Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 56-2 (2023) 3830–3835

Wi-Fi for Constrained Control

Matthias Pezzutto *

* *Service d'Automatique et d'Analyse des Systèmes, Université Libre de Bruxelles, Brussels, Belgium — e-mail: matthias.pezzutto@ulb.be*

Abstract: In this work, we assess the employment of Wi-Fi for constrained control applications. First, we carry out an experimental test to accurately characterize the network behavior in an office environment for different sampling periods and external interference conditions. Tests show that delays are increased for higher traffic loads on other co-existing networks, and that long blackouts can occur, largely limiting remote constrained control of unstable systems. Then, based on the experimental observations, we design a control strategy based on Reference Governor tailored for Wi-Fi networks. An experimental test involving a two-wheeled robot shows the feasibility and the effectiveness of the proposed algorithm compared to standard strategies.

Keywords: Control over Lossy Networks; Communication networks; Segway-like robot

1. INTRODUCTION

In the last years, control engineers have paid large attention to Networked Control Systems (NCSs). Informally, a NCS is a control system where the plant and the controller are connected through a communication network. Compared to standard control systems, NCSs come with enhanced flexibility in terms of installation, reconfigurability, and interoperability. Moreover, fewer computational capabilities can be allocated to the plant side, and more sophisticated control algorithms can be implemented. Unfortunately, differently from point-to-point dedicated cables, communication networks are affected by packet losses and delays. The resulting unreliability has been broadly investigated in the literature and we are now provided with solid theoretical knowledge and useful indications to deal with NCSs, as outlined e.g. by Zhang et al. (2012).

A first remarkable result on NCSs is the fundamental limitation for the stabilization of unstable systems. Schenato et al. (2007) show that there exists a critical threshold on the loss probability above which the system cannot be stabilized. A second established result is the relevance of the freshness of information. In fact, transmissions of the new output or the new control input are preferred over the retransmission of the lost packet, as pointed out by Azimi-Sadjadi (2003). Usually, the information is used as soon as delivered, e.g., the control input is applied at the sampling instant immediately after its reception.

These two aspects have their counterparts in the problem of constrained control. Differently from the unconstrained case, in the literature on model predictive control over communication networks, stable and unstable systems are usually treated analogously. Recursive feasibility and closed-loop stability are typically guaranteed if the number of consecutive packet losses is bounded (see e.g. Quevedo and Nešić (2010)). However, fundamental limitations of unstable systems have not been explicitly addressed. Moreover, the application of those strategies in practice is limited because the typical values of the number of consecutive packet losses for real networks are missing.

Also regarding the information freshness, existing works do not reach a consensus as in the unconstrained case. The main reason is the well-known problem of *input consistency*. More specifically, a control sequence will satisfy the constraints only if the estimate used by the controller and the actual state are *consistent*, that is when the past input trajectory used by the estimator and the actual applied input trajectory at the actuator are the same. It can be shown that the actual state and the estimate are always consistent if the link from the sensor to the controller is ideal (as considered by Quevedo and Nešić (2010)), or if a perfect acknowledgment mechanism is present for the link from the controller to the actuator (as considered by Li and Shi (2013)). In the other cases, input consistency has to be enforced in the control design. In the literature there exist mainly two solutions, corresponding to the two opposite ways of facing information freshness. On one hand, the first method, used e.g. by Findeisen and Varutti (2009), Pin and Parisini (2010), and Pin et al. (2020), proposes to introduce a fictitious fixed delay in the control loop. More specifically, the control input is computed a fixed number of steps in advance in order to zero the loss probability through retransmissions and/or longer allowed communication delays. In this way, the applied input and the computed control input are identical. However, the control system is specifically designed for the worst-case scenario and the information used is always outdated even in good channel conditions. On the other hand, the second solution proposes to compute the input for the next sampling instant. In good channel conditions, the solution boils down to a standard controller, without any additional conservativeness. However, in this way, input consistency is not automatically guaranteed due to dropouts, and the actuator must discard control inputs computed based on a wrong state estimate as done by Pezzutto et al. (2021a). Interestingly, hybrid solutions as Bemporad (1998) and Grüne et al. (2009) can be devised. However, so far, an effective comparison of the two approaches is not available. Additionally, the optimal solution depends on the network used and on the channel conditions, so an extensive experimental campaign is needed.

