

Optimal Cross-Layer Resource Allocation in Cellular Networks Using Channel- and Queue-State Information

Antonio G. Marques, *Member, IEEE*, Luis M. Lopez-Ramos, *Student Member, IEEE*, Georgios B. Giannakis, *Fellow, IEEE*, Javier Ramos, and Antonio J. Caamaño

Abstract—Algorithms that jointly allocate resources across different layers are envisioned to boost the performance of wireless systems. Recent results have revealed that two of the most important parameters that critically affect the resulting cross-layer designs are channel- and queue-state information (QSI). Motivated by these results, this paper relies on stochastic convex optimization to develop optimal algorithms that use instantaneous fading and queue length information to allocate resources at the transport (flow-control), link, and physical layers. Focus is placed on a cellular system, where an access point exchanges information with different users over flat-fading orthogonal channels. Both uplink and downlink setups are considered. The allocation strategies are obtained as the solution of a constrained utility maximization problem that involves average performance metrics. It turns out that the optimal allocation at a given instant depends on the instantaneous channel-state information (CSI) and Lagrange multipliers, which are associated with the quality-of-service (QoS) requirements and the operating conditions of the system. The multipliers are estimated online using stochastic approximation tools and are linked with the *window-averaged* length of the queues. Capitalizing on those links, queue stability and average queuing delay of the developed algorithms are characterized, and a simple mechanism is devised to effect delay priorities among users.

Index Terms—Congestion control, cross-layer design, dynamic resource management, network optimization, stochastic optimization.

I. INTRODUCTION

IN DYNAMIC wireless networks where the channel-state information (CSI) and the queue-state information (QSI) vary, nodes must adapt their transmission and reception parameters to these changes while adhering to power constraints and diverse

quality-of-service (QoS) requirements. In recent years, the design of optimal resource allocation schemes (ASs) that take into account cross-layer information has attracted the attention of information theory, signal processing, communication, adaptive control, and networking communities. Existing works can mainly be classified in one of the following two categories. In the first category, approaches rely on convex optimization and dual decomposition techniques; for example, see [4], [15]–[17], and [30]. In these works, the solution of a properly formulated optimization problem dictates how resources are allocated, whereas the structure of the solution typically suggests the design of signaling protocols. In the second category, approaches build on dynamic backpressure policies, pioneered by Tassiulas and Ephremides [32]; for example, see [6], [8], [9], [27], and [31]. These works rely on adaptive control tools (the so-called Lyapunov optimization) and aim at stabilizing queues of the wireless network. Such queues correspond not only to actual queues where packets are buffered before transmission but also to virtual queues whose role is to account for QoS requirements. Most existing works have focused on the impact of the cross-layer features on higher layers, and fewer works have addressed the design of cross-layer networking algorithms that take into account *instantaneous fading*; see [4], [7], [8], [15], [17], and [28] for some of the exceptions.

In this context, this paper aims at optimally designing cross-layer resource allocation algorithms for *fading* wireless *cellular* networks under the following operating conditions. At the transport layer, every node receives packets from higher layers. The packets of each user entail different utility levels, and nodes implement simple flow-control mechanisms to keep the network stable. At the link layer, users orthogonally share a set of parallel flat-fading channels. At the physical layer, nodes can adapt their power and rate loadings per channel realization. Both uplink (multiple-access) and downlink (broadcast) channels will be considered. The optimal cross-layer resource allocation (flows, powers, rates, and channels) is obtained as the solution of a constrained optimization problem, which naturally takes into account flow-specific utility functions, individual QoS requirements, and the network's operating conditions. The resulting optimum dynamic resource allocation is found in closed form, and it is shown to depend only on the current channel realization and user-specific prices (Lagrange multipliers). An adaptive dual stochastic algorithm is developed to estimate the optimal Lagrange multiplier values. The developed

Manuscript received June 10, 2011; revised November 2, 2011, January 17, 2012, and March 25, 2012; accepted April 4, 2012. Date of publication April 20, 2012; date of current version July 10, 2012. This work was supported in part by the Spanish Ministry of Science and Innovation under Grant TEC2009-12098, by the Spanish Ministry of Education, under FPU Grant AP2010-1050, and by the National Science Foundation under Grant 0830480, Grant 1016605, Grant 0824007, and Grant 1002180. This paper was presented in part at the 2010 IEEE Global Telecommunications Conference, Miami, FL. The review of this paper was coordinated by Prof. Y. Cheng.

A. G. Marques, L. M. Lopez-Ramos, J. Ramos and A. J. Caamaño are with the Department of Signal Theory and Communications, King Juan Carlos University, 28943 Madrid, Spain (e-mail: antonio.garcia.marques@urjc.es; luismiguel.lopez@urjc.es; javier.ramos@urjc.es; antonio.caamano@urjc.es).

G. B. Giannakis is with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: georgios@umn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVT.2012.2195732

algorithm is robust to nonstationarities, does not require knowledge of the channel distribution, and arbitrarily incurs minimal loss of performance (optimality). Furthermore, it is shown that, under very mild conditions, stochastic estimates of some of the Lagrange multipliers correspond to scaled versions of *window averages* of queue lengths. In other words, the developed schemes reveal one way in which QSI metrics that are related to the queue length can be used to optimize network performance. Building on the relationship between queues and multipliers, both stability and average queuing delay of the novel cross-layer resource allocation algorithms are characterized.

The main differences relative to the state of the art and the most relevant contributions of this paper are listed as follows.

- 1) The specific cross-layer resource allocation algorithms designed in this paper have not been considered.
- 2) The convergence results are stronger than the results available in the QSI literature. The convergence results are similar to the results in [28], which neither deals with queues nor considers window averages.
- 3) Approximations for the average queuing delay are provided, and a simple mechanism for effecting delay priorities is presented, which is attractive from a practical perspective.
- 4) The QSI corresponds to window averages of the instantaneous queue lengths.
- 5) The CSI/QSI adaptive schemes are designed by purely relying on convex optimization and dual stochastic algorithms, without *explicitly* including the queue lengths in the formulation.

Contributions 4 and 5 deserve further elaboration. From a practical perspective, contribution 4 is interesting, because the QSI variability will be smaller, and its time correlation is stronger; therefore, for example, more aggressive coding schemes can be employed for QSI feedback. Indeed, several actual congestion control and queue management protocols rely on window-averaged queue lengths [21]. From an analytical perspective, contribution 4 entails that the convergence (feasibility and optimality) proofs in the paper have to account for QSI updates being biased and outdated versions of the actual instantaneous queue lengths. As a result, the proofs can easily be adapted to deal with additional sources of QSI imperfections, which are relevant from a practical perspective (e.g., noise or delay in the signaling channels). The approach in contribution 5 opens the door to applying several known results to the problem at hand. For example, sensitivity analysis can be used to quantify tradeoffs between the considered constraints and the average queuing delay. The approach also allows for different versions of the dual stochastic updates, each potentially giving rise to different forms (metrics) of QSI; see [25], [28], and [33] for works that do not deal with QSI but use different stochastic updates to allocate resources in wireless networks.

The rest of this paper is organized as follows. The system setup and operating conditions are described in Section II. The constrained optimization problem is formulated in Section III. Its solution is presented in Section IV, where a stochastic method of estimating the optimum Lagrange multipliers is also developed, along with the connection between Lagrange multi-

pliers and queue lengths. Queue stability and the average delay are characterized in Section V, where a method of effecting delay priorities among users is also presented. Sections III–V deal with the uplink, whereas Section VI briefly presents their counterparts for the downlink. The signaling overhead required to implement the developed schemes is discussed in Section VII. Numerical results and conclusions in Sections VIII and IX wrap up this paper.¹

II. MODELING PRELIMINARIES

In this section, the system setup and channel model are first introduced. Then, the operation of the link and physical layers is presented. Finally, the flow-control mechanism and queue dynamics are described. To facilitate the readability of this paper, the most relevant notation introduced here and in the following sections is summarized in Table I.

Consider that M wireless terminals (users) are connected to an access point (AP). The overall bandwidth B is divided into K orthogonal channels, each with bandwidth B/K that is small enough to ensure that the fading per channel is flat, i.e., nonselective. The wireless link between the AP and user m on channel k is characterized by its random square magnitude h_m^k , which is assumed normalized by the receiver noise variance. The overall $MK \times 1$ gain vector is denoted by $\mathbf{h} := \{h_m^k, m = 1, \dots, M, k = 1, \dots, K\}$. The channel-fading process is assumed ergodic, stationary, and statistically independent across users and time. Moreover, it is assumed that \mathbf{h} remains invariant over a block of symbols but can vary from block to block (block-fading channel model). In other words, if n denotes the current block index (whose duration is dictated by the channel coherence interval), then $\mathbf{h}[n]$ remains constant, and $\mathbf{h}[n] \neq \mathbf{h}[n+1]$ with probability one (w.p.1). Notation $\mathbf{h}[n]$ will be used whenever the block time variance of the channel needs to be emphasized.

With regard to the link-layer operation, links at the outset are scheduled to simultaneously but orthogonally access in time (or frequency) any of the channels; for example, see [23] and [35]. To quantitatively describe access, let $w_m^k(\mathbf{h}) \in [0, 1]$ denote the nonnegative fraction of time (or channel bandwidth) that user m is scheduled to transmit over channel k during the channel realization \mathbf{h} . Since, individual user transmissions interfere with each other, it must hold that

$$\sum_{m=1}^M w_m^k(\mathbf{h}) \leq 1 \quad \forall k, \quad \forall \mathbf{h}. \quad (1)$$

This way, if $w_m^k(\mathbf{h}) = 0.9$ and $w_{m'}^{k'}(\mathbf{h}) = 0.1$, the traffic of user m is assigned to channel k during 90% of the duration of realization \mathbf{h} and to user m' during the remaining 10%,

¹As a notation, \star denotes the convolution operator, and x^* is the optimal value of variable x . $\mathbb{E}[\cdot]$ is the expectation, $\mathbb{1}_{\{\cdot\}}$ is the indicator function ($\mathbb{1}_{\{x\}} = 1$ if x is true; otherwise, it is zero), $[x]_a^b$ is the projection of x onto the $[a, b]$ interval, i.e., $[x]_a^b = \min\{\max\{a, x\}, b\}$, and $W(x[n], L) := 1/L \sum_{l=n-L+1}^n x[l]$ is the windowed average of length L of variable x at time $n \geq L$. Finally, for a function $f(\cdot)$, $(f)^{-1}(\cdot)$ denotes its inverse, and $\dot{f}(\cdot)$ denotes its derivative.

TABLE I
SUMMARY OF THE MOST SIGNIFICANT NOTATION

Symbol	Meaning
n	Time slot index
M, m	Number of users / User index
K, k	Number of channels / Channel index
\mathbf{h}, h_m^k	Channel gain vector / channel gain for user m on channel k
w_m^k	Scheduling variable (fraction of time user m occupies channel k)
$p_m^k(\mathbf{h})$	Power transmitted by user m over channel k
$r_m^k(\mathbf{h})$	Rate transmitted by user m over channel k
\check{p}^k	Maximum instantaneous power level in channel k
\check{p}_m	Maximum average power consumed by user m (Uplink)
\bar{p}	Maximum average power consumed by the BS (Downlink)
$C_m^k(\mathbf{h}, p_m^k(\mathbf{h}))$	Capacity (rate-power) function for user m over channel k
a_m	Instantaneous arrival rate of user m
\bar{a}_m	Average arrival rate of user m
q_m	Length of queue m
$U_m(\bar{a}_m)$	Utility function of user m 's service rate
\mathbf{P}^*	Sum-utility achieved by the optimal resource allocation [cf. (6)]
π_m	Lagrange multiplier for user m 's average power constraint
ρ_m	Lagrange multiplier for user m 's flow conservation constraint
α_m	Lagrange multiplier for user m 's average arrival rate constraint
$\boldsymbol{\lambda}$	Vector gathering all Lagrange multipliers
$\pi_m^*, \rho_m^*, \alpha_m^*, \boldsymbol{\lambda}^*$	(Optimal) Value of the <i>non-stochastic</i> Lagrange multipliers
L	Length of the averaging window
$\pi_{L,m}[n], \rho_{L,m}[n], \alpha_{L,m}[n], \boldsymbol{\lambda}_L[n]$	Stochastic estimate of the multipliers using an averaging window of length L at slot n
μ, μ_m	Stepsize / User-selective stepsize (for delay priorities)
$f(m, k, h, \boldsymbol{\lambda})$	Link Quality Indicator of channel k for user m

and traffic of no other user is scheduled. Clearly, if we restrict $w_m^k(\mathbf{h}) \in \{0, 1\}$, then $w_m^k(\mathbf{h})$ can readily be viewed as an indicator variable, which means that, if channel k is assigned to user m , then it holds that $w_m^k(\mathbf{h}) = 1$, whereas $w_{m'}^k(\mathbf{h}) = 0$ for all $m' \neq m$. No restriction is placed on the maximum number of channels that a user can access.

The resources adapted at the physical layer will be power and rate per user and channel. In particular, let $p_m^k(\mathbf{h})$ and $r_m^k(\mathbf{h})$ denote, respectively, the instantaneous power and rate user that m transmits over channel k during the channel realization \mathbf{h} if $w_m^k(\mathbf{h}) = 1$. Two types of power constraints are considered. Spectrum mask constraints are imposed to ensure that the *instantaneous* $p_m^k(\mathbf{h})$ does not exceed a maximum prespecified level \check{p}^k , i.e.,

$$p_m^k(\mathbf{h}) \leq \check{p}^k \quad \forall k, m, \mathbf{h}. \quad (2)$$

On the other hand, the maximum *average* power terminal that m can transmit is also bounded by \check{p}_m ; hence

$$\mathbb{E} \left[\sum_{k=1}^K p_m^k(\mathbf{h}) w_m^k(\mathbf{h}) \right] \leq \check{p}_m \quad \forall m \quad (3)$$

where the expectation was taken over the fading channel distribution. Since an uplink setup is considered here, users act as transmitters, and the power of *each individual user* has to be bounded [cf. (3)]. In the downlink, the AP acts as a transmitter, and thus, only the AP power must be constrained (see Section VI for details).

Under bit-error-rate or capacity constraints, power $p_m^k(\mathbf{h})$ is coupled with the corresponding rate $r_m^k(\mathbf{h})$. This rate–power coupling will be represented by the function $C_m^k(\mathbf{h}, p_m^k(\mathbf{h}))$. It is assumed throughout that the rate–power function $C_m^k(\mathbf{h}, p_m^k(\mathbf{h}))$ is increasing and strictly concave. For example,

if sufficiently strong coding schemes are used, $C_m^k(\mathbf{h}, p_m^k(\mathbf{h}))$ approaches Shannon's capacity formula $\log(1 + h_m^k p_m^k(\mathbf{h}))$, which is increasing and strictly concave. The one-to-one mapping between $p_m^k(\mathbf{h})$ and $r_m^k(\mathbf{h})$ implies that, when the optimization problem is formulated, it suffices to optimize over one of them.

With regard to the transport and network layers, the operation is described as follows. Packets are exogenously generated at higher layers. Packet streams will be referred to as flows, and there will be as many flows as users, each user having one flow. The packet arrival rate of flow m at a given instant n is a random variable denoted by $a_m[n]$. The average arrival rate of exogenous information of flow m is denoted by \bar{a}_m . Terminals are equipped with queues (buffers) that can store the incoming packets. Let $q_m[n]$ denote the queue size that corresponds to flow m at time slot n . Then, the queue obeys the recursion, i.e.,

$$q_m[n+1] = \left[q_m[n] + a_m[n] - \sum_{k=1}^K r_m^k(\mathbf{h}[n]) w_m^k(\mathbf{h}[n]) \right]_0^\infty \quad \forall m. \quad (4)$$

In practice, packet arrivals and departures vary at a time scale smaller than n . This case implies that definitions that are slightly different from the definition in (4) are also possible [8]. Such differences, however, are inconsequential for the subsequent analysis, and (4) has been chosen for mathematical simplicity.

