}$ and $q_{k,n}$) and the instantaneous attained service (a_k) of each user in order to take decisions. In each decision slot it allows the transmission of a single flow, and it is assumed to be preemptive (the service of a job can be interrupted at the beginning of a slot even if not completed). We refer to job completion or departure probability, $\mathbb{P}\left(a_k < X_k \leq a_k + r_{k,n} | X_k > a_k)\right)$, as $\mu_{k,(a,n)}$.

3. MDP Formulation

In this section we present a MDP formulation of the scheduling problem described in Section 2. First, we provide the MDP model of each job k. Then, we formulate the optimization problem for the joint MDP model, which takes into account all the jobs in the system.

3.1. MDP model of a job

In each time slot t, a user k that is in a state $s_k \in \mathcal{S}_k$ can be allocated either zero or full capacity. We refer to $\mathcal{B} := \{0, 1\}$ as the action space, in which action 0 means not serving and action 1 serving. Thus, the dynamics of user k is captured by the action process $b_k(.)$ and the state process $s_k(.)$. As a result of taking action $b_k(t)$ in state $s_k(t)$, the user k earns a reward, consumes the allocated capacity and evolves its state in the time slot $t+1$. In such a way, each user k is defined independently of other users by tuple $\left(\mathcal{S}_k, \left(\boldsymbol{R}^b_{k,s}\right)_{b\in\mathcal{B}}, \left(\boldsymbol{W}^b_{k,s}\right)_{b\in\mathcal{B}}, \left(\boldsymbol{P}^b_{k,s}\right)_{b\in\mathcal{B}}\right)$ as follows:

- $\mathcal{S}_k := (\mathcal{A}_k \times \{1, 2, \ldots, N_k\}) \cup \{*\}$ is the state space, where for $a \in \mathcal{A}_k$ and $n \in \mathcal{N}_k$ in each state (a, n) the job is uncompleted, and state * represents a flow already completed.
- $\boldsymbol{R}^b_k{:=}(R^b_{k,s})_{s\in\mathcal{S}_k}$, where $R^b_{k,s}$ is the expected one-slot reward earned by user k at state s if action b is decided at the beginning of a slot; it is defined as the expected cost of remaining in the system as:

$$R^0_{k,(a,n)} = -c_k, \qquad R^1_{k,(a,n)} = -c_k(1 - \mu_{k,(a,n)}), \qquad R^b_{k,*} = 0;$$

- $\boldsymbol{W}^b_k{:=}(W^b_{k,s})_{s\in\mathcal{S}_k}$, where $W^b_{k,s}$ is the expected one-slot capacity consumption or work required by user k at state s if action b is decided at the beginning of a slot, so that

$$W^0_{k,s} = 0, \qquad W^1_{k,s} = 1;$$

- $\boldsymbol{P}^b_k{:=}\left(p^b_k(s, s')\right)_{s,s'\in\mathcal{S}_k}$, where $p^b_k(s, s')$ is the probability of moving from state s to state s' under action b at the beginning of a slot, so that

$$p^0_k\left((a, n), (a, m)\right) = q_{k,m}, \qquad p^1_k\left((a, n), (a + r_n, m)\right) = q_{k,m}(1 - \mu_{k,(a,n)}),$$
$$p^1_k\left((a, n), *\right) = \mu_{k,(a,n)}, \qquad p^0_k(*, *) = 1;$$

Figure 1 summarizes and relates the aforedescribed MDP elements, in which a part of the state diagram of a user k MDP model is provided.

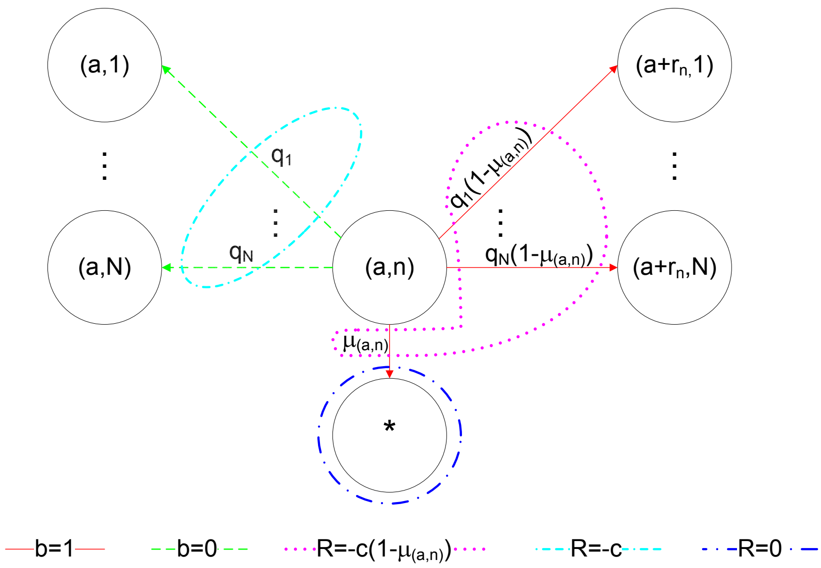

Figure 1: A part of a per-user state diagram.

3.2. Optimization problem

Now we can define the optimization problem for the joint or system MDP model associated with the single-job MDP model presented in the previous subsection. Let Π be the set of all admissible policies for the studied problem. We aim at finding a joint policy $\pi \in \Pi$ that maximizes the aggregate reward starting from the initial time epoch 0 subject to the allocation constraint of serving a single user. Thus, for a given discount factor β, we present the optimization problem we consider as:

$$\max_{\pi \in \Pi} \mathbb{E}_0^{\pi} \left[\sum_{t=0}^{\infty} \sum_{k \in \mathcal{K}} \beta^t R_{k,(a_k(t),n_k(t)) \cup *}^{b_k(t)} \right]$$
$$\text{subject to} \sum_{k \in \mathcal{K}} b_k(t) = 1 \qquad \forall t \tag{1}$$

It is known that the allocation constraint causes intractability in similar problems [19]. In this way, we will design an approximate solution for our problem in the next section.

4. Whittle Index-Based Solution

In this section we provide a closed-form solution for problem (1) based on Whittle methodology [19]. This method consists in obtaining a metric per job state that measures the dynamic service priority of a single user so as to achieve a simple scheduling index rule. With that aim, problem (1) can be relaxed by requiring to serve a job per slot on average as proposed in [19], and further approached by Lagrangian methods [17]. As a result, we obtain the following single-job price-based parametrized optimization problem:

$$\max_{\pi_k \in \Pi} \sum_{t=0}^{\infty} \mathbb{E}_0^{\pi} \beta^t \left[R_{k,(a_k(t),n_k(t))}^{b_k(t)} - v W_{k,(a_k(t),n_k(t))}^{b_k(t)} \right] \tag{2}$$

where the Lagrangian parameter v can be interpreted as the per-slot cost of serving. Note that we are not interested in obtaining an index value for the absorbing state $*$, since, once completed flow transmission, we will give the minimum priority to this state.

