Available online at www.sciencedirect.com

ScienceDirect

IFAC-PapersOnLine 48-22 (2015) 209–214

Decentralized Channel Access 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. (e-mail: {kgatsis, aribeiro, pappasg}@seas.upenn.edu).

Abstract: We consider a decentralized channel access mechanism for multiple sensors communicating over a shared wireless medium with corresponding actuators. Each sensor independently and iteratively adjusts the rate at which it accesses the channel, in response to packet collisions arising from simultaneous transmissions. We provide theoretical conditions under which the decentralized mechanism converges to an operating point where desired closed loop performance of all control loops is met. Control performance is abstracted as desired decrease rates of given Lyapunov functions for each loop, which translates to necessary packet success rates on each link. Numerical simulations of the proposed mechanism are presented.

Keywords: Wireless sensor-actuator systems. Decentralized sensor coordination. Random access protocols.

1. INTRODUCTION

The emergence of numerous wireless sensing devices in smart homes or modern industrial environments poses new design challenges in the interface of control and communication. Sharing efficiently the available wireless medium between such devices in a way that guarantees closed loop control performance and is easily implementable arises as an important problem.

The design of networked control systems (NCSs) sharing a wireless or wired communication medium is typically based on centralized approaches, i.e., scheduling – see Hespanha et al. (2007); Schenato et al. (2007) for a general introduction to NCSs. In static scheduling, sensors transmit in a periodic sequence that is predesigned in a centralized manner to meet control objectives – see, e.g., Zhang et al. (2001); Hristu-Varsakelis (2001); Le Ny et al. (2011); Alur et al. (2011). Deriving optimal such scheduling sequences is recognized as a hard combinatorial problem (Rehbinder and Sanfridson (2004); Gupta et al. (2006)). In dynamic scheduling, a central authority decides at each time step which device accesses the medium. This dynamic decision may be stochastic (Gupta et al. (2006)), or based on plant state information (Walsh et al. (2002); Donkers et al. (2011); Ramesh et al. (2013)), or on the wireless channel conditions (Gatsis et al. (2015a)).

Unlike centralized approaches, decentralized channel access mechanisms are attractive as they do not require coordination among the devices. For example in a random access mechanism, which is the subject of this paper, each sensor randomly and independently decides whether to transmit to its controller over the shared channel. The drawback is that collisions can occur from simultaneously transmitting sensors, resulting in lost packets. Hence control performance of the systems sharing the channel becomes convoluted. The more a control loop uses the wireless channel, the more the other control loops are affected and deteriorate. Hence we are interested in designing a mechanism for control loops to decide in a decentralized fashion the rate at which they should access the channel, so that desired performance of all control loops is guaranteed.

⋆ 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.

To the best of our knowledge, control under random access communication mechanisms has drawn limited attention. The focus had been on analyzing the impact of packet collisions for networked control systems and on comparing different medium access mechanisms, either numerically (Liu and Goldsmith (2004); Ramesh et al. (2013)) or analytically in simple cases (Rabi et al. (2010); Blind and Allgöwer (2011)). 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. Conditions for stability under packet collisions have also been examined via a hybrid system approach in Tabbara and Nesic (2008).

In contrast to the aforementioned work focused on analysis, our goal is to design the medium access mechanism, i.e., the rate at which sensors access the shared wireless medium, with control performance guarantees. Related work by Zhang (2003) considers instead the Aloha-like scheme and characterizes what retransmission policies lead to stability. For our random access mechanism we have recently shown that the optimal (minimum power) access rates can be decoupled among sensors (Gatsis et al. (2015b)). Particularly each sensor accesses the channel at a rate proportional to the desired control performance of its loop, and inverse proportional to its aggregate collision effect on all other control loops. Moreover, relying on a common access point, we have shown how the optimal access rates can be computed efficiently. Similar decoupled characterizations are known in the context of random access wireless networks (Lee et al. (2007); Hu and Ribeiro (2011)), where instead of control performance the relevant quantity of interest are data rates on links or general utility objectives. Besides closed loop control, optimal estimation from remote sensors over collision channels is considered recently by Vasconcelos and Martins (2014).

