Optimal Cognitive Transmission Exploiting Redundancy in the Primary ARQ process

Nicolò Michelusi[1*], **Osvaldo Simeone**[2†], **Marco Levorato**[§], **Petar Popovski**[3‡], **Michele Zorzi**[4*]

[*]Department of Information Engineering, University of Padova, Italy, {michelusi,zorzi}@dei.unipd.it
[†]CWCSPR, New Jersey Institute of Technology, New Jersey, USA, osvaldo.simeone@njit.edu
[‡]Department of Electronic Systems, Aalborg University, Denmark, petarp@es.aau.dk
[§]University of Southern California & Stanford University, USA, levorato@stanford.edu

***Abstract*—Cognitive radio technology enables the coexistence of Primary (PUs) and Secondary Users (SUs) in the same spectrum. In this work, it is assumed that the PU implements a retransmission-based error control technique (ARQ). This creates an inherent redundancy in the interference created by primary transmissions to the SU. We investigate secondary transmission policies that take advantage of this redundancy. The basic idea is that, if a Secondary Receiver (SR) learns the Primary Message (PM) in a given primary retransmission, then it can use this knowledge to cancel the primary interference in the subsequent slots in case of primary retransmissions, thus achieving a larger secondary throughput. This gives rise to interesting trade-offs in the design of the secondary policy. In fact, on the one hand, a secondary transmission potentially increases the secondary throughput but, on the other, causes interference to the reception of the PM at the Primary Receiver (PR) and SR. Such interference may induce retransmissions of the same PM, which plays to the advantage of the secondary user, while at the same time making decoding of the PM more difficult also at the SR and reducing the available margin on the given interference constraint at the PR. It is proved that the optimal secondary strategy prioritizes transmissions in the states where the PM is known to the SR, due to the ability of the latter to perform interference mitigation and obtain a larger secondary throughput. Moreover, when the primary constraint is sufficiently loose, the Secondary Transmitter should also trasmit when the PM is unknown to the SR. The structure of the optimal policy is found, and the throughput benefit of the proposed technique is shown by numerical results.**

***Index Terms*—Cognitive radio networks, dynamic resource allocation, Markov decision processes, interference, ARQ**

I. Introduction

Spectrum licensing has been traditionally used to protect wireless systems against mutual interference. While effective, this approach has also led to scarce utilization of the available spectrum resources [1]–[3]. Cognitive networks hold the promise to improve the spectral efficiency of wireless networks with respect to conventional licensing, by allowing the coexistence of Primary (licensed) and Secondary (unlicensed) Users (PUs and SUs, respectively) on the same radio band. To this end, cognitive radios are capable of sensing the radio environment, collecting information about the presence of active transmitters and potentially also further details about such transmissions such as the codebooks or even the specific messages. This information is used by the cognitive radios to make decisions and adapt their operation so as to optimize the network utilization, while limiting interference to the PUs [4].

There has been extensive research in the area of cognitive radio in the past few years, following different lines of inquiry. Among the most popular approaches, in the information theory community, cognitive network models have been studied by assuming a genie-aided SU that has non-causal access to the whole or part of the active Primary Message (PM) [5]–[7]. While this assumption allows for analysis of information-theoretic optimal stragies, it is not able to capture critical aspects of a cognitive network, such as imperfect sensing. Another line of inquiry is resource management where specific transmission strategies are considered and optimized using various tools from stochastic optimization or machine learning, see, e.g., [8] and references therein. This approach allows to take into account non-idealities in the system and to consider network constraints, such as delay or other QoS guarantees.

In this paper, we take the latter approach and consider the problem of optimizing the secondary transmission strategy in the presence of a primary system that employs a retransmission-based error control technique (ARQ). The SU wishes to maximize its own throughput while satisfying an interference constraint on the primary reception. Primary retransmissions give rise to redundancy in the interference generated by the PU to the SU, in the form of copies of the same primary packet transmitted over subsequent time-slots due to the use of ARQ. This can be exploited by the SU, by decoding the PM at a given primary transmission and exploiting this side information in the following primary ARQ rounds to cancel interference from the PU. This idea was first put forth in [9], where the authors propose several protocols in which the primary ARQ process is limited to one retransmission. In this work, we extend the analysis to multiple primary ARQ rounds. This gives rise to interesting trade-offs in the design of the secondary policy that do not appear

[1]N. Michelusi has been partially supported by a scholarship granted by "Fondazione Ing. Aldo Gini".

[2]The work of O. Simeone has been supported by the U.S. N.S.F. under grant CCF-0914899.

[3]P. Popovski has been partially supported by the EliteForsk program of the Danish Research Council.

[4]M. Zorzi has been partially supported by the European Commission within the Seventh Framework Program, under the SAPHYRE project (Grant Agreement no. 248001).

in a one-retransmission scenario. In fact, on the one hand, a secondary transmission potentially increases the secondary throughput but, on the other, it also causes interference to the reception of the PM at the Primary Receiver (PR) and SR. Such interference may induce retransmissions of the same PM in following rounds, which plays to the advantage of the secondary user, while at the same time making decoding of the PM more difficult also at the SR and reducing the available margin on the given interference constraint at the PR.

We formulate the optimization problem as a Markov Decision Process (MDP) following an approach similar to [10]. Reference [10] tackles a similar problem, in which, however, the SR is not allowed to perform interference cancellation based on decoding of the PM. This aspect plays instead a central role here. We prove that there are two operational regimes of the SU, depending on the level of interference that the SU is allowed to generate at the PU. Specifically, in the Low Interference Regime, where the SU is allowed to produce interference only below a certain threshold, the SU transmits exclusively when its receiver knows the PM, and stays idle otherwise. In fact, when the PM is known to the SR, the SU can achieve a larger secondary throughput by being able to cancel interference from the PU in such cases. In the High Interference Regime, where the SU is allowed to produce interference above the threshold, the SU transmits also in the states where its receiver does not know the PM, according to a specific structure that is explictly characterized.

The rest of this paper is organized as follows. Section II defines the system model under consideration, and the operation of PU and SU. In Section III we define the performance criteria for the system under consideration, along with the optimization problem. This is then reformulated in a MDP framework in Section IV. The main result, which characterizes the structure of the optimal secondary transmission policy, and an intuitive explanation of the optimal policy are given in Section V. In Section VI we present an algorithm to determine the optimal policy. In Section VII we present numerical results that validate the effectiveness of the proposed technique. Finally, Section VIII concludes the paper.

II. System Model

We consider the network depicted in Fig. 1, which represents a Primary and a Secondary Transmitters (PT and ST

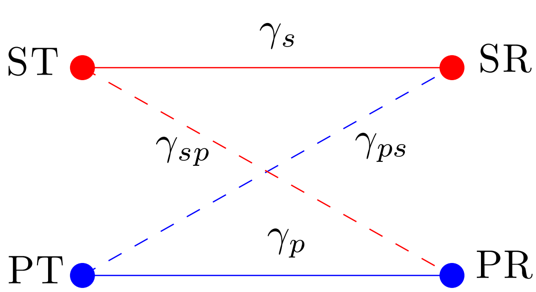

Figure 1. System model

respectively), transmitting packets to their respective destinations, PR and SR, through channels with instantaneous gain γ_p and γ_s, respectively. Transmission by a source interferes with the other ongoing communication through the interfering channels with gain γ_{ps} and γ_{sp}, respectively. The channel gains are modeled as quasi-static (i.e., time is divided into slots of fixed duration, during which the channel gains remain constant) and i.i.d. (independent identically distributed) processes.

We assume that no Channel State Information (CSI) is available at the PT/ST, so that the latter cannot adapt their transmission parameters, such as rate and power, to the current channel state. As a consequence, transmissions may suffer from outage when the selected rate is not supported by the current channel quality. We analyze in detail the operation of PU and SU and the corresponding outage performance at the PR/SR in the following two sub-sections.