Let us define the serving set $\mathcal{F} \subseteq \mathcal{S}_k$, which prescribes to serve a user k if $(a,n) \in \mathcal{F}$, while not to serve this user if $(a,n) \notin \mathcal{F}$. We will refer to states $(a,n) \in \mathcal{F}$ as active and $(a,n) \notin \mathcal{F}$ as passive. Therefore, it is possible to rewrite problem (2) as

$$\max_{\mathcal{F} \subseteq S_k} \mathbb{R}_{k,(a,n)}^{\mathcal{F}} - v \mathbb{W}_{k,(a,n)}^{\mathcal{F}} \tag{3}$$

where $\mathbb{R}_{(a,n)}^{\mathcal{F}}$ and $\mathbb{W}_{(a,n)}^{\mathcal{F}}$ are respectively the expected total reward and expected total work in (a,n) state.

The Whittle index is the break-even value of the Lagrangian parameter v, which measures the expected efficiency of serving in each state. Under Conjecture 1, we assume that the Whittle index exists for each (a,n) state of a user k. In this way, once concluded the main properties of the Whittle index in subsection 4.1, we will provide the methodology to obtain an analytically tractable index rule based on Whittle in subsection 4.2. From now on we omit user label k.

Conjecture 1. *Problem (2) is indexable.*

4.1. Analysis of Whittle index properties

As defined in [11] the Whittle index represents the rate between the marginal reward and the marginal work, where the marginal reward (work) is the difference of the expected reward earned (work required) by serving and not serving at initial state (a,n) and employing $\mathcal{F}$ afterwards. Thus, we can formally write the Whittle index, $v^*_{(a,n)}$, for problem (2) as:

$$v^*_{(a,n)} = v^{\mathcal{F}}_{(a,n)} = \frac{\mathbb{R}^{<1,\mathcal{F}>}_{(a,n)} - \mathbb{R}^{<0,\mathcal{F}>}_{(a,n)}}{\mathbb{W}^{<1,\mathcal{F}>}_{(a,n)} - \mathbb{W}^{<0,\mathcal{F}>}_{(a,n)}} \tag{4}$$

Lemma 1. *For any state (a,n) and under any policy $\mathcal{F}$ we obtain:*

$$v^*_{(a,n)} = v^{\mathcal{F}}_{(a,n)} = \frac{c\mu_{(a,n)} + \beta(1-\mu_{(a,n)}) \sum_{m\in\mathcal{N}} q_m \mathbb{R}^{\mathcal{F}}_{(a+r_n,m)} - \beta \sum_{m\in\mathcal{N}} q_m \mathbb{R}^{\mathcal{F}}_{(a,m)}}{1 + \beta(1-\mu_{(a,n)}) \sum_{m\in\mathcal{N}} q_m \mathbb{W}^{\mathcal{F}}_{(a+r_n,m)} - \beta \sum_{m\in\mathcal{N}} q_m \mathbb{W}^{\mathcal{F}}_{(a,m)}} \tag{5}$$

Proof. From the definition of reward and work we have

$$\mathbb{R}^{\mathcal{F}}_{(a,n)} = \begin{cases} -c(1-\mu_{(a,n)}) + \beta(1-\mu_{(a,n)}) \sum_{m\in\mathcal{N}} q_m \mathbb{R}^{\mathcal{F}}_{(a+r_n,m)} & (a,n)\in\mathcal{F} \\ -c + \beta \sum_{m\in\mathcal{N}} q_m \mathbb{R}^{\mathcal{F}}_{(a,m)} & (a,n)\notin\mathcal{F} \end{cases} \tag{6}$$

$$\mathbb{W}^{\mathcal{F}}_{(a,n)} = \begin{cases} 1 + \beta(1-\mu_{(a,n)}) \sum_{m\in\mathcal{N}} q_m \mathbb{W}^{\mathcal{F}}_{(a+r_n,m)} & (a,n)\in\mathcal{F} \\ \beta \sum_{m\in\mathcal{N}} q_m \mathbb{W}^{\mathcal{F}}_{(a,m)} & (a,n)\notin\mathcal{F} \end{cases} \tag{7}$$

Then we obtain expression (5) by substituting (6) and (7) in (4). □

Nevertheless, so as to achieve a closed-form characterization of the Whittle index (5) in a state (a,n) it is necessary to determine the optimal policy $\mathcal{F}$ of the future states that have influence on the index computation. In this way, first, we need to know if states (a,m) and $(a+r_n,m)$ are active or passive. However, the structure of the active set is not trivial for the problem considered since the combination of size properties and channel properties must be taken into account. We can suppose that a state with a better channel condition will be better than a state with a worse channel condition for certain attained service levels, resulting the first in a higher Whittle index value. Nonetheless, the attained service threshold value that satisfies the previous supposition is not straightforward.

We focus on the important class of size distributions with a Decreasing Hazard rate (DHR), particularly on the Pareto distribution with shape parameter $\alpha > 1$ and scale parameter $\gamma > 0$ whose density function for all $x \geq 0$ is provided in (8). Thus, for a Pareto distribution that belongs to DHR, in concordance to the previous supposition, we can assume that a state with a better channel condition will be better than a state with a worse channel condition until certain attained service due to the DHR property, but without prior knowledge of that attained service threshold.

$$f(x) = \frac{\gamma\alpha}{(1+\gamma x)^{\alpha+1}} \tag{8}$$

Therefore, finding the structure of the active set becomes challenging. In order to achieve fundamental properties of the Whittle index under study, which will be useful to determine the structure of the active set, we have employed an algorithm called $\mathcal{AG}$ (Adaptive Greedy, see [11] for a survey) that computes Whittle indices numerically. If the Whittle index value in the state (a,n) is higher than in the state (a',n'), state (a,n) will be active for $v^*_{(a',n')}$ computation, whereas state (a',n') passive for $v^*_{(a,n)}$ computation.