10.1016/j.ifacol.2023.10.1314

In this work, the two aforementioned aspects are studied for the specific case of Wi-Fi. Wi-Fi is an attractive solution for future control applications over wireless because of its high datarates, which allow low sampling periods, and its large spreading, which ensures a large availability of off-the-shelf devices and technical support. Unfortunately, packet losses and delays are dramatically emphasized because Wi-Fi is not specifically conceived for control applications but it favors the overall throughput over the timeliness of the communications. The overall behavior is difficult to be modeled because it depends on the complex communication system (e.g. stochastic back-off time, multiple antennas, frequency slot allocation) and on the uncontrolled channel conditions (e.g. channel access with competing devices, external interference). Attempts to close the feedback loop over Wi-Fi have been done by Wei et al. (2013) and Branz et al. (2021) but so far constrained control has not been tested.

The main contribution of this work can be summarized as follows: (1) we carry out a detailed experimental assessment of Wi-Fi when used in a typical control setup and we use it to point out the main limits of existing constrained control strategies over lossy communications, (2) we make available online a dataset containing Wi-Fi communication data for different sampling periods and disturbance conditions to help other researchers to test control strategies with realistic WiFi-in-the-loop simulations, (3) we enhance the constrained control strategy based on Reference Governor proposed by Pezzutto et al. (2021b) to better handle the typical behavior of Wi-Fi and we validate it through a full experiment involving a two-wheeled balancing robot.

2. SETUP AND PROBLEM FORMULATION

Consider a discrete-time linear system

$$x_{t+1} = Ax_t + Bu_t + w_t \tag{1}$$

where $x_t \in \mathbb{R}^n$ is the state, $u_t \in \mathbb{R}^m$ is the applied input, $w_t \in W \subset \mathbb{R}^n$ is an unknown disturbance, with W compact and containing the null vector. The system has to fulfill the set of constraints

$$H(Cx_t + Du_t) \leq h \quad \forall t \geq 0 \tag{2}$$

with $C \in \mathbb{R}^{p\times n}$, $D \in \mathbb{R}^{p\times m}$, $H \in \mathbb{R}^{q\times p}$, and $h \in \mathbb{R}^q$.

We assume that the plant and the control unit are connected through a wireless network introducing delays and packet losses in the loop, see Fig. 1. It follows that the applied input u_t cannot be readily manipulated and the control unit has not direct access to the state x_t of the plant. Let X_t denote the packet transmitted by the plant during the period $(t, t+1)$, and U_t denote the packet transmitted by the controller during the period $(t, t+1)$. The arrival process on the link from the plant to the controller, referred to as downlink, is represented by the binary variables

$$\gamma^t_{t-i} = \begin{cases} 1 \text{ if } X_{t-i} \text{ is available at the controller before sending } U_t \\ 0 \text{ otherwise} \end{cases}$$

Similarly, the arrival process on the link from the controller to the plant, referred to as uplink, is represented by the binary variables

$$\theta^t_{t-i} = \begin{cases} 1 \text{ if } U_{t-i} \text{ is available at the plant before sending } X_{t+1} \\ 0 \text{ otherwise} \end{cases}$$

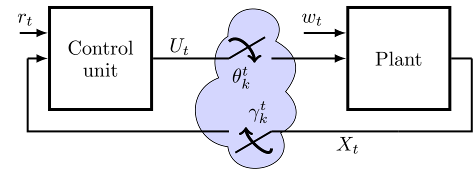

Fig. 1. Setup. Switches represent intermittent and delayed communications.

Note that, if a packet is available at a time instant k, then it will be available also for all future time instants $t > k$. We define the information available to the plant and the information available to the control unit as

$$\mathcal{I}^{\mathrm{P}}_t = \bigcup_{i=1}^{t} \gamma^t_{t-i} U_{t-i} \qquad \mathcal{I}^{\mathrm{C}}_t = \bigcup_{i=0}^{t} \theta^t_{t-i} X_{t-i}$$

respectively, where, with a little abuse of notation, if $\theta = 0$ then $\theta X = \emptyset$. The control objective is to design the laws

$$u_t = f(\mathcal{I}^{\mathrm{P}}_t) \qquad U_t = g(\mathcal{I}^{\mathrm{C}}_t)$$

so that the system (1) converges to the desired reference r_t while satisfying the constraints (2).

The aim of this work is twofold. First, we want to experimentally characterize the arrival processes γ^t_{t-i} and θ^t_{t-i} for a Wi-Fi network. Second, from the experimental data, we want to provide some meaningful guidelines for designing $f(\cdot)$ and $g(\cdot)$ when Wi-Fi is used.

3. EXPERIMENTAL CHARACTERIZATION OF WIFI

In this section, we provide an experimental assessment of Wi-Fi when used in a typical control setup. In view of the experimental outcomes, we draw a critical overview of the existing solutions and we provide some guidelines for constrained control design.