Different definitions of queue stability are available. In this paper, queues are deemed stable if $\lim_{n \rightarrow \infty} 1/n \times \sum_{l=1}^n \mathbb{E}[q_m[l]] < \infty$. This definition is referred to as “strong

stability” in [8]. For queues to be stable, the following necessary condition needs to be satisfied:

$$\bar{a}_m \leq \sum_{k=1}^K \mathbb{E} [r_m^k(\mathbf{h})w_m^k(\mathbf{h})] \quad \forall m. \quad (5)$$

The latter expression is typically known as the necessary average flow conservation condition. Equation (5), together with (1)–(3), is accounted for in the optimization problem presented in the next section.

III. PROBLEM FORMULATION AND DESIGN APPROACH

The optimal resource allocation will be viewed in this section as the solution of a constrained optimization problem. The *objective* will entail concave and increasing so-called *utility functions* $U_m(\cdot)$, which are commonly used in resource allocation tasks (not only restricted to communication systems) and account for the “social” utility (reward) to which a specific resource (here, \bar{a}_m) gives rise. Sum-utility maximization has frequently been employed by scheduling, medium access control layer, and networking algorithms; for example, see [8], [11], [18], and the references therein. Utility functions $U_m(\bar{a}_m)$ are chosen to be increasing (so that solutions that allow for higher arrival rates are promoted) and concave (so that the marginal utility for each user terminal diminishes as its rate increases), which offers a means of effecting fairness among different users. Typical utility functions $U_m(x)$ include $\alpha_m \log(\beta_m + x)$ and $\alpha_m(1 - \beta_m)^{-1}x^{1-\beta_m}$, where α_m and β_m are user-dependent positive constants. On the other hand, to effect QoS, a *minimum* average arrival rate \check{a}_m is guaranteed for certain users.

The optimal allocation is obtained by solving the following constrained sum-utility maximization:

$$\begin{aligned} \mathbf{P}^* := & \max_{\substack{\bar{a}_m \\ w_m^k(\mathbf{h}), p_m^k(\mathbf{h})}} \sum_{m=1}^M U_m(\bar{a}_m) & (6a) \\ \text{subject to} & \bar{a}_m \geq \check{a}_m, m = 1, \dots, M & (6b) \\ & (1)–(3), (5). & (6c) \end{aligned}$$

Recall that there is no need to optimize over $r_m^k(\mathbf{h})$, because $r_m^k(\mathbf{h})$ in (5) can be replaced with $C_m^k(\mathbf{h}, p_m^k(\mathbf{h}))$; thus, the optimum value of $p_m^k(\mathbf{h})$ will readily yield the optimum value of $r_m^k(\mathbf{h})$. Moreover, for “best effort” flows, the corresponding \check{a}_m in (6b) is set to zero. Both the cross-layer and the channel-adaptive attributes of the resource allocation problem are apparent, because variables of different layers are jointly optimized in (6), and several optimization variables and constraints are functions of \mathbf{h} . Finally, note that, in (6a), we have defined \mathbf{P}^* as the sum-rate utility value achieved by the optimum solution. This notation will be useful in the next sections to assess the potential loss of the performance of the algorithms to be developed.

The solution of (6) will be pursued in the next section, but several remarks are due before that. First, if the minimum rate requirements in (6b) are very high and the power budgets in (3) are very small, the optimization in (6) could be infeasible. This case, however, is readily detectable, because the Lagrange multipliers associated with some of the infeasible users would

grow unbounded. In such a case, the only option to stabilize the system is to simply drop users. Unfortunately, optimally selecting which users to drop (also known as admission control) often boils down to an NP-hard problem and goes beyond the scope of this paper. Second, the average rate constraints in (6b) guarantee that the *average* arrival rate of a flow remains above a given requirement (elastic traffic). However, there is no guarantee on the *instantaneous* transmit rate of real-time traffic. Last, note that queue dynamics were not explicitly taken into account in (6). Only the necessary condition for stability in (5) has explicitly been accounted for. However, it will be shown in the next section that queues can, in fact, be used as stochastic estimates of Lagrange multiplier values associated with (5). This result implies that the solution of (6) will also implicitly depend on QSI.

IV. OPTIMAL CROSS-LAYER ALLOCATION

The problem in (6) can readily be transformed to a convex one (see [23] or [35] for details), which can be solved using the dual approach. The optimal solution will first be presented as a function of the optimum Lagrange multipliers (dual variables). A stochastic scheme will be developed next for estimating the multipliers per time instant n . The last part of this section will be devoted to the relationship between stochastic multiplier estimates and queue lengths.

A. Optimal Allocation as a Function of the Multipliers

Let π_m , ρ_m , and α_m denote the Lagrange multipliers that are associated with the *average* constraints in (3), (5), and (6b), respectively. Collect all these multipliers in a vector $\boldsymbol{\lambda}$. There is no need to dualize the *instantaneous* constraints (1) and (2), because the solution will turn out to automatically satisfy them. Furthermore, let $(\dot{U}_m)^{-1}(\cdot)$ and $(\dot{C}_m^k)^{-1}(\mathbf{h}, \cdot)$ denote the inverse function of the derivative of $U_m(\cdot)$ and $C_m^k(\mathbf{h}, \cdot)$, respectively, and remember that x^* stands for the optimum value of a given variable x . Using these notational conventions and the Karush–Kuhn–Tucker (KKT) conditions [1] associated with (6), it is shown in Appendix A that the optimal cross-layer resource allocation is

$$\bar{a}_m^*(\boldsymbol{\lambda}^*) = \left[(\dot{U}_m)^{-1}(\rho_m^* - \alpha_m^*) \right]_0^\infty \quad (7)$$

$$p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}^*) = \left[(\dot{C}_m^k)^{-1}(\mathbf{h}, \pi_m^*/\rho_m^*) \right]_0^{\check{p}^k} \quad (8)$$

$$r_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}^*) = C_m^k(\mathbf{h}, p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}^*)). \quad (9)$$

For notational convenience, define $f(m, k, \mathbf{h}, \boldsymbol{\lambda}^*) := \rho_m^* r_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}^*) - \pi_m^* p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}^*)$. Then, Appendix A also shows that, for \mathbf{h} and $\boldsymbol{\lambda}^*$ given, only users m for which $f(m, k, \mathbf{h}, \boldsymbol{\lambda}^*)$ is maximum can transmit over channel k , i.e., users who satisfy $m_k = \arg \max_{m'} \{f(m', k, \mathbf{h}, \boldsymbol{\lambda}^*)\}$ [14], [23]. If the user who attains the maximum is unique, it readily follows that the optimum scheduling is

$$w_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}^*) = \mathbb{1}_{\{m = \arg \max_{m'} \{f(m', k, \mathbf{h}, \boldsymbol{\lambda}^*)\}\}}. \quad (10)$$

It will be shown later that, for the schemes developed in this paper, the user who attains the maximum is always unique, and thus, the optimum scheduling is always given by (10).

A close look at (7)–(9) and (10) reveals that the optimal management of resources depends only on the current channel realization \mathbf{h} and on the optimal Lagrange multipliers $\boldsymbol{\lambda}^*$. Although (7)–(9) are easy to derive and interpret, the scheduling in (10) deserves further elaboration. Equation (10) asserts that, based on *per-channel realization* \mathbf{h} , each channel k is uniquely assigned to the user index that maximizes the functional $f(m, k, \mathbf{h}, \boldsymbol{\lambda}^*)$; i.e., the scheduling follows a winner(s)-take-all strategy. With regard to ρ_m^* and π_m^* as prices for the rate and power, $f(m, k, \mathbf{h}, \boldsymbol{\lambda}^*)$ determines the net utility (rate reward minus power cost) of the m th user on channel k . This means that users with large ρ_m^* (or small π_m^*) are promoted for selection. Although, for a given \mathbf{h} , the winner is unique, the user who accesses the channel will vary as \mathbf{h} varies. Hence, long-term fairness is also accommodated. Last, in several practical systems (such as multicarrier systems), the channel gains in consecutive channels are expected to be very similar. This condition implies that the value of f would be similar for k and $k + 1$, and therefore, it is likely that the winner user in both channels is the same. Additional details on this issue are provided in Section VII, where different alternatives for reducing the signaling overhead required to implement our schemes are discussed. The winner-takes-all strategy (also known as max-weight scheduling) has been shown optimal for additional setups that deal with the orthogonal sharing of resources among users [13], [14], [16], [23], [31]. It is also worth noting that (7) dictates only the optimal value of the average arrival rate but not the specific distribution of the packet arrivals. In fact, from the point of view of (6), any arrival distribution with the average given by (7) is equally optimal. Nevertheless, the distribution of packet arrivals affects practical issues such as queue lengths or the delay packets experience. These issues and further conditions on the distribution of packet arrivals will be revisited when the stochastic algorithms are presented.

B. Estimating the Multipliers

In the previous section, the optimal resource allocation was characterized by the following two variables: 1) the current CSI \mathbf{h} and 2) the optimum Lagrange multipliers $\boldsymbol{\lambda}^*$. Various options are available to find $\boldsymbol{\lambda}^*$. Most options are iterative numerical methods that seek $\boldsymbol{\lambda}^*$ offline by capitalizing on the knowledge of the channel distribution [1, Ch. 4 and 6]. An

online approach is pursued here, under which the exact value of $\boldsymbol{\lambda}^*$ is never found. Instead, an estimate of $\boldsymbol{\lambda}^*$ is obtained using stochastic approximation iterations per time index n . This estimate, called $\boldsymbol{\lambda}_L[n]$ (the reason for using the subscript L will later be apparent), remains sufficiently close to $\boldsymbol{\lambda}^*$; for example, see [34]. The motivation behind this estimate is threefold: 1) The computational complexity of the stochastic schemes is lower than their offline counterparts; 2) stochastic schemes do not need to know the channel distribution and can cope with channel nonstationarities; and 3) connections between these multiplier estimates and queue lengths can be established, as will be shown in the next section.

Toward estimating the multipliers, let $a_m^*[\boldsymbol{\lambda}_L[n]]$ denote the instantaneous arrival of flow m during block n . The latter is a random variable that was drawn from a distribution with mean $\bar{a}_m^*(\boldsymbol{\lambda}_L[n])$, as given by (7). Moreover, with $x[n]$ denoting the value of a given variable x at time n , the sliding-window average of length L of variable x at time $n \geq L$ is² $W(x[n], L) := 1/L \sum_{l=n-L+1}^n x[l]$. Based on this average and with μ denoting a *constant* step size, the Lagrange multipliers are updated in (11)–(13), shown at the bottom of the page, where the primal variables, i.e., powers, rates, scheduling percentages, and flow rates, are found by substituting $\pi_{L,m}[n]$, $\rho_{L,m}[n]$, and $\alpha_{L,m}[n]$ into (7)–(10). Such primal will be referred to as stochastic primal variables.

Having described the stochastic schemes, it is appropriate to revisit the optimality of the winner-takes-all strategy in (10). It was stated in Section IV-A that the scheduling in (10) is, indeed, optimum, because for the schemes in this paper, the user m^* who maximizes $f(m, k, \mathbf{h}, \boldsymbol{\lambda}^*)$ is unique. To verify that this condition is true for the stochastic allocation, the user $m^*[n]$ who maximizes $f(m, k, \mathbf{h}[n], \boldsymbol{\lambda}_L[n])$ must be unique. Clearly, $f(m, k, \mathbf{h}[n], \boldsymbol{\lambda}_L[n])$ is a continuous random variable, because all terms that are involved in its definition are continuous random variables. Moreover, $f(m, k, \mathbf{h}[n], \boldsymbol{\lambda}_L[n])$ is uncorrelated across users, because, in practice, both arrivals and channel gains are independent across users. As a result, the probability of $f(m, k, \mathbf{h}[n], \boldsymbol{\lambda}_L[n]) = f(m', k, \mathbf{h}[n], \boldsymbol{\lambda}_L[n])$ has Lebesgue measure zero. For the winner to be not unique, two users

²For $n < L$, the sliding-window average is defined as $W(x[n], L) := 1/n \sum_{l=1}^n x[l]$. Although the definition for $n < L$ is needed for completeness, it will not be critical for the performance analysis, which holds because we will be interested in the long-term performance of our allocation schemes (more specifically, in the cumulative running average of the variables as $n \rightarrow \infty$). Clearly, if $x[n]$ is bounded, the behavior during the first $L - 1$ time slots (which is an interval of finite length) is not relevant for that purpose.

$$\pi_{L,m}[n+1] := \left[\pi_{L,m}[n] - \mu \left(\check{p}_m - W \left(\sum_{k=1}^K p_m^{k*}(\mathbf{h}[n], \boldsymbol{\lambda}_L[n]) w_m^{k*}(\mathbf{h}[n], \boldsymbol{\lambda}_L[n]), L \right) \right) \right]_0^\infty \quad (11)$$

$$\rho_{L,m}[n+1] := \left[\rho_{L,m}[n] + \mu \left(W \left(a_m^*[\boldsymbol{\lambda}_L[n]] - \sum_{k=1}^K r_m^{k*}(\mathbf{h}[n], \boldsymbol{\lambda}_L[n]) w_m^{k*}(\mathbf{h}[n], \boldsymbol{\lambda}_L[n]), L \right) \right) \right]_0^\infty \quad (12)$$

$$\alpha_{L,m}[n+1] := [\alpha_{L,m}[n] + \mu (\check{a}_m - W(a_m^*[\boldsymbol{\lambda}_L[n]], L))]_0^\infty \quad (13)$$

must attain the same value of f ; therefore, the previous finding readily implies that the winner is unique w.p.1.

The convergence, feasibility, and optimality of the stochastic schemes will rely on the fact that (11) and (13) are finite-length averages of stochastic subgradients of the dual function in (6); for example, see [1, Ch. 6]. Employing unbiased stochastic subgradients of the dual function for allocating resources in wireless fading networks has received some attention in recent years [24], [25], [28], [33], [34]. The focus and novelty here are on algorithms that use constant step size and combine (average) the L most recent stochastic subgradients. These two features will allow us to establish connections between some of the iterations in (11) and the (window average) length of the system queues. To be specific, assuming that the updates in (11)–(13) are bounded, the following result guarantees the feasibility and near optimality of the stochastic allocation iterates.

Proposition 1: The sample average of the stochastic resource allocation (i) is feasible and (ii) incurs arbitrarily small loss in performance relative to the average nonstochastic solution of the problem in (6). In particular, with $\delta_{\mathbf{P}}(\mu)$ denoting a positive number proportional to the step size μ and \mathbf{P}^* denoting the optimal value for the objective in (6), as $n \rightarrow \infty$, it holds w.p.1. that

i)

$$\begin{aligned} & \frac{1}{n} \sum_{l=1}^n \left(\sum_{k=1}^K r_m^{k*}(\mathbf{h}[l], \boldsymbol{\lambda}_L[l]) w_m^{k*}(\mathbf{h}[l], \boldsymbol{\lambda}_L[l]) \right) \\ & \geq \frac{1}{n} \sum_{l=1}^n a_m^*[\boldsymbol{\lambda}_L[l]] \geq \check{a}_m \end{aligned} \quad (14a)$$

$$\begin{aligned} & \frac{1}{n} \sum_{l=1}^n \left(\sum_{k=1}^K p_m^{k*}(\mathbf{h}[l], \boldsymbol{\lambda}_L[l]) w_m^{k*}(\mathbf{h}[l], \boldsymbol{\lambda}_L[l]) \right) \\ & \leq \check{p}_m. \end{aligned} \quad (14b)$$

ii)

$$\sum_m U_m \left(\frac{1}{n} \sum_{l=1}^n a_m^*[\boldsymbol{\lambda}_L[l]] \right) \geq \mathbf{P}^* - \delta_{\mathbf{P}}(\mu). \quad (15)$$

In other words, Proposition 1 guarantees the asymptotic optimality of the stochastic iterates, because they are feasible and achieve a value (performance) that is arbitrarily close to \mathbf{P}^* , which is the optimal value of the original (nonstochastic) solution of (6). The proof of the proposition is given in Appendix B and relies on the convergence of stochastic (epsilon) subgradient methods. For simplicity, the proofs are given for convergence in probability, but the extension to convergence w.p.1 is straightforward using the arguments in [2, Ch. 8], [3], and [28]. Moreover, the proof shows that, under very mild assumptions, the first inequality in (14a) and the inequality in (14b) hold with equality.