In this way, we have applied $\mathcal{AG}$-algorithm for our case. We have performed several numerical experiments considering a wide range of configurations. We have reduced experiments for three channel states since the algorithm becomes time-consuming. Being $A := \max_i |\mathcal{A}_i|$, it requires $\mathcal{O}(A^{N+1})$ elementary properties for computing all the index values for each configuration. We have paid attention to the undiscounted case ($\beta = 1$), using $\beta \approx 1$ to avoid indeterminacies; we have normalized Whittle index values by

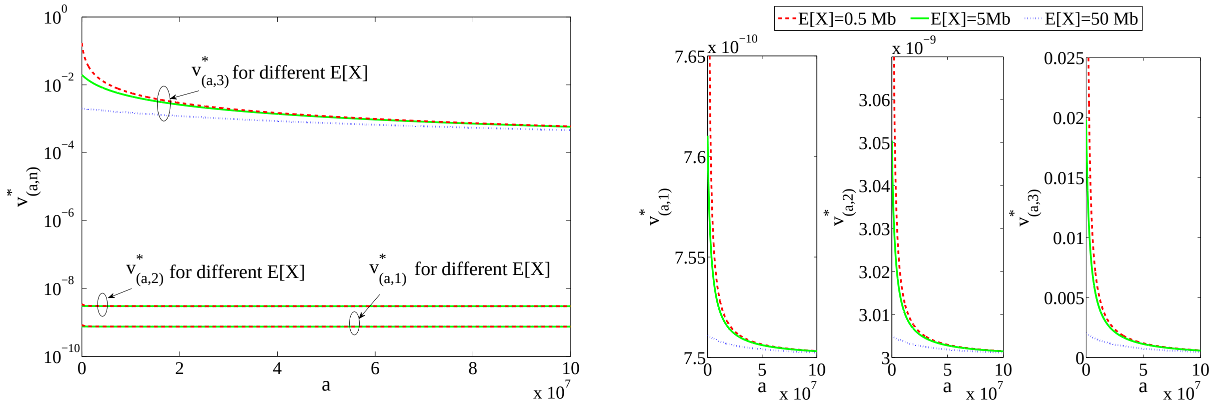

Figure 2: Normalized Whittle indices for a Pareto distribution with $\alpha = 1.5$, $q_1 = q_2 = q_3 = 1/3$, $r_1 = 8.4$ Kb, $r_2 = 16.8$ Kb and $r_3 = 33.6$ Kb.

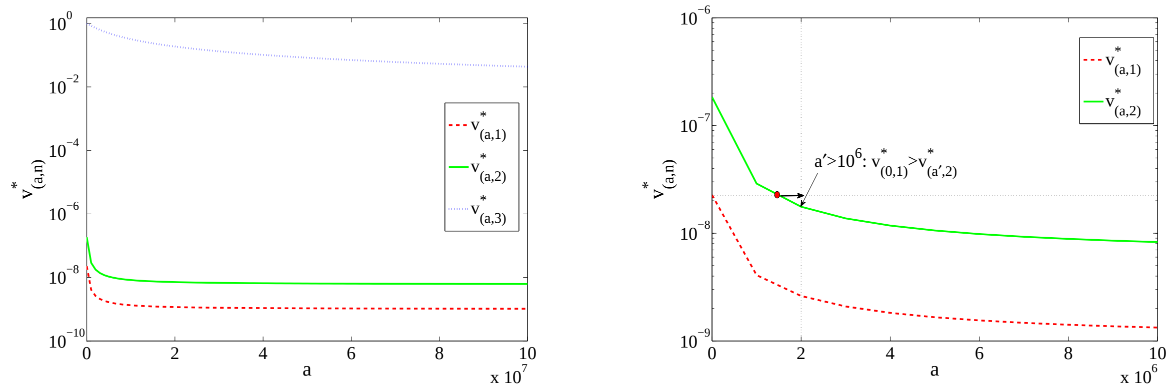

Figure 3: Normalized Whittle indices for a Pareto distribution with $\alpha = 1.5$ and E[X] $= 0.5$ Mb, $q_1 = q_2 = q_3 = 1/3$, $r_1 = 1$ Mb, $r_2 = 2$ Mb and $r_3 = 3$ Mb.

multiplying them by $1 - \beta$ to avoid large values and numerical computation instabilities. Thus, for the performed numerical experiments we conjecture the following properties:

(i) All the index values for the best channel condition are greater than those for the rest of channel conditions: $v^*_{(a,N)} > v^*_{(a',m)} \quad \forall a, a' \in \mathcal{A}, \forall m < N$ (see Figure 2 and Figure 3).

(ii) Generally, being in a better channel condition does not guarantee that the index value is greater than in a worse channel condition; $v^*_{(a,m_1)} > v^*_{(a',m_2)} \quad \forall a, a' \in \mathcal{A}, \forall m_1 < N, m_2 < N, m_1 > m_2$ is not always satisfied (illustrated in the right graph of Figure 3).

(iii) For the same attained service the Whittle index value in a better channel condition is greater: $v^*_{(a,m_1)} > v^*_{(a,m_2)}$ $\forall m_1 > m_2$ (shown in Figure 2 and Figure 3).

(iv) For the same channel condition Whittle index values are decreasing with attained service (observe the right graph of Figure 2).

(v) For the best channel condition the Whittle index is invariant with channel state probability (see the right graph of Figure 4).

(vi) For channel conditions that are not the best the Whittle index value decreases as long as the channel state probabilities of better channel conditions increase (illustrated in the left and middle graphs of Figure 4).

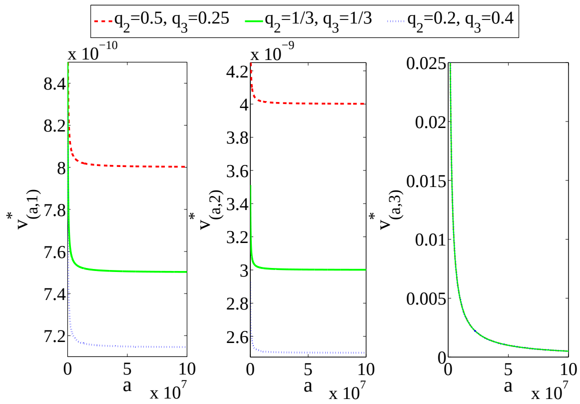

Figure 4: Normalized Whittle indices for different channel state probabilities. A Pareto distribution is used with $\alpha = 1.5$ and E[X] = 5 Mb, $r_1 = 8.4$ Kb, $r_2 = 16.8$ Kb and $r_3 = 33.6$ Kb.

