Exploiting Channel Memory for Wireless Scheduling with Limited Channel Probing: an Asymptotic Study

Paolo Giaccone, Emilio Leonardi

Exploiting Channel Memory for Wireless Scheduling with Limited Channel Probing: an Asymptotic Study

Paolo Giaccone, Emilio Leonardi
Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino (Italy)

***Abstract*— We consider a wireless downlink communication system with N parallel channels under the following two assumptions: (1) the channels states can be explicitly sensed at a limited rate λ_s, (2) channels are modeled as i.i.d., ON/OFF Markov chains with stationary distribution $[\pi_{OFF}, \pi_{ON}]$. The goal of this paper is to characterize the fundamental trade-off between system throughput and channel sensing rate λ_s in the challenging scenario in which $N \to \infty$ and jointly $\pi_{ON} \to 0$. To this end, we propose a simple scheduling algorithm that effectively exploits channel memory to maximize the successful transmission probability. We obtain sufficient and necessary conditions to achieve maximum (in order sense) system throughput. Finally, we show the performance of our policy under a non-asymptotic scenario.**

I. Introduction

Downlink communications from a wireless station to a large set of users are an important aspect that can be optimized to improve the performance and the scalability of the whole wireless network. We consider the specific case in which each user is associated with one channel and the channel state for each user evolves with the time. Throughput efficient scheduling algorithms for multi-user wireless download links are well known, when a perfect knowledge of channels states is assumed at all times. Indeed, the seminal work of Tassiulas and Ephremides [1] and its later extensions [2] permit to perfectly characterize both the achievable throughput as well as the optimal throughput schedulers under perfect channel information.

Realistic scenarios are much more challenging since either channel state information is not available, or it is only partially available.

Under the assumption that the state of channels evolves "slowly" with time, information about the current channels state can be inferred by past measurements. This is the idea exploited by recent papers [3], [4] where a preliminary investigation on the achievable capacity have been carried out in absence of channel state information. In particular both papers [3], [4] consider a download link with N users, whose associated channels evolve as independent ON/OFF Markov chains. Packets can be successfully transmitted only if the corresponding channels are in the ON state. The transmitter has to select a channel for transmission at every time so to maximize the overall throughput.

Unfortunately the problem of perfectly characterizing the achievable system throughput turns out to be very difficult. Even under the simplistic assumption that all transmission queues are constantly backlogged, the characterization of the achievable capacity requires the solution of a complex Markov Decision Process over an infinite number of states. Approximate solutions can be obtained only in toy cases (such as the case for $N = 2$) as shown [4]. In [3], instead, an inner and an outer bound on the capacity region are derived, but the obtained bounds are shown to be tight only when the number of users grows large and the traffic is symmetric.

In this paper we explore a different but tightly related problem. We characterize the fundamental trade-off between system throughput and channel sensing rate in the challenging scenario in which jointly the number of channels N tends to infinity and the probability of finding an individual channel in the ON state tends to zero. Note that the latter assumption makes the problem harder, indeed, if the probability for the channel to be in the ON state remains finite, then the number of ON channels tends to infinity, for $N \to \infty$, and the problem degenerates. In particular, we assume that our system is able to probe (equivalently, we say "sense") channels at a maximum rate λ_s, which grows sub-linearly with respect to N (i.e.,[1] $\lambda_s = o(N)$.) to avoid the complexity of sensing the whole, in order sense, set of channels. We propose a simple sensing and transmission scheduling algorithm, which effectively exploits channel memory to maximize the successful transmission probability. We obtain the sufficient and necessary conditions to achieve the maximum (in order sense) system throughput.

II. System description

We consider a wireless station transmitting data to N users through N independent Markovian channels. Each channel is modeled as a two-state ON/OFF Markov chain with transition rates α from OFF to ON state, and β from ON to OFF state. Hence, the average sojourn time in state ON is $E[T_{ON}] = 1/\beta$ and the average sojourn time in OFF is $E[T_{OFF}] = 1/\alpha$. A fixed non-null data rate μ is achieved when transmitting on any ON channel; the transmission rate is instead zero for any OFF channel. The wireless station is provided with N parallel queues, in which packets are queued according to their destined user. Queues, for simplicity, are assumed constantly backlogged. The station senses channels according to an homogeneous Poisson process at a rate λ_s; when a sensing event occurs, one channel is sampled. The transmission scheduling process occurs independently from the sensing process, even if it exploits some information from it. Our goal is to relate the system performance (i.e., the saturation throughput) to the following parameters: the sensing rate λ_s, the number of users/channels N and the channel dynamical behavior described by α and β.