In this paper we consider a random access architecture with multiple control loops (Fig. 1). We develop an iterative decentralized mechanism by which the sensors can adapt their channel access rates without the need of centralized coordination, as is the case, e.g., with the common access point in Gatsis et al. (2015b). We employ a Lyapunov-function control performance abstraction, motivated by our previous work (Gatsis et al. (2015a,b)). Each control system is abstracted via

10.1016/j.ifacol.2015.10.332

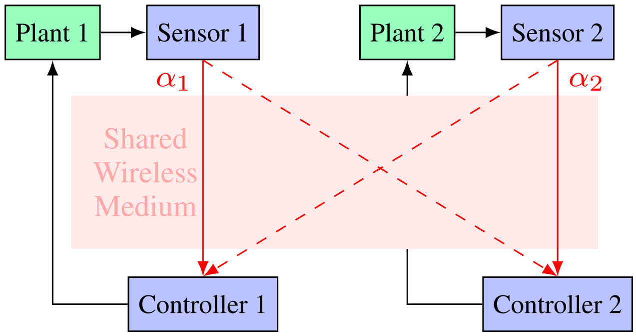

Fig. 1. Random access architecture over a shared wireless medium for $m = 2$ control loops. Each i sensor randomly decides with probability α_i whether to transmit to a corresponding controller computing the plant inputs. Packet collisions occur when both sensors transmit at the same slot. The goal of a decentralized channel access mechanism is to guarantee performance for all control loops.

a given Lyapunov function which is required to decrease at desired rates, stochastically due to the random packet losses and collisions on the shared medium (Section 2). These control requirements are shown to be equivalent to a minimum packet success rate on each link.

In Section 3 we develop our decentralized mechanism whereby each sensor iteratively responds to the discrepancy between desired and observed packet success rate locally on its link by appropriately adjusting its channel access rate. This mechanism is motivated by game-theoretic formulations of random access protocols (Jin and Kesidis (2002); Chen et al. (2010)). In Section 3 we provide technical conditions under which this mechanism converges locally to an operating point where control performance of all control loops is met. Moreover, for the special case of two control loops these conditions are shown to hold always, except for degenerate cases. In Section 4 we present numerical simulations of the decentralized mechanism, illustrating how it can also adapt to other changes in the environment, such as the introduction of new control systems over the shared channel. We conclude with some remarks in Section 5.

2. PROBLEM FORMULATION

We consider a wireless control architecture where m independent plants are controlled over a shared wireless medium. Each sensor i ($i = 1, 2, ..., m$) measures the output of plant i and transmits it to a corresponding controller i computing the plant control inputs. Packet collisions might arise on the shared medium between simultaneously transmitting sensors. The case for $m = 2$ control loops is shown in Fig. 1. We are interested in designing a decentralized mechanism for each sensor to decide whether to access the medium (random access) in a way that desirable control performance can be guaranteed for all control systems.

Communication takes place in time slots. At every time k each sensor i randomly and independently decides to access the channel with some constant probability $\alpha_i \in [0, 1]$, which is our design variable. If only sensor i transmits at a time slot, the message is not always successfully decoded at the access point/controller because of noise added to the transmitted signal on the wireless channel – see Gatsis et al. (2014). We assume that successful decoding occurs with some constant positive probability $q_i \in (0, 1]$.

To model the interference in the shared wireless medium, we suppose that if more than one sensors transmit at the same slot, a collision occurs on all sent packets. This model is usually employed in control literature (Zhang (2003); Tabbara and Nesic (2008)) and in wireless communication systems (Lee et al. (2007); Hu and Ribeiro (2011); Jin and Kesidis (2002); Chen et al. (2010)). Let us indicate with $\gamma_{i,k} \in \{0, 1\}$ the success of the transmission at time slot k for link/system i. This is a Bernoulli random variable with success probability

$$\mathbb{P}(\gamma_{i,k} = 1) = \alpha_i \, q_i \prod_{j \neq i} (1 - \alpha_j). \tag{1}$$

This expression states that the probability of system i 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.

Our goal is to design the communication aspects of the problem, hence we assume the dynamics for all m control systems are fixed, meaning that controllers have been already designed. We suppose the system evolution is described by a switched linear time invariant model,

$$x_{i,k+1} = \begin{cases} A_{c,i} \, x_{i,k} + w_{i,k}, & \text{if } \gamma_{i,k} = 1 \\ A_{o,i} \, x_{i,k} + w_{i,k}, & \text{if } \gamma_{i,k} = 0 \end{cases}. \tag{2}$$

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., Hespanha et al. (2007). 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 $W_i \succeq 0$. An example of such a networked control system model (2) is presented next.

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

$$\begin{aligned} x_{i,k+1} &= \lambda_{o,i} x_{i,k} + u_{i,k} + w_{i,k}, \\ y_{i,k} &= x_{i,k} + v_{i,k}, \end{aligned} \tag{3}$$

where $w_{i,k}$ and $v_{i,k}$ are i.i.d. Gaussian disturbance and measurement noise respectively. Each wireless sensor i transmits the output measurement $y_{i,k}$ to the controller. Consider a simple control law which applies a zero input $u_{i,k} = 0$ when no information is received, and upon receiving a measurement it applies an output feedback $u_{i,k} = f_i y_{i,k}$ leading to a stable closed loop mode $\lambda_{c,i} = \lambda_{o,i} + f_i$. The overall networked system dynamics are expressed as

$$x_{i,k+1} = \begin{cases} \lambda_{c,i} \, x_{i,k} + w_{i,k} + f_i v_{i,k}, & \text{if } \gamma_{i,k} = 1 \\ \lambda_{o,i} \, x_{i,k} + w_{i,k}, & \text{if } \gamma_{i,k} = 0 \end{cases}. \tag{4}$$

which is of the form (2). Dynamic control laws with local estimates of the plant state at the controller (e.g. Hespanha et al. (2007)) can be expressed similarly.