3.1 Experimental setup

The experimental setup consists of a host PC (Intel Core i5-6400, 2.70 GHz, 16GB RAM, Ubuntu 17.04 with wireless Qualcomm Atheros AR9227 chip and PC-link interface) and a computing board (Raspberry Pi 3 mod. B with Broadcom BCM2837 Wi-Fi module) connected by a Wi-Fi network (IEEE 802.11n standard, with the host PC as access point). For the purposes of the paper, the communication parameters are fixed during the experiments. In particular, at the Physical layer, MCS index is set to HT7 (64-QAM modulation, 5/6 coding rate) to trade-off between the datarate and the robustness of the communication. At the Data Link layer, the number of MAC transmission retries is set equal to 1 and, at the Transport layer, UDP protocol is selected with the aim of decreasing the traffic due to old information.

The experiments are carried out in an office environment where the two devices are 4 m apart and temporary obstructions might be given by objects or office users. The Wi-Fi network used for the experiment is set at channel 7 (2442 Mhz), while co-existing Wi-Fi networks are present at channels 5, 9, and 12. This setup is intended to recreate a typical scenario where control over wireless may be used. For instance, in modern factories or for mobile robotics, we expect that control tasks will take place in environments where co-existing Wi-Fi networks are present to provide Internet connection to other users or for IoT applications.

The experiment consists of the transmission of time-stamped packets from the host PC to the computing board and the other way around, with a fixed time-span T between two following transmissions from each machine. Transmission instants from the two machines are spanned approximately of a period $T/2$. This setup represents the standard architecture of NCSs as introduced in Sec. 2. In particular, the host PC plays the role of the control unit and the board represents the communication device at the plant side. Note that the time-span of $T/2$ is a common solution to allow the applied input at time t to exploit the information on the system state at time $t-1$. In this way, in the best case, only one step of delay, unavoidable to manage communications, is present in the control loop.

Using the time-stamps, we can obtain the arrival processes γ_{t-i}^t and θ_{t-i}^t occurred during the experiments. From the arrival processes, we compute the relevant quantities introduced in the following. For simplicity, the definitions are given for the downlink, but the same quantities can be introduced for the uplink.

The delay d_t with which the packet transmitted at time t is received can be computed as

$$d_t = \min d \ \text{ s.t. } \ \gamma_t^{t+d} = 1 \tag{3}$$

Note that, because of the time span of $T/2$ between the transmission of X_t and the transmission of U_t, based on the definition of γ_t^t, a null delay d_t is possible. We introduce the age of information a_t at time t defined as

$$a_t = \min a \ \text{ s.t. } \ \gamma_{t-a}^t = 1 \tag{4}$$

which represents how old is the last received packet at time t. Similarly, the blackout length b_t at time t is defined as

$$b_t = \min b \ \text{ s.t. } \ \gamma_{t-b}^{t-b} = 1 \tag{5}$$

which represents the number of consecutive packet losses at time t. The blackout measure is relevant because it adopts the most used packet loss definition in the literature, where a packet is regarded as lost if arrived later than the following transmission instant.

This setup has been used to test the network behavior for different sampling periods T and different channel conditions. In practice, we obtain different interference conditions by increasing the traffic load on one of the co-existing Wi-Fi networks. We make the sequences of the arrival processes available at `https://github.com/MatthiasPez/WiFi4Control.git`. They can be used to test control algorithms with real wireless communication.

3.2 Results

In this subsection, we report the results obtained for the downlink with $T = 5\,\text{ms}$ under high channel noise. More specifically, during the experiments, a constant traffic of 40 Mbit/s is present on one of the co-existing Wi-Fi networks. The results here reported are averaged over 10 runs of 30 s each, collected in different moments of the day.

Fig. 2 reports the empirical probability density function (pdf) and the cumulative distribution function (cdf) of the delays. As a comparison, we report the cdf of the case with low interference. We can see that the two distributions are substantially different. In the case with low interference, the great part of the packets arrives within the next transmission instant ($\Pr(d_t = 0) \simeq 0.95$), only a few packets arrive with one step of delay ($\Pr(d_t = 1) \simeq 0.03$), and the number of packets never arrived or arrived with more than 10 steps of delay, regarded as residual loss probability, is modest ($\Pr(d_t > 10) \simeq 0.02$). On the other hand, in the case of high noise, the number of packets arrived within the next transmission instant and the number of packets arrived with one step of delay are comparable ($\Pr(d_t = 0) \simeq 0.35$ and $\Pr(d_t = 1) \simeq 0.32$). The residual loss probability is relevant ($\Pr(d_t > 10) \simeq 0.08$).