A. Channel sensing policy

We consider a sensing policy with a limited storage available for the channels that have been recently sensed and found in the ON state. When a new packet must be transmitted, the scheduling policy access such storage and select the user/channel to whom/on which transmit. Note that the transmission policy may select an OFF channel, since this was sensed ON in the past but it has changed its state in the meanwhile.

We model the storage of the channels in the following way. Channels are dynamically divided into two classes: *tracked channels* and *untracked channels*. At time t, a channel is *tracked* if it was sensed ON during its last sensing event. Otherwise, a channel is *untracked*, since it was sensed OFF last time. We consider a double sensing process: tracked channels are sensed to check if they are still ON, whereas untracked channels are sensed to discover new ON channels. As soon as a tracked channel is sensed OFF, it becomes untracked. On the contrary, as soon as an untracked channel is sensed ON, it becomes tracked.

We denote with $k(t)$ the number of tracked channels at time t. To limit the maximum complexity of the sensing algorithm, we set a maximum number of tracked channels equal to K. Every tracked channel is periodically sensed at rate $\lambda_{TC} = \lambda_s/K$; while randomly selected untracked channels are sensed at rate $\lambda_s - k(t)\lambda_{TC}$.

B. Channel selection policy

We assume that the wireless station at the end of a packet transmission gets immediate feedback of the state of the corresponding channel. Channel selection for packet transmissions follows these rules. If the last packet has been correctly transmitted (i.e. the corresponding channel has not become OFF during the last packet transmission), at the end of its transmission a new packet is transmitted on the same channel. Otherwise the most recently sensed tracked channel is selected for the packet transmission.

III. Asymptotic analysis

In this section we analyze the dynamics of the number of tracked channels $k(t)$. Observe that such dynamics are complex to analyze due to the non-memoryless nature of sensed channels. Indeed, the probability that either a tracked or untracked channel n is found in ON state when sensed at time t, depends on τ, defined as the time elapsed since the last sensing event for the channel n. Indeed, it can shown that the probability of being ON for a tracked channel is

$$\hat{p}_{tr}^{ON}(\tau) = \pi_{ON} + (1 - \pi_{ON})e^{-(\alpha+\beta)\tau}$$

where $\pi_{ON} = \alpha/(\alpha + \beta)$ is the stationary distribution of ON state for each channel. Similarly, the probability of being OFF for an untracked channel is:

$$\hat{p}_{un}^{ON}(\tau) = \pi_{ON} - \pi_{ON}e^{-(\alpha+\beta)\tau}$$

Observe that $\hat{p}_{tr}^{ON}(\tau)$ is a decreasing function of its argument while $\hat{p}_{un}^{ON}(\tau)$ is increasing.

The dependence of $\hat{p}_{tr}^{ON}$ and $\hat{p}_{un}^{ON}$ on their arguments induces a complex structure in the temporal memory of $k(t)$, whose evolution in an interval $[t, t + \Delta t)$

[1]As reminder of Landau notation, given two functions $f(n) \geq 0$ and $g(n) \geq 0$: $f(n) = o(g(n))$ means $\lim_{n\to\infty} f(n)/g(n) = 0$; $f(n) = O(g(n))$ means $\limsup_{n\to\infty} f(n)/g(n) = c < \infty$; $f(n) = \omega(g(n))$ is equivalent to $g(n) = o(f(n))$; $f(n) = \Omega(g(n))$ is equivalent to $g(n) = O(f(n))$; $f(n) = \Theta(g(n))$ means that both $f(n) = O(g(n))$ and $g(n) = O(f(n))$ hold.

turns out to depend on the current number of tracked channels, $k(t)$, as well as on the times elapsed since the last measurements for all channels (both tracked and untracked ones). Furthermore the process describing the evolution of $k(t)$ is not reversible, thus, an exact analysis of $k(t)$ dynamics appears prohibitive.

We analyze the dynamics of $k(t)$ under the simplifying assumptions that the above mentioned dependencies are negligible, i.e., every tracked channel is sensed in ON with some given probability p_{tr}^{ON} which is a fixed value approximating (in order sense) $\hat{p}_{tr}^{ON}(\tau)$, and every untracked channel is sensed in ON state with some fixed probability p_{un}^{ON}, independently from the time elapsed since the last sensing measurement.