4.2. Closed-form Whittle index characterization

In this subsection we set out to derive a closed-form Whittle index, for which the properties concluded in the previous subsection have been useful for guessing the structure of the active set. If the optimal policy has an easy structure, we can exploit it to obtain a tractable index expression.

In this way, for Whittle index computation (expression (5)), first we need to determine if states (a, m) and $(a + r_n, m)$ are active or passive. According to concluded properties:

- On the one hand, for the same attained service in a better channel condition the index value is greater. Hence, in the index computation of (a, n) state, states with attained service a with a higher channel condition than n will be active, and if channel condition is lower, passive; that is, $(a, m) \in \mathcal{F} \quad \forall m > n$ and $(a, m) \notin \mathcal{F} \quad \forall m \leq n$ are satisfied.

- On the other hand, for a Pareto distribution that belongs to DHR class the Whittle index value is decreasing with attained service for each channel condition. Therefore, states with attained service a or higher than a and with channel condition n or lower will be passive; thus, $(a + r_n, m) \notin \mathcal{F} \quad \forall m \leq n$ is satisfied.

- However, it is not clear if states with channel condition higher than n and with attained service higher than a are active or passive; that is, for states $(a + r_n, m) \quad \forall m > n$ we do not know beforehand from which attained service on and from which channel condition on are active or passive.

Nevertheless, for the best channel condition the active set is totally specified. From concluded properties, states that have influence on the index computation of a state (a, N) are passive, because in these future states the attained service is equal or higher, without channel condition improvement. As shown in Proposition 1, the value of the Whittle index is infinite, which coincides with the result obtained in [3] for exponential sizes.

Proposition 1. *For the best channel condition and a Pareto size distribution the Whittle index for problem (2) in the undiscounted case is given by:*

$$v^*_{(a,N)} = \lim_{\beta \to 1} \frac{c\mu_{(a,N)}}{1 - \beta} = +\infty \tag{9}$$

Proof. For channel condition N $\mathcal{F} = \emptyset$ for all a. We use passive expressions of reward and work, (6) and (7) respectively, and by solving equation systems $\forall m \quad \mathbb{R}^{\mathcal{F}}_{(a,m)} = -c + \beta \sum_{m' \in \mathcal{N}} q_{m'} \mathbb{R}^{\mathcal{F}}_{(a,m')}$ and $\forall m \quad \mathbb{W}^{\mathcal{F}}_{(a,m)} = \beta \sum_{m' \in \mathcal{N}} q_{m'} \mathbb{W}^{\mathcal{F}}_{(a,m')}$ respectively, we achieve $\mathbb{R}^{\mathcal{F}}_{(a,m)} = \mathbb{R}^{\mathcal{F}}_{(a+r_n,m)} = \frac{-c}{1-\beta}$ and

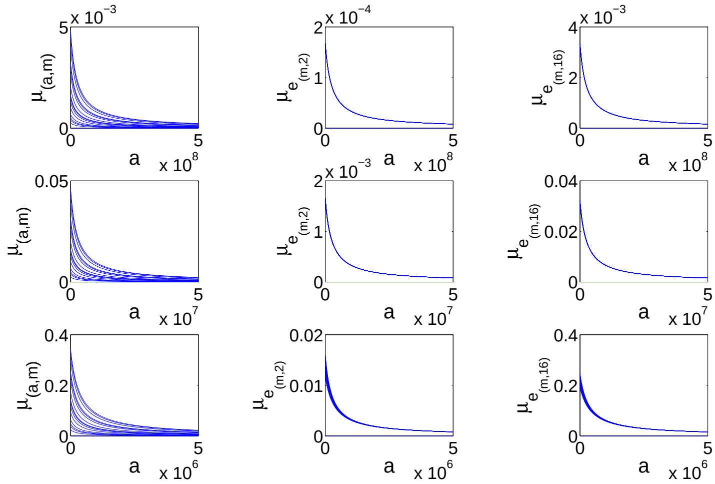

Figure 5: Evolution of μ with attained service (left), μ error for the worst non-null channel condition (middle) and μ error for the best channel condition (right), for different Pareto settings (mean size 50 Mbit, 5 Mbit and 0.5 Mbit top-down, for $\alpha = 1.5$) and different channel conditions.

$\mathbb{W}^{\mathcal{F}}_{(a,m)} = \mathbb{W}^{\mathcal{F}}_{(a+r_n,m)} = 0$. By substituting those reward and work terms in (5) we obtain expression (9). $\square$

Besides, due to the impossibility of obtaining a tractable Whittle index expression for any channel condition, caused by the unclear and difficult active set structure, in the next subsection we propose a heuristic based on an approximate Whittle approach.

4.2.1. An index rule proposal based on a Whittle index approximation

As previously stated, the Whittle index analytical expression for problem (2) is unachievable, except for the best channel condition. In order to obtain a closed-form index expression for any channel condition, we propose an approximate Whittle index solution in the following.

First of all, we will define an exact active set structure. From the previous statements, it is known that in the achievement of the Whittle index of (a, n) state, states $(a, m \leq n) \cup (a + r_n, m \leq n)$ are passive and states $(a, m > n)$ active, but the activity of $(a + r_n, m > n)$ states is unknown. Let's assume that for all $l = 1, 2, \ldots, N$ $\mu_{(a,m)} \approx \mu_{(a+r_l,m)}$ happens. In this way, knowing that $(a, m > n)$ states are active, being $\mu_{(a,m)} \approx \mu_{(a+r_n,m)}$, we can suppose that $(a + r_n, m > n)$ states are also active.

We define the completion probability error, $\mu_{e_{(m,l)}}$, as: $\mu_{e_{(m,l)}} = 1 - \mu_{(a+r_l,m)}/\mu_{(a,m)}$. Note that this error increases with r_l. We present the error analysis related to this approximation in Figure 5, for which for different mean sizes and for 16 channel conditions (rate values available in Table 1, which are typically used in 4G wireless settings), we show the evolution of μ with attained service (graphs on the left), the μ error for the worst non-null channel condition (graphs in the middle), and the μ error for the best channel condition (graphs on the right). As can be observed, for the first two size configurations good results are obtained; in the worst case about a 3% error is achieved. However, while the mean size decreases the error notably increases; for the case that introduces the highest error, $\mu_{e_{(2,16)}}$, when the attained service is the mean size about a 10% error is obtained. Nevertheless, we will show in the performance analysis section that the effects caused by the error of this approximation do not deteriorate the performance.

