Automatica 91 (2018) 1–9

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Random access design for wireless control systems[☆]

Konstantinos Gatsis [*], Alejandro Ribeiro, George J. Pappas

Department of Electrical and Systems Engineering, University of Pennsylvania, 200 South 33rd Street, Philadelphia, PA 19104, United States

ARTICLE INFO

Article history:
Received 21 April 2016
Received in revised form 3 June 2017
Accepted 17 December 2017

Keywords:
Networked control systems
Random access
Optimization
Decentralized systems
Communication channels

ABSTRACT

Interferences arising between wireless sensor–actuator systems communicating over shared wireless channels impact closed loop control performance. We design interference-aware channel access policies where the total transmit power of the sensors is minimized while desired control performance is guaranteed for each involved control loop. Control performance is abstracted as an expected decrease rate of a Lyapunov function for each loop. We prove that the optimal channel access policies are decoupled so that, intuitively, each sensor balances the gains from transmitting to its actuator with the negative interference effect on all other control loops. Moreover the optimal policies are of a threshold nature with respect to channel conditions, that is, a sensor transmits only under favorable local fading conditions. Finally, the optimal policies can be computed by a distributed iterative procedure which does not require coordination between the sensors.

1. Introduction

Wireless sensors are an essential part of modern smart communities where they are deployed to monitor and control physical processes in our homes, urban infrastructures, agriculture, and industrial plants. This abundance of wireless devices however also creates an increase in the wireless interferences arising between transmissions over the shared wireless medium. The development of decentralized communication mechanisms that can manage these interference effects and guarantee closed loop control performance arises as an important research direction.

The prevalent approach to the problem of sharing a wireless communication medium in networked control systems is centralized scheduling which guarantees no interferences. Static scheduling for example specifies that sensors transmit in some predefined periodically repeating sequence such as round-robin and this sequence is designed to meet control objectives, see, e.g., Hristu-Varsakelis (2001), Le Ny, Feron, and Pappas (2011) and Zhang, Branicky, and Phillips (2001). Deriving optimal scheduling sequences is recognized as a hard combinatorial problem (Gupta, Chung, Hassibi, & Murray, 2006; Rehbinder & Sanfridson, 2004). Scheduling can also be dynamic, where at each time step a central network coordination authority decides which device gets access to the medium. This dynamic decision may be stochastic (Gupta et al., 2006), based on plant state information (Donkers, Heemels, Van De Wouw, & Hetel, 2011; Walsh, Ye, & Bushnell, 2002), or based on the wireless channel conditions (Gatsis, Pajic, Ribeiro, & Pappas, 2015).

Besides scheduling, decentralized mechanisms where sensors independently decide access to the shared wireless medium also become practically useful for sensors that communicate information infrequently. Compared to centralized approaches they are easier to implement as they do not require predesigned sequences of how sensors access the medium, or a central authority to take scheduling decisions. The drawback of this decentralized approach however is that packet collisions can occur from simultaneously transmitting sensors, resulting in lost packets and control performance degradation. Hence sensor access policies need to be appropriately designed taking into account these effects. We consider specifically a random access mechanism where each sensor independently and randomly decides whether to transmit plant state measurements over a shared channel to an access point/controller (Fig. 1).

Control under random access communication mechanisms has drawn limited attention, to the best of our knowledge. Comparisons between different medium access mechanisms for networked control systems and the impact of packet collisions in stability and control performance have been considered either in numerical simulations (Liu & Goldsmith, 2004; Ramesh, Sandberg, & Johansson, 2013) or analytically in simple cases (Blind & Allgöwer, 2011; Rabi, Stabellini, Proutiere, & Johansson, 2010).

[☆] This work was supported in part by NSF CNS-0931239, and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA. The material in this paper was partially presented at the 54th IEEE Conference on Decision and Control, December 15-18, 2015, Osaka, Japan. This paper was recommended for publication in revised form by Associate Editor Claudio De Persis under the direction of Editor Christos G. Cassandras.

[*] Corresponding author.
E-mail addresses: kgatsis@seas.upenn.edu (K. Gatsis), aribeiro@seas.upenn.edu (A. Ribeiro), pappasg@seas.upenn.edu (G.J. Pappas).

https://doi.org/10.1016/j.automatica.2018.01.021
0005-1098/

These include random access mechanisms and related Aloha-like schemes, where after a packet collision the involved sensors wait for a random time interval and retransmit. Stability conditions under packet collisions were examined in Tabbara and Nesic (2008) and Zhang (2003). In contrast to these works, our goal is to directly design the medium access mechanism to guarantee control performance. Besides closed loop control, optimal remote estimation over collision channels is considered recently in Vasconcelos and Martins (2017).

We pose the design of channel access policies for multiple control loops over a shared wireless channel as an optimization problem (Section 2). The goal is to satisfy a control performance requirement for each control loop while minimizing the total expected transmit power expenditures. We adopt a Lyapunov-like control performance abstraction, motivated from our work on centralized scheduling (Gatsis, Pajic, et al., 2015). Each control system is abstracted via a given Lyapunov function which is required to decrease at a desired rate and in expectation over to the random packet losses and collisions on the shared medium.

Besides accounting for packet collisions, sensors can exploit channel fading state information. Fading refers to large unpredictable variations in wireless channel transferences (Goldsmith, 2005, Ch. 3,4), affecting the likelihood of successful packet decoding. This communication model has been used in estimation and control applications (Gatsis, Pajic, et al., 2015; Gatsis, Ribeiro, & Pappas, 2014; Quevedo, Ahlén, Leong, & Dey, 2012) but not under a random access mechanism. By adapting online to channel states sensors may, e.g., access the channel at higher rates under channel conditions with higher packet success. In preliminary work presented in Gatsis, Ribeiro, and Pappas (2015) we considered this random access problem but employing simpler policies that do not adapt to channel states online.

Based on Lagrange duality arguments we characterize the structure of the optimal sensor access policies (Section 4). We show that the optimal policies are of a threshold nature, that is, each sensor transmits only when its corresponding channel state is favorable enough and backs off otherwise. Moreover we reveal an intuitive decoupling among sensors; each sensor should select its threshold in a way that balances the control performance of its own closed loop with the collective negative effect it has on all other control loops due to collisions. Optimal decentralized policies are also known for general wireless random access networks targeting throughput (Adireddy & Tong, 2005; Hu & Ribeiro, 2011; Qin & Berry, 2006), but this is the first time this is shown for control performance.

In Section 5 we derive an iterative procedure to compute the optimal access policies. The procedure is easy to implement in our architecture as it does not require the sensors to coordinate among themselves, or to know what control performance the other sensors try to achieve. We conclude with a numerical example and some remarks (Sections 6, 7).

2. System description

We consider a wireless control architecture where m independent plants are controlled over a shared wireless medium. Each sensor i ($i = 1, 2, \ldots, m$) transmits measurements of plant i to an access point responsible for computing the plant control inputs. Packet collisions might arise on the shared medium between simultaneously transmitting sensors. See Fig. 1 for an illustration. We are interested in designing a mechanism for each sensor to independently decide whether to access the medium (random access) in a way that guarantees desirable control performance for all control systems.

The dynamics for each of the m control systems are assumed pre-designed independently of the communication policy, and are described by a switched model that depends on whether the controller manages to reach the access point or not. Thus, if we use $\gamma_{i,k} \in \{0, 1\}$ to indicate the success of the transmission at time k for link/system i and assume the system is linear and time invariant, we can model its evolution by the switched system

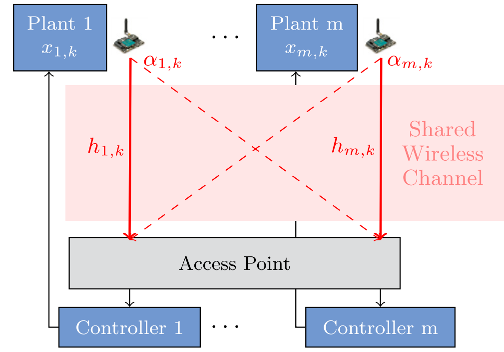

Fig. 1. Random access architecture for m control loops over a shared wireless medium. Each sensor i randomly transmits with probability $\alpha_{i,k}$ at time k to a common access point computing the plant control inputs. If only sensor i transmits, the successful decoding probability depends on local channel conditions $h_{i,k}$. If other sensors transmit at the same time a collision might occur at sensor i's transmission, rendering i's packet lost.

$$x_{i,k+1} = \begin{cases} A_{c,i}\, x_{i,k} + w_{i,k}, & \text{if } \gamma_{i,k} = 1, \\ A_{o,i}\, x_{i,k} + w_{i,k}, & \text{if } \gamma_{i,k} = 0. \end{cases} \tag{1}$$

Here $x_{i,k} \in \mathbb{R}^{n_i}$ denotes the state of control system i at each time k, which can in general include both plant and controller states — see, e.g., Example 1. At a successful transmission the system dynamics are described by the matrix $A_{c,i} \in \mathbb{R}^{n_i \times n_i}$, where 'c' stands for closed-loop, and otherwise by $A_{o,i} \in \mathbb{R}^{n_i \times n_i}$, where 'o' stands for open-loop. We assume that $A_{c,i}$ is asymptotically stable, implying that if system i successfully transmits at each slot the state evolution of $x_{i,k}$ is stable. The open loop matrix $A_{o,i}$ may be unstable. The additive terms $w_{i,k}$ model an independent (both across time k for each system i, and across systems) identically distributed (i.i.d.) noise process with mean zero and covariance $W_i \succeq 0$.