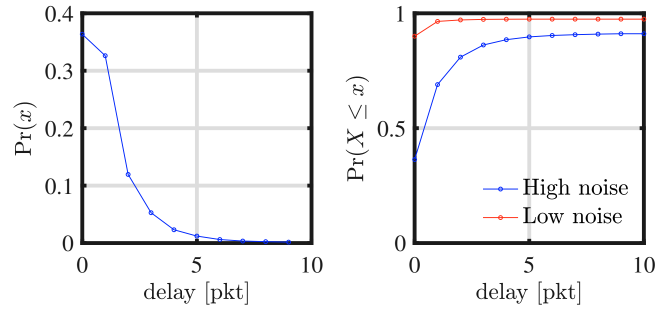

Fig. 2. Empirical pdf (left panel) and cdf (right panel) of delays over Wi-Fi. For higher noise, average delay and residual loss probability are increased.

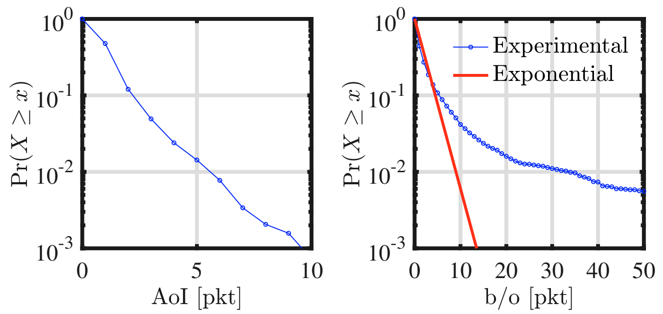

Fig. 3. Empirical ccdf of Age of Information (AoI, left panel) and Blackout (b/o, right panel) over Wi-Fi. Blackout distribution is heavy-tailed.

Fig. 3 reports the empirical complementary cumulative distribution function (ccdf) of the ages of information and of the blackouts on log scale. We can see that the probability of ages larger than 10 steps is small but not null ($\Pr(a_t > 10) \simeq 0.001$). On the other hand, the probability of blackouts longer than 10 steps is relevant ($\Pr(b_t > 10) \simeq 0.03$). As for comparison, we report the ccdf as if the blackout distribution would be exponential, corresponding to the well-known i.i.d. (independent and identically distributed) channel model with loss probability $\Pr(d_t > 0) = 0.65$. We see that longer blackouts have a higher probability with Wi-Fi and the ccdf in log scale has a sublinear convergence. This means that blackouts on Wi-Fi have a heavy-tailed distribution.

3.3 Discussion

The results presented above have some important practical consequences.

First, we can see that it is not possible to guarantee a bound on the blackout length in the case of Wi-Fi with high (but still realistic) channel noise. This limits the application of MPC schemes like those proposed by

Quevedo and Nešić (2010) and by Li and Shi (2013), which guarantee stability if the control horizon is longer than the maximum number of consecutive packet losses. It follows that, according to Fig. 3, even a control horizon equal to 50 steps is not sufficient to achieve a negligible probability of violating the prescribed stability bound. Moreover, during long blackouts, the control sequence would be applied in open-loop without any disturbance rejection. This would likely produce unsatisfactory performances, but poses also a technical issue: according to the standard robust approach (Marruedo et al. (2002)), also adopted by Quevedo and Nešić (2010) and Li and Shi (2013), the constraints are tightened along the control horizon to account for the possible effects of disturbances. However, if the system is unstable, the evolution under the worst-case disturbance diverges and the maximum control horizon to admit a non-empty terminal condition is limited (see (Rawlings and Mayne, 2012, Ch.3)). In turn, the set of network conditions that can be handled by the control algorithm is limited.

A possible way to overcome the aforementioned limits is to adopt the approach of Findeisen and Varutti (2009) and Pin and Parisini (2010) where a fictitious delay is introduced between the computation instant of the control packet and its applications at the actuator. With this approach, stability is guaranteed if the fictitious delay is longer than the maximum age of information, which is much smaller than the maximum blackout, see Fig. 3. Moreover, this approach can, in line of principle, limit the unstable open-loop propagation of the disturbance to the maximum age of information by adopting feedback-based MPC, such as Tube MPC proposed by Mayne et al. (2005). However, this solution requires controlling the system with a fixed long delay, which results into unsatisfactory evolutions in practice, especially with unstable systems. Moreover, robustness for any arbitrary channel condition can not be guaranteed for unstable systems since arbitrarily long fictitious delay cannot be admitted due to the open-loop propagation of the disturbance Even more, the solution might be conservative in several channel conditions since the algorithm is designed for the worst-case scenario. This can be particularly limiting with Wi-Fi networks, since it has substantially different behaviors in good and bad channel conditions, see Fig. 2.