The random packet success on link i modeled by (1) causes each control system i in (2) to switch in a random fashion between the two modes of operation (open and closed loop). As a result the access rate vector $\alpha \in [0, 1]^m$ to be designed affects the performance of *all* control systems. The following result characterizes, via a Lyapunov-like abstraction, a connection between control performance and the packet success rate.

Theorem 1. Consider a switched linear system i described by (2) with $\gamma_{i,k}$ being an i.i.d. sequence of 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 $P_i \in S_{++}^{n_i}$. Then the function decreases with an expected rate $\rho_i < 1$ at each step, i.e., we have

$$\mathbb{E}\left[V_i(x_{i,k+1})\,\middle|\,x_{i,k}\right] \leq \rho_i\, V_i(x_{i,k}) + Tr(P_i W_i) \qquad (5)$$

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

$$\mathbb{P}(\gamma_{i,k} = 1) \geq c_i, \qquad (6)$$

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

$$c_i = \min\{\theta \geq 0 : \; \theta A_{c,i}^T P_i A_{c,i} + (1-\theta) A_{o,i}^T P_i A_{o,i} \preceq \rho_i P_i\} \qquad (7)$$

Proof. The expectation over the next system state $x_{i,k+1}$ on the left hand side of (5) accounts via (2) for the randomness introduced by the process noise $w_{i,k}$ as well as the random success $\gamma_{i,k}$. In particular we have that

$$\mathbb{E}\left[V_i(x_{i,k+1})\,\middle|\,x_{i,k}\right] = \mathbb{P}(\gamma_{i,k} = 1)\; x_{i,k}^T A_{c,i}^T P_i A_{c,i} x_{i,k} + \mathbb{P}(\gamma_{i,k} = 0)\; x_{i,k}^T A_{o,i}^T P_i A_{o,i} x_{i,k} + Tr(P_i W_i), \qquad (8)$$

Here we used the fact that the random variable $\gamma_{i,k}$ is independent of the system state $x_{i,k}$. Plugging (8) at the left hand side of (5) we get for $x_{i,k} \neq 0$

$$\mathbb{P}(\gamma_{i,k} = 1) \geq \frac{x_{i,k}^T (A_{o,i}^T P_i A_{o,i} - \rho_i P_i) x_{i,k}}{x_{i,k}^T (A_{o,i}^T P_i A_{o,i} - A_{c,i}^T P_i A_{c,i}) x_{i,k}}. \qquad (9)$$

Since condition (5) needs to hold at any value of $x_{i,k} \in \mathbb{R}^{n_i}$, we can rewrite (9) as $\mathbb{P}(\gamma_{i,k} = 1) \geq c_i$ where

$$c_i = \sup_{y \in \mathbb{R}^{n_i}, y \neq 0} \frac{y^T (A_{o,i}^T P_i A_{o,i} - \rho_i P_i) y}{y^T (A_{o,i}^T P_i A_{o,i} - A_{c,i}^T P_i A_{c,i}) y}. \qquad (10)$$

This is equivalent to the semidefinite program (7).

The interpretation of the quadratic function $V_i(x_i)$ in this proposition is that it acts as a Lyapunov function for the control system. When the loop closes, the Lyapunov function of the system state decreases, while in open loop it increases. Condition (5) describes an overall decrease in expectation over the packet success. Consequently it guarantees control performance, in the sense that the variance of the plant state decreases exponentially at a rate ρ_i. 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. An example is shown next. In this paper the terms control performance (cf.(5)) and packet success rate (cf.(6)) are used interchangeably.

Example (continued). For the scalar wireless control system in (4) we can without loss of generality consider the quadratic function $V_i(x_i) = x_i^2$. Then for any desired decrease rate ρ_i in (5) the equivalent packet success rate by (7) becomes

$$c_i = \frac{\lambda_{o,i}^2 - \rho_i}{\lambda_{o,i}^2 - \lambda_{c,i}^2}. \qquad (11)$$

This illustrates the dependence of the packet success rate to the system dynamics (open and closed loop eigenvalues) and the control performance requirement, as can also be seen from the general form of (7). For example, higher desired control performance captured by smaller ρ_i requires a higher packet rate c_i, while a 'more stable' system with a smaller $\lambda_{c,i}$ requires a lower packet rate.