Example 1. Suppose each closed loop i consists of a linear plant and a linear output of the form

$$x_{i,k+1} = A_i x_{i,k} + B_i u_{i,k} + w_{i,k}, \tag{2}$$

$$y_{i,k} = C_i x_{i,k} + v_{i,k}, \tag{3}$$

where $\{w_{i,k},\ k \geq 0\}$ and $\{v_{i,k},\ k \geq 0\}$ are i.i.d. Gaussian disturbance and measurement noises respectively. Each wireless sensor i transmits the output measurement $y_{i,k}$ to the controller. A dynamic control law adapted to the packet drops keeps a local controller state $z_{i,k}$,

$$z_{i,k+1} = F_i\, z_{i,k} + \gamma_{i,k}\,(F_{c,i}\, z_{i,k} + G_i\, y_{i,k}) \tag{4}$$

which may for example represent a local estimate of the plant state (Hespanha, Naghshtabrizi, & Xu, 2007), and applies plant input $u_{i,k}$ as

$$u_{i,k} = K_i\, z_{i,k} + \gamma_{i,k}\,(K_{c,i}\, z_{i,k} + L_i\, y_{i,k}). \tag{5}$$

In other words, the controller updates appropriately the local state and input whenever a measurement is received. The overall closed loop system is obtained by joining plant and controller states into

$$\begin{bmatrix} x_{i,k+1} \\ z_{i,k+1} \end{bmatrix} = \begin{bmatrix} A_i + \gamma_{i,k} B_i L_i C_i & B_i K_i + \gamma_{i,k} B_i K_{c,i} \\ \gamma_{i,k} G_i C_i & F_i + \gamma_{i,k} F_{c,i} \end{bmatrix} \cdot \begin{bmatrix} x_{i,k} \\ z_{i,k} \end{bmatrix} + \begin{bmatrix} I & \gamma_{i,k} B_i L_i C_i \\ 0 & \gamma_{i,k} G_i \end{bmatrix} \begin{bmatrix} w_{i,k} \\ v_{i,k} \end{bmatrix} \tag{6}$$

which is of the form (1). □

The transmission success indicator variables $\gamma_{i,k}$ are random with a distribution that depends on the communication policy which here is supposed to be a slotted random access policy. Specifically, communication takes place in time slots generically indexed by k. At every slot k each sensor i transmits over the shared channel with some probability $\alpha_{i,k} \in [0, 1]$ to be designed. A sensor's transmission might fail due to two reasons, packet decoding errors and packet collisions. A collision is experienced on link i, rendering packet i lost, if some other sensor $j \neq i$ transmits in the same time slot (Tabbara & Nesic, 2008; Vasconcelos & Martins, 2017; Zhang, 2003) — see also Remark 1. Thus, the probability that sensor i's transmission is free of collisions, i.e., that no other sensor transmits and causes collisions on link i, equals $\prod_{j\neq i}[1-\alpha_{j,k}]$.

If sensor i transmits and has a collision-free time slot, the success of decoding the packet at the access point/receiver depends on the randomly varying channel conditions on link i. Denote then by $h_{i,k} \in \mathbb{R}_+$ the current channel fading conditions for link i at time k. We adopt a block fading model (Goldsmith, 2005, Ch. 4) whereby channel states $\{h_{i,k},\ k \geq 0\}$ are assumed constant during each transmission slot k, but i.i.d. across time with distribution ϕ_i. We also assume channel states are independent among systems i (Adireddy & Tong, 2005; Hu & Ribeiro, 2011; Qin & Berry, 2006), as well as independent of the plant process noise $w_{i,k}$. If sensor i transmits at a power level p_i, and no other sensor transmits, the power level of the received signal is the product $h_{i,k}p_i$ of the current channel fading gain and transmit power (Goldsmith, 2005, Ch. 3). During high channel fading gains for sensor i there is a higher received signal-to-noise ratio (SNR) at the access point/controller and consequently a higher chance to successfully decode the transmitted message. We let $q(h_{i,k})$ describe this relationship between the packet success and the channel state — for more details on this model the reader is referred to Gatsis et al. (2014). An illustration of this relationship is given in Fig. 2. Before deciding whether to transmit over the shared channel, each sensor i has access to the channel state $h_{i,k}$, e.g., channel may be estimated by a short pilot signal sent from the access point to the sensors at the beginning of each time slot (Gatsis et al., 2014; Hu & Ribeiro, 2011; Qin & Berry, 2006). The pilot signal may also be used for synchronization purposes of the slotted random access architecture (Hu & Ribeiro, 2011; Van de Beek et al., 1999). The function $q : \mathbb{R}_+ \to [0, 1]$ is assumed to be continuous and strictly increasing, i.e., higher channel fading states imply higher packet success probability.

Combining the effects of collisions and packet losses due to fading, the probability that a packet is successfully decoded at the access point can be written as

$$\mathbb{P}(\gamma_{i,k}=1) = \alpha_{i,k}\, q(h_{i,k}) \prod_{j\neq i}\Big[1-\alpha_{j,k}\Big]. \tag{7}$$

This expression states that the probability of system i in (1) closing the loop at time k equals the probability that transmission i is successfully decoded at the receiver, multiplied by the probability that no other sensor $j \neq i$ is causing collisions on ith transmission.

As channel states reveal information about how easy it is for each sensor to successfully communicate to the access point, we let the channel access decision $\alpha_{i,k}$ for each sensor i adapt to its local channel states $h_{i,k}$ respectively. Since channel states are i.i.d. over time we restrict attention to stationary policies, and drop the time index when not necessary. Hence we design policies that are measurable functions of the form $\alpha_{i,k} = \alpha_i(h_{i,k})$. The set of all access policies for sensor i is the set of functions

$$\mathcal{A}_i = \{\alpha_i : \mathbb{R}_+ \to [0, 1]\} \tag{8}$$

and the vector $\alpha(.)$ of access policies for all sensors belongs in the Cartesian product space $\mathcal{A} = \mathcal{A}_1 \times \cdots \mathcal{A}_m$. For fixed sensor access policies, the probability of successful transmission on link i can be expressed as

$$\mathbb{P}(\gamma_{i,k}=1) = \mathbb{E}_{h_i}[\alpha_i(h_i)\, q(h_i)] \prod_{j\neq i}\Big[1-\mathbb{E}_{h_j}[\alpha_j(h_j)]\Big]. \tag{9}$$

This expression follows from (7) by taking expectation with respect to the channel states and using the independence of channels among systems. The expectation is well-defined as both functions $\alpha(.)$, $q(.)$ are measurable and bounded in $[0, 1]$ hence integrable. We also make the following technical assumption on the probability distribution of channel states, which holds true for practically considered models (Goldsmith, 2005, Ch. 3).

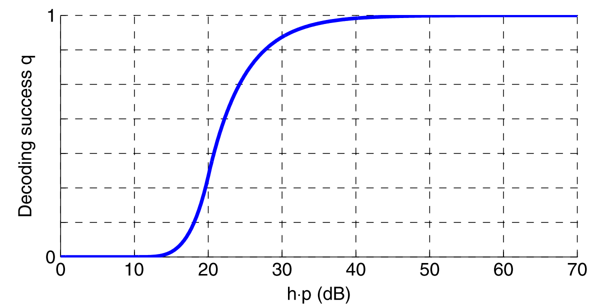

Fig. 2. Probability of successful decoding as a function of the received signal power level. Higher channel fading gains imply higher chance of packet success, hence more more favorable transmission opportunities.