A. PU operation

In order to improve reliability, the PU uses a retransmission-based error control technique, with a maximum number of transmissions equal to $T \geq 1$ (i.e., the same packet can be transmitted at most T times, after which it is dropped). After each transmission, the PR sends an acknowledgment back to the PT, in order to inform it about the transmission outcome. We define the *ARQ state* $t \in \mathbb{N}(1, T)$ as the retransmission index associated to the current packet transmitted by the PT, where $\mathbb{N}(n_0, n_1) = \{t \in \mathbb{N}, n_0 \leq t \leq n_1\}$ for $n_0 \leq n_1 \in \mathbb{N}$.

As a consequence of its activity, the ST generates interference over the channel link γ_{sp}, thus degrading the PU outage performance. We denote by $\rho_{p0}(R_p)$ and $\rho_{p1}(R_p)$ the outage probabilities at the PR when the ST is silent and transmits, respectively, as a function of the primary transmission rate R_p, measured in bits/s/Hz. Assuming that the PU is oblivious to the SU activity, so that secondary transmissions are treated as noise at the PR, we have:

$$\begin{cases} \rho_{p0}(R_p) = P\left(R_p > C\left(\gamma_p P_p\right)\right) \\ \rho_{p1}(R_p) = P\left(R_p > C\left(\frac{\gamma_p P_p}{1+\gamma_{sp} P_s}\right)\right) \end{cases} \quad (1)$$

where we have defined $C(x) = \log_2(1 + x)$ as the capacity of the Gaussian channel with SNR x at the receiver, P_p and P_s are the primary and secondary transmission powers, respectively, and the noise variance at the receiver is assumed without loss of generality to be equal to one. Notice that in general we have $\rho_{p1}(R_p) \geq \rho_{p0}(R_p)$.

B. SU operation

Along with the PU, we consider a ST transmitting packets to its respective destination SR. Unlike the PU (except in the limit case $T = 1$), it is assumed that the SU uses a "best effort" approach, by which a packet is discarded in case of transmission failure. Moreover, we assume that the SU has perfect knowledge of the PU parameters, such as length of the ARQ transmission window T and codebook used. This is coherent with the common characterization of the PU as a legacy system, and of the SU as an opportunistic system. By overhearing the feedback messages from the PR, the SU can thus track the primary ARQ state $t \in \mathbb{N}(1,T)^1$, and use this knowledge to perform decisions based on the current state of the PU. In particular, primary transmissions generate interference to the SU, thus degrading its performance. However, due to the use of ARQ, they introduce redundancy in the radio channel, in the form of copies of the same packet transmitted over subsequent time-slots. By decoding the PM, and by being able to track the primary ARQ state, the SR can thus exploit this redundancy in the subsequent primary ARQ rounds to achieve a larger secondary throughput via primary interference cancellation. We then assume that the SR tries to decode the PM, whenever the latter is unknown. In the following, $\phi \in \{0,1\}$ denotes the *SR state variable*, where $\phi = 1$ if the SR knows the PM, and $\phi = 0$ otherwise. Depending on whether the SR knows the PM or not, we can distinguish between two different operational regimes of the ST, from the perspective of transmission rate, outage behavior and accrued throughput.

1) PM Unknown to the SR ($\phi = 0$):

If the SR does not know the PM, it attempts to decode it in the present time-slot. The ability of the SR to decode the PM depends on the link quality and also on the secondary activity. In fact, secondary transmissions represent interference to the SR for the purpose of decoding the PM. Therefore, secondary activity must strike an ideal trade-off between secondary data transmission and decoding of the PM at the SR, which enables interference cancellation in the subsequent primary ARQ rounds. We define $\alpha_{s0}(R_p) = P\left[R_p \leq C(P_p\gamma_{ps})\right]$ and $\alpha_{s1}(R_{s0}, R_p)$ as the probability of decoding the PM at the SR in any time-slot, assuming the ST is silent and transmits, respectively, as a function of R_p and R_{s0}. Notice that the above implies that the SR does not exploit previous primary transmissions of the same packet to decode the PM at the current time-slot, using techniques such as maximum ratio combining [11]. Notice that in general we have $\alpha_{s0}(R_p) \geq \alpha_{s1}(R_{s0}, R_p)$.

We now analyze the throughput achievable in the SR state $\phi = 0$. We assume that secondary transmissions are performed with power P_s and rate R_{s0} (measured in bits/s/Hz). When the ST transmits, the accrued average throughput is denoted by $T_{s0}(R_{s0}, R_p) = R_{s0}\left(1 - \rho_{s0}(R_{s0}, R_p)\right)$, where $\rho_{s0}(R_{s0}, R_p)$ is the outage probability as a function of the primary and secondary rates R_p and R_{s0}, respectively [2].

2) PM Known by the SR ($\phi = 1$):

If the SR knows the PM (learned during previous primary ARQ rounds), secondary transmissions are performed with power P_s and rate R_{s1}. Due to the ability of the SR to completely cancel interference from the PU, the accrued average throughput is denoted by $T_{s1}(R_{s1}) = R_{s1}\left(1 - \rho_{s1}(R_{s1})\right)$, where $\rho_{s1}(R_{s1}) = P\left(R_{s1} > C(\gamma_s)\right)$ is the outage probability [13]. Since the choice of R_{s1} does not affect the outage behavior at the PR (1) and the evolution of the ARQ process (the PR treats secondary transmissions as noise), we can assume without loss of generality that the transmission rate R_{s1} is chosen in such a way as to maximize $T_{s1}(R_{s1})$. Notice that the same argument cannot be applied to R_{s0}, since the choice of this parameter reflects a trade-off between helping the SR to decode the PM and maximizing the throughput $T_{s0}(R_{s0}, R_p)$. These two objectives are conflicting, since the probability of decoding the PM at the SR is maximized by choosing $R_{s0} = 0$, which however would give $T_{s0}(0, R_p) = 0$. As a consequence of the optimization over R_{s1}, we then have the following inequality.

$$T_{s1}(R_{s1}) \geq T_{s1}(R_{s0}) \geq T_{s0}(R_{s0}, R_p) \tag{2}$$

The first inequality is due to the fact that R_{s1} is the rate which maximizes the secondary throughput $T_{s1}(R_s)$, the second to the fact that $\rho_{s1}(R_{s0}) \leq \rho_{s0}(R_{s0}, R_p)$.

In the following, we define $\tau = \frac{T_{s0}(R_{s0}, R_p)}{T_{s1}(R_{s1})}$. This represents the efficiency of the ST in delivering information to the SR when the latter does not know the PM, compared to when it knows it. From (2) we then have $\tau \leq 1$.

III. Performance Criteria and Problem Formulation

The ST follows a generic past-dependent policy μ, whose action set is denoted by $\mathcal{A} = \{0,1\}$, where actions 0 and 1 correspond to the ST staying silent or transmitting, respectively.

The following performance metrics will be used to formulate our optimization problem. We define the long term average throughput achieved by the PU under policy μ as

$$\mathcal{T}_p(\mu) = \lim_{N \to +\infty} \inf \frac{1}{N} \mathbb{E}\left[\sum_{n=0}^{N-1} \mathbf{1}\left(\Psi_P^n(\mu)\right) R_p\right] \tag{3}$$

where $\mathbf{1}(\cdot)$ is the indicator function, and $\Psi_P^n(\mu)$ is the event corresponding to a primary packet successfully delivered to its destination in time-slot n. Similarly, we define the long term average throughput achieved by the ST under policy $\mu \in \mathcal{U}$ as

$$\mathcal{T}_s(\mu) = \lim_{N \to +\infty} \inf \frac{1}{N} \mathbb{E}\left[\sum_{n=0}^{N-1} \mathbf{1}\left(\Psi_S^n(\mu)\right) R_{s\phi_n}\right] \tag{4}$$

where $\Psi_S^n(\mu)$ and ϕ_n are the event corresponding to a secondary packet successfully delivered to its destination, and the SR state variable in time-slot n, respectively. These metrics represent the average number of bits per second correctly delivered by the PT and ST to their respective destinations, when the SU is using policy μ. Finally, we define the long term average secondary transmission power under policy μ as

$$\mathcal{P}_s(\mu) = \lim_{N \to +\infty} \sup \frac{1}{N} \mathbb{E}\left[\sum_{n=0}^{N-1} \mathbf{1}\left(\Upsilon_S^n(\mu)\right) P_s\right] \tag{5}$$

where $\Upsilon_S^n(\mu)$ is the event corresponding to the ST transmitting in time-slot n.