Once determined the active set, and using the presented μ approximation, we will propose a simplified Whittle index expression. Considering that $\mu_{(a,m)} \approx \mu_{(a+r_n,m)}$, $R_{(a,m)} \approx R_{(a+r_n,m)}$ and $W_{(a,m)} \approx W_{(a+r_n,m)}$ are satisfied. Thus, we can assume that $R_a = \sum_m q_m R_{(a,m)} \approx \sum_m q_m R_{(a+r_n,m)}$ and

$W_a = \sum_m q_m W_{(a,m)} \approx \sum_m q_m W_{(a+r_n,m)}$. In this way, using these simplifications, the Whittle index formulation (5) is simplified to the following expression, denominated by $\tilde{v}^*_{(a,n)}$:

$$\tilde{v}^*_{(a,n)} = \frac{\beta\mu_{(a,n)}(c - R_a)}{1 - \beta\mu_{(a,n)}W_a} \tag{10}$$

Moreover, we could generalize the previous approximation for any attained service increase for any l channel condition, resulting in $R_a = \sum_m q_m R_{(a,m)} \approx \sum_m q_m R_{(a+r_l,m)}$ and $W_a = \sum_m q_m W_{(a,m)} \approx \sum_m q_m W_{(a+r_l,m)}$. Using this last approximation with the previous suppositions and approximations, we obtain the Whittle index approximation presented in Proposition 2, that we call *Attained Service dependant Potential Improvement* (ASPI). We provide the mathematical resolution to obtain the Whittle-index-based closed-form index expression in Appendix A.

Proposition 2. *The formulation of the ASPI index is (11):*

$$\boldsymbol{ASPI} = \tilde{\boldsymbol{v}}^*_{(\boldsymbol{a},\boldsymbol{n})} = \frac{\boldsymbol{c\mu_{(a,n)}}}{\sum\limits_{\boldsymbol{m>n}} \boldsymbol{q_m(\mu_{(a,m)} - \mu_{(a,n)})}} \tag{11}$$

As can be observed in expression (11), the obtained ASPI index is computationally tractable. This allocation strategy is a size-based and channel-aware policy, which depends on the size distribution, channel rates and channel state probabilities. The achieved expression is similar to the one obtained in [3] for the exponential distribution, called Potential Improvement (PI); but in our case the attained service is taken into account, which does not happen for the exponential approach. The index value equals the ratio between the instantaneous completion probability and the expected potential improvement of the instantaneous completion probability. Moreover, as in the original Whitte index, in the best channel condition the ASPI index is infinite, since the summation in the denominator is null. Thus, we summarize the proposed ASPI index rule in Definition 1.

Definition 1. *The ASPI index rule consists in: at every decision slot t,*

- *serving a user in its best channel condition with the highest value of $c\mu_{(a,N)}$;*
- *if there is no user in its best channel condition, serving the user with the highest value of $\frac{c\mu_{(a,n)}}{\sum_{m>n} q_m(\mu_{(a,m)} - \mu_{(a,n)})}$ (using (11)).*

Therefore, the ASPI index rule can be easily implemented. Furthermore, in the next section, we will show that the performance of this approximate Whittle index heuristic is adequate.

5. Performance Evaluation

In this section we analyze the performance of the ASPI index rule proposal, presented in Definition 1. To that end, we compare its achieved mean holding cost with the one obtained with well-known priority policies in several simulation scenarios. Note that, when all the holding costs are equal, the mean holding cost is equivalent to the mean delay, which is a relevant user-level performance metric. Furthermore, we will focus on wireless scenarios, and thus, we will use typical parameter values employed in this kind of networks. Nevertheless, the achieved results will be valid for any time-varying channel context.

Below we provide a brief description of the scheduling index rules used in our experiments:

- $c\mu$-rule, adapted to the attained service, denominated by cost and attained service dependant μ (cASμ): consists in serving the user k with the highest value of $v^{c\mu}_{k,(a,n)} = c_k\mu_{k,(a,n)}$.
- Max Rate (MR) scheduler: serves the user k with the highest value of $v^{\mathrm{MR}}_{k,n} = r_{k,n}$.

- Proportional Fair (PF) scheduler: consists in serving the user with the highest ratio of the current transmission rate and the attained throughput; that is, the highest value of $v^{\mathrm{PF}}_{k,(a,n,d)} = \frac{r_{k,n}}{a_k/d_k}$, where d_k is the time already spent in the system.

In case of ties, these are resolved randomly. As typical in current wireless networks, we use scheduling decisions and transmission slots of 1 ms.

It is well accepted that Internet traffic flows can be suitably modeled by Pareto distributions [16]. We use the Pareto distribution (8) in simulations, considering typical mean sizes employed in wireless settings: 0.5 Mb, 5 Mb and 50 Mb (see [7]). Even though in the analytical framework we do not consider arrivals of new users, we take into account arrivals of new users in simulations, so as to show that the performance of our scheduling proposal is satisfactory even in a more realistic case due to the presence of arrivals. Flows arrive according to a Poisson process with rate λ_k per slot. This arrival rate will determine the network load, ρ. We consider single-class settings, as well as scenarios with two classes of users. Thus, we define $\rho_k = \lambda_k \cdot \frac{E[X_k]}{r_{k,N}}$ per-class k and $\rho = \sum_k \rho_k$; in case of two classes, we have $\rho_1 = \rho_2$ for simplicity.

Moreover, we make use of transmission rates employed in 4G networks (adapted from [14]). The mapping between Channel Quality Indicator (CQI) indices and rates is provided in Table 1. Note that in each scenario we will only use several inputs (columns) from the previous table, depending on the number of channel conditions considered in each setting.

CQI	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
r	0	4.2	6.7	8.4	11.2	16.8	21.8	25.2	26.8	33.6	44.6	50.4	53.7	67.2	75.6	80.6

Table 1: CQIs and corresponding rates (Kb).

In order to guarantee that the achieved performance results are generally valid, we have carefully selected simulation scenarios. For that purpose, on the one hand, we have chosen a wide variety of configurations for different channel, size or/and cost characteristics. On the other hand, we have taken into account scenarios in which the μ error in the approximation of the achievement of ASPI is high. We have analyzed seven relevant settings, whose parameters are collected in Table 2.