We can conclude that constrained control of unstable systems over Wi-Fi is problematic and should be avoided if it is not possible to guarantee an ideal channel condition for the whole execution of the task. Hybrid solutions as Grüne et al. (2009) can be effective to trade-off between the residual loss probability and the promptness of the response. In particular, in conditions of high noise, introducing 1 or 2 steps of fictitious delay can be beneficial: as it can be seen in Fig. 2, if packets with 1 step of delay are not discarded, the residual loss probability is halved.

4. CONTROL DESIGN AND EXPERIMENTS

In this section, we propose a constrained control strategy for Wi-Fi. Based on the considerations above, we stick to considering stable or pre-stabilized systems and we propose to include a small fictitious delay in the control loop. The proposed strategy is based on Reference Governor (RG). Thanks to its capability in handling arbitrarily long blackouts and its low bandwidth demand, RG have been already used in the context of NCSs (see e.g. Casavola et al. (2006)). Along this line, in the following, we enhance the RG over Wireless studied by Pezzutto et al. (2021b) by including special features tailored for Wi-Fi.

4.1 Notation and preliminaries

Assume that A is asymptotically stable. We define

$$\hat{x}(k|x,u) = A^k x + \sum_{\ell=0}^{k-1} A^\ell B u \qquad F_k = C \bigoplus_{\ell=0}^{k} A^\ell W$$

where $\oplus$ represents a chain of Minkowsky set sum. Note that $\hat{x}(k|x,u)$ corresponds to the nominal prediction k steps ahead of the system state starting from initial condition x and under constant input u, while the set F_k is the reachable set in k steps starting from the origin for a (disturbance) input belonging to W. The constraints are conveniently reformulated as

$$Cx_t + Du_t \in Y \quad \forall t \geq 0$$

where

$$Y = \{y : Hy \leq h\}.$$

We recall the Maximal Output Admissible Set with Packet Loss introduced by Pezzutto et al. (2021b), defined as

$$O_\infty(i) = \{(x,u) : C\hat{x}(k|x,u) + Du \in Y \sim F_{k+i} \ \forall k \geq 0\}$$

where $\sim$ represents the Minkowsky set difference. For a given $\epsilon > 0$, we introduce the inner approximation

$$\tilde{O}_\infty(i) = O_\infty(i) \cap O_\epsilon$$

with

$$O_\epsilon = \{(x,u) : C(I-A)^{-1}Bu + Du \in (1-\epsilon)(Y \sim F_\infty)\}$$

As outlined by Pezzutto et al. (2021b), differently from $O_\infty(i)$, $\tilde{O}_\infty(i)$ consists of a finite number of equations and can be computed with an efficient iterative procedure. The set $\tilde{O}_\infty(i)$ is the key ingredient of the algorithm, since it includes the constraints required to be robust against arbitrarily long blackouts.

For ease of readability, we assume that the packet transmitted to the controller contains both the current state and the current input, i.e. $X_t = \{x_t, u_t\}$. This requires that the actuator and the sensor can communicate over a reliable link. This setup is more restrictive than the usual setup considered in the literature but it is reasonable in several cases, e.g., in mobile robotics, where sensors and actuators are co-located but sophisticated constrained control algorithms cannot be implemented onboard. However, the following strategy can be still implemented with simple modifications if the original system is stable and a lossy acknowledgment mechanism is included in the uplink.

4.2 Control algorithm

The modified RG over Wireless here presented has two main differences with respect to the original solution devised by Pezzutto et al. (2021b). First, the state estimate is computed starting from x_{t-a_t} instead of x_{t-b_t}. This is possible by avoiding discarding the packets with non-null delay and storing the most recent one. The second difference is the introduction of a fictitious delay $s \in \mathbb{N}$ in the control loop. In particular, at time instant t, the control input for the time instant $t + s$ is computed.

More formally, let v_{t+s} denote the input computed by the controller at time t to be applied at time $t + s$, so $U_t = \{v_{t+s}\}$. The underlying idea is to avoid discarding packets with non-null delay also in the uplink, along the line of Grüne et al. (2009) and Pin et al. (2020).

The actuator implements a modified zero-order hold

$$u_t = \theta_{t-s}^t\, v_t + (1 - \theta_{t-s}^t) u_{t-1}. \qquad (6)$$

It can be seen that the applied input is updated if $\theta_{t-s}^t = 1$, that is if the packet U_{t-s} is arrived at time instant t. In other words, the control packets are accepted if arrived with a delay of less than s steps, and discarded otherwise. By choosing a larger s, we decrease the discard rate at the price of using outdated information.