[1]We assume for simplicity that feedback messages are not subject to transmission failure, so that both the PU and the SU are able to perfectly track the ARQ state process.

[2]A closed form expression of $\rho_{s0}(R_{s0}, R_p)$, $\alpha_{s0}(R_p)$ and $\alpha_{s1}(R_{s0}, R_p)$ for Rayleigh Fading channels using joint decoding, as a function of the average SNR of the channel links can be found in [12]

In this work, we investigate the scenario where the goal of the ST is to determine the optimal secondary transmission policy maximizing its own average throughput, subject to a constraint on the maximum throughput loss at the PU, and on the maximum average secondary transmission power. This can be stated as:

$$\mu^*\left(\mathcal{T}_p^{(th)},\mathcal{P}_s^{(th)}\right) = \arg\max_\mu \mathcal{T}_s(\mu,\tau) \quad \text{subject to} \begin{cases} \mathcal{T}_p(\mu) \geq \mathcal{T}_p^{(th)} \\ \mathcal{P}_s(\mu) \leq \mathcal{P}_s^{(th)} \end{cases} \tag{6}$$

where $\mathcal{T}_p^{(th)} \in [R_p(1-\rho_{p1}), R_p(1-\rho_{p0})]$ and $\mathcal{P}_s^{(th)} \in [0, P_s]$ represent the primary throughput and the secondary power constraints, respectively.

The idea behind this problem is that the ST, as a consequence of its activity, which is governed by policy μ, induces a perturbation on the evolution of the primary ARQ process $\{t_n, n=0,\dots,+\infty\}$ and of the SR state $\{\phi_n, n=0,\dots,+\infty\}$. The ST can thus control to its own advantage the evolution of the ARQ and SR state processes by appropriately choosing policy μ so as to maximize its own throughput, while limiting the interference to the PU and the average secondary power consumption.

We use the following set of assumptions, which hold true in most practical scenarios:

$$\begin{cases} 0 < \rho_{p0}(R_p) < \rho_{p1}(R_p) < 1 \\ 0 < \alpha_{s1}(R_{s0}, R_p) < \alpha_{s0}(R_p) < 1 \\ \tau < 1 \\ T \geq 2 \end{cases} \tag{7}$$

The cases $\tau = 1$ or $T = 1$ are trivial. In fact, when $\tau = 1$ the SR does not experience a throughput loss when the PM is unknown, therefore any policy achieving the constraints with equality is optimal, and in particular the optimal policy described in this work. The case $T = 1$ corresponds to the PU using no ARQ: the SU transmits with a fixed transmission probability (this policy is a degenerate case of the optimal policy described in this work). As for the other parameters, the general case where the inequalities are not strict, while complicating the proofs, does not provide any additional insight to our analysis, and is thus omitted. We refer the interested reader to [14] for further details.

In the following, for the sake of notational convenience, we omit the dependence of the parameters defined above on the primary and secondary transmission rates.

IV. Stochastic Modeling of the Network

Let

$$\begin{aligned} \mathcal{S}_0 &= \{(t,0), t \in \mathbb{N}(1,T)\} \\ \mathcal{S}_1 &= \{(t,1), t \in \mathbb{N}(2,T)\} \end{aligned} \tag{8}$$

be the set of network states where the SR does not and does know the PM, respectively. The network is in state $(t,\phi) \in \mathcal{S}_\phi$ when the PU is in ARQ state t and the SR is in state ϕ. The state space of the network is then given by $\mathcal{S} = \mathcal{S}_0 \bigcup \mathcal{S}_1$. (The reason why state $(t,\phi) = (1,1) \notin \mathcal{S}$ is that ARQ state $t = 1$ corresponds to a fresh primary transmission, and therefore it is not possible for the SR to know the PM in advance).

We can view $\{(\Theta_n, a_n), n = 0,\dots,+\infty\}$ as a Markov Decision Process [15], where $\Theta_n \in \mathcal{S}$ is the network state in time-slot n, and $a_n \in \mathcal{A}$ is the action taken by the ST in time-slot n according to some policy μ. It can be shown that the solution of the problem (6) is a randomized stationary policy [16]. Therefore, $\mu : \mathcal{S} \mapsto [0,1]$ maps the network state $\Theta \in \mathcal{S}$ to the probability that the ST takes the actions in $\mathcal{A}$. Specifically, let $\mu_\phi(t)$ be the probability that the ST takes action $1 \in \mathcal{A}$ under state $(t,\phi) \in \mathcal{S}$, i.e., $\mu_\phi(t)$ represents the secondary transmission probability in state (t,ϕ), i.e., ARQ state t and SR state ϕ.

We define $\mathcal{U}_0$ and $\mathcal{U}_1$ as the sets of all the randomized stationary policies μ_0 and μ_1, i.e.

$$\begin{aligned} \mathcal{U}_0 &= \{\mu_0 : \mathbb{N}(1,T) \mapsto [0,1]\} \\ \mathcal{U}_1 &= \{\mu_1 : \mathbb{N}(2,T) \mapsto [0,1]\} \end{aligned} \tag{9}$$

The set of all the randomized stationary secondary policies $\mu = (\mu_0, \mu_1)$ is then defined as $\mathcal{U} \equiv \mathcal{U}_0 \times \mathcal{U}_1$.

Under the given class of stationary policies $\mathcal{U}$, the network state can be modeled as a Homogeneous Markov Process $\{\Theta_n, n = 0,\dots,+\infty\}$ taking values in the state space $\mathcal{S}$, where $\Theta_n = (t_n, \phi_n) \in \mathcal{S}$ corresponds to the PT performing the t_n-th transmission of a packet, with the SR in state ϕ_n in time-slot n.

Let $\pi_\mu : \mathcal{S} \mapsto [0,1]$ be the steady state distribution of the system under stationary policy $\mu \in \mathcal{U}$. We now restate the long term average primary and secondary throughputs, and the average secondary transmission power under policy μ, introduced in Section III, taking into account the stationarity assumption of μ. It is easy to prove, by averaging the instantaneous expected cost/reward in each state over the event of a secondary transmission, and weighting it by its steady state probability, that these metrics are given by

$$\begin{cases} \mathcal{T}_s(\mu) = T_{s0}\sum_{t=1}^T \pi_\mu(t,0)\mu_0(t) + T_{s1}\sum_{t=2}^T \pi_\mu(t,1)\mu_1(t) \\ \mathcal{P}_s(\mu) = P_s\left(\sum_{t=1}^T \pi_\mu(t,0)\mu_0(t) + \sum_{t=2}^T \pi_\mu(t,1)\mu_1(t)\right) \\ \mathcal{T}_p(\mu) = R_p\sum_{t=1}^T \pi_\mu(t,0)\left[1 - (1-\mu_0(t))\rho_{p0} - \mu_0(t)\rho_{p1}\right] \\ \qquad + R_p\sum_{t=2}^T \pi_\mu(t,1)\left[1 - (1-\mu_1(t))\rho_{p0} - \mu_1(t)\rho_{p1}\right] \end{cases}$$

Now, we define a function $\mathcal{W}(\mu,\nu), \mu \in \mathcal{U}, \nu \in [0,1]$ as

$$\mathcal{W}(\mu,\nu) = \nu\sum_{t=1}^T \pi_\mu(t,0)\mu_0(t) + \sum_{t=2}^T \pi_\mu(t,1)\mu_1(t) \tag{10}$$