Scenario	Channel	Size	c
1	CQI ={3,5}, $q_2 = 0.5$	$\{\alpha = 1.5, \gamma = 4 \cdot 10^{-7}, \mathbb{E}[X] = 5\text{Mb}\}$	1
2	CQI ={3,9}, $q_2 = 0.5$	$\{\alpha = 1.5, \gamma = 4 \cdot 10^{-7}, \mathbb{E}[X] = 5\text{Mb}\}$	1
3	CQI ={3,5}, $q_2 = 0.5$	$\{\alpha = 1.5, \gamma = 4 \cdot 10^{-6}, \mathbb{E}[X] = 0.5\text{Mb}\}$	1
4	See Table 3	$\{\alpha = 1.5, \gamma = 4 \cdot 10^{-7}, \mathbb{E}[X] = 5\text{Mb}\}$	1
5	See Table 4	$\{\alpha = 1.5, \gamma = 4 \cdot 10^{-7}, \mathbb{E}[X] = 5\text{Mb}\}$	1
6	Class 1: See Table 4	$\{\alpha_1 = 1.5, \gamma_1 = 4 \cdot 10^{-7}, \mathbb{E}[X_1] = 5\text{Mb}\}$	(1,1)
	Class 2: See Table 3	$\{\alpha_2 = 1.5, \gamma_2 = 4 \cdot 10^{-7}, \mathbb{E}[X_2] = 5\text{Mb}\}$	
7	See Table 5	$\{\alpha_1 = 1.5, \gamma_1 = 4 \cdot 10^{-7}, \mathbb{E}[X_1] = 5\text{Mb}\}$	(5,1)
		$\{\alpha_2 = 1.5, \gamma_2 = 4 \cdot 10^{-8}, \mathbb{E}[X_2] = 50\text{Mb}\}$	

Table 2: Parameter set in experimental study.

CQI	1	3	5	9	13
q	0.05	0.1	0.2	0.25	0.4

Table 3: CQIs and corresponding increasing channel state probabilities.

CQI	1	3	5	9	13
q	0.4	0.25	0.2	0.1	0.05

Table 4: CQIs and corresponding decreasing channel state probabilities.

Next we show the results obtained in the analyzed scenarios.

CQI	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
q	0.28	0.12	0.09	0.08	0.08	0.08	0.07	0.06	0.05	0.04	0.03	0.01	0.009	0.0005	0.0003	0.0002

Table 5: Channel state probabilities in a 4G network setting.

5.1. Scenario 1: Basic setting

We start analyzing scenarios with only two channel states so as to get fundamental insights into the performance of our scheduling proposal. In the first setting, we consider a typical wireless configuration that takes into account the equiprobable channel case and medium-sized flows. Figure 6 collects mean holding cost results for this setting for different scheduling algorithms under different network loads. As main conclusion, ASPI shows the lowest mean holding cost compared with the rest of policies.

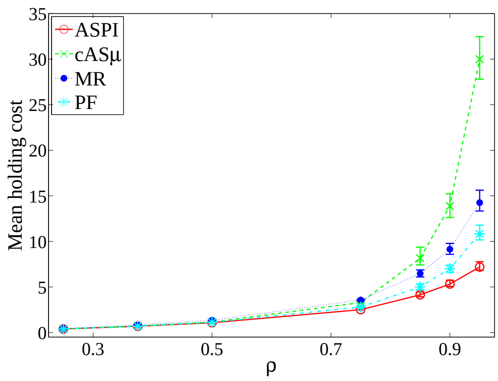

Figure 6: Mean holding cost for Scenario 1.

5.2. Scenario 2: A case with higher error due to rate

In this second family of simulations, the transmission rate of the best channel is four times higher than the one of the bad channel, whereas in the basic setting was the double. Thus, the error caused in the μ approximation for the ASPI index computation is higher. In spite of this higher error due to rate, as can be observed from Figure 7, ASPI still outperforms all the analyzed disciplines.

5.3. Scenario 3: A setting with higher error caused by size

Now we consider the smallest mean size, which introduces a higher μ error for the Whittle index approximation. Despite this higher error caused by size, as depicted in Figure 8, ASPI minimizes the mean holding cost with respect to the rest of the disciplines under study.

5.4. Scenario 4: Optimistic network case

In the following we analyze the first configuration in which the number of channel conditions is higher than two. We consider five channel states, whose transmission rates are multiples of the lowest rate taken into account; the rate of a certain channel condition is twice the one of its smallest contiguous channel condition. We assume that channel state probabilities are increasing with the improvement of channel quality. In this optimistic network situation, we can see from Figure 9 that ASPI achieves the best performance.

5.5. Scenario 5: Pessimistic network case

Contrary to the previous setting, we now assume that channel state probabilities are decreasing with rate in order to contrast the effects of opposite channel state probabilities. Under such a pessimistic network case, as can be appreciated in Figure 10, ASPI introduces the lowest mean holding cost. Note that we omit results from cASμ since it is unstable for this scenario.

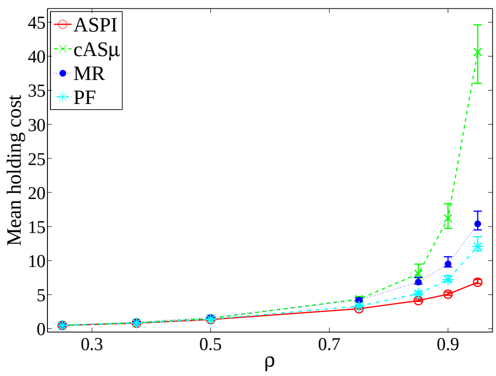

Figure 7: Mean holding cost for Scenario 2.

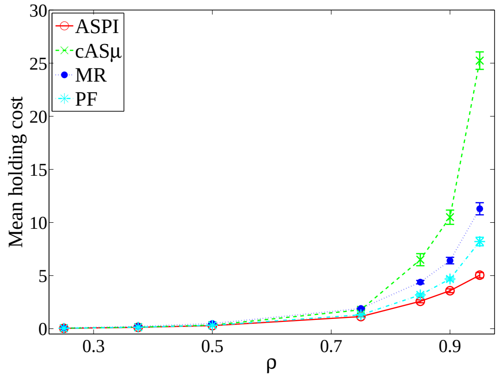

Figure 8: Mean holding cost for Scenario 3.

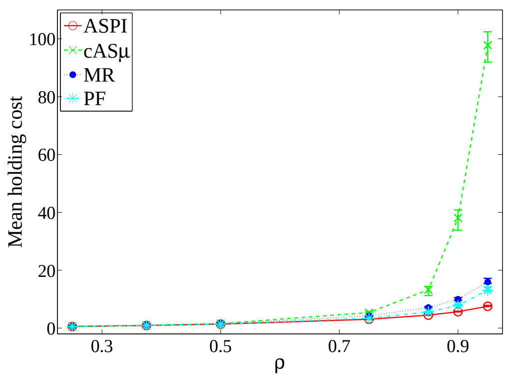

Figure 9: Mean holding cost for Scenario 4.