Assumption 1. The distributions ϕ_i of channel states $\{h_{i,k},\ k \geq 0\}$ for all $i = 1, \ldots, m$ are non-atomic, i.e., have a continuous distribution function on $\mathbb{R}_+$.

In the following section we formulate the problem of designing the optimal sensor access policies $\alpha \in \mathcal{A}$ such that a level of control performance is guaranteed for each involved control loop.

Remark 1. In practice multiple packets may be correctly decoded when multiple sensors access the channel simultaneously, e.g., due to capture (Adireddy & Tong, 2005; Luo & Ephremides, 2002). The probability of this multi-packet reception event depends on the received signal-to-interference-ratio (SINR) and hence on the channel fading conditions of all links. Instead, in our collision model the probability of multi-packet reception is zero and hence is a lower approximation of the actual probability. Similar approximate collision models are also common in wireless communications literature (Adireddy & Tong, 2005; Hu & Ribeiro, 2011; Qin & Berry, 2006). This introduces conservativeness in the design, i.e., the probability a system closes its loop in practice will be larger than in (9). However this simple model facilitates the design of optimal sensor access policies (Theorems 1, 2) which would become harder in the general multi-packet reception model and will be considered in future work. It is also worth noting that our model easily extends to the case where some sensors j do not interfere at all with some other sensor i by simply excluding those j from the product in (9).

3. Control performance and random access problem

The random packet success on link i modeled by (9) causes each control system i in (1) to switch in a random fashion between the two modes of operation (open and closed loop). As a result, the sensor access policies $\alpha(.)$ to be designed affect the performance of all control systems. In this paper we account for control performance via Lyapunov functions. The following proposition establishes a connection between Lyapunov control performance abstractions and the packet success rates over wireless links.

Proposition 1 (*Control Performance Abstraction*). *Consider a switched linear system i described by (1) and $\gamma_{i,k}$ being a sequence of i.i.d. Bernoulli random variables, and a quadratic function $V_i(x_i) = x_i^T P_i x_i,\ \ x_i \in \mathbb{R}^{n_i}$ with a positive definite matrix $P_i \succ 0$. Then the function $V_i(x_i)$ decreases with an expected rate $\rho_i < 1$ at each step, i.e., we have*

$$\mathbb{E}\left[V_i(x_{i,k+1}) \mid x_{i,k}\right] \leq \rho_i\, V_i(x_{i,k}) + Tr(P_i W_i) \tag{10}$$

for all $x_{i,k} \in \mathbb{R}^{n_i}$, if and only if

$$\mathbb{P}(\gamma_{i,k} = 1) \geq c_i, \tag{11}$$

where $c_i \geq 0$ is computed by the semidefinite program

$$c_i = \min\{\ \theta \geq 0 : \\ \theta\ A_{c,i}^T P_i A_{c,i} + (1-\theta) A_{o,i}^T P_i A_{o,i} \preceq \rho_i P_i\}. \tag{12}$$

Proof. The expectation over the next system state $x_{i,k+1}$ on the left hand side of (10) accounts via (1) for the randomness introduced by the process noise $w_{i,k}$ and the random success $\gamma_{i,k}$. In particular we have that

$$\mathbb{E}\left[V_i(x_{i,k+1}) \mid x_{i,k}\right] = \mathbb{P}(\gamma_{i,k} = 1)\ x_{i,k}^T A_{c,i}^T P_i A_{c,i} x_{i,k} \\ + \mathbb{P}(\gamma_{i,k} = 0)\ x_{i,k}^T A_{o,i}^T P_i A_{o,i} x_{i,k} + Tr(P_i W_i). \tag{13}$$

Here we used the fact that the random variable $\gamma_{i,k}$ is independent of the system state $x_{i,k}$. Plugging (13) at the left hand side of (10) we get

$$\mathbb{P}(\gamma_{i,k} = 1)\ x_{i,k}^T A_{c,i}^T P_i A_{c,i} x_{i,k} \\ + \mathbb{P}(\gamma_{i,k} = 0)\ x_{i,k}^T A_{o,i}^T P_i A_{o,i} x_{i,k} \leq\ \rho_i x_{i,k}^T P_i x_{i,k}. \tag{14}$$

Since condition (10) needs to hold for all $x_{i,k} \in \mathbb{R}^{n_i}$ we can rewrite (14) as a linear matrix inequality (Boyd & Vandenberghe, 2009)

$$\mathbb{P}(\gamma_i = 1) A_{c,i}^T P_i A_{c,i} + (1 - \mathbb{P}(\gamma_i = 1)) A_{o,i}^T P_i A_{o,i} \preceq \rho_i P_i, \tag{15}$$

where we dropped the time indices from $\gamma_{i,k}$ since they are i.i.d. by assumption. The values $\mathbb{P}(\gamma_i = 1)$ that satisfy this linear matrix inequality belong in some closed convex set. As a result there is a minimum value c_i, given by the semidefinite program (12), such that condition (15) is equivalent to $\mathbb{P}(\gamma_i = 1) \geq c_i$. ■

The interpretation of this result is as follows. When the loop closes, the quadratic function $V_i(x_i)$ of the system state decreases due to the stable closed loop dynamics of (1), while during packet drops and collisions it increases due to the potentially unstable dynamics. Condition (10) requires that $V_i(x_i)$ acts as a Lyapunov function and decreases in expectation over the packet success regardless of the current plant state. In this paper we assume that quadratic Lyapunov functions $V_i(x_i)$ and desired expected decrease rates ρ_i are given for each control system. They present a control interface for communication design over a shared wireless medium. We design the sensor access policies so that *all Lyapunov functions* (10) decrease at the desired rates $\rho_i < 1$ at any time k *in expectation*. By the above proposition, these control performance requirements are equivalent to minimum packet success rates (11) for each link i, computed by (12). Hence we need to ensure that (11) holds for all links i. An implicit assumption for Proposition 1 and throughout the paper is that if system i always remains in closed loop then the desired decrease rate is met, that is, $A_{c,i}^T P_i A_{c,i} \preceq \rho_i P_i$.

Besides control performance, it is desired that the sensors' channel access mechanism makes an efficient use of their power resources. We assume that when sensor i decides to access the channel it transmits with a fixed power $p_i > 0$. The total expected power consumption at each slot is given by $\sum_{i=1}^m \mathbb{E}_{h_i} \alpha_i(h_i) p_i$ summing up the transmit power of each system i if the system decides to transmit. We pose then the design of the sensor access policies α that minimize the total expected power consumption subject to the desired control performance (10) (equivalently (11)) for all plants as

$$\underset{\alpha \in \mathcal{A}}{\text{minimize}}\quad \sum_{i=1}^m \mathbb{E}_{h_i} \alpha_i(h_i) p_i \tag{16}$$

$$\text{subject to}\ \ c_i \leq \mathbb{E}_{h_i}[\alpha_i(h_i) q(h_i)] \prod_{j \neq i} \Big[1 - \mathbb{E}_{h_j}[\alpha_j(h_j)]\Big] \\ i = 1, \ldots, m. \tag{17}$$

Technically we assume that the problem is strictly feasible, a common constraint qualification assumption that will allow us to examine the Lagrange dual problem in the proof in Appendix A and in particular in Proposition 2 therein. Sufficient conditions for this can be found by checking any given channel access policy. For example, considering all sensors to transmit with a constant equal probability $\alpha_i(h_i) = \frac{1}{m}$, a sufficient condition for (17) to be strictly feasible becomes $c_i < \mathbb{E}_{h_i} q(h_i) \frac{1}{m} (1 - \frac{1}{m})^{m-1}$ for all $i = 1, \ldots, m$. This condition depends jointly on the required control performance c_i, on the channel distribution $\phi_i(h_i)$, as well as on the transmission characteristic $q(h_i)$. More generally we assume the following.

Assumption 2. There exists $\alpha' \in \mathcal{A}$ that satisfies constraints (17) with strict inequality.

We make a final technical comment. In order to meet the constraints (17) it is obvious that the successful transmission rate $\mathbb{E}_{h_i}[\alpha_i(h_i) q(h_i)]$ for each sensor i must be strictly positive. Similarly the channel access rate $\mathbb{E}_{h_j}[\alpha_j(h_j)]$ for any sensor j must be strictly less than one. We assume a sufficiently small positive parameter $\zeta > 0$ is known such that implicit constraints $\mathbb{E}_{h_i}[\alpha_i(h_i) q(h_i)] > \zeta$ and $\mathbb{E}_{h_j}[\alpha_j(h_j)] < 1 - \zeta$ hold without loss. A sufficient choice for example is $\zeta \leq \min_i c_i$.