Notice that, since $\nu \in [0,1]$, $\sum_{(t,\phi)\in\mathcal{S}} \pi_\mu(t,\phi) = 1$ and $\mu_\phi(t) \in [0,1]$, then $\mathcal{W}(\mu,\nu) \in [0,1]$. It is then easy to prove the following:

$$\begin{cases} \mathcal{T}_s(\mu) = T_{s1}\mathcal{W}(\mu,\tau) \\ \mathcal{P}_s(\mu) = P_s\mathcal{W}(\mu,1) \\ \mathcal{T}_p(\mu) = R_p\left[1 - (1-\mathcal{W}(\mu,1))\rho_{p0} - \mathcal{W}(\mu,1)\rho_{p1}\right] \end{cases} \tag{11}$$

To avoid confusion, since we express the primary and secondary throughputs and the secondary power as a function of $\mathcal{W}(\mu,\nu)$, for a specific value of $\nu \in [0,1]$, in the following we refer to $\mathcal{W}(\mu,\tau)$ as the *secondary throughput* (normalized to T_{s1}), and to $\mathcal{W}(\mu,1)$ as the *secondary access rate*. In fact, $\mathcal{W}(\mu,1)$ represents the long term average number of secondary transmissions per time-slot. When convenient, we will also make explicit the dependence of $\mathcal{W}(\mu,\nu)$ on the secondary policy $\mu=(\mu_0,\mu_1),\mu_0\in\mathcal{U}_0,\mu_1\in\mathcal{U}_1$, by rewriting $\mathcal{W}(\mu,\nu)\equiv\mathcal{W}(\mu_0,\mu_1,\nu)$.

The equivalence (11) is due to the fact that, for any secondary transmission, the ST transmits with power P_s, whereas the PU experiences a throughput loss equal to $R_p(\rho_{p1}-\rho_{p0})$, due to an increase of the outage probability. The average throughputs and secondary power are then obtained by averaging over the secondary access rate.

As a consequence of (11) we have:

Lemma 1. The optimization problem (6) can be equivalently restated as

$$\mu^{*(\epsilon)}=\arg\max_{\mu\in\mathcal{U}}\mathcal{W}(\mu,\tau) \quad (12)$$
$$\text{s.t.}\ \mathcal{W}(\mu,1)\leq\min\left\{\frac{R_p(1-\rho_{p0})-\mathcal{T}_p^{(th)}}{R_p(\rho_{p1}-\rho_{p0})},\frac{\mathcal{P}_s^{(th)}}{P_s}\right\}\equiv\epsilon$$

Remark 1. The above Lemma states that maximizing the secondary throughput with a constraint on the maximum secondary power and primary throughput loss is equivalent to maximizing the secondary throughput, with a constraint on the maximum secondary access rate.

Remark 2. Notice that, since $\mathcal{T}_p^{(th)}\in[R_p(1-\rho_{p1}),R_p(1-\rho_{p0})]$ and $\mathcal{P}_s^{(th)}\in[0,P_s]$, then $\epsilon\in[0,1]$. This also agrees with the fact that the secondary access rate $\mathcal{W}(\mu,1)\in[0,1]$.

V. Analysis and Network Optimization

In this Section we overview the structure of the optimal policy $\mu^{*(\epsilon)}=\left(\mu_0^{*(\epsilon)},\mu_1^{*(\epsilon)}\right)\in\mathcal{U}$, as a function of the constraint on the secondary access rate $\epsilon\in[0,1]$.

We define $\bar{0}_\phi\in\mathcal{U}_\phi$ and $\bar{1}_\phi\in\mathcal{U}_\phi$ as the secondary policies in state ϕ with the following structure

$$\begin{cases}\bar{0}_\phi(t) &= 0 \quad \forall t\in\mathbb{N}(1,T)\\ \bar{1}_\phi(t) &= 1 \quad \forall t\in\mathbb{N}(1,T)\end{cases} \quad (13)$$

Therefore, if the SU uses policy $\bar{0}_\phi$ ($\bar{1}_\phi$) in state $\phi\in\{0,1\}$, then it always stays silent (it always transmits) when its receiver is operating in state ϕ.

Moreover, we define $\mu_0(n_l,n_r)\in\mathcal{U}_0$ as a policy, parameterized by $(n_l,n_r)\in\{(x,y)\in[0,T]^2:x+y\leq T\}$, with the following structure.

$$\mu_0(n_l,n_r,t)=\begin{cases}1 & t\leq\lfloor n_l\rfloor\\ n_l-\lfloor n_l\rfloor & t=\lfloor n_l\rfloor+1\\ 1 & t>T-\lfloor n_r\rfloor\\ n_r-\lfloor n_r\rfloor & t=T-\lfloor n_r\rfloor\\ 0 & \text{otherwise}\end{cases} \quad (14)$$

where $\lfloor x\rfloor$ denotes the floor operator (largest integer less than or equal to x). With a slight abuse of notation, in the following we refer to $\mu_0(n_l,n_r)$ and $\mu_0(n_l,n_r,t)$ as the policy with the structure defined above and the corresponding transmission probability in state $(t,0)$, respectively, and to μ_0 and $\mu_0(t)$ as a generic randomized stationary policy in $\mathcal{U}_0$ and the corresponding transmission probability in state $(t,0)$, respectively. As an example, the structure of policy $\mu_0(n_l,n_r)$ is depicted in Fig. 2 for $T=7$. The parameters n_l and n_r correspond to the areas covered by the leftmost and rightmost grey regions, respectively. The transmission probability in each ARQ state then corresponds to the area covered by the grey region in each bin, relative to the total bin area. The integer part of n_l (respectively, n_r) represents the total number of states in which the ST always transmits, concentrated in the initial (final) primary ARQ states where the SR does not know the PM, whereas the residual part $n_l-\lfloor n_l\rfloor$ ($n_r-\lfloor n_r\rfloor$) represents the transmission probability in the ARQ state immediately succeeding (preceding) the sequence of states in which the ST always transmits. For example, in Fig. 2, where $n_l=2$ and $n_r=2.5$, the ST always transmits in the two initial and final ARQ states when its receiver does not know the PM, transmits with probability 0.5 in ARQ state $t=5$ and never transmits in ARQ state $t=3$.

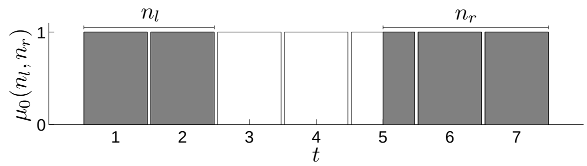

Figure 2. Structure of policy $\mu_0(n_l,n_r)$, $T=7$, $n_l=2$, $n_r=2.5$

We are now ready to state the main result of this work in the following sub-section.

A. Optimal Policy

According to the optimal policy, we distinguish two operational regimes of the PU/SU network, depending on the level of interference the SU is allowed to generate to the PU. Specifically, if $\epsilon>\mathcal{W}(\bar{0}_0,\bar{1}_1,1)$, i.e., the PU allows a secondary access rate above a threshold given by $\mathcal{W}(\bar{0}_0,\bar{1}_1,1)$, then secondary transmissions occur in both the states where the SR does not and does know the PM, according to a specific structure. We call this mode of operation the *High Interference Regime* (*HIR*), since in this case the secondary access constraint allows the SU to generate a relatively large amount of interference to the PU.

Otherwise ($\epsilon\leq\mathcal{W}(\bar{0}_0,\bar{1}_1,1)$), secondary transmissions are executed only in the states where the SR knows the PM. We call this mode of operation the *Low Interference Regime* (*LIR*), since in this case the secondary access constraint is such that the SU can only generate a relatively small amount of interference to the PU. This is formalized in the following Proposition.

Proposition 1. *In the* LIR *the optimal policy is given by*

$$\mu^{*(\epsilon)}=\left(\bar{0}_0,\mu_1^{*(\epsilon)}\right) \quad (15)$$

where $\mu_1^{*(\epsilon)}\in\mathcal{U}_1$ is an arbitrary policy satisfying the constraint with equality, i.e., $\mathcal{W}(\bar{0}_0,\mu_1^{*(\epsilon)},1)=\epsilon$.