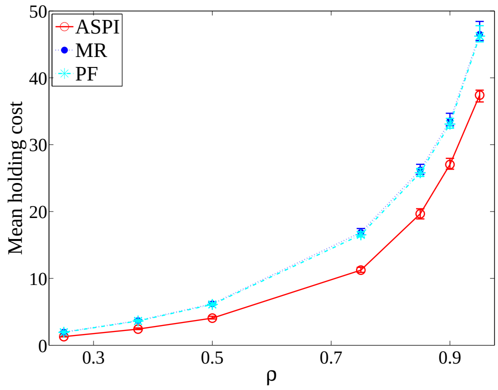

Figure 10: Mean holding cost for Scenario 5.

5.6. Scenario 6: A mixture of optimistic and pessimistic network cases

This setting mixes the previous optimistic and pessimistic single-class cases, resulting in a multiclass scenario with users that differ in channel configuration. From Figure 11 and Figure 12 we conclude that not only ASPI is superior for the aggregate of classes (see Figure 11), but also inside each class (see Figure 12).

5.7. Scenario 7: A 4G network case

Finally, we present the performance analysis of a setting that resembles a real 4G network context. To that end, we employ channel traces obtained from a system-level radio access simulator [9]. This configuration leads to channel state probabilities that decrease with the improvement of channel rate. Apart from that, we consider two classes of users, which differ in size and cost; we give priority to the small-sized class by augmenting its cost value. Results in Figure 13 and Figure 14 indicate that, in this emerging scenario, ASPI outperforms all the policies considered in the mixture of classes (from Figure 13), as well as per-class (from Figure 14).

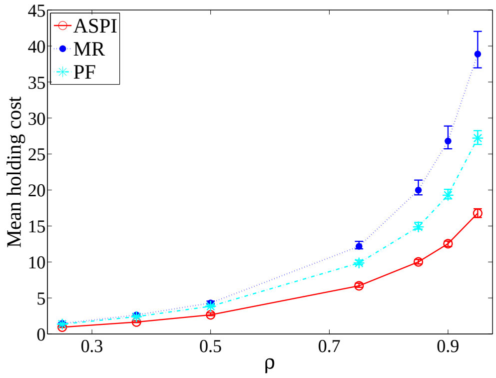

Figure 11: Mean holding cost for Scenario 6.

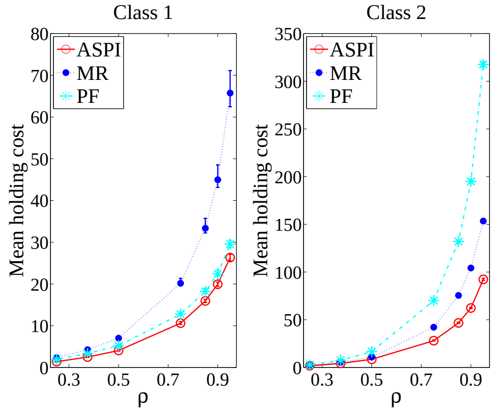

Figure 12: Mean holding cost per class for Scenario 6.

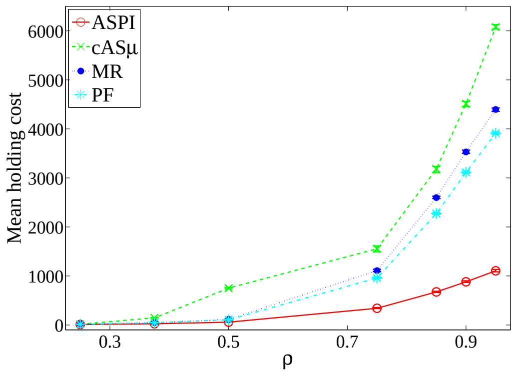

Figure 13: Mean holding cost for Scenario 7.

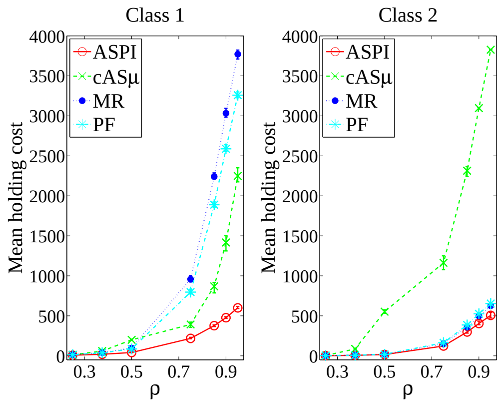

Figure 14: Mean holding cost per class for Scenario 7.

6. Conclusions

In this paper we provide a Whittle-index-based methodology so as to achieve a simple solution for the scheduling problem aimed at minimizing the average holding cost considering flows with general size distribution in randomly time-varying channels. Particularly, we obtain the opportunistic and non-anticipating size-aware ASPI index rule scheduler for realistic DHR size distributions, which is analytically tractable and performs well in relevant simulation scenarios.

Therefore, the results of this work will be useful for network providers in order to guarantee a better quality of service to their customers in time-varying wireless networks. Apart from that, the simplification applied in the Whittle method would be useful to other problems when the gradient of the system components is considerably small among consecutive slots.

Besides, even though simulation results illustrate that the proposed scheduler is stable for $\rho < 1$, it would be interesting to prove its stability.

References

[1] S. Aalto and P. Lassila. Flow-level stability and performance of channel-aware priority-based schedulers. In *Next Generation Internet (NGI), 2010 6th EURO-NF Conference on*, pages 1–8. IEEE, 2010.

[2] S. Aalto, A. Penttinen, P. Lassila, and P. Osti. On the optimal trade-off between SRPT and opportunistic scheduling. In *Proceedings of ACM Sigmetrics*, 2011.

[3] U. Ayesta, M. Erausquin, and P. Jacko. A modeling framework for optimizing the flow-level scheduling with time-varying channels. *Performance Evaluation*, 67:1014–1029, 2010.