By definition, when computing v_{t+s}, the last known applied input is u_{t-a_t-1}. Thereafter, the computed inputs v_k, $k = t-a_t, \ldots, t+s-1$, may be applied depending on the realization of the random sequence θ_{k+s}^k. It follows that we have 2^{a_t+s} possible input sequences in the period from $t - a_t$ to $t + s - 1$. Let $\Theta = (\vartheta_1, \ldots, \vartheta_i)$ be an arbitrary binary vector of length i and denote

$$U_{k-i:k} = \left[\, v_{k-1}',\ \cdots,\ v_{k-i}',\ u_{k-i-1}' \,\right]'. \qquad (7)$$

Introduce the estimator

$$\hat{x}(t+s, \Theta) = A^{a_t+s+1} x_{t-a_t-1} + \mathcal{R}_{a_t+s} \Lambda_\Theta U_{t-a_t:t+s} \qquad (8)$$

where $\mathcal{R}_i = [B,\ AB,\ \cdots,\ A^{i-1}B,\ A^iB]$ is the reachability matrix in i steps and Λ_Θ is the selection matrix given by the sequence Θ

$$\Lambda_\Theta = \begin{bmatrix} \vartheta_1 \mathrm{I}_m & (1-\vartheta_1)\vartheta_2 \mathrm{I}_m & \ldots & \prod_{\ell=1}^{i}(1-\vartheta_\ell)\mathrm{I}_m \\ 0 & \vartheta_2 \mathrm{I}_m & \ldots & \prod_{\ell=2}^{i}(1-\vartheta_\ell)\mathrm{I}_m \\ 0 & 0 & \ldots & \ldots \\ 0 & 0 & \vartheta_i \mathrm{I}_m & (1-\vartheta_i)\mathrm{I}_m \\ 0 & 0 & 0 & \mathrm{I}_m \end{bmatrix}.$$

Finally, the control input is obtained as

$$v_{t+s} = \arg\min_v \|r_{t+s} - v\|^2 \qquad (9)$$

$$\text{s.t. } (\hat{x}(t+s, \Theta), v) \in \tilde{O}_\infty(a_t + s),\ \Theta \in \{0,1\}^{a_t+s} \qquad (10)$$

If the problem is infeasible, v_{t+s} does not exist. In that case, in order to keep the formalism introduced for the arrival process, the remote control unit transmits an empty packet which, even if received, does not produce any effect at the actuator. If $\|r_{t+s} - v_{t+s}\| > \|r_{t+s} - u_{t-a_t-1}\|$ we assume that v_{t+s} is discarded and an empty packet is sent.

Note that, in order to manage the uncertainty on the initial state of the optimization problem due to the disturbances occurred since the last received packet, we use the Maximal Output Admissible Set with Packet Loss $\tilde{O}_\infty(a_t + s)$. Moreover, in order to manage the uncertainty on the initial state of the optimization problem due to the unknown applied inputs, we impose that the new computed input v_{t+s} satisfies the constraints for any possible sequence of applied inputs. This is expressed by the condition on $\hat{x}(t + s, \Theta)$, which, by varying Θ, returns any possible state compatible with any possible applied input sequence. Note that, at the best case, $a_t = 0$. It follows that, in ideal channel conditions, there is a conservativeness due to tighter constraints in $\tilde{O}_\infty(s)$ with respect to $\tilde{O}_\infty(0)$ and the condition imposed on $\hat{x}(t+s, \Theta)$ for $\Theta \in \{0, 1\}^s$ instead of on $Ax_t + Bu_t$. The following proposition summarizes the properties of the proposed solution.

Proposition 1. Assume $(x_0, u_0) \in \tilde{O}_\infty(s)$. Then, constraints are satisfied with probability 1 for any $t \geq 0$. Additionally, assume that $\mathbf{P}(\bigcap_{k>0}\{\theta_{t_k}^{t_k}\gamma_{t_k}^{t_k+s} = 0\}) = 0$ for any infinite sequence $\{t_k\}$ with $t_{k+1} > t_k$. If $W = \{0\}$ and $r_t = r$, the applied input u_t reaches the reference r^* and the state converges to $(I - A)^{-1}Br^*$, where $r^* = \arg\min_v \|r - v\|^2$ s.t. $((I - A)^{-1}Bv, v) \in \tilde{O}_\infty(s)$.