In the latter case the dynamic behavior of $k(t)$ degenerates into a simple birth-and-death Markov Chain, whose steady state solution can be expressed as, for $h = 0, \dots, K$:

$$\Pr\{k(t) = h\} = \frac{\left(\frac{\lambda_{TC}\, p_{un}^{ON}}{\lambda_{TC}(1 - p_{tr}^{ON})}\right)^h \prod_{j=1}^{h} \frac{(K-j)}{j}}{\sum_{l=1}^{K} \left[\left(\frac{\lambda_{TC}\, p_{un}^{ON}}{\lambda_{TC}(1 - p_{tr}^{ON})}\right)^l \prod_{j=1}^{l} \frac{(K-j)}{j}\right]} \tag{1}$$

From this simplified model, it is possible to estimate in order sense the asymptotic behavior of the performance in our system when $N \to \infty$. One of the key steps in the reasoning is that it can be shown that the number of tracked channels can be bounded by below, and by above, using two simplified systems, with different parameters, whose performance are the same (in order sense) than the original one we are investigating. However, we skip the detailed proof.

A. A challenging scenario

The first asymptotic scenario we consider, is a challenge scenario in which the number of channels in ON does not scale with N. This implies that

$$\liminf_{N\to\infty} N\pi_{ON} > 0 \tag{2}$$

and

$$\liminf_{N\to\infty} N\pi_{ON} < \infty \tag{3}$$

Observe that (2) expresses the minimal requirement to have a finite throughput even for the ideal condition in which the state of every channel is perfectly known by the scheduler. In this scenario we will evaluate the minimum sensing rate that guarantees a non vanishing throughput, exploiting channel memory. This minimum channel rate will critically come to depend on the channel persistence β^{-1}. In general we are interested to the case in which $\beta \to 0$, i.e. channel persistence is very large with respect to packet transmission rate. This condition is typically met in networks characterized by marginal mobility.

In the simplified birth-and-death chain model, first, for tracked channels we set:

$$p_{tr}^{ON} = \lambda_{TC} \int p_{tr}^{ON}(\tau) e^{-\lambda_{TC}\tau} d\tau$$

i.e. p_{tr}^{ON} represents the probability that a sampled tracked channel is found in the ON state, unconditionally over the time since the last measurement. Instead, as mentioned before, for untracked channels we set $p_{un}^{ON} = \pi_{ON}$. This approximation is reasonable when the typical lag between two successive sensing events at the same channel is equal or larger than the average cycle time associated to channel evolution.

Observe that when $N \to \infty$, a necessary condition to have a non-vanishing throughput in the simplified system is that:

$$\liminf_{N\to\infty} \Pr\{k(t) > 0\} > 0$$

This can be obtained by setting $K = \Theta(1/(N\pi_{ON})) = \Theta(1)$.

Indeed, observe that the probability of successfully transmitting a new packet at time t, given that $k(t) > 0$ is, by construction, lower bounded, in order sense, by p_{tr}^{ON}. In case the last packet transmission was successful, the next packet is transmitted on the same channel; the resulting probability that the new packet is successfully transmitted can be evaluated as the probability that the tracked channel remains in ON for the whole transmission of the considered new packet. Let p_o denote such probability. Typically $p_o \approx 1 \gg p_{tr}^{ON}$, since the transmission time of a packet is very short with respect to the dynamics of the channels, under the limited mobility considered in this work. In the case the last packet transmission was not successful, the most recently sensed tracked channel is selected for packet transmission. In this case, the transmission of the new packet is successful if the sampled tracked channel is sampled ON and the channel remains ON during the packet transmission; thus, the probability of successful transmission is $p_{tr}^{ON} p_o$.

As a conclusion, to sustain a non-vanishing throughput, it must be $p_{tr}^{ON} p_o > 0$. If λ_{TC} is sufficiently large, in particular it is faster than the ON-OFF transition rate (i.e., $\lambda_{TC} = \Omega(\beta)$), then

$$p_{tr}^{ON} > 0$$

and the throughput is finite. As a consequence, the overall sensing rate $\lambda_s = K\lambda_{TC}$ must be asymptotically lower bounded as $\Omega(\beta/\pi_{ON})$.