Our goal is to select sensor access rates α so that the Lyapunov functions *for all* control loops $i = 1, \ldots, m$ decrease *in expectation* at the desired rates $\rho_i < 1$ at any time k. By the above theorem, such control performance requirements are equivalent to minimum packet success rates c_i in (6) for each link i, computed by (7).

Without loss of generality and to simplify notation in the rest of the paper we consider the decoding probabilities on each link i

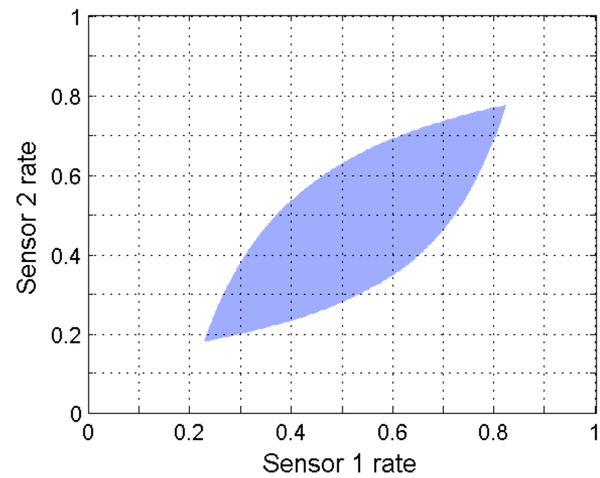

Fig. 2. Feasible set of sensor access rates for $m = 2$ loops. This is an example of the set or rates α_1, α_2 satisfying inequalities of the form (12) for some c_1, c_2, meeting equivalently the desired control performance of the two loops.

are $q_i = 1$. So the control requirement (5)-(6) that each sensor i needs to satisfy is equivalent to

$$c_i \leq \alpha_i \prod_{j \neq i} \left(1 - \alpha_j\right). \qquad (12)$$

If $q_i < 1$ it is immediate that that the requirement (6) can be expressed as in (12) with the left hand side modified to c_i/q_i. Intuitively by (12) a sensor i needs to select an access rate α_i high enough in order to overcome the collisions caused by other sensors and achieve control performance. On the other hand it cannot select a very high access rate otherwise too many collisions occur and other sensors j cannot meet their performance c_j – see Fig. 2 for an example of the set of feasible rates.

In the following section we present our decentralized mechanism for sensors to select appropriate access rates without a need for coordination. It relies on the fact that each sensor can observe locally the impact that other sensors have on its own control performance due to collisions, and react by adjusting its channel access rate. After establishing theoretical convergence to desired control performance for all systems, numerical simulations of the mechanism are presented in Section 4.

3. DECENTRALIZED CHANNEL ACCESS MECHANISM

The proposed mechanism is iterative. Suppose at iteration t each sensor selects an access rate $\alpha_i(t)$. Sensor i then observes in a decentralized manner on link i a discrepancy between the current packet success rate $\alpha_i(t) \prod_{j \neq i}(1 - \alpha_j(t))$ and the desired one c_i for control performance – see Remark at the end of this section for an implementation. We denote this discrepancy by

$$d_i(t) = c_i - \alpha_i(t) \prod_{j \neq i} \left(1 - \alpha_j(t)\right). \qquad (13)$$

If $d_i(t) > 0$ sensor i intuitively can increase its access rate to meet the control performance. Otherwise sensor i can decrease its rate to mitigate collisions on other sensors. Hence consider a simple decentralized update rule

$$\alpha_i(t+1) = \left[\alpha_i(t) + \frac{\varepsilon(t)}{\alpha_i(t)} d_i(t)\right]_{\mathcal{A}}. \qquad (14)$$

Here $\varepsilon(t) > 0$ is a suitably small positive step size. The term $\alpha_i(t)$ is added at the denominator to simplify notation, but is not required for the convergence results to follow. Since the access rate always needs to be in the unit interval, here $[\ ]_{\mathcal{A}}$ denotes the projection on a subset $\mathcal{A} = [\alpha_{\min}, 1]$ of the unit interval, where $\alpha_{\min} > 0$ is a small positive value to guarantee the denominator in (14) is well-behaved.

In the following theorem we analyze the convergence of this decentralized mechanism. It follows the arguments of Jin and Kesidis (2002); Chen et al. (2010) where a similar random access adaptation mechanism is developed – see also the remark at the end of this section.