Equally relevant, convergence results can also be obtained when the cumulative running average in Proposition 1 is replaced by a finite-size sliding-window averaging or with an exponentially decaying window averaging. For these cases, convergence in the distribution of the modified left-hand side

of (15) to a Gaussian whose mean is larger than the right-hand side of (15) can be shown. A rigorous proof can be obtained following the methodology in [12, Ch. 11].

C. Relating Multipliers With Queue Lengths

On top of being optimal, the stochastic schemes are also meaningful, because they reveal the cross-layer attribute of our algorithm. To see how, temporarily set $L = 1$ and compare (12) to (4). It is clear that $\rho_m[n]$ and $q_m[n]$ are related in a way such that the stochastic Lagrange multipliers can be interpreted as a scaled version of the queue sizes. A very similar result is also true for $L > 1$. For this case, the following proposition holds.

Proposition 2: If there exists n_0 such that $\rho_{L,m}[n] > 0$ and $q_m[n] > 0$ for $n \geq n_0$, then $\rho_{L,m}[n] = \mu W(q_m[n], L) + \delta_m(n_0)$, where $\delta_m(n_0) := \rho_{L,m}[n_0] - \mu q_m[n_0]$.

Proof: The proof for $L = 1$ is straightforward, because the only difference between (12) and (4) is the step size μ . To prove the equivalence for $L > 1$, we will view $\rho_{L,m}[n]$ as the output $y[n]$ of a discrete linear time-invariant filter with input $x[n] = a_m^*[n] - \sum_k r_m^{k*}(\mathbf{h}[n]) w_m^{k*}(\mathbf{h}[n])$, where the dependence of variables on $\boldsymbol{\lambda}_L[n]$ has been dropped for notational brevity. Upon examining (12) and with \star denoting convolution, $y[n]$ can be written as $y[n] = h_{IIR}[n] \star (h_{FIR}[n] \star (\mu x[n]))$. In the latter expression, $h_{FIR}[n]$ implements the window averaging of length L , whereas $h_{IIR}[n]$ implements an autoregressive filter of order one, i.e., if $y_{IIR}[n]$ and $x_{IIR}[n]$ denote, correspondingly, the output and the input of filter $h_{IIR}[n]$, then $y_{IIR}[n] = y_{IIR}[n-1] + x_{IIR}[n]$. Using properties of the convolution operator, $y[n] = h_{IIR}[n] \star (h_{FIR}[n] \star (\mu x[n]))$ can be rewritten as $y[n] = \mu (h_{FIR}[n] \star (h_{IIR}[n] \star x[n]))$. Relying on the expression in (4) and taking into account that $x[n] = a_m^*[n] - \sum_k r_m^{k*}(\mathbf{h}[n]) w_m^{k*}(\mathbf{h}[n])$, it readily follows that $q_m[n] = (h_{IIR}[n] \star x[n])$. For this condition to be true, the projection operator in (4) needs to be transparent, which is the reason for requiring the queues not to be empty.³ Substituting the latter expression into $y[n] = \mu (h_{FIR}[n] \star (h_{IIR}[n] \star x[n]))$ yields $y[n] = \mu (h_{FIR}[n] \star q_m[n])$. Because $h_{FIR}[n]$ implements the window averaging and $y[n] = \rho_{L,m}[n]$, it follows that $\rho_{L,m}[n] = \mu W(q_m[n], L)$, which is the claim of the proposition. ■

As aforementioned, the result of Proposition 2 is meaningful, because it shows one way in which queues can be used to allocate resources in the network. In fact, Propositions 1 and 2 establish that, if ρ_m^* in (7)–(10) is approximated by $\mu W(q_m[n], L)$, the resulting stochastic resource allocation is optimum, as long as μ is sufficiently small [cf. (15)]. Although the term $\delta_m(n_0)$ introduces a discrepancy between $\rho_{L,m}[n]$ and $\mu W(q_m[n], L)$, this is not a concern for the following two reasons: 1) In practice, $\delta_m(n_0) \approx 0$, and 2) from an implementation perspective, the value of $\rho_{L,m}[n]$ at instant n_0 can always be redefined as $\rho_{L,m}[n_0] := \rho_{L,m}[n_0] - \delta_m(n_0)$ without affecting the long-term performance of the algorithm. Similarly, the presence of n_0 is not critical from a practical

³Intuitively, this assumption is reasonable, because the optimization in (6) aims at obtaining as high arrival rates as possible, which require the system to work close to its saturation point.

perspective either and has been assumed to simplify the proof. Nevertheless, exhaustive simulations have shown that a nonzero n_0 is present for small-medium values of L and μ . As an alternative, the iterations in (11)–(13) can slightly be modified so that the correspondence $\rho_{L,m}[n] = \mu W(q_m[n], L)$ holds $\forall n$. Although the proof is a bit more tedious, it can be shown that Proposition 1 also holds in this case.

In addition to revealing how QSI can optimally be used for allocating resources, Proposition 2 is also relevant for the following three extra reasons: 1) to analyze the stability of resource allocation algorithms; 2) to estimate the queuing delay from which packets will suffer from; and 3) to establish connections with other well-known cross-layer resource allocation algorithms such as the dynamic backpressure algorithm generalizations reported in [8]. The next section deals with reasons 1 and 2. With regard to reason 3, the stochastic iterates with $L = 1$ allocate resources in a manner related to those in [8] and [36] as follows: the role of $\rho_L[n]$ is played by the actual queues, the roles of $\pi_L[n]$ and $\alpha_L[n]$ are played by the so-called “virtual” queues, and some tuning parameters in [8] correspond to the inverse of the step size μ . Nevertheless, the results presented so far are distinct from [8] not only because different channel realizations are used to update the multipliers (because $L > 1$) but also because the analytical framework (Lyapunov optimization versus stochastic convex optimization), the CSI model, and the convergence claims are all different. Interestingly, works such as [36] also established relationships between the queue lengths and the optimal value of the multipliers of a related optimization problem. More specific details on this issue will be given in the next section after presenting the results on the queue stability of our schemes.

V. QUEUE STABILITY AND DELAY ANALYSIS

The previous section established that, although queue dynamics have not explicitly been accounted for in the formulation of (6), they naturally emerge as *scaled* stochastic estimates of the Lagrange multipliers associated with (5). In this section, the queue stability of the developed resource allocation algorithms is first characterized. Next, estimates of the average queuing delay that the flows experience are obtained and used to outline a mechanism for effecting delay priorities among users.

The stability analysis of the stochastic resource allocation in (7)–(13) relies on Proposition 2 to establish the following result with regard to the convergence of queue lengths.

Proposition 3: If $\bar{q}_m[n] := 1/n \sum_{l=1}^n q_m[l]$ denotes the sample average of the queue length $q_m[n]$ and $\delta_q(\mu)$ is a finite number that satisfies $\delta_q(\mu) \rightarrow 0$ as $\mu \rightarrow 0$, it then holds that

$$|\mu \bar{q}_m[n] - \rho_m^*| < \delta_q(\mu) \text{ as } n \rightarrow \infty \text{ w.p.1.} \quad (16)$$

As a corollary, (16) yields $\lim_{n \rightarrow \infty} \bar{q}_m[n] < \infty$, provided that $\rho_m^* < \infty$, which implies that the iterations are stable, as long as the original problem in (6) is feasible. The proof is given in Appendix B2. Note that Proposition 3 focuses on the convergence of the sample average of the queues and not on the convergence of the instantaneous multipliers. This is the case, because we are interested in the long-term queue stability and

not in issues such as queue length (or delay) moments. As simulations will confirm, the time trajectory of the stochastic estimates consists of two phases. During the first phase, the multipliers will move from the initialization point toward the optimum value. Grossly speaking, the speed of convergence is expected to be linear, because our algorithms are modified versions of a stochastic first-order (gradient) iteration. During the second phase, the estimates will hover around their optimum value. The estimates will not converge to a fixed point (because they are continuously updated based on a continuous random variable), and the “hovering noise” will be proportional to the step size considered.

Before moving to the next section, note that, for an algorithm related to the one proposed in this paper, a relationship between the *expected* queue length and the optimum value of a Lagrange multiplier has also been established in [36]. In particular, [36] relies on adaptive control tools to propose a cross-layer algorithm that relies on CSI and QSI to allocate resources. The authors then analyze the performance that such an algorithm achieves and show that the expected queue length approximately corresponds to the value of a Lagrange multiplier of a deterministic optimization problem.

A. Average Delay and Delay Priorities

The relationship between queues and Lagrange multipliers can also be leveraged to estimate the average queuing delay of the proposed stochastic resource allocation. To this end, Little’s result asserts that, with stable queues, the long-term average delay is given by the long-term average queue length divided by the long-term average arrival rate, which means that the delay of a specific flow is $\bar{d}_m = \bar{q}_m / \bar{a}_m$. Using Proposition 3, it follows that $|\mu \bar{d}_m - \rho_m^* / (\mu \bar{a}_m^*)| < \delta_q(\mu) / \bar{a}_m$ w.p.1. This case readily implies that the average delay for the stochastic resource algorithms developed in this paper can be *approximated* as

$$\bar{d}_m \approx \rho_m^* / (\mu \bar{a}_m^*) \quad \forall m. \quad (17)$$

In other words, the average delay of our stochastic algorithm can be estimated based on the optimal solution of (6) and the step size of the proposed iterations [24]. The KKT conditions associated with (6) can readily be used to show that $\rho_m^* = \dot{U}_m(\bar{a}_m^*) + \alpha_m^*$. Substituting the latter expression into (17), it follows that the delay can be written as the sum of two non-negative terms $\bar{d}_m \approx \dot{U}_m(\bar{a}_m^*) / (\mu \bar{a}_m^*) + \alpha_m^* / (\mu \bar{a}_m^*)$. Clearly, the first term depends only on the exogenous arrival rates \bar{a}_m^* . This condition implies that the arrival rates can be used to estimate the value of the average delay for a given flow (when the minimum rate constraint of such a flow is not active and, thus, $\alpha_m^* = 0$) or a lower bound on that value (when the minimum rate constraint of the flow is active and, thus, $\alpha_m^* > 0$).

Upon examining (17), it is also apparent that changes in the step size induce changes in the average delay. In particular, the larger the step size, the smaller the average queuing delay. The intuition is that large step sizes accelerate convergence and improve the ability of the iterates to react against events that would otherwise increase the queuing delay. However, large

step size values also lead to more severe hovering in the dual domain and, thus, more pronounced loss of optimality. Equally interesting, (17) can also be used to effect different delay priorities. Key for this purpose is the fact that the iterations in (11)–(13) converge not only if the step size is the same for all entries of λ but if it is different for each entry as well. This way, flows (users) with stricter delay constraints can employ a larger step size. In other words, by allowing the step size to be user dependent, it is possible to control the average delay performance of individual users.

VI. DOWNLINK SETUP

In this section, we briefly elaborate on the generalization of the developed schemes for use in a downlink setup (fading broadcast channel). The first step to address the optimal design is to reinterpret some of the notation. In particular, in this section, the entries of \mathbf{h} represent the channel from the AP (transmitter) to the users (receivers), \bar{a}_m is the average exogenous rate at the AP destined for user m , and $p_m^k(\mathbf{h})$ and $r_m^k(\mathbf{h})$ are the power and rate transmitted by the AP to user m over channel k . With \check{p} denoting the average power budget of the AP, the second step is to replace the M individual power constraints in (3) with the single constraint

$$\sum_{k,m} p_m^k(\mathbf{h}) w_m^k(\mathbf{h}) \leq \check{p}. \quad (18)$$

Now, we are ready to formulate the new optimization problem, which gives rise to the optimal cross-layer allocation for the downlink setup, i.e.,

$$\mathbf{P}^* := \max_{\bar{a}_m} \sum_{m=1}^M U_m(\bar{a}_m) \quad (19a)$$

$$\text{subject to } \bar{a}_m \geq \check{a}_m \forall m \quad (19b)$$

$$(1), (2), (5), (18). \quad (19c)$$

Similar to (6), the problem in (19) can be transformed into a convex problem and solved in the dual domain. The main difference is that, instead of M multipliers $\{\pi_m\}_{m=1}^M$, now, only a single multiplier π is involved. Clearly, $\{\pi_m\}_{m=1}^M$ were associated with the M constraints in (3), whereas π is associated with the single constraint in (18). With this change, the optimal allocation of resources for the broadcast channel is

$$\bar{a}_m^*(\lambda^*) = \left[(\dot{U}_m)^{-1} (\rho_m^* - \alpha_m^*) \right]_0^\infty \quad (20)$$

$$p_m^{k*}(\mathbf{h}, \lambda^*) = \left[\left(\dot{C}_m^k \right)^{-1} (\mathbf{h}, \pi^* / \rho_m^*) \right]_0^{\check{p}^k} \quad (21)$$

$$r_m^{k*}(\mathbf{h}, \lambda^*) = C_m^k(\mathbf{h}, p_m^{k*}(\mathbf{h}, \lambda^*)) \quad (22)$$

$$w_m^{k*}(\mathbf{h}, \lambda^*) = \mathbb{1}_{\{m=\arg \max_{m'} \{f(m', k, \mathbf{h}, \lambda^*)\}\}} \quad (23)$$

where $f_{DL}(m, k, \mathbf{h}, \lambda^*) := \rho_m^* r_m^{k*}(\mathbf{h}, \lambda^*) - \pi^* p_m^{k*}(\mathbf{h}, \lambda^*)$. Comparing (20)–(23) with (7)–(10), we indeed observe changes in (8)–(10), which are the resource allocation variables that are dependent on π_m^* .

With regard to the stochastic iterates, the only modification required is to replace the M updates in (11) with the following single update, in (24), shown at the bottom of the page.

The results in Propositions 1–3, as well as the results in (17), also hold for this case. Together with the changes in the formulation, consideration of a downlink setup entails changes in the signaling schemes and feedback overhead. This issue is discussed in more detail in the next section.

VII. SIGNALING AND COMPUTATIONAL OVERHEAD

So far, algorithms that yield the optimal allocation of resources as a function of several variables have been developed, and their performance has been analyzed. This section briefly discusses the signaling (and computational) overhead required to implement the developed algorithms and points out means of alleviating it. We assume that the algorithms are run in a centralized fashion so that the AP (which acts as a scheduler) carries out most of the required tasks. The operation of the central scheduler proceeds into the following four steps (phases):

- S1: gathering of the information required to find the optimum resource allocation;
- S2: calculation of the optimum resource allocation;
- S3: notification of the resulting allocation to the transmitters;
- S4: update of the corresponding multipliers.

Carrying out these tasks is challenging, because the variables involved are available at different locations. For example, the channel gains $\mathbf{h}[n]$ are available at the receiver side through training but not at the transmitter side (unless channels are reciprocal). On the other hand, the instantaneous values of the queues $\{q_m[n]\}_{m=1}^M$ are only known at the transmitter side, because they depend on the instantaneous arrivals. If convenient for reducing the signaling overhead, the multipliers can become available at both sides, provided that the corresponding stochastic iterations are run both at the transmitter and the receiver. The overhead associated with each of the aforementioned steps depends on the system setup (downlink versus uplink); for this reason, the two setups are separately analyzed.