In the following section we prove that the optimal access policies α^* follow a simple and intuitive decoupled structure. Each sensor independently accesses the channel in a way that trades off its own control loop with the effect of collisions on all other control loops collectively. In Section 5 we develop a procedure to find these optimal policies.

4. Channel-aware random access design

Our main result is the following characterization of the optimal access policies for the sensors.

Theorem 1 (*Optimal Sensor Access*). *Consider a random access architecture with m control loops of the form (1), communication modeled by (9), and control performance abstracted by (10)–(11) for each loop $i = 1, \ldots, m$. Consider the design of optimal sensor access policies (16)–(17), and let Assumptions 1, 2 hold. Then there exist non-negative vectors $\nu^* \in \mathbb{R}_+^m, \mu^* \in \mathbb{R}_+^m$ such that the optimal sensor access policy is written as*

$$\alpha_i^*(h_i) = \begin{cases} 1 & \text{if}\ \ \nu_i^*\, q(h_i) \geq\ p_i + \mu_i^* \\ 0 & \text{otherwise.} \end{cases} \tag{18}$$

for each sensor $i = 1, \ldots, m$.

We observe the following interesting facts. First note that the optimal policies are deterministic, that is, given current channel conditions each sensor either transmits or not. Second by the strict monotonicity of the packet success function $q(.)$, the optimal sensor access policies in (18) are threshold policies. That is, a sensor attempts to close its loop only when its corresponding channel quality is above some threshold, and backs off otherwise. Third and more importantly, given the vectors ν^*, μ^* the optimal policies are decoupled among the sensors. The policy α_i^* (or equivalently the threshold for sensor i) in (18) only depends on parameters pertinent to system i, i.e., its transmit power p_i, and the elements ν_i^*, μ_i^*. Decentralized threshold-based policies have also been shown to be optimal for general wireless communication networks (Adireddy & Tong, 2005; Hu & Ribeiro, 2011; Qin & Berry, 2006). The context differs however, since in these works the objective is throughput-based utility functions in contrast to control system performance here.

Technically the vectors ν^*, μ^* correspond to the optimal Lagrange multipliers of an appropriately defined equivalent problem (cf. (A.5)–(A.8)). Intuitively each element ν_i^* can be interpreted as the importance of control performance of system i, and each element μ_i^* as the collision effect that sensor i has on all other systems $j \neq i$. The optimal access policy for sensor i in (18), or equivalently the optimal channel threshold, trades off the requirement on loop i and the collective negative effect on all other control loops. A larger value ν_i^* corresponds to a lower threshold so that sensor transmits more often, while a larger value μ_i^* corresponds to a higher threshold so that sensor backs off more often. Moreover, a high transmit power p_i in (18) also implies that sensor i should access the channel less often to limit expenditures.

The decoupling of the optimal sensor access policies in Theorem 1 relies on knowing the values ν^*, μ^*. In the following section we develop a distributed iterative procedure to obtain this knowledge.

Remark 2. In our previous work in Gatsis, Ribeiro, et al. (2015) we consider simpler random access policies not taking into account channel state information. In particular at every time k each sensor i randomly and independently transmits with some constant probability $\tilde{\alpha}_i \in [0, 1]$ to be designed. Similarly to (9) the probability of successfully closing each loop is given by $\mathbb{P}(\gamma_{i,k} = 1) = \tilde{\alpha}_i \prod_{j \neq i} \Big[1 - \tilde{\alpha}_j\Big]$. It turns out (Gatsis, Ribeiro, et al., 2015, Theorem 2) that the optimal access rates solving a problem equivalent to (16)–(17) can be expressed as

$$\tilde{\alpha}_i = \frac{\tilde{\nu}_i}{p_i + \tilde{\mu}_i} \quad \text{for each } i = 1, \ldots, m \tag{19}$$

for some vectors $\tilde{\nu} \in \mathbb{R}_+^m, \tilde{\mu} \in \mathbb{R}_+^m$, which have the same interpretation as the ones of Theorem 1 (but are not equal as they correspond to different problems). Hence for the non-channel-aware case the sensors need to randomize ($0 < \tilde{\alpha}_i < 1$). In contrast, conditioned on channel state information being available the optimal policies for the sensors are deterministic, exploiting favorable channel conditions to transmit.

5. Computation of channel-aware random access policies

In the previous section the optimal sensor access policies that guarantee control performance of all closed loop systems are characterized in terms of some appropriate vectors ν^*, μ^* (Theorem 1). In this section we capitalize on this result and develop an iterative procedure to determine the optimal sensor access policies by computing ν^*, μ^*. The procedure is easily implemented in the architecture of Fig. 1 and distributed in the sense that the common access point/controller is responsible for finding ν^*, μ^* and the sensors do not need to directly coordinate among themselves.

Technically as we argued in the proof of Theorem 1 the values ν^*, μ^* are the optimal Lagrange dual variables of an appropriately defined problem (cf. (A.5)–(A.8)). Algorithm 1 corresponds to a dual subgradient algorithm (Bertsekas, Nedić, & Ozdaglar, 2003, Ch. 8) to find these optimal dual variables. This procedure can be implemented in the wireless control architecture of Fig. 1 as follows. At each period t the access point maintains tentative vectors $\nu(t)$, $\mu(t)$ and some additional vectors $\lambda(t)$. At the beginning of each period, the access point (AP) sends to each sensor i the values $\nu_i(t)$, $\mu_i(t)$ via the reverse channel (Step 3). For the rest of the period t each sensor uses a random access policy $\alpha(h_i; t)$ as if the received values $\nu_i(t)$, $\mu_i(t)$ corresponded to the optimal ones (Step 4). Here $\alpha(h_i; t)$ denotes the valuation of the policy during period t at any channel state $h_i \in \mathbb{R}_+$. Then the AP measures the gap between desired and current control performance of each system during this period and updates the vectors $\nu(t)$, $\mu(t)$ to $\nu(t+1)$, $\mu(t+1)$ to prepare for the next period, along with an update of the additional variables $\lambda(t)$ (Step 7). To perform this update the AP needs to compute[1] the average transmission and packet success rates for each system during this period (Step 5) and also perform some auxiliary computations (Step 6).

Algorithm 1 Random access design algorithm

1: Initialize $\lambda(0) \in \mathbb{R}_+^m$, $\nu(0) \in \mathbb{R}_+^m$, $\mu(0) \in \mathbb{R}_+^m$ at the AP

2: **loop** At period $t = 0, 1, \ldots$

3: AP sends $\nu_i(t)$, $\mu_i(t)$ to each sensor i.

4: During the period each sensor i accesses the channel according to policy

$$\alpha_i(h_i; t) \leftarrow \begin{cases} 1 & \text{if } \nu_i(t)\, q(h_i) \geq p_i + \mu_i(t) \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

5: AP measures $\mathbb{E}_{h_i}[\alpha_i(h_i; t)\, q(h_i)]$, $\mathbb{E}_{h_i}[\alpha_i(h_i; t)]$ for all sensors $i = 1, \ldots, m$ during the period.

6: AP computes the auxiliary variables

$$\beta_i(t) \leftarrow \max\{\zeta, \frac{\lambda_i(t)}{\nu_i(t)}\}, \tag{21}$$

$$\delta_i(t) \leftarrow 1 - \max\{\zeta, \frac{\sum_{j \neq i} \lambda_j(t)}{\mu_i(t)}\} \tag{22}$$

for all $i, j \in \{1, \ldots, m\}$, where $\zeta > 0$ is a small positive constant introduced at the end of Section 3.

7: AP updates the new dual variables

$$\begin{aligned} \nu_i(t+1) &\leftarrow \Big[\nu_i(t) + \epsilon(t) \left(\beta_i(t) - \mathbb{E}_{h_i}[\alpha_i(h_i; t) q(h_i)] \right) \Big]_+ \\ \mu_i(t+1) &\leftarrow \Big[\mu_i(t) + \epsilon(t) \left(\mathbb{E}_{h_i}[\alpha_i(h_i; t)] - \delta_i(t) \right) \Big]_+ \\ \lambda_i(t+1) &\leftarrow \Big[\lambda_i(t) + \epsilon(t) \Big(\log(c_i) - \log(\beta_i(t)) \\ &\qquad - \sum_{j \neq i} \log(1 - \delta_j(t)) \Big) \Big]_+ \end{aligned} \tag{23}$$

for all $i \neq j \in \{1, \ldots, m\}$, where $[\ ]_+$ denotes the projection to the non-negatives $\mathbb{R}_+$.