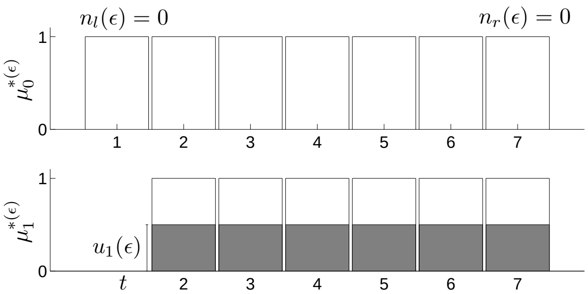

Figure 3. Structure of optimal policy $\mu^{*(\epsilon)}$ in the *LIR*

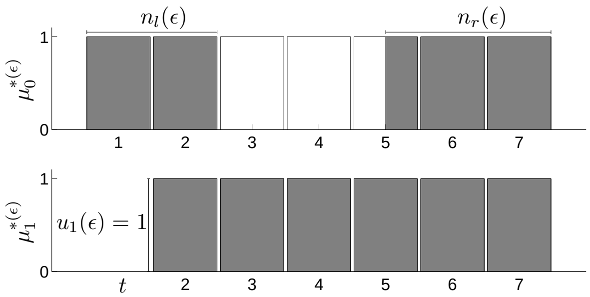

Figure 4. Structure of optimal policy $\mu^{*(\epsilon)}$ in the *HIR*

In the HIR *the optimal policy is given by*

$$\mu^{*(\epsilon)} = \left(\mu_0\left(n_l\left(\epsilon\right), n_r\left(\epsilon\right)\right), \bar{1}_1\right) \tag{16}$$

where $(n_l(\epsilon), n_r(\epsilon)) \in \left\{(x,y) \in [0,T]^2 : x+y \leq T\right\}$ *are derived using Algorithm VI-A in Section VI. Moreover, the optimal policy is randomized in at most one state, i.e., at least one of* $n_l(\epsilon)$ *and* $n_r(\epsilon)$ *must be an integer.*

Remark 3. Notice that in the *LIR*, since the optimal policy $\mu_1^{*(\epsilon)} \in \mathcal{U}_1$ is any solution of $\mathcal{W}\left(\bar{0}_0, \mu_1^{*(\epsilon)}, 1\right) = \epsilon$, we may exploit this degree of freedom and choose $\mu_1^{*(\epsilon)}$ so as to optimize other metrics of the network, such as the primary failure probability. In the following, we choose $\mu_1^{*(\epsilon)} \equiv u_1(\epsilon)\bar{1}_1$, where $u_1(\epsilon) \in [0,1]$ is the unique solution of $\mathcal{W}\left(\bar{0}_0, u_1(\epsilon)\bar{1}_1, 1\right) = \epsilon$, i.e., the ST transmits with a fixed transmission probability $u_1(\epsilon)$ in all states where its receiver knows the PM.

The structure of the optimal policy in the *LIR* and *HIR* is depicted in Figs. 3 and 4, respectively.

Remark 4. In the *LIR*, the ST transmits with a fixed transmission probability only when the PM is known to the SR. In this case, the optimal policy is described by a single parameter, namely the transmission probability $u_1(\epsilon)$.

In the *HIR*, the ST always transmits when the PM is known to the SR. Otherwise, it concentrates transmissions in the initial and final primary ARQ states. In this case, the optimal policy is described by two parameters, namely $n_l(\epsilon)$ and $n_r(\epsilon)$.

The overall policy is then described by three parameters, namely $u_1(\epsilon)$, $n_l(\epsilon)$ and $n_r(\epsilon)$. In particular, in the *LIR* we have $n_l(\epsilon) = 0$ and $n_r(\epsilon) = 0$ (i.e., no transmissions in the states where the PM is unknown to the SR); in the *HIR*, we have $u_1(\epsilon) = 1$ (i.e., the ST always transmits when the SR knows the PM).

B. Discussion of the structure of the optimal policy

An intuitive explanation of the optimal policy as a function of the constraint on the secondary access rate $\epsilon \in [0,1]$ is given as follows. When $\epsilon = 0$, the ST never transmits, and therefore the optimal policy is $(\bar{0}_0, \bar{0}_1)$. As ϵ gets larger, the ST can augment its access rate, by allocating secondary transmissions in the states according to a specific structure. Namely, it first fills transmissions in the states where the PM is known to the SR, since the latter can cancel interference generated by the PU, thus accrueing a better decoding performance and a larger secondary throughput.

As long as $\epsilon \leq \mathcal{W}\left(\bar{0}_0, \bar{1}_1, 1\right)$ (*LIR*), it is sufficient to transmit in the states where the SR knows the PM to achieve the constraint on the secondary access rate, and therefore the optimal policy is $\mu^{*(\epsilon)} = (\bar{0}_0, u_1(\epsilon)\bar{1}_1)$. I.e., the ST stays silent when its receiver does not know the PM ($n_l(\epsilon) = n_r(\epsilon) = 0$), and transmits with a fixed transmission probability $u_1(\epsilon)$ otherwise.

As ϵ gets larger, and the system enters the *HIR*, transmitting only in the states where the PM is known to the SR satisfies the constraint on the secondary access rate loosely. Therefore, under this policy the PU can accept further interference from the SU, i.e., there are additional opportunities left for the ST to transmit with respect to the previous case, by increasing its access rate and its interference to the primary, so as to improve the secondary throughput. It might then be beneficial for the ST to transmit also in the states where its receiver does not know the PM, other than in the states where it knows it. Specifically, the ST concentrates transmissions in those states in the initial and in the final primary ARQ states by an amount $n_l(\epsilon)$ and $n_r(\epsilon)$, respectively (Fig. 4). The reason behind this result is quite subtle, and can be explained by observing that secondary transmissions in the initial ARQ states induce primary retransmissions, thus augmenting the steady state probability associated to the states characterized by a larger ARQ index and where the SR knows the PM ($\phi = 1$). Therefore, by concentrating transmissions in the initial ARQ states, in the long term the network spends more time in the states where the SR knows the PM, and can thus perform interference cancellation, thus getting a larger long term throughput reward. In other words, by concentrating transmissions in the initial ARQ states, the SU induces redundancy in the ARQ process, which can be exploited once the SR knows the PM to enhance the secondary throughput. In the final ARQ states, it might not be beneficial for the ST to stay silent. In fact, the SU has only few opportunities left to exploit the redundancy in the ARQ process before the deadline T is reached, i.e., there are too few opportunities to exploit the knowledge of the PM in the subsequent ARQ rounds, to justify the SU to stay idle, so as to help the SR to decode the PM. Therefore, the SU is incentivized to transmit, thus maximizing its instantaneous throughput reward, although this impairs the ability of the SR to decode the PM.

VI. Algorithm to Determine $n_l(\epsilon)$ and $n_r(\epsilon)$ in *HIR*

In Prop. 1 we have stated that in the *HIR* the ST concentrates transmissions in the states where its receiver knows the PM in the initial and final ARQ states, by an amount $n_l(\epsilon)$ and $n_r(\epsilon)$, respectively, providing an intuitive explanation of this result in Section V-B. These two parameters reflect a trade-off between concentrating transmissions in the initial primary ARQ states, with the objective of inducing redundancy in the ARQ process, which can be exploited in the following ARQ rounds to mitigate interference and improve the secondary throughput, and concentrating them in the final ARQ states, in order to exploit the few available primary retransmissions before the primary transmission cycle ends. However, in order to determine the optimal policy, we need to optimize this trade-off, by determining the value of the optimal parameters $n_l(\epsilon)$ and $n_r(\epsilon)$ as a function of the constraint on the secondary access rate $\epsilon \in [0,1]$.