Consider first the downlink case. To find the optimum resource allocation, the AP needs to know $(q_m[n], \alpha_{L,m}[n]) \forall m$, $\pi_L[n]$, and $\mathbf{h}[n]$. In this setup, the AP acts as a transmitter so that $q_m[n]$, $\alpha_{L,m}[n]$, and $\pi_L[n]$ are available at the AP, and the only problem is to get $\mathbf{h}[n]$. In time-division duplexing systems, uplink and downlink channels are reciprocal. Hence, the AP can acquire $\mathbf{h}[n]$ by estimating the channel in the reverse link.

$$\pi_L[n+1] = \left[\pi_L[n] - \mu W \left(\check{p} - \sum_{k,m} p_m^{k*}(\mathbf{h}[n], \lambda_L[n]) w_m^{k*}(\mathbf{h}[n], \lambda_L[n]), L \right) \right]_0^\infty \quad (24)$$

However, in frequency-division duplexing systems, the forward and reverse channels are nonreciprocal, and therefore, $\mathbf{h}[n]$ is not available at the AP. In this case, during S1, every user has to send the estimation of its own channels to the AP. If the channel varies sufficiently slow and the number of channels is not very high, the users could send an analog estimation of their channels. However, if the channel varies very fast or the number of channels is very high, the feedback rate needs to be reduced. Effective methods of feeding back the values of $\mathbf{h}[n]$ are to rely on quantized CSI [20], [22] and exploit correlation among channel gains (e.g., by grouping channels [19]). S2 is locally run at the AP and only entails computational overhead. For the schemes developed in this paper, the required complexity to carry out such an operation is affordable. In particular, finding the optimum arrival flows, powers, and rates requires the evaluation of closed-form expressions [cf. (7)–(9)]. Finding the optimum scheduling requires the evaluation of the closed-form functional $f(m, k, \mathbf{h}[n], \boldsymbol{\lambda}_L[n])$ for every user m and channel k and computing a maximum per channel [cf. (10)]. During S3, the scheduler needs to notify the optimum allocation to the terminals. This approach requires broadcasting the index of the winner user in each of the (group of) channels. If the users can detect the modulation/coding mode that the AP is using, no further action is required. If not, the AP needs to notify the modulation/coding scheme to be used. If only a finite set of adaptive modulation and coding modes is used (which is the typical case in quantized CSI systems), then only the index of the corresponding mode has to be identified. For this setup, users do not need to know any multipliers; therefore, S4 is run only at the AP, and no feedback is required. With regard to the computational complexity, S4 requires running only one iteration of the updates in (11)–(13), which is trivial, because the optimum allocation has been found.

The uplink case is more challenging. In this case, we can safely assume that $\mathbf{h}[n]$, $\{\alpha_{L,m}[n]\}_{m=1}^M$, and $\{\pi_{L,m}[n]\}_{m=1}^M$ are available at the AP (the AP acts as receiver, and therefore, the channel values are known through training). However, the values of $\{q_m[n]\}_{m=1}^M$ are known only by the users; therefore, during S1, such values have to be fed back to the AP. If the duration of a slot n is long enough, it is reasonable to assume that every user can send its own $q_m[n]$ to the AP through a signaling (feedback) control channel. If not, schemes for reducing the signaling are required. As in the case of the CSI, quantizing the queue length or exploiting the correlation that the queue lengths exhibit across time is an alternative that is worth exploring. A simpler alternative consists of sending the (window average) queue length only one out of L slots. Clearly, this approach would reduce the feedback rate by a factor of L . The proofs in Appendix B can be adapted to show that the optimality/convergence results in Proposition 1 would also hold for such a case. The computational complexity to run S2 is the same as in the downlink setup. During S3, the AP broadcasts the index of the winner user in each of the channels. Moreover, the power and rate that the winner users need to load also have to be broadcast. For fast-fading channels, the most reasonable alternative for reducing this overhead is to use quantized CSI. In such a case, the AP only needs to identify the index of the rate–power pair to be used.

We close this section by pointing out that lines of research that are worth looking at are developing schemes that work in a fully distributed scenario, are robust to communication errors, and are suboptimal (imperfect) but require much lower signaling overhead [16]. Although all these lines are certainly of interest, they go beyond the scope of this paper and are left as future work.

VIII. NUMERICAL RESULTS

In this section, numerical results are presented to assess the analytical claims. Unless explicitly mentioned, the default simulation setup will be an orthogonal frequency-division multiplexing (OFDM) uplink system with $M = 4$ users and $K = 64$ parallel channels. Channels are time and frequency selective, and the corresponding signal-to-noise ratios (SNRs) follow an exponential distribution (Rayleigh fading channel model). The average power gains for users 1–4 are 6, 4.5, 3, and 1.5 dB, respectively. Arrivals are randomly drawn from a binomial distribution. Simulations are run by assuming that the signaling channels were ideal; therefore, both CSI and QCI are assumed to be perfectly known.

Test Case 1: Optimality and Feasibility: First, we illustrate that the developed algorithms are optimal and feasible, and we compare their performance with existing alternatives. Values of the average transmit power, transmit rate, sum rate, and sum utility for the five ASs tested are given as follows.

- AS1: Utilities are linear (therefore, the objective corresponds to the sum rate), and there are no rate restrictions.
- AS2: The utility functions are $U_m(a) = \log(1 + a) \forall m$, and there are no rate restrictions.
- AS3: This AS is similar to AS2, but with $\tilde{a}_1 = 300$.
- AS4: This AS is a simplified version of AS2, but with fixed subcarrier allocation.
- AS5: This algorithm is compatible with the Worldwide Interoperability for Microwave Access (WiMAX) standard, where the subcarrier and power allocation are fixed, and the SNR per subcarrier is used to select the transmit rate from a codebook with eight prespecified values.

In all five schemes, the average power constraints are $\check{\mathbf{p}} := [\check{p}_1, \check{p}_2, \check{p}_3, \check{p}_4]^T = [40, 20, 12, 7]^T$. To reduce the number of schemes simulated, in this test case, we only simulate the nonstochastic version of our algorithms. Extensive simulations that assess the optimality loss of the stochastic algorithms relative to their stochastic counterparts will be carried out in test case 3.

The results listed in Table II illustrate that the proposed algorithm can provide fairness and guarantees on minimum average rate requirements, whereas it clearly outperforms other suboptimal schemes, which fail to satisfy these requirements. With regard to optimality, we see how indeed AS1 outperforms AS2 in terms of the sum rate (sum rate is the metric maximized by AS1) and how AS2 does better than AS1 in terms of the sum utility (sum utility is the metric maximized by AS2). Moreover, we see that, to satisfy the minimum rate requirement for user 1, AS3 needs to sacrifice some utility. (This case is because the unconstrained problem in AS2 does not fulfill the requirement \tilde{a}_1 , and therefore, the constraint needs to be activated.) Recall that the optimality of the nonstochastic

TABLE II
SUM UTILITY (HERE, $P^* := \sum_m U_m(\bar{a}_m)$) AND SUM RATE:
INDIVIDUAL AVERAGE RATES AND POWERS FOR DIFFERENT TEST CASES
(TRANSIENTS HAVE BEEN REMOVED WHEN FORMING AVERAGES). THE
AVERAGE POWER \bar{p} IS [40.0, 20.0, 12.0, 7.0] FOR ALL FIVE TEST CASES

	(AS1)	(AS2)	(AS3)	(AS4)	(AS5)
P^*	17.354	17.632	17.097	14.057	8.968
$\sum_m \bar{a}_m$	467.06	442.76	451.96	157.72	82.49
\bar{a}_1	282.46	239.81	300.00	76.14	53.39
\bar{a}_2	115.54	119.16	85.82	43.16	21.74
\bar{a}_3	49.36	58.31	45.00	25.09	6.68
\bar{a}_4	19.70	25.48	21.14	13.33	0.68
\bar{p}	[40.0, 20.0, 12.0, 7.0]				

schemes has theoretically been proved and the results here try to illustrate the gains with respect to (w.r.t.) suboptimal schemes. Finally, note that the small differences in terms of sum-utility values among AS1–AS3 are due to the specific setup selected. Had the utilities been different or had the rates been smaller, the difference in terms of the sum utility would have been larger. However, for the simulated configuration, the transmit rates are large enough, and therefore, the selected utility is basically flat.

Test Case 2: Dynamic Behavior of the Stochastic Schemes:

To gain further insight into the convergence of our algorithms, we analyze the dynamic behavior of the stochastic iterates. We focus on the test case in AS3, with $L = 1$ and $\mu = 10^{-4}$. Fig. 1 comprises six subplots, each depicting the evolution over time of a different subset of variables. Fig. 1(a) is the evolution of the average transmit rate $\bar{a}_m[n] := 1/n \sum_{l=1}^n a_m^*(\lambda_L[l])$, whereas Fig. 1(b) is the evolution of the average power consumption $\bar{p}_m[n] := 1/n \sum_{l=1}^n (\sum_k p_m^{k*}(\mathbf{h}[l], \lambda_L[l]) w_m^{k*}(\mathbf{h}[l], \lambda_L[l]))$. In Fig. 1(a) and (b), dashed lines correspond to the performance when the optimal multipliers are known, whereas continuous lines correspond to the stochastic algorithms proposed in this paper. Fig. 1(c), (d), and (f) plot the instantaneous value of the Lagrange multipliers $\rho_{L,m}[n]$, $\alpha_{L,m}[n]$, and $\pi_{L,m}[n]$, respectively. Finally, Fig. 1(e) depicts the time trajectory of the instantaneous value of the queue lengths $q_m[n]$.

Starting with the trajectories of the running averages $\bar{a}_m[n]$ and $\bar{p}_m[n]$ in Fig. 1(a) and (2), we observe that the considered constraints are satisfied (with equality) and the stochastic allocation converges in some hundred (correspondingly, a couple of thousand) iterations. The behavior of the Lagrange multipliers plotted in Fig. 1(c)–(e) is slightly different. Similar to Fig. 1(a) and (b), there is an initial phase during which the multipliers go from the initial point (which was randomly chosen) toward the optimal value. However, once the multipliers approach their optimal value, they do not strictly converge but hover around it. This case is not unreasonable, because $\rho_{L,m}[n]$, $\alpha_{L,m}[n]$, and $\pi_{L,m}[n]$ are instantaneous variables, whereas $\bar{a}_m[n]$ and $\bar{p}_m[n]$ are running averages. The numerical results also reveal that, for the simulated test case, $\alpha_{L,1}$ is nonzero. This case means that the rate constraint for $m = 1$ needs to be enforced, as mentioned in test case 1. We also observe that all $\pi_{L,m}$ and $\rho_{L,m}$ are also nonzero $\forall m$, confirming that the constraints in (3) and (5) are always active. Finally, the trajectories of queue lengths are depicted in Fig. 1(e). Note that the queues are stable and the relationship between $\rho_{L,m}[n]$ and $q_m[n]$ holds in practice.

Test Case 3: Changing the Step Size and the Window Length:

The numerical results for this test case are listed in Table III. The results correspond to a setup where the power gains for users 1–4 are 6, 5, 4, and 3 dB, respectively; the maximum average power constraints are $\check{\mathbf{p}} = [40, 20, 16, 10]^T$. The minimum average rate constraints in the upper half of the table are $\check{\mathbf{a}} := [\check{a}_1, \check{a}_2, \check{a}_3, \check{a}_4]^T = [0, 0, 0, 55]^T$, whereas the average rate constraints in the lower half of the table are $\check{\mathbf{a}} = [150, 0, 0, 0]^T$. All other parameters are set as in the default case. The values for μ and L vary, as indicated in the entries of the table. The numerical results reveal that larger step sizes give rise to larger gaps from the optimum. This trend is reasonable, because the bounds on the loss of optimality in Proposition 1 are proportional to μ . Note, however, that the bounds were derived using a worst case approach so that the actual loss does not have always to follow such a trend. The variation of the losses w.r.t. L is more difficult to explain. On the one hand, the bound on the optimality loss grows if the value of L increases (this case was not established in the statement of Proposition 1, but it was revealed in its proof). On the other hand, larger values of L render the stochastic updates of the multipliers closer to the original (nonstochastic) updates. This contradictory behavior may be the cause for the results in Table III. For $\mu = 10^{-5}$, larger values of L incur more sizeable losses; for $\mu = 10^{-4}$, the trend is the opposite; and for $\mu = 3 \cdot 10^{-5}$, the variation of the losses is not monotonic.

To illustrate the behavior of our algorithm in a scenario with a larger number of users, additional numerical results (that confirm the previous findings) are presented in Table IV. The simulation setup is similar to the previous setup, but in this case, there are $M = 8$ users instead of 4. The average power gains are $[6, 6, 4.5, 4.5, 3, 3, 1.5, 1.5]^T$, and the maximum average power constraints are $\check{\mathbf{p}} = [80, 40, 40, 20, 30, 15, 20, 10]^T$. The following two different rate requirements are considered: 1) $\check{\mathbf{a}} = [100, 0, 0, 0, 0, 0, 0, 0]^T$ for the configurations listed in Table 4, upper half, and 2) $\check{\mathbf{a}} = [0, 70, 0, 40, 0, 0, 0, 0]^T$ for the configurations listed in the lower half of Table IV.

Test Case 4: Evaluating the Average Delay: Finally, we analyze the queuing delay associated with the developed schemes. We consider a downlink setup with $M = 4$, $L = 1$, $K = 64$, $\check{\mathbf{a}} = [200, 80, 40, 30]^T$, and $\check{p} = 80$. The following six test cases are considered.

- ASa: our stochastic scheme with $\mu = 10^{-5}$ and $\lambda_L[0] = 0$;
- ASb: our stochastic scheme with $\mu = 10^{-6}$ and $\lambda_L[0] = 0$;
- ASc: the scheme in ASa, but with $\lambda_L[0] = 0.9\lambda^*$;
- ASd: the counterpart of ASc for ASb;
- ASe: the nonstochastic scheme, which assumes that λ^* is known and uses these values to adapt resources;
- ASf: a modified version of the nonstochastic scheme in ASe, where the average arrival rate $\bar{a}^*(\lambda^*)$ is slightly backed off so that the delay is decreased (in particular, simulations are run with $\bar{a} = 0.97\bar{a}^*(\lambda^*)$).

In all six cases, packet arrival rates follow a binomial distribution. The corresponding average queue length, average delay, and utility loss are listed in Table V.