Theorem 2. Consider a random access architecture with m control loops (2), communication modeled by (1), and control performance abstracted by (5)-(6) for each loop $i = 1, \dots, m$. Then an access rate vector $\alpha^* \in [0, 1]^m$ that satisfies all control performance constraints with equality is

(1) an equilibrium of the decentralized access mechanism (14),
(2) locally stable for mechanism (14) if $\varepsilon(t) > 0$ is sufficiently small, and if the square matrix $M(\alpha^*)$ defined for $1 \leq i \neq j \leq m$ as
$$M_{ii} = -\frac{c_i}{(\alpha_i^*)^2}, \quad M_{ij} = \prod_{\ell \neq i,j} (1 - \alpha_\ell^*) \tag{15}$$
is negative definite.

Proof. Proof of the first part is straightforward. Define the functions $f_i(\alpha) = \alpha_i \prod_{j \neq i}(1 - \alpha_j)$, so that the feasible set of access rates is defined as $\{\alpha \in [0, 1]^m : f(\alpha) \geq c\}$. The point α^* in the statement satisfies $f(\alpha^*) = c$, so that $\alpha_t = \alpha^*$ implies $d_i(t) = 0$ in (13) for all $i = 1, \dots, m$, and therefore $\alpha_{t+1} = \alpha_t = \alpha^*$.

To prove the second part, we begin by defining the function
$$H(\alpha) = \sum_{i=1}^{m} c_i \log(\alpha_i) + \prod_{j=1}^{m}(1 - \alpha_j). \tag{16}$$
Note that the partial derivatives equal
$$\frac{\partial H}{\partial \alpha_i} = \frac{c_i}{\alpha_i} - \prod_{j \neq i}(1 - \alpha_j), \tag{17}$$
from which we conclude that mechanism (13)-(14) is equivalently written as
$$\alpha_{t+1} = \left[\alpha_t + \varepsilon_t \nabla H(\alpha_t)\right]_{\mathcal{A}} \tag{18}$$
where the projection $[\ \ ]_{\mathcal{A}}$ is element-wise. So the update mechanism moves along the gradients of the function $H(\alpha)$.

Then we argue that for a sufficiently small value ε_t the function $H(\alpha)$ increases along the trajectory. In particular consider the Taylor expansion at the point α_{t+1}
$$\begin{aligned} H(\alpha_{t+1}) =& H(\alpha_t) + \nabla H(\alpha_t)^T(\alpha_{t+1} - \alpha_t) \\ &+ 1/2(\alpha_{t+1} - \alpha_t)^T \nabla^2 H(\tilde{\alpha})(\alpha_{t+1} - \alpha_t) \end{aligned} \tag{19}$$
for some $\tilde{\alpha}$ that is a convex combination of α_t, α_{t+1}. We can bound the term involving the gradient in this expression as follows. By the projection theorem (see Prop. 2.2.1 by Bertsekas et al. (2003)) we have that for any point $w \in \mathcal{A}^m$ and any point z it holds that $(z - [z])^T(w - [z]) \leq 0$. Applying this inequality for $z = \alpha_t + \varepsilon_t \nabla H(\alpha_t)$ and $w = \alpha_t$ we get that
$$\|\alpha_{t+1} - \alpha_t\|^2 - \varepsilon_t \nabla H(\alpha_t)(\alpha_{t+1} - \alpha_t) \leq 0 \tag{20}$$
Substituting (20) in (19) we have
$$\begin{aligned} &H(\alpha_{t+1}) \geq H(\alpha_t) + \\ &+ 1/2(\alpha_{t+1} - \alpha_t)^T \left[1/\varepsilon_t I + \nabla^2 H(\tilde{\alpha})\right](\alpha_{t+1} - \alpha_t). \end{aligned} \tag{21}$$

For a sufficiently small ε_t the matrix $1/\varepsilon_t I$ dominates the Hessian $\nabla^2 H(\tilde{\alpha})$, and the matrix in the brackets above becomes positive definite. So we have $H(\alpha_{t+1}) > H(\alpha_t)$.

We have thus shown that for a sufficiently small ε_t the mechanism is essentially a gradient ascent. This will converge to some local maximum. A sufficient condition for the point α^* to be a local maximum is negative definiteness of the Hessian $\nabla^2 H(\alpha^*) \prec 0$. By differentiating (17) we have that
$$\frac{\partial^2 H}{\partial \alpha_i^2} = -\frac{c_i}{(\alpha_i^*)^2}, \quad \frac{\partial^2 H}{\partial \alpha_i \partial \alpha_j} = \prod_{\ell \neq i,j} (1 - \alpha_\ell^*), \tag{22}$$
which shows that $\nabla^2 H(\alpha^*)$ equals the matrix $M(\alpha^*)$ defined in the statement of the theorem and completes the proof. ∎

The theorem states that the iterative decentralized channel access mechanism will converge to an operating point where all control specifications are exactly met, as long as the starting point is close enough. Otherwise, it is possible that the sensors do not converge to feasible access rates. For example if two sensors initially transmit at very high rates, causing many collisions to each other, they will keep on increasing their access rates at every iteration in an attempt to meet their control performance, and a deadlock is reached – see also a numerical example of this situation in Section 4. On the other hand, the theorem provides explicit conditions under which the mechanism converges locally to a desirable operating point. For the special case of $m = 2$ control loops we will show next that except for degenerate cases these conditions always hold true. Hence the mechanism can be expected to converge in cases of practical interest.