Proof. The critical point of the proof is to show that the sequence $\{\tau_k\}$ of time instants where the optimization problem is feasible is infinitely long. This can be proved as follows: suppose $\{\tau_k\}$ is not infinite, i.e., the optimization problem is not feasible from a certain τ onward. So $u_{\tau+i} = u_\tau$ for $i \geq 0$. By assumption on the arrival processes, there exists with probability 1 a time instant t where $\theta_t^t = 1$ and the problem is feasible because $v_{t+s} = u_\tau$ is an admissible solution. We reach an absurd and we conclude that $\{\tau_k\}$ is infinitely long. The rest of the proof can be obtained following the proof of Proposition 3.3 in Pezzutto (2022).

Interestingly, we can see that constraints are satisfied without any requirement on the network. This is possible because the system is stable, so that there always exists an input which can be kept constant for any arbitrary long period and such that the constraints are enforced. Also, the convergence to the best admissible approximation r^* of the desired setpoint r requires only mild assumptions on the arrival processes. Essentially, it is sufficient that any infinite sequence of instants has two consecutive successful receptions, one per link. Clearly, theoretical properties are guaranteed also if some control packets arrive after s time steps, differently e.g. from Pin and Parisini (2010).

4.3 Experimental setup

The experimental setup consists of a remote PC and a Segway-like robot connected through a Wi-Fi network. Details on the mechanics, the electronics, and the mathematical model are given in Pezzutto (2022). The proposed control architecture comprises a stabilizing LQR controller and a simple estimator, implemented on board, and the proposed RG with $s = 2$, implemented on the PC. The control objective is to track the desired position reference $r_t = 0\,\mathrm{m}$ for $t < 2\,\mathrm{s}$ and $r_t = 3\,\mathrm{m}$ for $t \geq 2\,\mathrm{s}$ while keeping the tilt angle in the interval $[-0.1\,\mathrm{rad}, 0.1\,\mathrm{rad}]$. The Wi-Fi network is set as in the previous section. The experiment is conducted under mild channel interference.

We use the communication data obtained during the experiment to compare the proposed strategy with the Networked MPC by Pin and Parisini (2010) through an accurate WiFi-in-the-loop simulation on Simulink. For the Networked MPC, we propose to set $\tau_{ac} = 10$, where τ_{ac} is the fictitious delay introduced between the computation instant of the control packet and its applications at the actuator. Following the original idea by Pin and Parisini (2010), no inner controller is present. The same estimator used with the proposed strategy is employed. More details can be found in Pezzutto (2022).

4.4 Results

The channel evolution is represented in Fig. 4 in terms of age of information. The results are in line with the network assessment presented in the previous section. Remarkably,

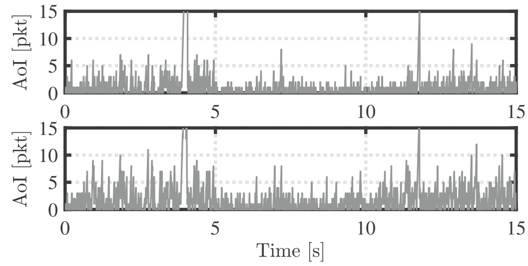

Fig. 4. Age of Information evolution during the experiment. Top panel: uplink. Bottom panel: downlink.

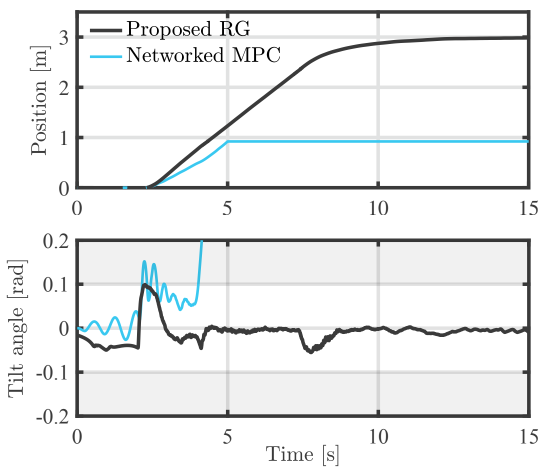

Fig. 5. System responses. Proposed RG succeeds in the task while Networked MPC violates the constraints.

we can see that the behavior on the two links is correlated, as higher interference conditions affect both the links. Note that AoI is large ($\simeq 30$ steps in the uplink) around $t = 4$ s.