8: **end loop**

This algorithm is guaranteed to converge to the optimal sensor access policies as we state next.

Theorem 2 (*Sensor Access Optimization*)**.** *Consider the setup of Theorem 1. The iterations of Algorithm 1 with stepsizes $\varepsilon(t) > 0$ in Step 7 satisfying $\sum_{t \geq 0} \varepsilon(t)^2 < \infty$, $\sum_{t \geq 0} \varepsilon(t) = \infty$ converge to the optimal sensor access policies, i.e.,*

$$c_i \leq \lim_{t \to \infty} \mathbb{E}_{h_i}[\alpha_i(h_i; t) q(h_i)] \prod_{j \neq i} \Big[1 - \mathbb{E}_{h_j} \alpha_j(h_j; t)\Big], \tag{24}$$

for all $i = 1, \ldots, m$, and

$$\lim_{t \to \infty} \sum_{i=1}^m \mathbb{E}_{h_i}[\alpha_i(h_i; t)] p_i = \sum_{i=1}^m \mathbb{E}_{h_i}[\alpha_i^*(h_i)] p_i. \tag{25}$$

[1] Here we assume that even when collisions arise the AP can identify which sensor transmits at each time slot. Hence it can measure the average rate $\mathbb{E}_{h_i}[\alpha_i(h_i; t)]$ at which each sensor i accesses the channel, as well as the packet success ratio $\mathbb{E}_{h_i}[\alpha_i(h_i; t)\, q(h_i)]$ when only sensor i transmits.

One caveat of this distributed implementation is that it requires information sent from the access point to the sensors, hence it introduces some communication overhead. This overhead however burdens mainly the access point which typically has more capabilities compared to the simple wireless sensors. In the following section we present numerical simulations of this random access design.

6. Numerical simulations

We first present a numerical example of the random access design for two ($m = 2$) asymmetric control systems. We assume the first system is two-dimensional, with one input and one sensor output described as in (2)–(3) by the state transition, input, and output matrices respectively

$$A_{o,1} = \begin{bmatrix} -1.0 & -0.4 \\ -0.5 & 0.3 \end{bmatrix},\ B_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix},\ C_1 = \begin{bmatrix} 2, & 1 \end{bmatrix}. \tag{26}$$

We consider a simple output-based controller which whenever receives the sensor measurement $y_{1,k} = C_1 x_{1,k} + v_{1,k}$ over the wireless channel applies a feedback of the form $u_{1,k} = K_1 y_{1,k}$, while applies zero input $u_{1,k} = 0$ when no sensor measurement is received. This system is a special case of the one considered in Example 1 with open and closed loop matrices $A_{o,1}$ and $A_{c,1} = A_{o,1} + B_1 K_1 C_1$ respectively. The open loop has a largest unstable eigenvalue of -1.14 and by picking $K = 0.3$ the largest closed loop eigenvalue is 0.59.

To obtain a communication requirement for the described system we require a Lyapunov function $V_1(x_1) = x_1^T P_1 x_1$. We pick this function such that the perfect closed loop system decreases at a rate 0.7, i.e., the matrix P_1 solves the linear matrix equality $A_{c,1}^T P_1 A_{c,1} = 0.7 P_1 + I$. To account for the possibility of packet losses and collisions, we require that this Lyapunov function decreases over the shared wireless channel at a rate $\rho_1 = 0.95$ (cf. (10)), which is slower that the perfect closed loop rate 0.7 but still guarantees stability.

The second system is asymmetric from the first one, in particular it has scalar integrator open loop dynamics $A_{o,i} = 1$ and stable closed loop dynamics $A_{c,2} = 0.2$. The Lyapunov function $V_2(x_2) = x_2^2$ ($P_2 = 1$) is required to decrease with expected rate $\rho_2 = 0.9$ (cf. (10)).

Both systems are perturbed by zero-mean unit-variance Gaussian noises. Both channel states $h_{1,k}, h_{2,k}$ are i.i.d. exponential with mean 1. In isolation each sensor faces a packet success probability modeled by the function $q(h_{i,k})$, $i = 1, 2$ shown in Fig. 2. The transmit powers are taken equal $p_1 = p_2 = 1$. By Proposition 1 the control performance requirements for the two systems are equivalent to required packet success rates $c_1 \approx 0.42$, $c_2 \approx 0.10$ for the two sensors, computed by (12). Hence System 1, which is more unstable, requires a higher packet success rate.

We solve the random access design problem (16)–(17) using Algorithm 1 . At each iteration expectations with respect to the channel state distributions need to be used at Steps 5 and 7. In the simulation we approximate these expectations with averages from a large number of samples readily drawn from the exponential channel distributions. The iterates of the (dual) variables $\lambda(t)$, $\mu(t)$, $\nu(t)$ converge to their optimal values, as also shown in the proof of Theorem 2. We plot the evolution of the sensor access policies $\alpha_i(h_i; t)$, or equivalently the thresholds of these policies during the iterations of the algorithm in Fig. 3. As also established in Theorem 2 the channel thresholds converge to their optimal values in the limit. We observe that the sensor 1 has a lower threshold, meaning that it transmits more often, which is reasonable since it corresponds to the more demanding control system.

Next we simulate the random access architecture with the obtained channel access policies. In Fig. 4 we plot the empirical

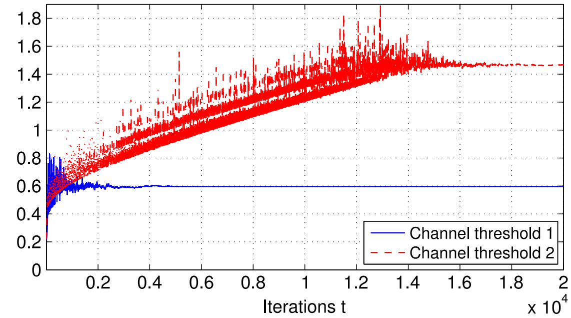

Fig. 3. Channel access thresholds during the optimization algorithm. The thresholds for both sensors converge to their optimal values. Sensor 1 has a lower threshold, i.e., transmits more often, since it supports a more demanding control performance.

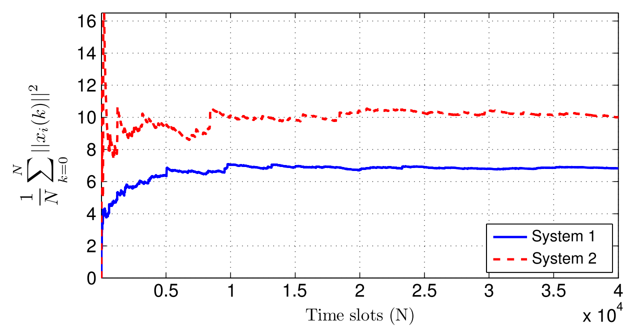

Fig. 4. Evolution of control systems using the optimal random access policies. Both systems remain stable despite collisions and packet drops.

average quadratic cost of the systems and verify that both systems remain stable despite packet collisions over the shared channel. The empirical transmit rates $1/N \sum_{k=1}^{N} \alpha_{i,k}$ at which each sensor transmits equal 0.55 and 0.23 for $i = 1, 2$ respectively capturing again the asymmetry in the design of the shared wireless control architecture.

Finally we discuss the scalability of the wireless random access architecture for multiple control systems. For simplicity we consider multiple identical scalar systems with integrator open loop dynamics $A_{o,i} = 1$ and reset-to-zero closed loop dynamics $A_{c,i} = 0$. Each system i's Lyapunov function $V(x_i) = x_i^2$ is required to decrease at a common rate ρ. For different rate values we search for the maximum number of systems that can be supported over the shared random access channel. In particular for each value ρ we search for the maximum number m such that the Algorithm 1 converges, implying that the design is feasible. This search is heuristic as we do not verify that the algorithm diverges to infinity but instead observe the (dual) iterates $\lambda(t)$, $\mu(t)$, $\nu(t)$ grow very large. The results are summarized in Table 1. With strict control performance (fast decrease rate $\rho \ll 1$) only few systems can be supported, while under more lenient control performance the architecture can admit many more systems. By symmetry all systems have the same optimal channel access threshold. As shown in Table 1 the threshold increases as the number of systems grows, i.e., systems access the channel at a lower rate $\mathbb{E}_{h_i}\alpha_i(h_i)$ but more opportunistically.