Besides iterative adaptation to the other sensors' access rates, a sensor may use the decentralized mechanism to adapt to other changes in the environment. This can be for example the case where control performance specifications are varying over time, which due to Theorem 1 would correspond to an equivalent varying packet success rate $c_i(t)$ in (13)-(14). Moreover we show in simulations in Section 4 how the mechanism adapts when new control systems are introduced in the shared wireless medium.

Remark. For a practical implementation of the proposed mechanism, a sensor needs to observe the discrepancy between current and desired packet success rate in (13). A sensor can measure this discrepancy using acknowledgments sent from its corresponding receiver/controller in Fig. 1 for every transmitted packet. This way the sensor can measure empirically over a period of time the packet success rate on its link/loop. This procedure will include some errors which are neglected in the analysis of this paper.

Remark. The decentralized mechanism has a game theoretic interpretation according to Jin and Kesidis (2002); Chen et al. (2010). At each round t of the game, each sensor/agent i selects an action $\alpha_i(t)$ and all agents observe corresponding outcomes, which here can be thought as the observed packet success rates. At the next round each agent i responds to the actions of the other agents by adjusting its action to a new value $\alpha_i(t+1)$ according to (14). This action is selected to improve i's utility assuming that all other agents will retain their previous actions. Such a policy is called better response, or gradient play (Chen et al. (2010)). Utilities are not explicitly formulated in our paper, but intuitively each sensor is satisfied with the smallest access rate that meets its control performance. The proof of Theorem 2 is based on a common potential function, similar to the game-theoretic framework of Jin and Kesidis (2002); Chen et al. (2010).

3.1 Special case: Random access for two control systems

Consider the setup with $m = 2$ sensors trying to achieve control performance captured by packet success rates c_1, c_2 respectively. A pair of sensor access rates α_1, α_2 that meet exactly the requirements satisfy

$$\alpha_1 (1 - \alpha_2) = c_1, \qquad \alpha_2 (1 - \alpha_1) = c_2. \tag{23}$$

According to Theorem 2, a solution α^* of these equations will be a locally stable operating point of the decentralized channel access mechanism (14) if it satisfies the negative definiteness condition of the matrix $M(\alpha^*)$ given by

$$\begin{bmatrix} -\frac{c_1}{(\alpha_1^*)^2} & 1 \\ 1 & -\frac{c_2}{(\alpha_2^*)^2} \end{bmatrix} = \begin{bmatrix} -\frac{1-\alpha_2^*}{\alpha_1^*} & 1 \\ 1 & -\frac{1-\alpha_1^*}{\alpha_2^*} \end{bmatrix} \tag{24}$$

where the equality follows by substituting (23). Examining the characteristic polynomial, it follows that $M(\alpha^*) \prec 0$ holds if and only if $\alpha_1^* + \alpha_2^* < 1$.

We then directly solve (23) to check this condition. From the first equation we get $\alpha_1 = c_1/(1 - \alpha_2)$, which if we plug in the second equation yields a quadratic equation

$$(\alpha_2)^2 - (1 + c_2 - c_1)\alpha_2 + c_2 = 0 \tag{25}$$

This equation has a (real) solution if and only if the discriminant is non-negative, i.e.,

$$\Delta = (1 - c_1 - c_2)^2 - 4c_1c_2 \geq 0. \tag{26}$$

We note in passing that given the relationship between packet success rate c_i and control performance of system i (cf. Theorem 1) condition (26) describes the control performance specifications that can be supported by our random access architecture.

In general (25) might have two solutions, as seen e.g., in Fig. 2. The minimum solution is of the form $\alpha_2^* = (1 + c_2 - c_1 - \sqrt{\Delta})/2$. A symmetric argument shows that $\alpha_1^* = (1 + c_1 - c_2 - \sqrt{\Delta})/2$. From these two expressions we verify that the minimum equilibrium point satisfies $\alpha_1^* + \alpha_2^* = 1 - 2\sqrt{\Delta}$, which is always less than unity except for the degenerate case $\Delta = 0$. To sum up, for $m = 2$ sensors in almost all cases the access rates meeting the control performance requirements are a locally stable equilibrium of the decentralized channel access mechanism. In the following section we present numerical simulations of the decentralized channel access mechanism.