To this end, we now define an iterative Algorithm which generates a sequence of policies corresponding to increasing values of the secondary access rate, by activating at each iteration the state which, when activated, gives the highest increase of the secondary throughput per unit increase of the secondary access rate. This process continues until either the secondary throughput starts decreasing, and therefore allocating further transmissions is sub-optimal, or the constraint on the secondary access rate is violated, and therefore further transmissions would excessively impair the PU.

We first present the Algorithm in Section VI-A, followed by an explanation of the steps involved in Section VI-B. To this end, we redefine $\mathcal{W}(n_l, n_r, \nu) \equiv \mathcal{W}(\mu_0(n_l, n_r), \bar{1}_1, \nu), \nu \in [0,1]$ as a function of only the parameters n_l, n_r associated to the optimal policy. Moreover, we define the *transmission efficiency* $\eta\left(\mu^*, t, \phi\right)$ in state $(t, \phi) \in \mathcal{S}$ under policy $\mu^* \in \mathcal{U}$ as

$$\eta\left(\mu^*, t, \phi\right) = \left.\frac{\frac{\mathbf{d}\mathcal{W}(\mu,\tau)}{\mathbf{d}\mu_\phi(t)}}{\frac{\mathbf{d}\mathcal{W}(\mu,1)}{\mathbf{d}\mu_\phi(t)}}\right|_{\mu=\mu^*} \tag{17}$$

which represents how fast the secondary throughput increases per unit increase of the secondary access rate, due to augmenting the transmission probability in state (t, ϕ). In particular, we define the transmission efficiency under policy $\mu = (\mu_0(n_l, n_r), \bar{1}_1)$ calculated at the leftmost and rightmost idle states as follows.

$$\begin{aligned} \eta_l(n_l, n_r) &= \eta(\mu, \lfloor n_l \rfloor + 1, 0) \\ \eta_r(n_l, n_r) &= \eta(\mu, T - \lfloor n_r \rfloor, 0) \end{aligned} \tag{18}$$

These quantities represent how fast the secondary throughput increases, per unit increase of the secondary access rate, by activating respectively the leftmost or the rightmost idle state.

A. Algorithm

- ○ **Initialization:** $n_l^{(0)} = 0$, $n_r^{(0)} = 0$, $i = -1$
- ○ **Main:**

1) $i := i + 1$
2) **Case** $\eta_l(n_l^{(i)}, n_r^{(i)}) > 0 \bigcup \eta_r(n_l^{(i)}, n_r^{(i)}) > 0$
   a) **Case** $\eta_l(n_l^{(i)}, n_r^{(i)}) > \eta_r(n_l^{(i)}, n_r^{(i)})$
      - **Set** $n_l^{(i+1)} = n_l^{(i)} + 1$ and $n_r^{(i+1)} = n_r^{(i)}$

      i) **If** $\mathcal{W}(n_l^{(i+1)}, n_r^{(i+1)}, 1) \geq \epsilon$
         - **Set** $n_r(\epsilon) = n_r^{(i)}, n_l(\epsilon) \in \left(n_l^{(i)}, n_l^{(i+1)}\right]$ as the unique solution of $\mathcal{W}(n_l(\epsilon), n_r(\epsilon), 1) = \epsilon$
         - **Return** $(n_l(\epsilon), n_r(\epsilon))$

      ii) **If** $\mathcal{W}(n_l^{(i+1)}, n_r^{(i+1)}, 1) < \epsilon$
         - **Repeat from** Step 1)

   b) **Case** $\eta_r(n_l^{(i)}, n_r^{(i)}) \geq \eta_l(n_l^{(i)}, n_r^{(i)})$
      - **Set** $n_l^{(i+1)} = n_l^{(i)}$ and $n_r^{(i+1)} = n_r^{(i)} + 1$

      i) **If** $\mathcal{W}(n_l^{(i+1)}, n_r^{(i+1)}, 1) \geq \epsilon$
         - **Set** $n_l(\epsilon) = n_l^{(i)}, n_r(\epsilon) \in \left(n_r^{(i)}, n_r^{(i+1)}\right]$ as the unique solution of $\mathcal{W}(n_l(\epsilon), n_r(\epsilon), 1) = \epsilon$
         - **Return** $(n_l(\epsilon), n_r(\epsilon))$

      ii) **If** $\mathcal{W}(n_l^{(i+1)}, n_r^{(i+1)}, 1) < \epsilon$
         - **Repeat from** Step 1)
3) **Case** $\eta_l(n_l^{(i)}, n_r^{(i)}) \leq 0 \bigcap \eta_r(n_l^{(i)}, n_r^{(i)}) \leq 0$
   - **Set** $n_l(\epsilon) = n_l^{(i)}$ and $n_r(\epsilon) = n_r^{(i)}$
   - **Return** $(n_l(\epsilon), n_r(\epsilon))$

- ○ **Set the optimal policy:** $\mu^{*(\epsilon)} = \left(\mu_0\left(n_l(\epsilon), n_r(\epsilon)\right), \bar{1}_1\right)$

B. Comments and Explanation of the Algorithm

The Algorithm is initialized by allocating no transmissions in the states where the SR does not know the PM, i.e., $n_l^{(0)} = n_r^{(0)} = 0$. It then determines a sequence of policies, labeled by the counter i, giving increasing values of the secondary access rate, by allocating transmissions to either the leftmost or the rightmost idle state, according to which one maximizes the transmission efficiency, i.e., the state which, if activated, most improves the secondary throughput per unit increase of the secondary access rate, or alternatively, the state which maximizes the throughput reward to the SU, while minimizing the cost in terms of throughput loss of the PU.

Specifically, in the ith iteration, given the current values of the parameters $n_l^{(i)}$ and $n_r^{(i)}$ defining policy $\mu_0(n_l^{(i)}, n_r^{(i)})$ (see (14) and Fig. 2), if the transmission efficiency at both the leftmost and rightmost idle states is negative (Case 3), then activating any further state would determine a throughput loss for both the SU and the PU (this is a consequence of Theorem 2 in the Appendix and of the definition of transmission efficiency), thus resulting in a sub-optimal policy. The current structure of the policy is then returned, since it maximizes the secondary throughput.

Otherwise (Case 2), activating either the leftmost or the rightmost idle state determines an increase of the secondary throughput. In this case, the Algorithm allocates secondary transmissions to the state which maximizes the transmission efficiency, i.e., the state which determines the steepest increase of the secondary throughput per unit increase of the secondary access rate, as a consequence of activating that state. We have two cases: if the transmission efficiency is larger in the lefmost (respectively the rightmost) than in the rightmost (leftmost) idle state, then it is more efficient to allocate transmissions in the former than in the latter. The parameter $n_l^{(i)}$ ($n_r^{(i)}$) is then increased by one unit for the next iteration.

Notice that, as a consequence of activating the most efficient state, both the secondary access rate and the secondary throughput increase. This might incur a violation of the constraint on the secondary access rate (Cases 2(a)i and 2(b)i). In this case, the optimal state cannot be entirely activated, but the transmission probability is reduced until the constraint on the secondary access rate is attained with equality. A larger transmission probability in that state would incur a violation of the secondary access rate, whereas a smaller one would decrease the secondary throughput, thus resulting in a sub-optimal policy. Otherwise (Cases 2(a)ii and 2(b)ii), activating the most efficient state does not violate the constraint on the secondary access rate, and therefore it might be possible to further improve the secondary throughput by allocating more secondary transmissions, thus increasing the secondary access rate. In this last case, the Algorithm proceeds with a new iteration.