As a summary, we have provided the main insights for which the following property holds:

Property 1: In the wireless downlink scenario in which $\pi_{ON} \to 0$ and $\liminf_{N\to\infty} N\pi_{ON} > 0$, a storage size $K = \Theta(1/(N\pi_{ON})) = \Theta(1)$ and a sensing rate $\lambda_s = \Omega(\beta/\pi_{ON})$ are sufficient to obtain a non-vanishing throughput.

B. A more favorable scenario

Now, we consider the more favorable scenario in which the average number of ON channels grows to infinity, i.e.

$$\liminf_{N\to\infty} N\pi_{ON} \to \infty$$

even if we still assume that the probability of finding a channel in the ON state vanishes, i.e.,

$$\limsup_{N\to\infty} \pi_{ON} \to 0$$

Also in this case, we assume that $\beta \to 0$.

In this case, if the storage is enough large as $K = \omega(1/(N\pi_{ON}))$ (but still $K = o(N)$), then we obtain:

$$\liminf_{N\to\infty} \Pr\{k(t) > 0\} \to 1$$

while

$$p_{tr}^{ON} = \lambda_{TC} \int \hat{p}_{tr}^{ON}(\tau) e^{-\lambda_{TC}\tau} d\tau \to 1$$

iff $\lambda_{TC} = \omega(\beta)$.

As a consequence, also under $\lambda_{TC} = \omega(\beta)$ and $K = \omega(1/\pi_{ON})$ the maximum throughput is asymptotically achieved by our sensing and scheduling policy. Under the weaker conditions $\lambda_{TC} = \Omega(\beta)$ and $K = \Omega(1/\pi_{ON})$ the throughput system throughput is also non-vanishing.

IV. Numerical analysis

Throughout some simulations, we have investigated the performance of the proposed sensing and transmission schedule policy in a discrete-time system, in which the state of each channel evolves according to a discrete-time Markov chain. Our simulated system is representative of the continuous-time Markov chain considered in the analytical model.

Similarly to the algorithm described in Sections II-A and II-B, the wireless station senses in total λ_s channels per timeslot according to the following rules: at each timeslot, the station starts to sense the tracked channels and then eventually senses other random untracked channels, until λ_s channels are sensed. If a channel is sensed ON, it becomes/remains tracked and its state is stored in the station. Otherwise, it becomes/remains untracked. Note that λ_s is an indicator of the storage requirement needed to track the channels since, by construction, at most λ_s channels are tracked at the same time.

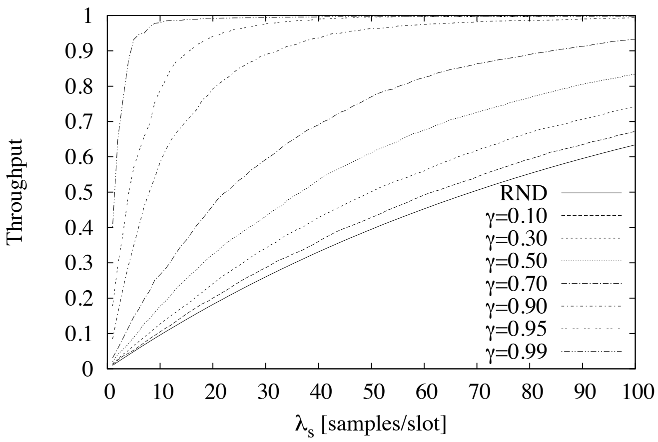

Fig. 1. Throughput for $N = 1000$ users/channels, under $\pi_{ON} = 0.01$ and different values of channel auto-covariance γ

TABLE I
Auto-covariance γ of the channel and corresponding duration of ON and OFF periods

γ	$E[T_{ON}]$	$E[T_{OFF}]$
0.10	1.122	111.2
0.30	1.443	142.9
0.50	2.020	200.0
0.70	3.367	333.3
0.90	10.10	1000
0.95	20.2	2000
0.99	101	10000

We model the channel memory through the auto-covariance coefficient of the Markov chain, which is an indicator of the "persistency" of the state of the channel. Let $p_{s1,s2}$ be the transition probability from state $s1 \in \{ON, OFF\}$ to $s2 \in \{ON, OFF\}$. It can be shown [3] that the auto-covariance coefficient γ is equal to:

$$\gamma = \mathbb{E}[(S(t) - \mathbb{E}S(t))(S(t+1) - \mathbb{E}S(t+1))] \\ = p_{ON,ON} - p_{OFF,ON}$$