4. NUMERICAL SIMULATIONS

We first consider two ($m = 2$) scalar control systems. Using the notation of the example of Section 2, we assume both systems have identical dynamics, unstable open loop $\lambda_{o,1} = \lambda_{o,2} = 1.05$ and stable closed loop $\lambda_{c,1} = \lambda_{c,2} = 0.1$. We suppose that the two control systems have different control performance requirements given by the desired Lyapunov decrease rates $\rho_1 = 0.9$ and $\rho_2 = 0.95$, corresponding to different packet success rates on the two links according to Theorem 1. The set of feasible sensor access rates α_1, α_2 that meet both control performances is shown in Fig. 2.

We employ the decentralized update rule (14) for each sensor to attempt to reach a global operating point. In Fig. 3 we show the evolution of the two sensor access rates for different initial points. On one case the feasible point where both control performances are exactly satisfied is reached after some iterations. Intuitively, sensor 1 access the medium more often because it has a more demanding desired control performance captured by $\rho_1 < \rho_2$. On the other case, both sensors initially transmit at a very high rate, and no feasible operating point is reached. Each sensor tries to respond to the packet collisions inflicted by the other by increasing its access rate higher and higher. The iteration reaches a deadlock where both sensors transmit at full rate $\alpha_1 = \alpha_2 = 1$. In practice one would have to enforce some arbitration rule to overcome such deadlocks. For example, sensors may restart the algorithm by transmitting at some fixed low rates till the desired operation point is reached.

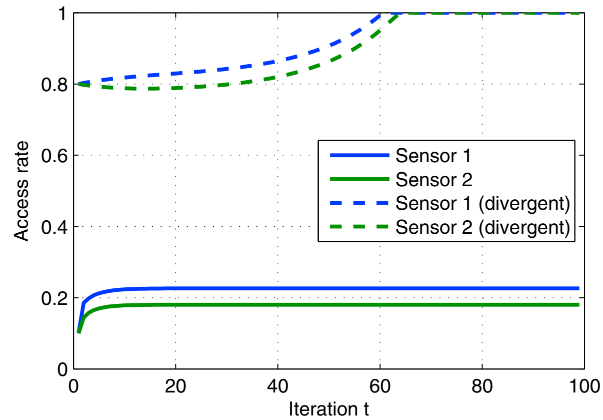

Fig. 3. Evolution of the decentralized channel access evolution for different starting points. Both convergence to a desired operating point, and divergence to the non-feasible region can occur.

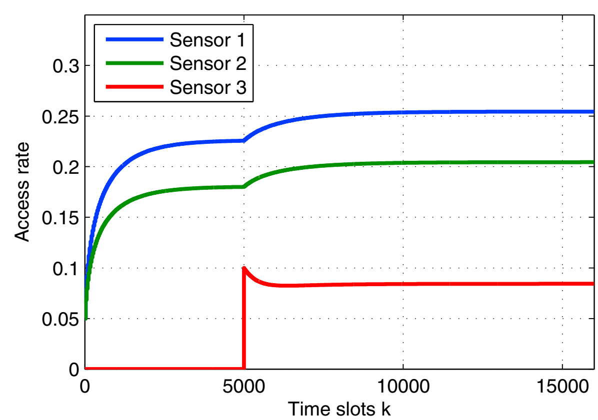

Fig. 4. Evolution of the access rates after an introduction of a third control system in the wireless medium. The sensors need to increase their access rates to mitigate the increased collisions and meet control performance.

Next we turn our attention to how the mechanism can adapt to other changes in the environment, in particular to an introduction of a new control system in the shared wireless medium. After the two systems are in a steady state, a third system with parameters $\lambda_{o,3} = 1$, $\lambda_{c,3} = 0$, $\rho_3 = 0.95$ is introduced. We assume that an iteration of the sensor access update rule (14) is performed every 50 time slots. Each sensor i measures the discrepancy $d_i(t)$ in (13) during this period, for example using acknowledgments from the corresponding receiver/controller, and adjusts its access rate. In practice this measurement of the discrepancy has some error which is however neglected in our simulations. In Fig. 4 we plot the trajectory of the sensor access rates. We observe that after the introduction of the new system there is a period of adaptation till the new operating point is found. The access rates are increased to mitigate the counteract the increased collisions on the medium due to the introduction of the new system.