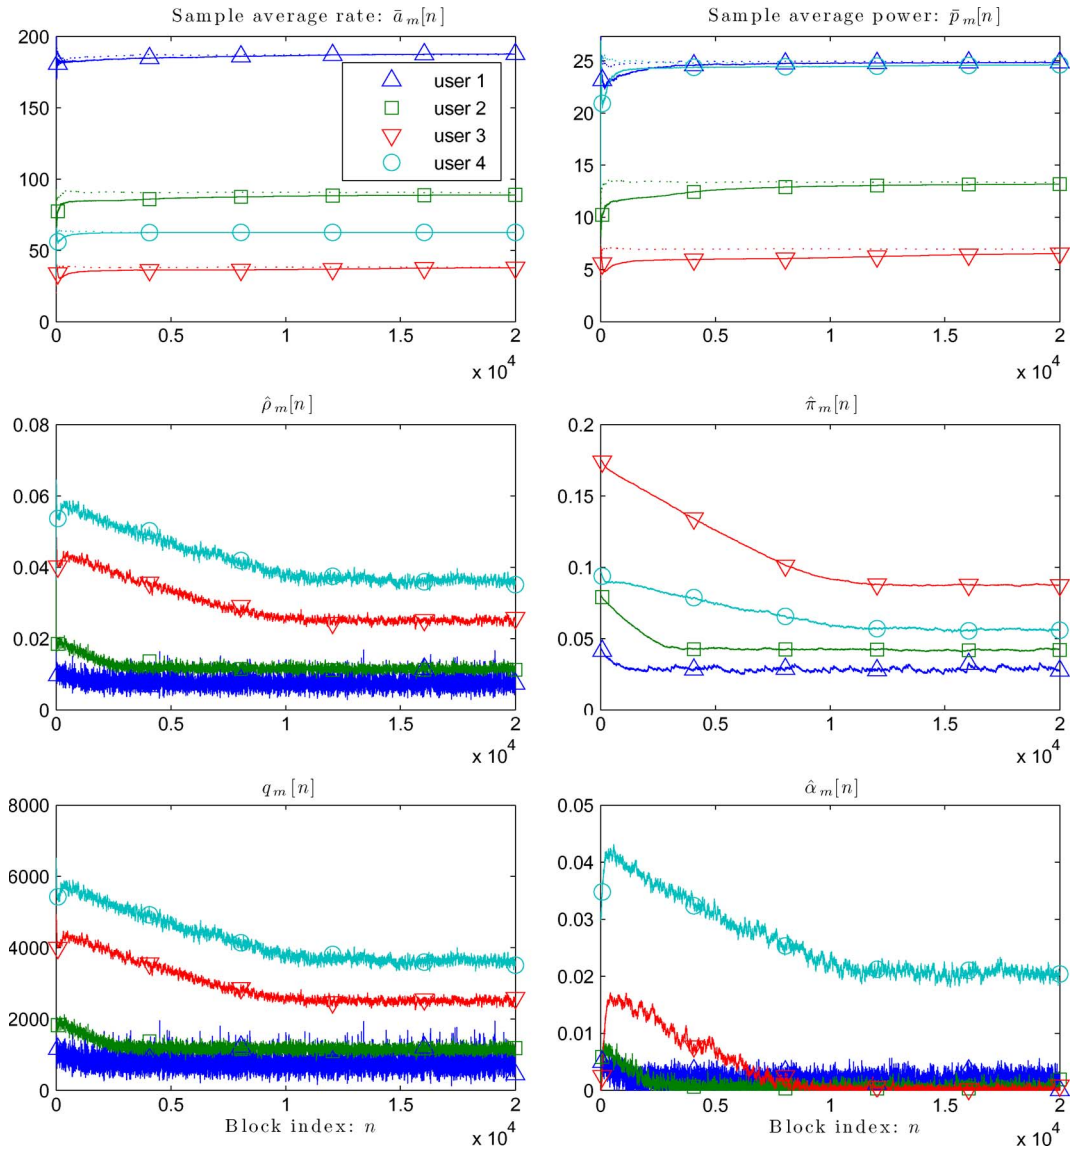


Fig. 1. Trajectories of different primal and dual variables in test case 2. (a) Sample average rate $\bar{a}_m[n]$. (b) Sample average power $\bar{p}_m[n]$. (c) Flow conservation multiplier $\hat{\rho}_m[n]$. (d) Estimate of the average power multiplier $\hat{\pi}_m[n]$. (e) Queue length $q_m[n]$. (f) Estimate of the average arrival rate multiplier $\hat{\alpha}_m[n]$.

The main observation from the numerical results in Table V is that the average delay is indeed inversely proportional to the step size. Moreover, the results also show that the delay of the stochastic schemes ASa–ASd is considerably lower than the nonstochastic optimal solution in ASe. In addition, the comparison of the simulated average delays for the schemes ASf and ASe reveals that the delay of the nonstochastic schemes can considerably be reduced, at the expense of a small loss in terms of optimality (sum utility). Differently, stochastic schemes ASc and ASd also achieve a delay reduction relative to their counterparts ASa and ASb but without incurring a sum-utility loss (relative to ASa and ASb). The delay improvement observed in stochastic schemes that use a “smart” initialization is in agreement with the results in [36]. In particular, [36] develops different algorithms that explicitly consider the instantaneous queue length and shows that, if the actual queue lengths are offset by a positive number (which is referred to as placeholder bits), the delay optimality tradeoffs considerably improve. The effect of such placeholder bits is very similar to the nonzero

initialization considered in ASc and ASd. Characterizing the optimality versus delay tradeoff for the nonstochastic schemes, their stochastic counterparts (relying on the results in [36]), and stochastic versus nonstochastic alternatives are interesting research lines but go beyond the scope of this paper and are left as future work.

To gain insight into the sensitivity of the average delay w.r.t. the step size, we run a last experiment. In this experiment, the value of the step size is modified online, and the effect on the size of the queues and the average delay is analyzed. In particular, the step size remains constant during a period of 10 000 time iterations, and then, it is reduced by a factor of 1/4. The simulation comprises five periods; therefore, the total simulated time is 50 000 instants. The corresponding results are plotted in Fig. 2, where Fig. 2(a) represents the queue lengths of all users, Fig. 2(b) depicts the delay experienced by the users, and Fig. 2(c) represents the product delay times step size. Results in Fig. 2(a) and (b) show that, every time the step size changes, the queue length and delay notably increase. However,

TABLE III
SUM UTILITY (P^*) AND SUM RATE FOR DIFFERENT
VALUES OF μ AND WINDOW SIZES

$\tilde{\mathbf{a}} = [0, 0, 0, 55]$				
CASE	Stepsize μ	Window size L	P^*	$\sum_m \bar{a}_m$
OFF-LINE	N/A	N/A	18.765	490.38
ONLINE	1.0E-5	1	18.722	483.87
		3	18.712	482.44
		8	18.571	464.98
	3.0E-5	1	18.424	435.13
		3	18.474	452.21
		8	18.432	453.49
	1.0E-4	1	17.592	337.50
		3	18.085	399.57
		8	18.371	449.02
$\tilde{\mathbf{a}} = [250, 0, 0, 0]$				
CASE	Stepsize μ	Window size L	P^*	$\sum_m \bar{a}_m$
OFF-LINE	N/A	N/A	18.784	513.21
ONLINE	1.0E-5	1	18.736	508.85
		3	18.662	503.79
		8	18.389	483.58
	3.0E-5	1	18.354	483.53
		3	18.271	477.65
		8	18.062	462.59
	1.0E-4	1	17.551	424.10
		3	17.871	451.05
		8	17.735	445.18

TABLE IV
SUM UTILITY (P^*) AND SUM RATE FOR DIFFERENT VALUES OF μ
AND WINDOW SIZES. $M = 8$ USERS, AND $K = 10$ CHANNELS

$\tilde{\mathbf{a}}_1 = 100$				
CASE	Stepsize μ	Window size L	P^*	$\sum_m \bar{a}_m$
OFF-LINE	N/A	N/A	24.095	211.20
ONLINE	3.0E-5	1	23.956	208.69
		4	23.951	208.64
	3.0E-6	4	24.046	210.44
		10	24.042	210.37
$\tilde{\mathbf{a}}_2 = 70; \tilde{\mathbf{a}}_4 = 40$				
CASE	Stepsize μ	Window size L	P^*	$\sum_m \bar{a}_m$
OFF-LINE	N/A	N/A	24.010	197.09
ONLINE	3.0E-5	1	23.845	193.74
		4	23.846	193.81
	3.0E-6	4	23.974	195.98
		10	23.967	195.85

TABLE V
QUEUE LENGTHS, AVERAGE DELAYS [cf. (17)], AND SUM UTILITY
(P^*) FOR THE ALLOCATION SCHEMES IN TEST CASE 4

	(ASa)	(ASb)	(ASc)	(ASd)	(ASe)	(ASf)
\bar{q}_1	960	7791	115	726	143608	463
\bar{q}_2	1051	9436	119	974	70832	333
\bar{q}_3	1480	13382	104	1449	46116	217
\bar{q}_4	2162	19543	220	2077	32603	203
\bar{d}_1	4.5	39.0	0.6	3.6	717.8	2.4
\bar{d}_2	10.6	89.8	1.1	9.3	673.6	3.3
\bar{d}_3	21.8	181.2	1.4	19.6	623.4	3.0
\bar{d}_4	47.0	389.5	4.5	41.3	649.1	4.2
P^*	18.050	18.218	18.153	18.220	18.222	18.101

after a transient period, the stochastic schemes can drive the queues to a stable level (note that this transient seems to also increase as the step size decreases). Moreover, the product delay times the step size is stabilized around a steady-state value, confirming that the average delay is inversely proportional to the step size (as predicted by Proposition 3).

TABLE VI
SUM UTILITY (P^*), SUM RATE, AND AVERAGE DELAYS [cf. (17)]
FOR THE ALLOCATION SCHEMES IN TEST CASE 5

	(A1)	(A2)	(A3)
P^*	23.954	24.077	24.055
$\sum_m \bar{a}_m$	208.65	210.89	210.78
$\sum_m \bar{d}$	1531	12314	1526

Test Case 5: Comparison With Other Optimal Resource Algorithms: The last test case is devoted to compare the performance of our algorithms with existing state-of-the-art alternatives. The metrics to be compared are the sum average delay and sum utility. The following three algorithms are compared:

- A1: our stochastic optimization algorithm with $L = 1$ and $\mu = 10^{-5}$;
- A2: a stochastic algorithm, with the time-diminishing step size similar to the step size proposed in [34] (the step size considered is $\mu[n] = \mu[0]/(1 + 100n/N)$, where $N = 50\,000$ is the total simulation time);
- A3: the optimal queue control algorithm proposed in [36].

Neither [34] nor [36] consider the exact same operating conditions as in the conditions in this paper; therefore the solutions proposed in these papers require some adaptation. Because [36] considers an unconstrained optimization problem, we will assume that the value of the *optimum* values of the multipliers associated with the average rate and power constraints (i.e., α_m^* and $\pi_m^* \forall m$) are perfectly known. To ensure fair comparison, the initialization of the queues (multipliers) in both algorithms is the same, and the parameter V in [36] is set to $V = 1/\mu$.

The scenario considered is the same as in the scenario considered in test case 3 ($M = 8$, and $\tilde{\mathbf{a}} = [100, 0, 0, 0, 0, 0, 0, 0]^T$). The results are listed in Table VI. The first observation is that A2 achieves the best sum-utility value. This result was expected, because the schemes in [34] are optimal as $n \rightarrow \infty$. On the other hand, A2 also exhibits the highest delay. This result was also expected, because the steady state of the algorithms in [34] is, in fact, the same as the optimal nonstochastic solution (recall that the results in test case 4 confirmed that the queues for the nonstochastic algorithm can unboundedly grow). Comparing A3 with A1, the results show that the sum-utility loss is slightly higher for our scheme whereas the achieved delays are similar. This result is also reasonable, because A3 capitalizes on the perfect knowledge of α_m^* and π_m^* (which amounts to having knowledge of the channel distribution), whereas A1 is a purely stochastic (online) algorithm.

IX. CONCLUSION

Cross-layer algorithms have been designed in this paper to allocate resources (flows, channel access, power, and rates) in a cellular system where users transmit over a set of orthogonal channels. Both uplink and downlink setups have been addressed. The developed resource allocation strategy depends on the instantaneous fading, queue lengths, and user-specific weights. The user-specific weights correspond to Lagrange multipliers and are estimated online using stochastic approximation iterations. Capitalizing on a relationship between




Fig. 2. Trajectories of (a) queue length, (b) average delay, and (c) product delay step sizes for a slightly varying step size μ . The variation of the step size corresponds to a (decreasing) piecewise constant function.

multiplier estimates and the window-averaged queue lengths, stability, and average delay performance were analyzed. Finally, a mechanism for effecting delay priorities among users by tuning the step size of individual users was also discussed.

APPENDIX A OPTIMAL SOLUTION FOR (6)

To solve for the optimal resource allocation, the constraints in (6) will be split into two groups. The first group is formed by (3), (5), and (6b). These average constraints are generally difficult to handle, and dual techniques will be used to deal with them. The second group is formed by (1) and (2) plus the nonnegative constraints that were not explicitly written in (6) but confine all $w_m^k(\mathbf{h})$, $p_m^k(\mathbf{h})$ and \bar{a}_m to be nonnegative. The constraints in the second group are easy to handle, and there is no need to dualize them.

To find the optimal solution for (6), we will do the following three tasks: 1) form the corresponding Lagrangian; 2) minimize the Lagrangian while guaranteeing that the constraints in the second group are satisfied; and 3) substitute for the optimum value of the Lagrange multipliers to recover the primal variables.

We start by introducing a notation to form the Lagrangian that is associated with the problem in (6). Let \mathbf{x} be a vector that contains all primal variables $(\bar{a}_m, w_m^k(\mathbf{h}), p_m^k(\mathbf{h}) \forall m, k, \mathbf{h})$. Note that \mathbf{x} has infinite length, because \mathbf{h} takes infinite values. Using these notational conventions, the Lagrangian is

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) = & \sum_{m=1}^M \left(-U_m(\bar{a}_m) + \alpha_m(\bar{a}_m - \check{a}_m) \right. \\ & + \rho_m \left(\mathbb{E}_{\mathbf{h}} \left[\bar{a}_m - \sum_{k=1}^K w_m^k(\mathbf{h}) C_m^k(\mathbf{h}, p_m^k(\mathbf{h})) \right] \right) \\ & \left. + \pi_m \left(\mathbb{E}_{\mathbf{h}} \left[\sum_{k=1}^K w_m^k(\mathbf{h}) p_m^k(\mathbf{h}) - \check{p}_m \right] \right) \right). \end{aligned} \quad (25)$$

For a given $\boldsymbol{\lambda}$, we need to minimize $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ w.r.t. \mathbf{x} . As it will be apparent next, the structure of $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ allows the minimization w.r.t. \bar{a}_m , $p_m^k(\mathbf{h})$, and $\{w_m^k(\mathbf{h})\}_{m=1}^M$ to be separately performed. Because $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ is strictly convex and differentiable w.r.t. \bar{a}_m and $p_m^k(\mathbf{h})$, minimization w.r.t. these variables

amounts to equating the corresponding partial derivatives of $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ to zero. Differently, the minimization w.r.t. $w_m^k(\mathbf{h})$ is slightly more complicated, mainly because $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ is linear in $w_m^k(\mathbf{h})$.