7. Concluding remarks

We design a random access mechanism for sensors transmitting measurements of multiple plants over a shared wireless channel to a controller. The goal of the sensors is to guarantee control performance for all control systems by mitigating the effect of packet collisions from simultaneous transmissions as well as by

Table 1
Random access scalability of wireless control systems.

ρ	.75	.8	.85	.9	.92	.94	.96	.98
Max # systems	1	2	2	3	4	6	9	18
Channel threshold	1.37	1.02	1.69	1.40	1.54	1.86	2.23	2.90
Access rate	0.25	0.39	0.19	0.35	0.27	0.17	0.12	0.06

adapting online to randomly varying channel conditions. Via a Lyapunov function abstraction, control performance is transformed to required packet success rates of each closed loop. We show that the optimal random access policies can be decoupled between the sensors and are of a threshold form with respect to channel states. Moreover we develop a distributed procedure to compute the optimal policies. In future research we will explore decentralized online sensor adaptation to plant measurements (Gatsis, Ribeiro, & Pappas, 2016), as well as adaptation to changes in the environment, for example, the admission of new control loops in the architecture.

Appendix A. Proof of Theorem 1

First we convert problem (16)–(17) into an equivalent auxiliary optimization problem in (A.5)–(A.8) which has zero duality gap. We begin by removing the product of the expectations appearing in the constraints (17). Taking the logarithm at each side of (17) preserves the feasible set of variables by monotonicity. Then the logarithm of the product at the right hand side of (17) becomes a sum of logarithms, and we can rewrite the constraints equivalently as

$$\log(c_i) \leq \log(\mathbb{E}_{h_i}[\alpha_i(h_i)\, q(h_i)]) + \sum_{j\neq i} \log(1 - \mathbb{E}_{h_j}[\alpha_j(h_j)]),\ i = 1, \ldots, m. \tag{A.1}$$

The logarithms in (A.1) are well-defined and finite by the implicit assumption that the terms $\mathbb{E}_{h_i}[\alpha_i(h_i)\, q(h_i)]$ and $\mathbb{E}_{h_j}[\alpha_j(h_j)]$ are bounded away from 0 and 1 by some small positive quantity $\zeta > 0$ respectively — see end of Section 3. Next, we replace the term $\mathbb{E}_{h_i}[\alpha_i(h_i)\, q(h_i)]$ in constraint (A.1) by an auxiliary variable β_i for $i = 1, \ldots, m$, and the terms $\mathbb{E}_{h_j}[\alpha_j(h_j)]$ in (A.1) by variables δ_j for $j \neq i$. Hence we rewrite (A.1) as

$$\log(c_i) \leq \log(\beta_i) + \sum_{j\neq i} \log(1 - \delta_j). \tag{A.2}$$

To force the auxiliary variables to behave like the expectations we introduce additional constraints of the form

$$\beta_i \leq \mathbb{E}_{h_i}[\alpha_i(h_i)\, q(h_i)] \tag{A.3}$$

$$\delta_j \geq \mathbb{E}_{h_j}[\alpha_j(h_j)] \tag{A.4}$$

for all $i, j \in \{1, \ldots, m\}$. By the aforementioned implicit assumption these vector variables β, δ take values in the sets $\mathcal{B} := [\zeta, +\infty)^m$, $\mathcal{D} := (-\infty, 1-\zeta]^m$ respectively. Overall we formulate the auxiliary optimization problem

$$\underset{\alpha,\beta,\delta\in\mathcal{A}\times\mathcal{B}\times\mathcal{D}}{\text{minimize}} \quad \sum_{i=1}^m \mathbb{E}_{h_i}\alpha_i(h_i)p_i \tag{A.5}$$

$$\text{subject to} \quad \log(c_i) \leq \log(\beta_i) + \sum_{j\neq i} \log(1 - \delta_j), \tag{A.6}$$

$$\beta_i \leq \mathbb{E}_{h_i}[\alpha_i(h_i)\, q(h_i)], \tag{A.7}$$

$$\delta_i \geq \mathbb{E}_{h_i}[\alpha_i(h_i)], \quad i = 1, \ldots, m. \tag{A.8}$$

We argue that this auxiliary problem is equivalent to the original one in (16)–(17), i.e., a feasible solution of one problem corresponds to a feasible solution with the same objective value for the other problem. Indeed let us start with a feasible solution α for (16)–(17). Let us define variables β, δ that make (A.7), (A.8) hold with equality. Then α, β, δ is feasible for problem (A.5)–(A.8) and has the same objective. Reversely, consider a feasible α, β, δ for problem (A.5)–(A.8). Without loss we can assume that all constraints (A.7)–(A.8) hold with equality. Otherwise if, say, an inequality i in (A.7) is strict, we can increase the value of variable β_i till equality in (A.7) is reached without loss of feasibility in (A.6) and without changing the objective value in (A.5). A similar procedure can be performed if some inequality (A.8) is strict, leading to a new feasible point satisfying (A.7)–(A.8) with equalities. Then it is immediate that α is also feasible for (17) and has the same objective.

Based on the established equivalence, in the rest of the proof it suffices to show that (18) describes an optimal solution for the auxiliary problem (A.5)–(A.8). The advantage of formulating this auxiliary problem is that it has zero duality gap (Ribeiro, 2012). To formally state this result, let us denote the optimal value of this problem by P^* (finite by feasibility Assumption 2) and let us define the Lagrange dual problem. We associate dual variables $\lambda_i \geq 0$ with inequalities (A.6), $\nu_i \geq 0$ with (A.7), and $\mu_i \geq 0$ with (A.8), for $i \in \{1, \ldots, m\}$. We write the Lagrangian function as

$$\begin{aligned} L(\alpha, \beta, \delta, \lambda, \nu, \mu) &= \sum_{i=1}^m \mathbb{E}_{h_i}\alpha_i(h_i)p_i \\ &+ \sum_{i=1}^m \lambda_i \Big\{\log(c_i) - \log(\beta_i) - \sum_{j\neq i}\log(1-\delta_j)\Big\} \\ &+ \sum_{i=1}^m \nu_i(\beta_i - \mathbb{E}_{h_i}[\alpha_i(h_i)\, q(h_i)]) \\ &+ \sum_{j=1}^m \mu_j(\mathbb{E}_{h_j}[\alpha_j(h_j)] - \delta_j). \end{aligned} \tag{A.9}$$

We can reorder the Lagrangian terms to the form

$$\begin{aligned} L(\alpha, \beta, \delta, \lambda, \nu, \mu) = \sum_{i=1}^m \Big\{ &\mathbb{E}_{h_i}\, \alpha_i(h_i)\Big[p_i + \mu_i - \nu_i q(h_i)\Big] \\ &- \big(\sum_{j\neq i}\lambda_j\big)\log(1-\delta_i) - \mu_i\delta_i + \nu_i\beta_i - \lambda_i\log(\beta_i) \\ &+ \lambda_i \log(c_i)\Big\}. \end{aligned} \tag{A.10}$$

This form is useful because each primal variable $\alpha_i(h_i)$, β_i, δ_i for each i is decoupled from the others, a fact we will exploit next. Then we can define the Lagrange dual function

$$g(\lambda, \nu, \mu) = \inf_{\alpha,\beta,\delta\in\mathcal{A}\times\mathcal{B}\times\mathcal{D}} L(\alpha, \beta, \delta, \lambda, \nu, \mu), \tag{A.11}$$

as well as the Lagrange dual problem whose optimal value we denote by D^* as

$$D^* = \sup_{\lambda,\nu,\mu\in\mathbb{R}_+^m\times\mathbb{R}_+^m\times\mathbb{R}_+^m} g(\lambda, \nu, \mu). \tag{A.12}$$

Then we can establish the following zero duality property about the auxiliary problem (A.5)–(A.8).

Proposition 2 (*Strong Duality*)**.** *Let Assumptions 1 and 2 hold. Then the problem (A.5)–(A.8) has zero duality gap, i.e.,* $P^* = D^*$*. Moreover if* $\alpha^*, \beta^*, \delta^*$ *are optimal solutions, and* λ^*, ν^*, μ^* *are optimal solutions for the dual problem (A.12), then*

$$\alpha^*, \beta^*, \delta^* \in \operatorname*{argmin}_{\alpha,\beta,\delta \in \mathcal{A}\times\mathcal{B}\times\mathcal{D}} L(\alpha, \beta, \delta, \lambda^*, \nu^*, \mu^*). \tag{A.13}$$

The result follows from Ribeiro (2012, Theorems 1 and 4) where general stochastic optimization problems of the form (A.5)–(A.8) are examined under non-atomic probability measures (Assumption 1) and strict feasibility (Assumption 2). The proof is omitted due to space limitations. The above proposition suggests that we can recover the optimal variables α^* of our problem by just minimizing the unconstrained Lagrangian function in (A.13). A technical caveat of (A.13) is that it describes an inclusion only, implying that in general there might be Lagrangian minimizers that are not optimal. The following lemma excludes such cases by establishing that the functions α that minimize the Lagrangian are unique up to a set of measure zero and given by an explicit expression. The proof can be found in Appendix B.