In a typical scenario, $\gamma > 0$ since each state is positively correlated. Note that $\gamma = 0$ means that the states are i.i.d. over the timeslots and the channel state is not persistent; on the contrary, when $\gamma \to 1$ the channel state becomes

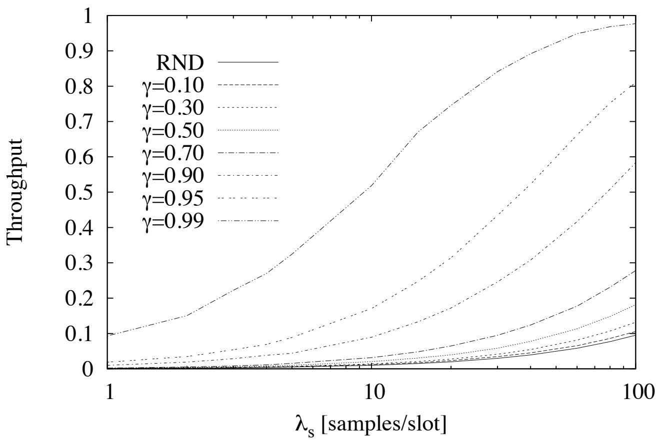

Fig. 2. Throughput for $N = 10000$ users/channels, under $\pi_{ON} = 0.001$ and different values of channel auto-covariance γ

more and more persistent since transitions between states are rare. As a reference, Table I reports the values of average duration of ON and OFF states for the specific values of γ considered in the following simulations. For simplicity, we assume the same value of γ for all the channels.

We have considered two scenarios, with increasing number of users/channels N while decreasing the probability π_{ON} of being ON for each channel, in order to keep the average number of ON channels available at each time equal to the constant value 10. Fig.s 1-2 show the throughput obtained for different values of sampling rate λ_s and different values of channel persistency. In both figures we added the curves RND corresponding to a random sampling of λ_s channels per slot and without storing any tracked channel. These curves were obtained by evaluating $1 - (1 - \pi_{ON})^{\lambda_s}$.

For high channel persistency, i.e. higher values of γ, Fig. 1 shows that our policy achieves high throughput even if the number of sampled channels is very limited, being around 1-10% of all the channels, whereas the average number of channels that are ON at the same time is just 1% of all the channels. When the channel state is less persistent, the effect of the storing the tracked channel is less effective, since the channel states tend to become i.i.d. and the past history of the channel is less meaningful. As expected, in the worst case, our policy approaches the random sampling policy RND for i.i.d. channels for which storing the tracked channels is useless.

Fig. 2 shows a similar qualitative behavior for a larger number of users/channels. Notably, in this scenario the average number of ON channels is just 0.1% of all the channels, and some relevant throughput can be obtained even by sampling only 10 channels per slot (i.e., around 0.1% of all the channels), when the channel persistence is enough large.

V. Conclusions

We have investigated a downlink wireless scenario in which a station communicates with a large number of users and in which the channels' states evolve with the time, and they can be partially probed by the station, which is provided with some limited sensing capability. We showed some fundamental tradeoff between the dynamics of the channels and the sensing rate to achieve an non-negligible throughput in an asymptotic scenario. We propose a sensing policy that leverages the channels persistency to track the best channel candidates for transmission and optimize (in order sense) the throughput. This result is achieved through a limited storage and capability for channel tracking.

Finally, we have simulated the performance of our policy in some scenarios, with a large (but finite) number of users/channels. Our results highlight the potentials of our policy also in non-asymptotic scenarios.

VI. Acknowledgements

The research leading to these results has funded by the European Union Seventh Framework Programme under the FET project "STAMINA".

References

[1] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks", IEEE Trans. Automatic Control, vol. 37, no. 12, pp. 1936-1948, Dec. 1992

[2] M. Neely, E. Modiano and C. Rohrs, "Power Allocation and Routing in Multi-Beam Satellites with Time Varying Channels", IEEE Transactions on Networking, Feb. 2003.

[3] C.P. Li and M. Neely, "Exploiting channel memory for multiuser wireless scheduling without channel measurement: Capacity regions and algorithms", IEEE WiOpt, Avignone, France, May 2010.

[4] K. Jagannathan, S. Mannor, I. Menache, E. Modiano, "A State Action Frequency Approach to Throughput Maximization over Uncertain Wireless Channels", IEEE Infocom (Mini-conference), Shanghai, China, April 2011.