VII. Numerical Results

In this section, we discuss some numerical results demonstrating the performance improvement achievable by exploiting the redundancy introduced by the primary ARQ process to achieve a larger secondary throughput via primary interference cancellation, over traditional techniques which do not take advantage of this side-information. We recall that the system under consideration is the one depicted in Fig. 1, where a pair of Primary and Secondary Transmitters (PT and ST, respectively) transmit to their respective receivers PR and SR over the links γ_p and γ_s, thus mutually interfering over the links γ_{ps} and γ_{sp}. Each channel link is modeled as i.i.d. Rayleigh fading with power $\Gamma_x = E\left[|\gamma_x|^2\right]$, $x \in \{s, p, sp, ps\}$ (i.e., x is the label associated to any of the active links), with zero mean-unit variance circular Gaussian noise at each receier. Therefore, letting without loss of generality $P_p = 1$ and $P_s = 1$ be the primary and secondary transmission powers, respectively, the Signal to Noise Ratio (SNR) at the output of each link has exponential distribution with mean Γ_x. In particular, we assume that the SNRs assume a common value, i.e., $\Gamma_s = \Gamma_p = \Gamma_{ps} = \Gamma_{sp} = \Gamma$. In fact, this represents a worst case scenario, since interference from the PU to the SU is neither strong enough to make correct decoding of the PM likely, nor weak enough to be comparable with the noise level, thus creating very little interference. On the other hand, since the PR treats secondary transmissions as noise, this value of the SNR is strong enough to severely impair the PU, in case of secondary transmission. Moreover, we choose $\Gamma = 5$ as the common value of the SNRs. In fact, using values too close to unity would make the signal power comparable with the noise power, thus making the effect of interference at each receiver of little significance.

We compare the performance of four different secondary transmission policies, in terms of the achievable primary and secondary throughputs as a function of the secondary access rate constraint, $\epsilon \in [0, 1]$. The first policy is the optimal one described in this paper, and is denoted by $\mu^{*(\epsilon)}$. Its structure is derived using the algorithm described in Section VI-A.

Then, we consider two sub-optimal policies which do not exploit the redundancy introduced by the ARQ process. In the first case the SR performs multi-user detection of primary and secondary messages to enhance secondary decoding performance, but, if the PM is correctly decoded, it does not exploit this knowledge in the following ARQ rounds, if primary retransmissions occur. The achievable throughput in case of secondary transmission is given by $T_{s0}(R_s, R_p) = R_s\left(1 - \rho_{s0}(R_s, R_p)\right)$ (see the system model in Section II). This is maximized with respect to R_s in order to optimize the performance. We then let $T_{s,\textit{M-User}} = \max_{R_s \in \mathbb{R}_+} T_{s0}(R_s, R_p)$ be the achievable optimized throughput, and $R_{s,\textit{M-User}}$ be the optimal transmission rate for this scenario.

In the second case, the SR decodes the secondary message by treating primary transmissions as noise. We denote the achievable throughput in case of secondary transmission by $T_{s,\textit{Noise}}$, which is again maximized with respect to the secondary transmission rate.

Finally, we also compare the optimal policy with an *oracle policy*, which assumes perfect knowlege of the PM by the SR in advance. This agrees with some Information Theoretical models in the literature [5]–[7]. Under this policy, the achievable throughput in case of secondary transmission, after maximizing it with respect to the transmission rate is given by $T_{s,\textit{Oracle}} = \max_{R_s} T_{s1}(R_s)$, where $T_{s1}(R_{s1})$ was defined in the system model, Section II. For the three cases above, the optimal secondary transmission policy for a given constraint on the secondary access rate $\epsilon \in [0, 1]$, consists in transmitting with a fixed transmission probability ϵ, thus achieving a long term secondary throughput $\epsilon T_{s,i}$, where $i \in \{\textit{M-User,Noise,Oracle}\}$ is a label which refers to either the sub-optimal policy with Multi-User detection, to the sub-optimal policy which considers the PM as noise, and to the Oracle policy.

Table I
System parameters with common SNR $\Gamma = 5$

SU (optimal policy)		**PU**	
α_{s0}	0.5748	R_p	1.9141
α_{s1}	0.4431	ρ_{p0}	0.4252
R_{s0}	1.3214	ρ_{p1}	0.8475
R_{s1}	1.9141	**SU (oracle policy)**	
T_{s0}	0.7427	$R_{s,\textit{Oracle}}$	1.9141
T_{s1}	1.1002	$T_{s,\textit{Oracle}}$	1.1002
SU (PM as noise)		**SU (multiuser detection)**	
$R_{s,\textit{Noise}}$	1.0253	$R_{s,\textit{M-User}}$	1.3214
$T_{s,\textit{Noise}}$	0.4095	$T_{s,\textit{M-User}}$	0.7427

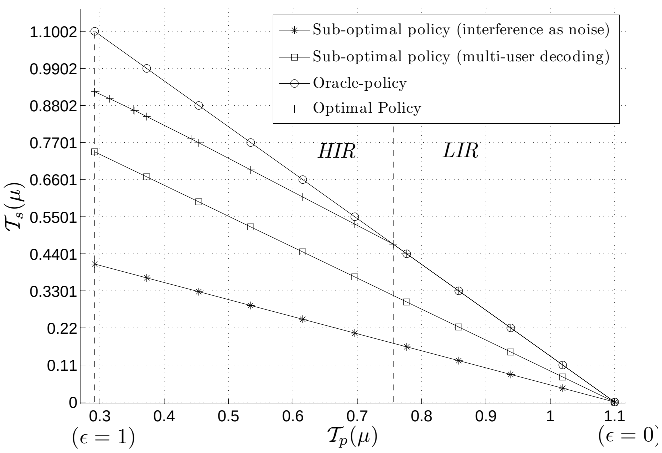

Figure 5. Primary vs secondary throughputs, length of primary ARQ transmission window $T = 5$

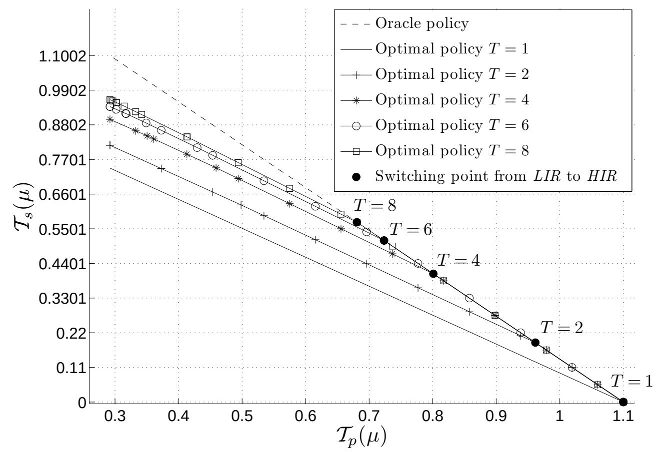

Figure 6. Primary vs secondary throughputs for different values of the length of the primary ARQ window $T \in \{1, 2, 4, 6, 8\}$

In Table I we summarize the main parameters used in the numerical results, after optimization of the throughput terms as described above. These are calculated using the results derived in [12]. We refer to the system model, Section II for details on their meaning.

Fig. 5 depicts the primary/secondary throughput pairs achievable using the four policies described above, as a function of $\epsilon \in [0, 1]$. The size of the ARQ transmission window chosen is $T = 5$. In agreement with intuition, the sub-optimal policy treating primary transmissions as noise performs the worst. In fact, according to this policy, the SR uses no side-information about the PU to enhance its own performance. A better performance is achieved by the sub-optimal policy which performs multi-user decoding of primary and secondary messages. Finally, the optimal policy described in this paper lies in between the sub-optimal policy performing multi-user detection and the oracle policy which always assumes knowledge of the PM at the SR. In general, the performance improves as more side-information about the PU is exploited by the SR. Remarkably, although the assumption that the SR knows the PM in advance is not realistic, in the *LIR* the optimal policy developed in this paper, which does not assume any prior knowledge of the PM, attains the performance of the oracle policy. The reason behind this result lies in the fact that, even if the SR does not know the PM in advance, in the *LIR* secondary transmissions occur only in the states where the SR knows the PM, thus mimicking the scenario where the PM is known in advance by the SR. As the system approaches the threshold between *LIR* and *HIR*, the optimal policy described in this paper allocates more and more transmissions in the states where the PM is known to the SR, until transmissions always occur in those states (at the threshold, marked by a vertical dashed line). At this point, if the SU is allowed to create further interference to the PU, it should start transmitting even in the states where the PM is unknown to the SR, thus entering the *HIR*. However, by transmitting in those states the transmission efficiency starts decreasing, due to interference from the PU and to the inability to perform primary interference cancellation in those states, and therefore the optimal policy starts performing worse than the oracle policy.