Consider first the optimum arrival flows. Differentiating $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ w.r.t. \bar{a}_m and equating the derivative to zero yields $\dot{U}_m(\bar{a}_m^*) + \alpha_m - \rho_m = 0$. Solving the latter expression w.r.t. \bar{a}_m yields

$$\bar{a}_m^*(\boldsymbol{\lambda}) = (\dot{U}_m)^{-1}(\rho_m - \alpha_m). \quad (26)$$

To guarantee that the latter expression satisfies the nonnegative constraints, we simply need to project \bar{a}_m^* onto the feasible set (nonnegative orthant). Thus, the solution is

$$\bar{a}_m^*(\boldsymbol{\lambda}) = \left[(\dot{U}_m)^{-1}(\rho_m - \alpha_m) \right]_0^\infty. \quad (27)$$

Proceeding similar to the optimum power allocation, we set the partial derivative of $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ w.r.t. $p_m^k(\mathbf{h})$ to zero. This approach yields $\rho_m w_m^k(\mathbf{h}) \dot{C}_m^k(\mathbf{h}, p_m^k(\mathbf{h})) - \pi_m w_m^k(\mathbf{h}) = 0$, which can be rewritten as $[\rho_m \dot{C}_m^k(\mathbf{h}, p_m^k(\mathbf{h})) - \pi_m] w_m^k(\mathbf{h}) = 0$. Clearly, the equality is satisfied if either of the following two conditions hold: 1) $w_m^k(\mathbf{h}) = 0$ or 2) $\rho_m \dot{C}_m^k(\mathbf{h}, p_m^k(\mathbf{h})) - \pi_m = 0$. Supposing that $w_m^k(\mathbf{h}) \neq 0$, condition 2 needs to hold. The root of condition 2 is $p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}) = (\dot{C}_m^k)^{-1}(\mathbf{h}, \pi_m / \rho_m)$. As aforementioned, we must guarantee that $p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda})$ is nonnegative and satisfies (3); hence, the optimal power allocation is

$$p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}) = \left[(\dot{C}_m^k)^{-1}(\mathbf{h}, \pi_m / \rho_m) \right]_0^{\check{p}^k}. \quad (28)$$

If $w_m^k(\mathbf{h}) = 0$, then any finite value of $p_m^k(\mathbf{h})$ is equally optimum. In fact, $w_m^k(\mathbf{h})$ being zero means that channel k is not assigned to user m ; therefore, the effective transmit power (and rate) is zero for any finite value of $p_m^k(\mathbf{h})$, including the one in (28). For this reason, the power allocation in (28) is optimum, regardless of the value $w_m^k(\mathbf{h})$. Interestingly, if $C_m^k(\cdot)$ corresponds to Shannon's capacity, the solution in (28) reduces to

$$p_m^k(\mathbf{h}, \boldsymbol{\lambda}) = \left[\frac{\rho_m}{\pi_m} - \frac{1}{h_m^k} \right]_0^{\check{p}^k} \quad (29)$$

which corresponds to the celebrated *waterfilling* solution (with ρ_m/π_m being the waterfilling level), which maximizes the sum capacity of a set of parallel orthogonal channels [5].

As aforementioned, minimization w.r.t. $w_m^k(\mathbf{h})$ has to more carefully be handled. The first reason is that $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ is linear in $w_m^k(\mathbf{h})$; hence, the optimal scheduling cannot be found by differentiating w.r.t. $w_m^k(\mathbf{h})$. The second reason is that (1), which was not dualized, couples the scheduling across users, and therefore, the mere projection in (27) and (28) cannot be used here. The third reason is that, to find the optimum $w_m^{k*}(\mathbf{h}, \boldsymbol{\lambda})$, the optimum values $\{p_{m'}^{k*}(\mathbf{h}, \boldsymbol{\lambda})\}_{m'=1}^M$ need to be known. To solve for $w_m^{k*}(\mathbf{h}, \boldsymbol{\lambda})$, we first note that the terms in $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ that depend on $w_m^k(\mathbf{h}, \boldsymbol{\lambda})$ are [cf. (25)]

$$-\sum_{m=1}^M \rho_m \mathbb{E}_{\mathbf{h}} \left[\sum_{k=1}^K w_m^k(\mathbf{h}) C_m^k(\mathbf{h}, p_m^k(\mathbf{h})) \right] + \sum_{m=1}^M \pi_m \mathbb{E}_{\mathbf{h}} \left[\sum_{k=1}^K w_m^k(\mathbf{h}) p_m^k(\mathbf{h}) \right]. \quad (30)$$

Substituting for the optimum $p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda})$, rearranging terms, and defining the functional $f(m, k, \mathbf{h}, \boldsymbol{\lambda}) := \rho_m C_m^k(\mathbf{h}, p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda})) - \pi_m p_m^{k*}(\mathbf{h}, \boldsymbol{\lambda})$ (cf. Section IV-A), it is possible to rewrite (30) as

$$\mathbb{E}_{\mathbf{h}} \left[\sum_{k=1}^K \left(- \sum_{m=1}^M w_m^k(\mathbf{h}) f(m, k, \mathbf{h}, \boldsymbol{\lambda}) \right) \right]. \quad (31)$$

Based on the previous expression, it readily follows that finding the optimal $w_m^k(\mathbf{h}, \boldsymbol{\lambda})$ that minimizes (25) amounts to solving, for each \mathbf{h} and k , the following problem:

$$\min_{\{w_m^k(\mathbf{h}, \boldsymbol{\lambda})\}_{m=1}^M} - \sum_{m=1}^M w_m^k(\mathbf{h}) f(m, k, \mathbf{h}, \boldsymbol{\lambda}) \quad (32a)$$

$$\text{s.t.} : \sum_{m=1}^M w_m^k(\mathbf{h}, \boldsymbol{\lambda}) \leq 1 \quad (32b)$$

$$w_m^k(\mathbf{h}, \boldsymbol{\lambda}) \geq 0 \quad \forall m \quad (32c)$$

where (32b) and (32c) are the constraints that involve $w_m^k(\mathbf{h})$ and were not dualized.

Assuming that $f(m, k, \mathbf{h}, \boldsymbol{\lambda})$ is nonnegative for at least one user,⁴ the solution of the previous problem is straightforward and consists of setting $w_m^k(\mathbf{h}, \boldsymbol{\lambda}) = 1$ for the user m that maximizes $f(m, k, \mathbf{h}, \boldsymbol{\lambda})$ while setting $w_m^k(\mathbf{h}, \boldsymbol{\lambda}) = 0$ for all other users. This policy can be written in closed form using the indicator function as

$$w_m^{k*}(\mathbf{h}, \boldsymbol{\lambda}) = \mathbb{1}_{\{m = \arg \max_{m'} \{f(m', k, \mathbf{h}, \boldsymbol{\lambda})\}\}}. \quad (33)$$

Substituting $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*$ into (27), (28), and (33) yields the optimal solution in (7), (8), and (10), respectively.

⁴The convexity of $C_m^k(\cdot)$ can be used to rigorously show that $f(m, k, \mathbf{h}, \boldsymbol{\lambda})$ is always nonnegative.

APPENDIX B CONVERGENCE AND OPTIMALITY OF STOCHASTIC SCHEMES

We begin by briefly reviewing some concepts from the optimization theory and introducing notation. Then, we present a result on the boundedness of Lagrange multipliers (Lemma 1) to be used in subsequent proofs. Using Lemma 1, we prove Proposition 1 and, subsequently, Proposition 3. For brevity, several of the proofs presented here guarantee only convergence in the mean (hence, in probability), but more sophisticated tools can be used to establish convergence w.p.1 [2, Sec. 8.2.2], [3], [28].

The definition of a stochastic subgradient is first reviewed. Let f be a scalar real-valued function of vector \mathbf{x} . Vector $\boldsymbol{\partial}(\mathbf{x}_0)$ is an unbiased stochastic subgradient of f at point \mathbf{x}_0 if and only if $f(\mathbf{x}) - f(\mathbf{x}_0) \geq \mathbb{E}[\boldsymbol{\partial}(\mathbf{x}_0)]^T(\mathbf{x} - \mathbf{x}_0)$ w.p.1. Going back to our problem, we will use $D(\boldsymbol{\lambda})$ to denote the dual function of (6) [1, Sec. 5.1.2]. Key to the subsequent proofs is the fact that the updates in (11)–(13) are finite-length averages of the stochastic subgradients of $D(\boldsymbol{\lambda})$. To be more precise, let $\boldsymbol{\partial}[n]$ denote a $3M \times 1$ vector whose i th entry is $\partial_i[n] := \check{p}_i - \sum_k p_i^{k*}(\mathbf{h}[n], \boldsymbol{\lambda}_L[n]) w_i^{k*}(\mathbf{h}[n], \boldsymbol{\lambda}_L[n])$ if $1 \leq i \leq M$, $\partial_i[n] := \sum_k r_{i-M}^{k*}(\mathbf{h}[n], \boldsymbol{\lambda}_L[n]) w_{i-M}^{k*}(\mathbf{h}[n], \boldsymbol{\lambda}_L[n]) - a_{i-M}^*[\boldsymbol{\lambda}_L[n]]$ if $M+1 \leq i \leq 2M$, and $\partial_i[n] := a_{i-2M}^*[\boldsymbol{\lambda}_L[n]] - \check{a}_{i-2M}$ if $2M+1 \leq i \leq 3M$. It is easy to verify that, for $L=1$, $\boldsymbol{\partial}[n]$ comprises the correction terms in (11)–(13); therefore, (11)–(13) can collectively be written as the following first-order vector recursion: $\boldsymbol{\lambda}_L[n+1] = [\boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}[n]]_0^\infty$. For the general case $L > 1$, define $\boldsymbol{\partial}_L[n] := W(\boldsymbol{\partial}[n], L)$ and write (11)–(13) as $\boldsymbol{\lambda}_L[n+1] = [\boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}_L[n]]_0^\infty$, where the (11)–(13) correction terms are window averages of $\boldsymbol{\partial}[n]$.

Using results from the optimization theory, it readily follows that $\boldsymbol{\partial}[n]$ is an unbiased stochastic gradient of $D(\boldsymbol{\lambda})$ at point $\boldsymbol{\lambda}_L[n]$ (cf. [1, Ch. 6]). This case implies, by definition, that $D(\boldsymbol{\lambda}) - D(\boldsymbol{\lambda}_L[n]) \geq \mathbb{E}[\boldsymbol{\partial}[n]]^T(\boldsymbol{\lambda}_L[n] - \boldsymbol{\lambda})$ w.p.1, which is an important property to be used later. Note, however, that $\boldsymbol{\partial}_L[n]$ for $L > 1$ may not be an *unbiased* stochastic subgradient of $D(\boldsymbol{\lambda})$. In fact, the subsequent proofs will reveal that $\boldsymbol{\partial}_L[n]$ is, in general, a *biased* estimate of the subgradient of $D(\boldsymbol{\lambda})$ and therefore results from the convergence of ϵ -subgradients that need to be used [2, Sec. 4.3]. This case will slightly complicate the convergence proofs. Finally, let $\tilde{\boldsymbol{\partial}}_L[n]$ denote the normalized multiplier error $\tilde{\boldsymbol{\partial}}_L[n] := (\boldsymbol{\lambda}_L[n+1] - \boldsymbol{\lambda}_L[n])/\mu$. In words, $\tilde{\boldsymbol{\partial}}_L[n] = \boldsymbol{\partial}_L[n]$ if $\boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}_L[n]$ is nonnegative, and therefore, the projection operator in $\boldsymbol{\lambda}_L[n+1] = [\boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}_L[n]]_0^\infty$ is not needed. Otherwise, the entries of $\boldsymbol{\partial}_L[n]$, which are large enough to render $\boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}_L[n]$ negative, are partially clipped in $\tilde{\boldsymbol{\partial}}_L[n]$. Based on this definition, we can write $\boldsymbol{\lambda}_L[n+1] = \boldsymbol{\lambda}_L[n] - \mu \tilde{\boldsymbol{\partial}}_L[n]$. This notation will be used at several points in the proofs to bypass the nonlinearity of the projection operator in $\boldsymbol{\lambda}_L[n+1] = [\boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}_L[n]]_0^\infty$.

Now, we are ready to present a result on the boundedness of the multipliers.

Lemma 1: If the problem in (6) is strictly feasible and the constraint set in (6b) and (6c) is bounded, then there exists a finite constant c such that $\mathbb{E}[\|\boldsymbol{\lambda}_L[n]\|] < c \forall n$.

This result is standard for dual (sub)gradient algorithms, and its proof will only be sketched here. The proof follows the steps in, for example, [26, Lemma 3], after writing $\boldsymbol{\partial}_L[n]$ as a sum of L terms and taking expectations. Intuitively speaking, this claim is reasonable because of the following three reasons: 1) Feasibility implies that the optimum multipliers $\boldsymbol{\lambda}^*$ are finite; 2) the subgradient iterations drive $\mathbb{E}[\boldsymbol{\lambda}_L[n]]$ closer to $\boldsymbol{\partial}^*$; and 3) the maximum correction term is bounded. Because convergence in the mean implies convergence in probability, it is guaranteed that $\|\boldsymbol{\lambda}_L[n]\|$ is bounded in probability. As aforementioned, the following appendices will prove convergence in probability but not convergence w.p.1. One of the key steps in showing convergence w.p.1 is to prove that $\|\boldsymbol{\lambda}_L[n]\| < c \forall n$ w.p. 1. This step can be done using a slight generalization of the supermartingale convergence theorem; see [2, Props. 8.2.10 and 8.2.11] and [2, Prop. 8.2.11] for details or [28] for a related argument.

Proof of Proposition 1:

Proof of Proposition 1.(i): Recall that the updates in (11)–(13) can be written in a single vector recursion as $\boldsymbol{\lambda}_L[n+1] = [\boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}_L[n]]_0^\infty$. Because the stochastic iteration is projected over the nonnegative orthant, it follows that $\boldsymbol{\lambda}_L[n+1] \geq \boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}_L[n]$. Applying this inequality recursively yields $\boldsymbol{\lambda}_L[n+1] \geq \boldsymbol{\lambda}_L[0] - \mu \sum_{l=1}^n \boldsymbol{\partial}_L[l]$. Dividing both sides by n and initializing with $\boldsymbol{\lambda}_L[0] \geq \mathbf{0}$ yields $(n^{-1})\boldsymbol{\lambda}_L[n+1] \geq -\mu(n^{-1}) \sum_{l=1}^n \boldsymbol{\partial}_L[l]$. Because the multipliers are bounded (cf. Lemma 1), as $n \rightarrow \infty$, the left-hand side of the last inequality goes to zero, which implies that $0 \geq -\lim_{n \rightarrow \infty} (n^{-1}) \sum_{l=1}^n \boldsymbol{\partial}_L[l]$. Now, take into account the following two facts: 1) in (6), all the constraints are convex (in fact, linear) and 2) the constraint violation averaged over the last L slots is represented by $-\boldsymbol{\partial}_L[l]$ (which follows from the definitions of $\boldsymbol{\partial}[l]$ and $\boldsymbol{\partial}_L[l]$). Based on these facts, $-\lim_{n \rightarrow \infty} (n^{-1}) \sum_{l=1}^n \boldsymbol{\partial}_L[l]$ amounts to the time-averaged constraint violation, and therefore, the inequality $0 \geq -\lim_{n \rightarrow \infty} (n^{-1}) \sum_{l=1}^n \boldsymbol{\partial}_L[l]$ implies that the stochastic solution is feasible, thus proving (i).

If it holds that, after a time instant n_0 , the multipliers of all active constraints are nonnegative, then the projection operator becomes transparent, which readily implies that $0 = \lim_{n \rightarrow \infty} (n - n_0)^{-1} \sum_{l=n_0}^n \boldsymbol{\partial}_L[l] = \lim_{n \rightarrow \infty} n^{-1} \sum_{l=1}^n \boldsymbol{\partial}_L[l]$. The latter expression together with the linearity of the constraints in (6) implies that the bounds in (14a) and (14b) associated with the *active constraints* are satisfied with equality. For the problem in (6), it is easy to verify that the flow conservation and average power constraints are always active. However, the constraints in (6b) may or may not be active, depending on individual user rate requirements and the system's operating conditions (average channel gain and power budgets).

Proof of proposition 1.(ii): To show the bound in (15), we need first to prove the following lemma.

Lemma 2: If the entries of $\boldsymbol{\partial}[n]$ are bounded, it holds that $\lim_{n \rightarrow \infty} 1/n \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L^T[l] \mathbb{E}[\boldsymbol{\partial}[l]]] \leq \mathcal{O}(\mu)$.

Proof: We first present the proof for $L = 1$. Based on this proof, we present the proof for $L = 2$, which can easily be generalized for $L > 2$.