Lemma 1. *Let Assumption 1 hold. Consider any dual variables* $\lambda, \nu, \mu \in \mathbb{R}^m_+ \times \mathbb{R}^m_+ \times \mathbb{R}^m_+$*. Then the functions* $\alpha \in \mathcal{A}$ *that minimize the Lagrangian* $L(\alpha, \beta, \delta, \lambda, \nu, \mu)$ *are uniquely defined except for a set of arguments* $h \in \mathbb{R}^m_+$ *of* ϕ*-measure zero, and are given by*

$$\alpha_i(h_i) = \begin{cases} 1 & \text{if} \quad \nu_i q(h_i) \geq p_i + \mu_i \\ 0 & \text{otherwise.} \end{cases} \tag{A.14}$$

for each $i = 1, \ldots, m$ *and for every value* $h_i \in \mathbb{R}_+$.

To sum up we have shown in (A.13) that the optimal solution $\alpha(.)$ to problem (A.5)–(A.8) belongs in the set of Lagrange minimizers at λ^*, ν^*, μ^*, and by Lemma 1 these minimizers are unique up to a set of measure zero. As a result, all these minimizers will have the same objective and constraint slack in problem (A.5)–(A.8), and they will all be optimal for this problem. In particular, the specific minimizer defined by $\alpha_i(h_i)$ given in (A.14) at the values λ^*, ν^*, μ^* will be optimal for the problem, and corresponds exactly to the one given in (18) at the statement of the theorem. ■

Appendix B. Proof of Lemma 1

Consider the problem of minimizing the Lagrangian $L(\alpha, \beta, \delta, \lambda, \nu, \mu)$ over variables $\alpha \in \mathcal{A}$. Due to the separability of the Lagrangian given in the form (A.10) over variables α, β, δ, we can separate the problem into subproblems

$$\operatorname*{argmin}_{\alpha_i \in \mathcal{A}_i} \mathbb{E}_{h_i} \alpha_i(h_i) \Big[p_i + \mu_i - \nu_i q(h_i) \Big] \tag{B.1}$$

for $i \in \{1, \ldots, m\}$. Next we need to verify that (A.14) is optimal for (B.1). Note that without loss of generality we can exchange the expectation operator $\mathbb{E}_{h_i}$ and the minimization over functions $\alpha_i : \mathbb{R}_+ \to [0, 1]$ to equivalently solve

$$\operatorname*{argmin}_{\alpha_i(h_i) \in [0,1]} \alpha_i(h_i) \Big[p_i + \mu_i - \nu_i q(h_i) \Big] \tag{B.2}$$

pointwise at all values $h_i \in \mathbb{R}_+$. This is valid because any function α_i that minimizes (B.1) can differ from the minimizer in (B.2) at a set of values $h_i \in \mathbb{R}_+$ with measure at most zero. Then we can verify that (A.14) is the minimizer in (B.2). That is because the right hand side in (B.2) is a linear expression of $\alpha_i(h_i) \in [0, 1]$. Hence the minimizer $\alpha_i(h_i)$ is uniquely defined, and takes values either 0 or 1 except for the values h_i where $p_i + \mu_i - \nu_i q(h_i) = 0$. In the latter case the minimizer is not uniquely defined. However due to the strict monotonicity assumption for $q(h_i)$ this case occurs for at most one value h_i, hence it is a measure zero event since measure ϕ_i is non-atomic by Assumption 1. This completes the proof.

Appendix C. Proof of Theorem 2

A sufficient condition for (24) and (25) is that the policies $\alpha_i(h_i; t)$ at Step 4 of the algorithm converge to the optimal solution α^* of problem (16)–(17) in the sense that

$$\lim_{t\to\infty} \mathbb{E}_{h_i}[\alpha_i(h_i; t)] = \mathbb{E}_{h_i}[\alpha_i^*(h_i)] \tag{C.1}$$

$$\lim_{t\to\infty} \mathbb{E}_{h_i}[\alpha_i(h_i; t) q(h_i)] = \mathbb{E}_{h_i}[\alpha_i^*(h_i) q(h_i)] \tag{C.2}$$

hold for all $i = 1, \ldots, m$. In the proof of Theorem 1 we argued that problem (16)–(17) is equivalent to the auxiliary problem (A.5)–(A.8). Hence it suffices to show that the algorithm converges to the optimal solution of this auxiliary problem in the sense of (C.1)–(C.2).

Recall the Lagrange dual function $g(\lambda, \nu, \mu)$ of the auxiliary problem (A.5)–(A.8), defined in (A.11). We will argue that the variables $\lambda(t), \nu(t), \mu(t)$ selected by the algorithm converge to the optimal dual variables λ^*, ν^*, μ^* that maximize this dual function. The reason is that at each iteration t of the algorithm, the dual variables $\lambda(t), \nu(t), \mu(t)$ according to Step 7 move towards a subgradient direction of the dual function. To show this note that the variable $\alpha(t)$ selected by the algorithm at step (20) minimizes the Lagrangian $L(\alpha, \beta, \delta, \lambda(t), \nu(t), \mu(t))$ with respect to the variable $\alpha \in \mathcal{A}$. This follows directly from Lemma 1. Similarly the variables $\beta(t), \delta(t)$ at steps (22) minimize the Lagrangian function $L(\alpha, \beta, \delta, \lambda(t), \nu(t), \mu(t))$ with respect to the variables $\beta, \delta \in \mathcal{B} \times \mathcal{D}$. This fact follows by examining the first order conditions for the Lagrangian given in the form (A.10), subject to the constraints $\beta, \delta \in \mathcal{B} \times \mathcal{D}$. Then by standard arguments (Bertsekas et al., 2003, p. 457), the constraint slack of the minimizers of the Lagrangian is a subgradient of the dual function. To sum up, at each iteration of the algorithm the variables $\lambda(t), \nu(t), \mu(t)$ move towards a subgradient direction of the dual function by (23) . Additionally the subgradients are bounded — in particular the logarithms at the right hand side of (23) are finite by the restriction $\beta(t), \delta(t) \in \mathcal{B} \times \mathcal{D}$. Under the bounded subgradient condition, convergence of $\lambda(t), \nu(t), \mu(t)$ to the optimal dual variables λ^*, ν^*, μ^* for stepsizes as in the statement of the theorem is a standard subgradient method result (Bertsekas et al., 2003, Prop. 8.2.6).

In the rest of the proof, based on the established convergence of the dual variables to the optimal ones, we will show that the same holds for the primal variable $\alpha(.; t)$ in the sense of (C.1)–(C.2). Note that at any iteration t the function $\alpha_i(.; t)$ takes the value 1 when $\nu_i(t) q(h_i) \geq p_i + \mu_i(t)$ and 0 otherwise. Due to the strict monotonicity of the function $q(.)$ this is a threshold-like function taking the value 1 when $h_i \geq \bar{h}_i(t) = q^{-1}((p_i + \mu_i(t))/\nu_i(t))$. Since we have established that $\nu(t) \to \nu^*, \mu(t) \to \mu^*$, and since the function $q(.)$ is continuous hence its inverse too, we conclude that the threshold $\bar{h}_i(t)$ converges to $\bar{h}_i^* = q^{-1}((p_i + \mu_i^*)/\nu_i^*)$. By Theorem 1 this limit value equals the threshold of the optimal access policy.

Hence we conclude that $\alpha_i(.; t) \to \alpha_i^*(.)$ pointwise for all $h_i \in \mathbb{R}_+$ except perhaps for the point $\bar{h}_i^*$, i.e., the optimal threshold point. Since the probability measure ϕ_i is non-atomic the point $\bar{h}_i^*$ has a probability measure zero. Hence $\alpha(.; t) \to \alpha^*(.)$ almost everywhere. Also both sequences of functions $\alpha_i(.; t)$ and $\alpha_i(.; t)q(.)$ are uniformly bounded in [0, 1]. By the bounded convergence theorem (Billingsley, 1995, Theorem 16.5) we conclude that convergence in expectation, i.e., (C.1) and (C.2), also holds. This completes the proof. ■

References

Adireddy, S., & Tong, L. (2005). Exploiting decentralized channel state information for random access. *IEEE Transactions on Information Theory*, *51*(2), 537–561.