Fig. 6 depicts the primary/secondary throughput pairs achievable using the optimal policy, for $\epsilon \in [0, 1]$ and for different values of the length of the primary ARQ transmission window $T \in \{1, 2, 4, 6, 8\}$. It is worth noticing that for small values of T the performance of the optimal policy approaches the lower bound where the SU does not exploit the redundancy in the primary ARQ process and performs multi-user decoding of primary and secondary messages (this is attained with equality in the case $T = 1$). The reason is that, for such values of T, there is very little redundancy introduced in the system due to the primary ARQ process (no redundancy when $T = 1$). As T gets larger, the SU can potentially benefit from a longer sequence of primary retransmissions, i.e., a larger level of redundancy in the system due to the primary ARQ process, thus resulting in improved performance. However, the performance improvement diminishes as T gets larger. In fact, the states characterized by increasingly larger values of the ARQ state t are accessed less and less frequently, thus giving a smaller and smaller contribution to the performance of the SU. Notice also that the switching point between *LIR* and *HIR* is achieved at increasing values of the secondary access rate (or equivalently, decreasing values of the primary throughput) for increasing T. This is a consequence of the fact that, for increasing values of T, there is an increasing number of states in which the SR knows the PM.

VIII. Concluding Remarks

Coexistence between secondary and primary systems is based on the premise that the activity of primary users typically leaves some margin in the usage of radio resources for secondary users. For instance, this may take the form of a cap on the maximum additional interference that secondary systems are allowed to create to the primary system. In this paper, following [9], we have argued that the presence of a primary user that performs retransmissions offers relevant margins to the activity of secondary users. In particular, secondary receivers may estimate the primary packet from a given (re)transmission of the primary packet and then use this side information to perform interference cancellation on the following retransmission. We have tackled the problem of maximizing the secondary throughput under primary interference constraints and characterized in full the optimal transmission strategy of the secondary user. Specifically, we have shown that, if the SU is allowed to produce interference only below a certain threshold (Low Interference Regime), then the SU should transmit exclusively after its receiver has decoded the primary message from previous (re)transmissions; instead, if the interference constraint allows the secondary system to produce interference above the threshold, the SU should transmit also in the states where its receiver does not know the PM, according to a specific structure that is explictly characterized. Numerical results have validated and quantified the performance advantages derived by exploiting the proposed interference mitigation strategy.

APPENDIX

In this appendix, we state some useful properties related to the randomized stationary policies. These are explicitly used in the proof of Prop. 1, which is not provided here, but can be found in [14]. However, they are presented here to help the reader understand the main results presented in this work. In particular, Theorem 2 states that if we increase the transmission probability in any state $(t, 0)$, the secondary access rate $\mathcal{W}(\mu_0, \mu_1, 1)$ increases. Theorem 3, on the other hand, states that if we increase the transmission probability in any state $(t, 1)$, both the secondary access rate $\mathcal{W}(\mu_0, \mu_1, 1)$ and the secondary throughput $\mathcal{W}(\mu_0, \mu_1, \tau)$ increase. For the proofs of the theorems, we refer the interested reader to [14].

Theorem 2. *Under the set of assumptions (7), for all* $t \in \mathbb{N}(1, T)$, $\mathcal{W}(\mu, 1)$ *is a strictly increasing function of* $\mu_0(t)$.

Theorem 3. *Under the set of assumptions (7), for all* $t \in \mathbb{N}(1, T)$, $\mathcal{W}(\mu, \nu)$ *is a strictly increasing function of* $\mu_1(t)$.

Notice that, although these results seem rather intuitive, increasing the transmission probability in state (t, ϕ), while augmenting the expected cost/reward accrued in state (t, ϕ), also causes a perturbation of the steady-state distribution of the system, by affecting the outage behavior at the primary and secondary systems. The two theorems then state that the increase of the expected cost/reward outweighs the loss due to the perturbation of the steady-state distribution induced by augmenting the transmission probability in state (t, ϕ). Notice also that from Theorem 2 the secondary throughput might not be an increasing function of $\mu_0(t)$. In fact, observe that augmenting the transmission probability in state $(t, \phi = 0)$ degrades the outage performance at the SR, and thus its ability to successfully decode the PM. This in turns affects the ability of the SR to exploit the knowledge of the PM in the subsequent primary ARQ rounds, so as to achieve a larger secondary throughput via primary interference cancellation. This, in the long term, might degrade the overall secondary throughput.

REFERENCES

[1] FCC Spectrum Policy Task Force, "Report of the spectrum efficiency working group," FCC, Tech. Rep., Nov. 2002.

[2] J. Peha, "Approaches to spectrum sharing," *IEEE Communications Magazine*, vol. 43, no. 2, pp. 10 –12, Feb. 2005.

[3] ——, "Sharing spectrum through spectrum policy reform and cognitive radio," *Proceedings of the IEEE*, vol. 97, no. 4, pp. 708 –719, April 2009.

[4] J. Mitola and G. Maguire, "Cognitive radio: making software radios more personal," *IEEE Personal Communications*, vol. 6, no. 4, pp. 13 –18, Aug. 1999.

[5] A. Goldsmith, S. Jafar, I. Maric, and S. Srinivasa, "Breaking spectrum gridlock with cognitive radios: An information theoretic perspective," *Proceedings of the IEEE*, vol. 97, no. 5, pp. 894 –914, May 2009.

[6] I. Maric, A. Goldsmith, G. Kramer, and S. Shamai, "On the capacity of interference channels with a partially-cognitive transmitter," in *IEEE International Symposium on Information Theory, ISIT*, June 2007, pp. 2156 –2160.

[7] A. Jovicic and P. Viswanath, "Cognitive radio: An information-theoretic perspective," in *IEEE International Symposium on Information Theory, ISIT*, July 2006, pp. 2413 –2417.

[8] S. Ahmad, M. Liu, T. Javidi, Q. Zhao, and B. Krishnamachari, "Optimality of myopic sensing in multichannel opportunistic access," *IEEE Transactions on Information Theory*, vol. 55, no. 9, pp. 4040 –4050, Sept. 2009.

[9] R. Tannious and A. Nosratinia, "Cognitive Radio Protocols Based on Exploiting Hybrid ARQ Retransmissions," *IEEE Transactions on Wireless Communications*, vol. 9, no. 9, pp. 2833 –2841, Sep. 2010.

[10] M. Levorato, U. Mitra, and M. Zorzi, "Cognitive interference management in retransmission-based wireless networks," in *47th Annual Allerton Conference on Communication, Control, and Computing*, Sep. 2009, pp. 94 –101.

[11] P. Frenger, S. Parkvall, and E. Dahlman, "Performance comparison of HARQ with Chase combining and incremental redundancy for HSDPA," in *IEEE 54th Vehicular Technology Conference (VTS)*, 2001.

[12] R. Narasimhan, "Individual Outage Rate Regions for Fading Multiple Access Channels," in *IEEE International Symposium on Information Theory, ISIT*, Jun. 2007, pp. 24–29.

[13] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. John Wiley & Sons, Inc., New York, 2006.

[14] N. Michelusi, M. Levorato, P. Popovski, O. Simeone, and M. Zorzi, "Optimal Cognitive Transmission utilizing Redundancy in the Primary ARQ process," Tech. Rep., 2011, www.dei.unipd.it/~michelusi/papers/techrep_cognitive2011.pdf.

[15] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Athena Scientific, 2000.

[16] K. W. Ross, "Randomized and past-dependent policies for markov decision processes with multiple constraints," *Operations Research*, vol. 37, no. 3, pp. 474–477, 1989.