The first step in proving the lemma for $L = 1$ consists of using the nonexpansive property of the projection operator $[\cdot]_0^\infty$

in $\boldsymbol{\lambda}_L[n+1] = [\boldsymbol{\lambda}_L[n] - \mu \boldsymbol{\partial}_L[n]]_0^\infty$ to write

$$\|\boldsymbol{\lambda}_L[n+1]\|^2 \leq \|\boldsymbol{\lambda}_L[n]\|^2 + \mu^2 \|\boldsymbol{\partial}_L[n]\|^2 - 2\mu \boldsymbol{\lambda}_L^T[n] \boldsymbol{\partial}_L[n]. \quad (34)$$

Supposing that the “path” (history) of the system up to instant n , i.e., $\boldsymbol{\lambda}_L[0], \dots, \boldsymbol{\lambda}_L[n]$, is given, consider taking expectations on both sides of (34) over the future of the path to obtain

$$\mathbb{E} \left[\|\boldsymbol{\lambda}_L[n+1]\|^2 \right] \leq \|\boldsymbol{\lambda}_L[n]\|^2 + \mu^2 \mathbb{E} \left[\|\boldsymbol{\partial}_L[n]\|^2 \right] - 2\mu \boldsymbol{\lambda}_L^T[n] \mathbb{E} [\boldsymbol{\partial}_L[n]]. \quad (35)$$

Now, taking expectations over any possible path $\boldsymbol{\lambda}_L[0], \dots, \boldsymbol{\lambda}_L[n]$ yields

$$\mathbb{E} \left[\|\boldsymbol{\lambda}_L[n+1]\|^2 \right] \leq \mathbb{E} \left[\|\boldsymbol{\lambda}_L[n]\|^2 \right] + \mu^2 \mathbb{E} \left[\|\boldsymbol{\partial}_L[n]\|^2 \right] - 2\mu \mathbb{E} \left[\boldsymbol{\lambda}_L^T[n] \mathbb{E} [\boldsymbol{\partial}_L[n]] \right]. \quad (36)$$

Successively applying the last inequality and with Δ_L^{\max} upper bounding $\mathbb{E}[\|\boldsymbol{\partial}_L[n]\|^2]$, we arrive at

$$\mathbb{E} \left[\|\boldsymbol{\lambda}_L[n+1]\|^2 \right] \leq \mathbb{E} \left[\|\boldsymbol{\lambda}_L[0]\|^2 \right] + \mu^2 \sum_{l=1}^n \Delta_L^{\max} - 2\mu \sum_{l=1}^n \mathbb{E} \left[\boldsymbol{\lambda}_L[l]^T \mathbb{E} [\boldsymbol{\partial}_L[l]] \right]. \quad (37)$$

Dividing by n and rearranging the terms yields

$$\frac{2\mu}{n} \sum_{l=1}^n \mathbb{E} \left[\boldsymbol{\lambda}_L[l]^T \mathbb{E} [\boldsymbol{\partial}_L[l]] \right] \leq \frac{\mu^2}{n} \sum_{l=1}^n \Delta_L^{\max} + \frac{1}{n} \left(\mathbb{E} \left[\|\boldsymbol{\lambda}_L[0]\|^2 \right] - \mathbb{E} \left[\|\boldsymbol{\lambda}_L[n+1]\|^2 \right] \right). \quad (38)$$

Because the multipliers are bounded (cf. Lemma 1), as $n \rightarrow \infty$, the last inequality reduces to

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mathbb{E} \left[\boldsymbol{\lambda}_L^T[l] \mathbb{E} [\boldsymbol{\partial}_L[l]] \right] \leq \frac{\Delta_L^{\max}}{2} \mu. \quad (39)$$

For $L = 1$, it holds by definition that $\boldsymbol{\partial}_L[l] = \boldsymbol{\partial}[l]$. Substituting $\boldsymbol{\partial}_L[l] = \boldsymbol{\partial}[l]$ into (39), the bound in Lemma 2 readily follows.

The proof for $L = 2$ will rely on the definition of $\boldsymbol{\partial}_L$ and the bound in (39), which holds, regardless of the value of L . Let us rewrite the left-hand side of (39) as follows:⁵

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mathbb{E} \left[\boldsymbol{\lambda}_L^T[l] \mathbb{E} [\boldsymbol{\partial}_L[l]] \right] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mathbb{E} \left[\boldsymbol{\lambda}_L^T[l] \left(\frac{1}{2} \mathbb{E} [\boldsymbol{\partial}[l-1]] + \frac{1}{2} \mathbb{E} [\boldsymbol{\partial}[l]] \right) \right] \\ &\stackrel{(a)}{=} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=2}^{n+1} \mathbb{E} \left[\frac{1}{2} (\boldsymbol{\lambda}_L[l-1] + \boldsymbol{\lambda}_L^T[l]) \mathbb{E} [\boldsymbol{\partial}[l-1]] \right] \end{aligned}$$

⁵To simplify the proof, we will not take into account that the window-averaging operator renders the definition of $\boldsymbol{\partial}_L[l]$ different for $l = 1$. As explained in Footnote 2, such a difference is not relevant when we look at the long-term average and take $n \rightarrow \infty$.

$$\begin{aligned}
&\stackrel{(b)}{=} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=2}^{n+1} \mathbb{E} \left[\frac{1}{2} \left(\boldsymbol{\lambda}_L[l-1] + \boldsymbol{\lambda}_L[l-1] - \mu \tilde{\boldsymbol{\theta}}_L[l-1] \right)^T \mathbb{E}[\boldsymbol{\theta}[l-1]] \right] \\
&\stackrel{(c)}{=} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L^T[l-1] \mathbb{E}[\boldsymbol{\theta}[l-1]]] \\
&\quad - \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mu \frac{1}{2} \mathbb{E}[\tilde{\boldsymbol{\theta}}_L^T[l-1] \mathbb{E}[\boldsymbol{\theta}[l-1]]] \quad (40)
\end{aligned}$$

where, in (40)(a), we have rearranged the terms; in (40)(b), we have used that $\boldsymbol{\lambda}_L[l] = \boldsymbol{\lambda}_L[l-1] - \mu \tilde{\boldsymbol{\theta}}_L[l-1]$ (cf. the definition of $\tilde{\boldsymbol{\theta}}_L[n]$); and in (40)(c), we have shifted the summation index. Substituting (40) into (39) yields

$$\begin{aligned}
&\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L^T[l-1] \mathbb{E}[\boldsymbol{\theta}[l-1]]] \\
&\quad - \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mu \frac{1}{2} \mathbb{E}[\tilde{\boldsymbol{\theta}}_L^T[l-1] \mathbb{E}[\boldsymbol{\theta}[l-1]]] \leq \frac{\Delta_L^{\max}}{2} \mu. \quad (41)
\end{aligned}$$

Because we have assumed that the multiplier updates are bounded, $|\mathbb{E}[\tilde{\boldsymbol{\theta}}_L^T[l-1] \mathbb{E}[\boldsymbol{\theta}[l-1]]]|$ is bounded by, for example, $\tilde{\Delta}^{\max}$. Upon substituting this bound into (41) and rearranging the terms, we have $\lim_{n \rightarrow \infty} (n^{-1}) \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L^T[l-1] \mathbb{E}[\boldsymbol{\theta}[l-1]]] \leq (\Delta_L^{\max}/2 + \tilde{\Delta}^{\max}/2)\mu$, which completes the proof of Lemma 2. ■

Note that, although the lemma assumed that the entries of $\boldsymbol{\theta}[n]$ are bounded, the proof relied on the fact that the second-order statistics of $\boldsymbol{\theta}[n]$ are bounded. Nevertheless, the assumption that the entries of $\boldsymbol{\theta}[n]$ are bounded does not entail a loss of generality from a practical perspective (the arrival rates and the transmit powers and rates are always bounded in real systems), and therefore, it was assumed through.

Having proved Lemma 2, we are ready to establish the bound in (15). Taking expectations on the left-hand side of (15), we have

$$\begin{aligned}
&\sum_{m=1}^M U_m \left(\frac{1}{n} \sum_{l=1}^n \mathbb{E}[a_m^*(\boldsymbol{\lambda}_L[l])] \right) \\
&\stackrel{(a)}{=} \sum_{m=1}^M U_m \left(\frac{1}{n} \sum_{l=1}^n \mathbb{E}[\bar{a}_m^*(\boldsymbol{\lambda}_L[l])] \right) \\
&\stackrel{(b)}{\geq} \frac{1}{n} \sum_{l=1}^n \mathbb{E} \left[\sum_{m=1}^M U_m (\bar{a}_m^*(\boldsymbol{\lambda}_L[l])) \right] \\
&\stackrel{(c)}{=} \frac{1}{n} \sum_{l=1}^n \mathbb{E} \left[\left(\sum_{m=1}^M U_m (\bar{a}_m^*(\boldsymbol{\lambda}_L[l])) \right) + \boldsymbol{\lambda}_L^T[l] \mathbb{E}[\boldsymbol{\theta}[l]] \right] \\
&\quad - \frac{1}{n} \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L^T[l] \mathbb{E}[\boldsymbol{\theta}[l]]]
\end{aligned}$$

$$\begin{aligned}
&\stackrel{(d)}{=} \frac{1}{n} \sum_{l=1}^n \mathbb{E}[D(\boldsymbol{\lambda}_L[l])] \\
&\quad - \frac{1}{n} \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L^T[l] \mathbb{E}[\boldsymbol{\theta}[l]]] \\
&\stackrel{(e)}{\geq} D(\boldsymbol{\lambda}^*) - \frac{1}{n} \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L^T[l] \mathbb{E}[\boldsymbol{\theta}[l]]] \quad (42)
\end{aligned}$$

where (a) is due to the definition of $a_m^*(\boldsymbol{\lambda}_L[l])$, (b) holds because $U_m(\cdot)$ are concave, (c) follows by adding and subtracting the same term, (d) is due to the definition of the dual function [1, Ch. 5], and (e) relies on the fact that $D(\boldsymbol{\lambda}_L[l]) \geq D(\boldsymbol{\lambda}^*)$, which holds, because the dual function is minimized at $\boldsymbol{\lambda}^*$.

Because the problem in (6) is convex and feasible, it has zero duality gap, and thus, $D(\boldsymbol{\lambda}^*) = \mathbf{P}^*$. After substituting $D(\boldsymbol{\lambda}^*) = \mathbf{P}^*$ into (42) and letting $n \rightarrow \infty$, it holds that

$$\begin{aligned}
&\sum_{m=1}^M U_m \left(\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mathbb{E}[a_m^*(\boldsymbol{\lambda}_L[l])] \right) \\
&\geq \mathbf{P}^* - \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L^T[l] \mathbb{E}[\boldsymbol{\theta}[l]]]. \quad (43)
\end{aligned}$$

Using the bound provided in Lemma 2, the statement in (ii) follows. ■

Proof of Proposition 3: The sketch of the proof proceeds in two steps. First, we show that the averaged iterates of the Lagrange multipliers converge to a neighborhood of the optimum Lagrange multipliers. With $\bar{\boldsymbol{\lambda}}_L := \lim_{n \rightarrow \infty} 1/n \sum_{l=1}^n \mathbb{E}[\boldsymbol{\lambda}_L[l]]$, this approach will be accomplished by capitalizing on the fact that the dual function $D(\boldsymbol{\lambda})$ is continuous and $D(\bar{\boldsymbol{\lambda}}_L) - D(\boldsymbol{\lambda}^*) \leq \delta_D(\mu)$, where $\delta_D(\mu)$ is proportional to μ . This case will be shown in Lemma 3, which is presented at the end of this appendix. Because the bound is proportional to μ , $\bar{\boldsymbol{\lambda}}_L$ belongs to the (level) set $\mathcal{B}_\lambda := \{\boldsymbol{\lambda} : D(\boldsymbol{\lambda}) - D(\boldsymbol{\lambda}^*) \leq \delta_D(\mu)\}$. Because $D(\boldsymbol{\lambda})$ is continuous, the size of \mathcal{B}_λ shrinks⁶ as $\mu \rightarrow 0$, and therefore, $\|\bar{\boldsymbol{\lambda}}_L - \boldsymbol{\lambda}^*\| \rightarrow 0$ as $\mu \rightarrow 0$. The second step consists of using Proposition 2 to write $\rho_{m,L}[n] = \mu W(q_m[n], L) + \delta_m(n_0)$. Using the first step, it then follows that $|\mu \lim_{v \rightarrow \infty} v^{-1} \sum_{n=1}^v \mathbb{E}[W(q_m[n], L)] - \rho_{m,L}^*| \rightarrow 0$ as $\mu \rightarrow 0$. Swapping the order of $\mathbb{E}[\cdot]$ and $W(\cdot)$, we deduce that $|\mu \lim_{v \rightarrow \infty} v^{-1} \sum_{n=1}^v \mathbb{E}[q_m[n]] - \rho_{m,L}^*| \rightarrow 0$ as $\mu \rightarrow 0$, which is the claim of the proposition.

Lemma 3: If the entries of $\boldsymbol{\theta}[n]$ are bounded, it holds that $D(\bar{\boldsymbol{\lambda}}_L) - D(\boldsymbol{\lambda}^*) \leq \mathcal{O}(\mu)$.

Proof: For brevity, we consider $L = 2$, but the proof can easily be extended for $L > 2$. Taking into account that

⁶Expressions (or bounds) for $\delta_D(\mu)$ and the size of \mathcal{B}_λ (thus, for the value of $\|\bar{\boldsymbol{\lambda}}_L - \boldsymbol{\lambda}^*\|$) can be derived; for example, for $L = 2$, $\delta_D(\mu)$ is expressed as in (52). With regard to the size of \mathcal{B}_λ , we can capitalize on the expression of $D(\boldsymbol{\lambda})$ associated with (6). Because a closed-form expression for $D(\boldsymbol{\lambda})$ may be difficult to obtain, bounds on $D(\boldsymbol{\lambda})$ (which, in general, are easier to find) can be used to bound the size of \mathcal{B}_λ . On the other hand, generic bounds on \mathcal{B}_λ that do not exploit the specific form of (6) but only the fact that $D(\boldsymbol{\lambda})$ is a dual function can also be found along the lines of, for example, [26].