Bertsekas, D. P., Nedić, A., & Ozdaglar, A. E. (2003). *Convex analysis and optimization*. Athena Scientific.

Billingsley, P. (1995). *Probability and measure*. John Wiley & Sons.

Blind, R., & Allgöwer, F. (2011). Analysis of networked event-based control with a shared communication medium: part I-pure ALOHA. In *IFAC world congress*.

Boyd, S., & Vandenberghe, L. (2009). *Convex optimization*. Cambridge University Press.

Donkers, M., Heemels, W., Van De Wouw, N., & Hetel, L. (2011). Stability analysis of networked control systems using a switched linear systems approach. *IEEE Transactions on Automatic Control*, *56*(9), 2101–2115.

Gatsis, K., Pajic, M., Ribeiro, A., & Pappas, G. J. (2015). Opportunistic control over shared wireless channels. *IEEE Transactions on Automatic Control*, *60*(12), 3140–3155.

Gatsis, K., Ribeiro, A., & Pappas, G. J. (2014). Optimal power management in wireless control systems. *IEEE Transactions on Automatic Control*, *59*(6), 1495–1510.

Gatsis, K., Ribeiro, A., & Pappas, G. J. (2015). Control with random access wireless sensors. In *54th IEEE conference on decision and control* (pp. 318–323).

Gatsis, K., Ribeiro, A., & Pappas, G. J. (2016). Control-aware random access communication. In *ACM/IEEE 7th int. conference on cyber-physical systems*.

Goldsmith, A. (2005). *Wireless communications*. Cambr. Univ. Press.

Gupta, V., Chung, T. H., Hassibi, B., & Murray, R. M. (2006). On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage. *Automatica*, *42*(2), 251–260.

Hespanha, J., Naghshtabrizi, P., & Xu, Y. (2007). A survey of recent results in networked control systems. *Proceedings of the IEEE*, *95*(1), 138–162.

Hristu-Varsakelis, D. (2001). Feedback control systems as users of a shared network: Communication sequences that guarantee stability. In *Proc. of the 40th IEEE conf. on decision and control, 2001*, Vol. 4, (pp. 3631–3636).

Hu, Y., & Ribeiro, A. (2011). Adaptive distributed algorithms for optimal random access channels. *IEEE Transactions on Wireless Communications*, *10*(8), 2703–2715.

Le Ny, J., Feron, E., & Pappas, G. J. (2011) (2011). Resource constrained lqr control under fast sampling. In *14th int. conf. on hybrid systems: computation and control* (pp. 271–280).

Liu, X., & Goldsmith, A. (2004). Wireless medium access control in networked control systems. In *Proc. of the 2004 American control conference*, Vol. 4, (pp. 3605–3610).

Luo, W., & Ephremides, A. (2002). Power levels and packet lengths in random multiple access. *IEEE Transactions on Information Theory*, *48*(1), 46–58.

Qin, X., & Berry, R. (2006). Distributed approaches for exploiting multiuser diversity in wireless networks. *IEEE Transactions on Information Theory*, *52*(2), 392–413.

Quevedo, D. E., Ahlén, A., Leong, A. S., & Dey, S. (2012). On Kalman filtering over fading wireless channels with controlled transmission powers. *Automatica*, *48*(7), 1306–1316.

Rabi, M., Stabellini, L., Proutiere, A., & Johansson, M. (2010). Networked estimation under contention-based medium access. *International Journal of Robust and Nonlinear Control*, *20*(2), 140–155.

Ramesh, C., Sandberg, H., & Johansson, K. H. (2013). Design of state-based schedulers for a network of control loops. *IEEE Transactions on Automatic Control*, *58*(8), 1962–1975.

Rehbinder, H., & Sanfridson, M. (2004). Scheduling of a limited communication channel for optimal control. *Automatica*, *40*(3), 491–500.

Ribeiro, A. (2012). Optimal resource allocation in wireless communication and networking. *EURASIP Journal on Wireless Communications and Networking*, *2012*(1), 1–19.

Tabbara, M., & Nesic, D. (2008). Input–output stability of networked control systems with stochastic protocols and channels. *IEEE Transactions on Automatic Control*, *53*(5), 1160–1175.

Van de Beek, J.-J., Borjesson, P. O., Boucheret, M.-L., Landstrom, D., Arenas, J. M., Odling, P., et al. (1999). A time and frequency synchronization scheme for multiuser OFDM. *IEEE Journal on Selected Areas in Communications*, *17*(11), 1900–1914.

Vasconcelos, M. M., & Martins, N. C. (2017). Optimal estimation over the collision channel. *IEEE Transactions on Automatic Control*, *62*(1), 321–336.

Walsh, G. C., Ye, H., & Bushnell, L. G. (2002). Stability analysis of networked control systems. *IEEE Transactions on Control Systems Technology*, *10*(3), 438–446.

Zhang, W. (2003). Stabilization of networked control systems over a sharing link using ALOHA. In *Proc. of the 42nd IEEE conference on decision and control*, Vol. 1, (pp. 204–209).

Zhang, W., Branicky, M., & Phillips, S. (2001). Stability of networked control systems. *IEEE Control Systems Magazine*, *21*(1), 84–99.

Konstantinos Gatsis received the Diploma degree in electrical and computer engineering from the University of Patras, Patras, Greece in 2010, and the Ph.D. degree in electrical and systems engineering from the University of Pennsylvania, Philadelphia, PA, USA, in 2016. Currently, he is a Postdoctoral Researcher in the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia. His research interests include cyber–physical systems, networked control systems, as well as security and resource optimization problems arising in them. Dr. Gatsis received the Joseph, D'16, and Rosaline Wolf Award for Best Doctoral Dissertation from the department of Electrical and Systems Engineering at the University of Pennsylvania. He also received the 2014 O. Hugo Schuck Best Paper Award, the Student Best Paper Award at the 2013 American Control Conference, and was a Best Paper Award Finalist at the 2014 ACM/IEEE International Conference on Cyber–Physical Systems.

Alejandro Ribeiro received the B.Sc. degree in electrical engineering from the Universidad de la Republica Oriental del Uruguay, Montevideo, in 1998 and the M.Sc. and Ph.D. degree in electrical engineering from the Department of Electrical and Computer Engineering, the University of Minnesota, Minneapolis in 2005 and 2007. From 1998 to 2003, he was a member of the technical staff at Bellsouth Montevideo. After his M.Sc. and Ph.D. studies, in 2008 he joined the University of Pennsylvania (Penn), Philadelphia, where he is currently the Rosenbluth Associate Professor at the Department of Electrical and Systems Engineering. His research interests are in the applications of statistical signal processing to the study of networks and networked phenomena. His focus is on structured representations of networked data structures, graph signal processing, network optimization, robot teams, and networked control. Dr. Ribeiro received the 2014 O. Hugo Schuck best paper award, and paper awards at the 2016 SSP Workshop, 2016 SAM Workshop, 2015 Asilomar SSC Conference, ACC 2013, ICASSP 2006, and ICASSP 2005. His teaching has been recognized with the 2017 Lindback award for distinguished teaching and the 2012 S. Reid Warren, Jr. Award presented by Penn's undergraduate student body for outstanding teaching. Dr. Ribeiro is a Fulbright scholar class of 2003 and a Penn Fellow class of 2015.

George J. Pappas is the Joseph Moore Professor and Chair of the Department of Electrical and Systems Engineering at the University of Pennsylvania. He also holds a secondary appointment in the Departments of Computer and Information Sciences, and Mechanical Engineering and Applied Mechanics. He is member of the GRASP Lab and the PRECISE Center. He has previously served as the Deputy Dean for Research in the School of Engineering and Applied Science. His research focuses on control theory and in particular, hybrid systems, embedded systems, hierarchical and distributed control systems, with applications to unmanned aerial vehicles, distributed robotics, green buildings, and biomolecular networks. He is a Fellow of IEEE, and has received various awards such as the Antonio Ruberti Young Researcher Prize, the George S. Axelby Award, the O. Hugo Schuck Best Paper Award, the National Science Foundation PECASE, and the George H. Heilmeier Faculty Excellence Award.