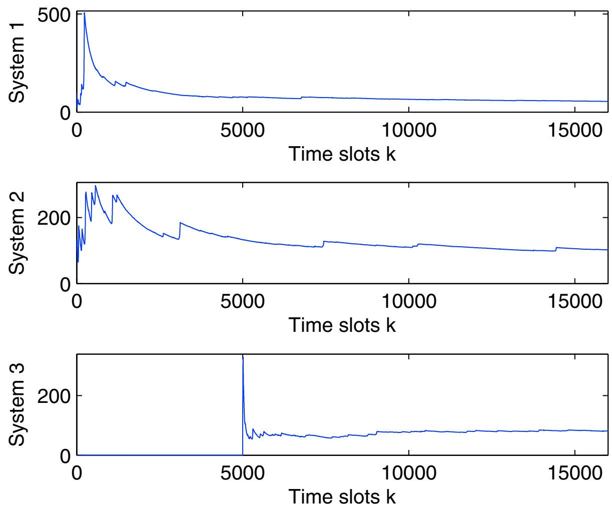

Fig. 5. Evolution of the mean square magnitude $1/N\sum_{k=1}^{N} x_{i,k}^2$ of the plant states during the introduction of a third system. All systems remain stable.

We are also interested in how the plant states of each control system evolve. As an indication of the latter we plot the average square magnitude of the state $1/N\sum_{k=1}^{N} x_{i,k}^2$ over time for all systems in Fig. 5. After some transients the mean square magnitudes reach a steady state corresponding to desired control performance. Even as the third system is introduced no significant deviations are observed, since convergence to the new necessary access rates is relatively fast (Fig. 4).

5. CONCLUSION

We consider a random access architecture for multiple control loops sharing a common wireless channel. Each sensor randomly decides whether to access the channel, at a rate that needs to be appropriately selected to both meet control performance and mitigate the effect of packet collisions from simultaneous transmissions. We develop a decentralized mechanism that reaches a desired operating point without the need to coordinate among the sensors.

Our work is a starting point for exploring efficient, easily implementable, and distributed mechanisms for control systems over shared wireless channels. Future work includes theoretical analysis of the decentralized mechanism under varying or asymmetric channel conditions. Further exploration is also required for mechanisms adapting to the varying plant conditions, similar to the state-based approaches of e.g., Gatsis et al. (2014); Donkers et al. (2011); Ramesh et al. (2013), as well as random access protocols for estimation and linear quadratic control (Schenato et al. (2007)).

REFERENCES

Alur, R., D'Innocenzo, A., Johansson, K.H., Pappas, G.J., and Weiss, G. (2011). Compositional modeling and analysis of multi-hop control networks. *IEEE Trans. on Automatic control*, 56(10), 2345–2357.

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

Blind, R. and Allgöwer, F. (2011). Analysis of networked event-based control with a shared communication medium: Part I-Pure ALOHA. In *IFAC World Congress*.

Chen, L., Low, S.H., and Doyle, J.C. (2010). Random access game and medium access control design. *IEEE/ACM Trans. on Networking*, 18(4), 1303–1316.

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

Gatsis, K., Pajic, M., Ribeiro, A., and Pappas, G.J. (2015a). Opportunistic control over shared wireless channels. *IEEE Trans. on Automatic Control*. To appear.

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

Gatsis, K., Ribeiro, A., and Pappas, G.J. (2015b). Control with random access wireless sensors. In *54th IEEE Conference on Decision and Control*. To appear.

Gupta, V., Chung, T.H., Hassibi, B., and 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., and Xu, Y. (2007). A survey of recent results in networked control systems. *Proc. 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*, volume 4, 3631–3636.

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

Jin, Y. and Kesidis, G. (2002). Equilibria of a noncooperative game for heterogeneous users of an aloha network. *IEEE Communications Letters*, 6(7), 282–284.

Le Ny, J., Feron, E., and Pappas, G.J. (2011). Resource constrained lqr control under fast sampling. In *14th Int. Conf. on Hybrid Systems: computation and control*, 271–280.

Lee, J.W., Chiang, M., and Calderbank, A.R. (2007). Utility-optimal random-access control. *IEEE Trans. on Wireless Communications*, 6(7), 2741–2751.

Liu, X. and Goldsmith, A. (2004). Wireless medium access control in networked control systems. In *Proc. of the 2004 American Control Conference (ACC)*, volume 4, 3605–3610.

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

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

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

Schenato, L., Sinopoli, B., Franceschetti, M., Poolla, K., and Sastry, S.S. (2007). Foundations of control and estimation over lossy networks. *Proc. of the IEEE*, 95(1), 163–187.

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

Vasconcelos, M. and Martins, N. (2014). Remote estimation games over shared networks. In *52nd Annual Allerton Conf. on Communication, Control, and Computing*, 12–18.

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

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

Zhang, W. (2003). Stabilization of networked control systems over a sharing link using ALOHA. In *Proc. of the 42nd IEEE Conf. on Decision and Control (CDC)*, volume 1, 204–209.