[4] U. Ayesta, M. Erausquin, and P. Jacko. Resource Sharing in a Single Server with Time-Varying Capacity. In *Allerton Conference*, 2011.
[5] U. Ayesta, M. Erausquin, M. Jonckheere, and I. M. Verloop. Scheduling in a random environment: Stability and asymptotic optimality. *IEEE/ACM Transactions on Networking*, 21(1):258–271, 2013.
[6] S. Borst. User-level performance of channel-aware scheduling algorithms in wireless data networks. *IEEE/ACM Transactions on Networking*, 13(3):636–647, 2005.
[7] F. Cecchi and P. Jacko. Scheduling of users with Markovian time-varying transmission rates. In *Proceedings of the ACM SIGMETRICS/international conference on Measurement and modeling of computer systems*, pages 129–140. ACM, 2013.
[8] J. C. Gittins, K. Glazebrook, and R. Weber. *Multi-Armed Bandit Allocation Indices.* Wiley-Blackwell, 2011.
[9] J.C Ikuno, M. Wrulich, and M. Rupp. System level simulation of LTE networks. In *Proc. 2010 IEEE 71st Vehicular Technology Conference*, Taipei, Taiwan, May 2010.
[10] J. Kim, B. Kim, J. Kim, and Y. H. Bae. Stability of flow-level scheduling with Markovian time-varying channels. *Performance Evaluation*, 70(2):148–159, 2013.
[11] J. Niño Mora. Dynamic priority allocation via restless bandit marginal productivity indices. *TOP*, 15(2):161–198, 2007.
[12] C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of optimal queueing network. *Mathematics of Operations Research*, 24(2):293–305, 1999.
[13] B. Sadiq and G. de Veciana. Balancing SRPT prioritization vs opportunistic gain in wireless systems with flow dynamics. In *ITC 22*, 2010.
[14] Stefania Sesia, Issam Toufik, and Matthew Baker. *LTE: the UMTS long term evolution.* Wiley Online Library, 2009.
[15] I. Taboada, J. O. Fajardo, F. Liberal, and B. Blanco. Size-based and channel-aware scheduling algorithm proposal for mean delay optimization in wireless networks. In *Communications (ICC), 2012 IEEE International Conference on*, pages 6596–6600. IEEE, 2012.
[16] K. Thompson, G.J. Miller, and R. Wilder. Wide-area Internet traffic patterns and characteristics. *Network, IEEE*, 11(6):10–23, 1997.
[17] V. Visweswaran. Decomposition techniques for MILP: Lagrangian relaxation decomposition techniques for milp: Lagrangian relaxation. In *Encyclopedia of Optimization*, pages 632–638. Springer, 2009.
[18] R. Weber and G. Weiss. On an index policy for restless bandits. *Journal of Applied Probability*, 27(3):637–648, 1990.
[19] P. Whittle. Restless bandits: Activity allocation in a changing world. *A Celebration of Applied Probability, J. Gani (Ed.), Journal of Applied Probability*, 25A:287–298, 1988.

Appendix A. Proof of Proposition 2

Lemma 2. *Suppose that states $(a, m \leq n) \cup (a + r_n, m \leq n)$ are passive and that states $(a, m > n) \cup (a + r_n, m > n)$ are active, and*

$$\tilde{v}^*_{(a,n)} = \frac{\mu_{(a,n)}(c - \beta R_a)}{1 - \beta \mu_{(a,n)} W_a} \tag{A.1}$$

Along with this, if we assume that $R_a = \sum_m q_m R_{(a,m)} \approx \sum_m q_m R_{(a+r_l,m)}$ and $W_a = \sum_m q_m W_{(a,m)} \approx \sum_m q_m W_{(a+r_l,m)}$ for any $l = 1, 2, \ldots N$ value, for the undiscounted case:

$$\tilde{v}^*_{(a,n)} = \frac{c\mu_{(a,n)}}{\sum_{m>n} q_m(\mu_{(a,m)} - \mu_{(a,n)})} \tag{A.2}$$

Proof. Referring to reward elements, using (6):

$$\begin{aligned} &R_{(a,m\leq n)} = -c + \beta R_a \\ &R_{(a,m>n)} = -c(1 - \mu_{(a,m)}) + \beta(1 - \mu_{(a,m)}) \sum_m q'_m R_{(a+r_m,m')} = \\ &= -c(1 - \mu_{(a,m)}) + \beta(1 - \mu_{(a,m)}) R_a \end{aligned} \tag{A.3}$$

And by (A.3),

$$R_a = (1 - \sum_{m>n} q_m)(-c + \beta R_a) + \sum_{m>n} q_m(-c(1-\mu_{(a,m)}) + \beta(1-\mu_{(a,m)})R_a)$$

And by isolating R_a,

$$R_a = \frac{-c(1 - \sum\limits_{m>n} q_m \mu_{(a,m)})}{1 - \beta + \beta \sum\limits_{m>n} q_m \mu_{(a,m)}} \tag{A.4}$$

Analogously, for work elements, using (7):

$$\begin{aligned} W_{(a,m\leq n)} &= \beta W_a \\ W_{(a,m>n)} &= 1 + \beta(1-\mu_{(a,m)}) \sum_m q'_m W_{(a+r_m,m')} = 1 + \beta(1-\mu_{(a,m)}) W_a \end{aligned} \tag{A.5}$$

And by (A.5),

$$W_a = (1 - \sum_{m>n} q_m)\beta W_a + \sum_{m>n} q_m(1 + \beta(1-\mu_{(a,m)})W_a)$$

And isolating W_a,

$$W_a = \frac{\sum\limits_{m>n} q_m}{1 - \beta + \beta \sum\limits_{m>n} q_m \mu_{(a,m)}} \tag{A.6}$$

Substituting (A.4) and (A.6) in (A.1), and simplifying,

$$\begin{aligned} \tilde{v}^*_{(a,d)} &= \frac{c\mu_{(a,n)} - \beta\mu_{(a,n)} \frac{c(-1+\sum_{m>n} q_m \mu_{(a,m)})}{1-\beta+\beta \sum_{m>n} q_m \mu_{(a,m)}}}{1 - \beta\mu_{(a,n)} \frac{\sum_{m>n} q_m}{1-\beta+\beta\sum_{m>n} q_m\mu_{(a,m)}}} = \\ &= \frac{c\mu_{(a,n)}(1 - \beta + \beta \sum\limits_{m>n} q_m \mu_{(a,m)}) - \beta\mu_{(a,n)} c(-1 + \sum\limits_{m>n} q_m \mu_{(a,m)})}{1 - \beta + \beta \sum\limits_{m>n} q_m \mu_{(a,m)} - \beta \mu_{(a,n)} \sum\limits_{m>n} q_m} = \\ &= \frac{c\mu_{(a,n)}}{1 - \beta + \beta \sum\limits_{m>n} q_m(\mu_{(a,m)} - \mu_{(a,n)})} \end{aligned} \tag{A.7}$$

And thus, for $\beta = 1$, expression (A.7) equals (A.2).

□