$\lambda_L[n+1] = [\lambda_L[n] - \mu \partial_L[n]]_0^\infty$, it follows that

$$\begin{aligned} \|\lambda_L[n+1] - \lambda^*\|^2 &\leq \|\lambda_L[n] - \lambda^*\|^2 \\ &+ \mu^2 \|\partial_L[n]\|^2 - 2\mu \partial_L^T[n] (\lambda_L[n] - \lambda^*) \end{aligned} \quad (44)$$

where the inequality is due to the projection $[\cdot]_0^\infty$ operator. Considering that the path of the system up to instant n is given, taking expectations on both sides of (44) yields

$$\begin{aligned} \mathbb{E} \left[\|\lambda_L[n+1] - \lambda^*\|^2 \right] &\leq \mathbb{E} \left[\|\lambda_L[n] - \lambda^*\|^2 \right] + \mu^2 \mathbb{E} \left[\|\partial_L[n]\|^2 \right] \\ &- 2\mu \mathbb{E} \left[\partial_L^T[n] (\lambda_L[n] - \lambda^*) \right]. \end{aligned} \quad (45)$$

For $L=2$, we have that $\partial_L[n] = (1/2)(\partial[n] + \partial[n-1])$, whereas in Lemma 2, we ignore that the equality does not hold for $n=1$ (cf. Footnote 5). Based on this case, the last term in (45) can be rewritten as

$$\begin{aligned} \mathbb{E} \left[\partial_L^T[n] (\lambda_L[n] - \lambda^*) \right] &= \frac{1}{2} \mathbb{E} \left[\partial^T[n] (\lambda_L[n] - \lambda^*) \right] \\ &+ \frac{1}{2} \mathbb{E} \left[\partial^T[n-1] (\lambda_L[n-1] - \mu \tilde{\partial}_L[n-1] - \lambda^*) \right]. \end{aligned} \quad (46)$$

Using the fact that $\partial[n]$ and $\partial[n-1]$ are subgradients of the dual function at points $\lambda_L[n]$ and $\lambda_L[n-1]$, respectively, (46) can be bounded as

$$\begin{aligned} \mathbb{E} \left[\partial_L^T[n] (\lambda_L[n] - \lambda^*) \right] &\leq \frac{1}{2} (D(\lambda[n]) - D(\lambda^*)) + \frac{1}{2} \\ &\times (D(\lambda[n-1]) - D(\lambda^*)) - \frac{\mu}{2} \mathbb{E} \left[\partial^T[n-1] \tilde{\partial}[n-1] \right]. \end{aligned} \quad (47)$$

Substituting (47) into (45) yields

$$\begin{aligned} \mathbb{E} \left[\|\lambda_L[n+1] - \lambda^*\|^2 \right] &\leq \mathbb{E} \left[\|\lambda_L[n] - \lambda^*\|^2 \right] + \mu^2 \mathbb{E} \left[\|\partial_L[n]\|^2 \right] \\ &- 2\mu \frac{1}{2} (D(\lambda[n]) - D(\lambda^*)) - 2\mu \frac{1}{2} \\ &\times (D(\lambda[n-1]) - D(\lambda^*)) \\ &+ \mu^2 \mathbb{E} \left[\partial^T[n-1] \tilde{\partial}[n-1] \right]. \end{aligned} \quad (48)$$

Taking expectations in (48) over any possible path $\lambda_L[0], \dots, \lambda_L[n]$ and with Δ_L^{\max} and $\tilde{\Delta}_L^{\max}$ denoting upper bounds on $\mathbb{E}[\|\partial_L[n]\|^2]$ and $|\mathbb{E}[\partial_L^T[n] \partial_L[n]]|$, respectively, it follows that

$$\begin{aligned} \mathbb{E} \left[\|\lambda_L[n+1] - \lambda^*\|^2 \right] &\leq \mathbb{E} \left[\|\lambda_L[n] - \lambda^*\|^2 \right] + \mu^2 \Delta_L^{\max} \\ &- \frac{\mu}{2} \mathbb{E} [D(\lambda[n]) - D(\lambda^*)] \\ &- \frac{\mu}{2} \mathbb{E} [D(\lambda[n-1]) - D(\lambda^*)] + \mu^2 \tilde{\Delta}_L^{\max}. \end{aligned} \quad (49)$$

Successively applying (49) yields

$$\begin{aligned} \mathbb{E} \left[\|\lambda_L[n+1] - \lambda^*\|^2 \right] &\leq \mathbb{E} \left[\|\lambda_L[0] - \lambda^*\|^2 \right] \\ &+ \mu^2 \sum_{l=1}^n \left(\Delta_L^{\max} + \tilde{\Delta}_L^{\max} \right) - 2\mu \sum_{l=1}^n \left(\mathbb{E} [D(\lambda_L[l])] - D(\lambda^*) \right). \end{aligned} \quad (50)$$

Dividing by n and rearranging terms, we arrive at

$$\begin{aligned} \frac{\mu^2}{n} \sum_{l=1}^n \left(\Delta_L^{\max} + \tilde{\Delta}_L^{\max} \right) &\geq \frac{2\mu}{n} \sum_{l=1}^n \left(\mathbb{E} [D(\lambda_L[l])] - D(\lambda^*) \right) \\ &+ \frac{1}{n-1} \left(\mathbb{E} \left[\|\lambda_L[n+1] - \lambda^*\|^2 \right] - \mathbb{E} \left[\|\lambda_L[0] - \lambda^*\|^2 \right] \right). \end{aligned} \quad (51)$$

Because the multipliers are bounded (cf. Lemma 1), as $n \rightarrow \infty$, the last inequality yields

$$\begin{aligned} \frac{\mu \left(\Delta_L^{\max} + \tilde{\Delta}_L^{\max} \right)}{2} &\geq \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \left(\mathbb{E} [D(\lambda_L[l])] - D(\lambda^*) \right) \\ &\geq D \left(\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=1}^n \mathbb{E} [\lambda_L[l]] \right) - D(\lambda^*) \end{aligned} \quad (52)$$

where, for the last inequality, we have used that the dual function is convex, which concludes the proof. ■

REFERENCES

- [1] D. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.
- [2] D. Bertsekas, A. Nedic, and A. E. Ozdaglar, *Convex Analysis and Optimization*. Belmont, MA: Athena Scientific, 2003.
- [3] S. Boyd and A. Mutapcic, "Stochastic subgradient methods," Notes for EE364b, Stanford, CA: Stanford Univ., 2008.
- [4] L. Chen, S. H. Low, M. Chiang, and J. C. Doyle, "Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks," in *Proc. IEEE INFOCOM*, Barcelona, Spain, Apr. 2006, pp. 1–13.
- [5] T. Cover and J. Thomas, *Elements of Information Theory*. Hoboken, NJ: Wiley-Interscience, 1991.
- [6] A. Eryilmaz and R. Srikant, "Joint congestion control, routing, and MAC for stability and fairness in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1514–1524, Aug. 2006.
- [7] N. Gatsis, A. Ribeiro, and G. B. Giannakis, "A class of convergent algorithms for resource allocation in wireless fading networks," *IEEE Trans. Wireless Commun.*, vol. 9, no. 5, pp. 1808–1823, May 2010.
- [8] L. Georgiadis, M. Neely, and L. Tassiulas, "Resource allocation and cross-layer control in wireless networks," *Found. Trends Netw.*, vol. 1, no. 1, pp. 1–144, Apr. 2006.
- [9] A. Giannoulis, K. P. Tsoukatos, and L. Tassiulas, "Lightweight cross-layer control algorithms for fairness and energy efficiency in CDMA ad hoc networks," in *Proc. Int. Symp. Model. Optim. Mobile, Ad hoc, Wireless Netw.*, Boston, MA, Apr. 2006, pp. 1–8.
- [10] L. Huang and M. J. Neely, *Delay Reduction via Lagrange Multipliers in Stochastic Network Optimization*, Ithaca, NY, 2009. [Online]. Available: <http://arxiv.org/abs/0904.3795v1>
- [11] Z. Jiang, Y. Ge, and Y. Li, "Max-utility wireless resource management for best effort traffic," *IEEE Trans. Wireless Commun.*, vol. 4, no. 1, pp. 100–111, Jan. 2005.
- [12] H. Kushner and G. Yin, *Stochastic Approximation Algorithms and Applications*, 2nd ed. Berlin, Germany: Springer-Verlag, 2003.

- [13] J.-W. Lee, R. R. Mazumdar, and N. B. Shroff, "Opportunistic power scheduling for dynamic multiserver wireless systems," *IEEE Trans. Wireless Commun.*, vol. 5, no. 6, pp. 1506–1515, Jun. 2006.
- [14] L. Li and A. J. Goldsmith, "Capacity and optimal resource allocation for fading broadcast channels—Part I: Ergodic capacity," *IEEE Trans. Inf. Theory*, vol. 47, no. 3, pp. 1083–1102, Mar. 2001.
- [15] X. Lin and N. B. Shroff, "The impact of imperfect scheduling on cross-layer congestion control in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 14, no. 2, pp. 302–315, Apr. 2006.
- [16] X. Lin, N. B. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1452–1463, Aug. 2006.
- [17] Y.-H. Lin and R. L. Cruz, "Opportunistic link scheduling, power control, and routing for multihop wireless networks over time-varying channels," in *Proc. Allerton Conf. Commun., Control, Comput.*, Monticello, IL, Sep. 2005, pp. 976–985.
- [18] P. Liu, R. A. Berry, and M. L. Honig, "A fluid analysis of a utility-based wireless scheduling policy," *IEEE Trans. Inf. Theory*, vol. 52, no. 7, pp. 2872–2889, Jul. 2006.
- [19] L. M. Lopez-Ramos, A. G. Marques, J. Ramos, and A. Caamano, "Cross-layer resource allocation for downlink access using instantaneous fading and queue length information," in *Proc. IEEE GLOBECOM Workshops*, Miami, FL, Dec. 6–10, 2010, pp. 1212–1216.
- [20] D. J. Love, R. W. Heath, V. K. Lau, D. Gesbert, B. Rao, and M. Andrews, "An overview of limited feedback in wireless communication systems," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 8, pp. 1341–1365, Oct. 2008.
- [21] S. H. Low, "A duality model of TCP and queue management algorithms," *IEEE/ACM Trans. Netw.*, vol. 11, no. 4, pp. 525–536, Aug. 2003.
- [22] A. G. Marques, G. B. Giannakis, F. F. Digham, and F. J. Ramos, "Power-efficient wireless OFDMA using limited-rate feedback," *IEEE Trans. Wireless Commun.*, vol. 7, no. 2, pp. 685–696, Feb. 2008.
- [23] A. G. Marques, G. B. Giannakis, and J. Ramos, "Optimizing orthogonal multiple access based on quantized channel-state information," *IEEE Trans. Signal Process.*, vol. 59, no. 10, pp. 5023–5038, Oct. 2011.
- [24] A. G. Marques, G. B. Giannakis, and J. Ramos, "Stochastic cross-layer resource allocation for wireless networks using orthogonal access: Optimality and delay analysis," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Dallas, TX, Mar. 14–19, 2010, pp. 3154–3157.
- [25] A. G. Marques, X. Wang, and G. B. Giannakis, "Dynamic resource management for cognitive radios using limited-rate feedback," *IEEE Trans. Signal Process.*, vol. 57, no. 9, pp. 3651–3666, Sep. 2009.
- [26] A. Nedic and A. Ozdaglar, "Approximate primal solutions and rate analysis for dual subgradient methods," *SIAM J. Optim.*, vol. 19, no. 4, pp. 1757–1780, Feb. 2009.
- [27] M. J. Neely, E. Modiano, and C. E. Rohrs, "Dynamic power allocation and routing for time-varying wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 89–103, Jan. 2005.
- [28] A. Ribeiro, "Ergodic stochastic optimization algorithms for wireless communication and networking," *IEEE Trans. Signal Process.*, vol. 58, no. 12, pp. 6369–6386, Dec. 2010.
- [29] A. Ribeiro and G. B. Giannakis, "Optimal FDMA over wireless fading ad hoc networks," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Las Vegas, NV, Mar./Apr. 2008, pp. 2765–2768.
- [30] P. Soldati, B. Johansson, and M. Johansson, "Proportionally fair allocation of end-to-end bandwidth in STDMA wireless networks," in *Proc. 7th ACM Int. Symp. Mobile Ad Hoc Netw. Comput.*, Florence, Italy, May 2006, pp. 286–297.
- [31] A. Stolyar, "Maximizing queuing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Syst.*, vol. 50, no. 4, pp. 401–457, Aug. 2005.
- [32] L. Tassiulas and A. Ephremides, "Stability properties of constrained queuing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.
- [33] X. Wang and N. Gao, "Stochastic resource allocation in fading multiple-access and broadcast channels," *IEEE Trans. Inf. Theory*, vol. 56, no. 5, pp. 2382–2391, May 2010.
- [34] X. Wang, G. B. Giannakis, and A. G. Marques, "A unified approach to QoS-guaranteed scheduling for channel-adaptive wireless networks," *Proc. IEEE*, vol. 95, no. 12, pp. 2410–2431, Dec. 2007.
- [35] C. Y. Wong, R. S. Cheng, K. B. Lataief, and R. D. Murch, "Multiuser OFDM with adaptive subcarrier, bit, and power allocation," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 10, pp. 1747–1758, Oct. 1999.
- [36] L. Huang and M. J. Neely, "Delay reduction via Lagrange multipliers in stochastic network optimization," *IEEE Trans. Autom. Control*, vol. 56, no. 4, pp. 842–857, Apr. 2011.



Antonio G. Marques (M'07) received the B.Sc.–M.Sc. degree (with highest honors) in telecommunications engineering and the Ph.D. degree (with highest honors) in telecommunications engineering from the Carlos III University of Madrid, Madrid, Spain, in 2002 and 2007, respectively.

In 2003, he joined the Department of Signal Theory and Communications, King Juan Carlos University, Madrid, where he currently develops his research and teaching activities as an Assistant Professor. Since 2005, he has also been a Visiting Researcher with the Department of Electrical Engineering, University of Minnesota, Minneapolis. His research interests include communications theory, signal processing, and networking. His current research is focused on stochastic cross-layer resource allocation, cognitive radios, and signal processing for networks.

Dr. Marques has received various awards at in several conferences, including the 2007 IEEE International Conference on Acoustics, Speech, and Signal Processing.



Luis M. Lopez-Ramos (S'10) received the B.Sc. degree (with highest honors) in telecommunications engineering from the King Juan Carlos University of Madrid, Madrid, Spain, in 2010. He is currently working toward the M.Sc. degree in multimedia and communications with Carlos III University of Madrid.

Since then, he has been a Research Assistant at the Department of Signal Theory and Communications, King Juan Carlos University. His research interests include signal processing for wireless networks, cognitive radio, and computer vision.



Georgios B. Giannakis (F'97) received the Diploma degree in electrical engineering in 1981 from the National Technical University of Athens, Athens, Greece, and the M.Sc. degree in electrical engineering in 1983, the M.Sc. degree in mathematics in 1986, and the Ph.D. degree in electrical engineering in 1986 from the University of Southern California, Los Angeles.

Since 1999, he has been a Professor with the University of Minnesota, Minneapolis, where he currently holds the ADC Chair in Wireless Telecommunications with the Department of Electrical and Computer Engineering and serves as the Director of the Digital Technology Center. He is a coholder of 21 patents issued. His research interests include communications, networking, and statistical signal processing, subjects on which he has published more than 325 journal papers, 525 conference proceedings, 20 book chapters, two edited books, and two research monographs. His current research is focused on compressive sensing, cognitive radios, cross-layer designs, wireless sensors, and social and power grid networks.

Dr. Giannakis is a Fellow of the European Association for Signal Processing (EURASIP) and has served the IEEE in a number of posts, for example, as a Distinguished Lecturer for the IEEE Signal Processing (SP) Society. He has coreceived eight Best Paper Awards from the IEEE SP Society and the IEEE Communications Society, including the G. Marconi Prize Paper Award in Wireless Communications. He also received the Technical Achievement Award from the IEEE SP Society in 2000 and from EURASIP in 2005, the Young Faculty Teaching Award, and the G. W. Taylor Award for Distinguished Research from the University of Minnesota.



Javier Ramos received the B.Sc. and M.Sc. degrees in telecommunications engineering from the Polytechnic University of Madrid, Madrid, Spain, and the Ph.D. degree in 1995.

Between 1992 and 1995, he was involved with several research projects with Purdue University, West Lafayette, IN, working on signal processing for communications. In 1996, he was a Postdoctoral Research Associate with Purdue University. From 1997 to 2003, he was an Associate Professor with Carlos III University of Madrid. Since 2003, has

been with the King Juan Carlos University, Madrid, where he is currently a Professor and the Dean of the School of Telecommunications Engineering. His research interests include broadband wireless services and technologies and distributed sensing.

Dr. Ramos received the Ericsson Award for the Best Ph.D. Dissertation on Mobile Communications in 1996.



Antonio J. Caamaño received the B.Sc. and M.Sc. degrees in theoretical physics from the Autonomía University of Madrid, Madrid, Spain, in 1995 and the Ph.D. degree in telecommunications engineering from the Carlos III University of Madrid in 2003.

In 2003, he was an Associate Professor with the King Juan Carlos University of Madrid, where he has been the Head of the Department of Signal Theory and Communications since 2005. He has directed or cooperated in more than 40 private and publicly funded research projects. He has published research

results in field emission devices for communications, blind source separation, multiuser detection, and processing of biological signals. His research interests include large-scale Mobile Ad Hoc Network optimization, vehicular communications, distributed computing, and neuromorphic processing.