System evolution is reported in Fig. 5. In the top panel, we can see that the reference is reached by the proposed strategy. As we can see in the bottom panel, constraints are satisfied despite packet loss and the tilt angle evolution is quite smooth. Conversely, evolutions obtained by Networked MPC are not satisfactory. In the first part of the simulation, we can see that the tilt angle oscillates noticeably and, when the robot starts moving, constraints are violated. Through extensive simulations, we have observed that the oscillations are caused by the noisy velocities obtained by discrete differentiation but they are drastically emphasized by the fictitious delay τ_{ac}. During the long blackout, constraints are violated and the robot falls.

5. CONCLUSION

In this work, we have shown experimentally the main challenges of constrained control over Wi-Fi and how to tackle them. Future works will study how to effectively use Wi-Fi for multi-agent constrained control applications.

REFERENCES

Azimi-Sadjadi, B. (2003). Stability of networked control systems in the presence of packet losses. In *Proc. 42nd IEEE International Conference on Decision and Control*, volume 1, 676–681. IEEE.

Bemporad, A. (1998). Predictive control of teleoperated constrained systems with unbounded communication delays. In *Proc. 37th IEEE Conference on Decision and Control*, volume 2, 2133–2138. IEEE.

Branz, F., Antonello, R., Pezzutto, M., Vitturi, S., Tramarin, F., and Schenato, L. (2021). Drive-by-Wi-Fi: Model-based control over wireless at 1 khz. *IEEE Trans. Control Syst. Technol.*, 30(3), 1078–1089.

Casavola, A., Papini, M., and Franzé, G. (2006). Supervision of networked dynamical systems under coordination constraints. *IEEE Trans. Autom. Control*, 51(3).

Findeisen, R. and Varutti, P. (2009). Stabilizing nonlinear predictive control over nondeterministic communication networks. In *Nonlinear model predictive control*, 167–179. Springer.

Grüne, L., Pannek, J., and Worthmann, K. (2009). A prediction based control scheme for networked systems with delays and packet dropouts. In *Proc. 48h IEEE Conference on Decision and Control*, 537–542. IEEE.

Li, H. and Shi, Y. (2013). Network-based predictive control for constrained nonlinear systems with two-channel packet dropouts. *IEEE Trans. Ind. Electron.*, 61(3), 1574–1582.

Marruedo, D.L., Alamo, T., and Camacho, E. (2002). Input-to-state stable MPC for constrained discrete-time nonlinear systems with bounded additive uncertainties. In *Proc. 41st IEEE Conference on Decision and Control*, volume 4, 4619–4624. IEEE.

Mayne, D.Q., Seron, M.M., and Raković, S. (2005). Robust model predictive control of constrained linear systems with bounded disturbances *Automatica*, 41(2), 219–224.

Pezzutto, M., Farina, M., Carli, R., and Schenato, L. (2021a). Remote MPC for tracking over lossy network. *IEEE Contr. Syst. Lett.*, 6, 1040–1045.

Pezzutto, M. (2022). Predictive constrained control over wireless: from stability to safety. *Ph.D. Thesis*, `https://www.research.unipd.it/handle/11577/3438844`.

Pezzutto, M., Garone, E., and Schenato, L. (2021b). Constrained Control with Communication Blackouts: Theory and Experimental Validation over Wi-Fi. In *Proc. 29th Mediterranean Conference on Control and Automation*, 238–243. IEEE.

Pin, G., Fenu, G., Casagrande, V., Zorzenon, D., and Parisini, T. (2020). Robust stabilization of a class of nonlinear systems controlled over communication networks. *IEEE Trans. Autom. Control*, 66(7).

Pin, G. and Parisini, T. (2010). Networked predictive control of uncertain constrained nonlinear systems: Recursive feasibility and input-to-state stability analysis. *IEEE Trans. Autom. Control*, 56(1), 72–87.

Quevedo, D.E. and Nešić, D. (2010). Input-to-state stability of packetized predictive control over unreliable networks affected by packet dropouts. *IEEE Trans. Autom. Control*, 56(2), 370–375.

Rawlings, J.B. and Mayne, D.Q. (2012). *Model Predictive Control: Theory and Design*. Nob Hill Publishing.

Schenato, L., Sinopoli, B., Franceschetti, M., Poolla, K., and Sastry, S.S. (2007). Foundations of control and estimation over lossy networks. *Proc. IEEE*, 95(1).

Wei, Y.H., Leng, Q., Han, S., Mok, A.K., Zhang, W., and Tomizuka, M. (2013). RT-WiFi: Real-time high-speed communication protocol for wireless cyber-physical control applications. In *Proc. 34th Real-Time Systems Symposium*, 140–149. IEEE.

Zhang, L., Gao, H., and Kaynak, O. (2012). Network-induced constraints in networked control systems—a survey. *IEEE Trans. Ind. Inf.*, 9(1), 403–416.