

Energy-Optimal Scheduling with Dynamic Channel Acquisition in Wireless Downlinks

Chih-ping Li, Michael J. Neely

Abstract—We consider a wireless base station serving L users through L time-varying channels. It is well known that opportunistic scheduling algorithms with full channel state information (CSI) can stabilize the system and achieve the full capacity region. However, opportunistic scheduling algorithms with full CSI may not be energy efficient when the cost of channel acquisition is high and traffic rates are low. In particular, under the low traffic rate regime, it may be sufficient and more energy-efficient to transmit data with no CSI, i.e., to transmit data blindly, since no power for channel acquisition is consumed. In general, we show that purely channel-aware or purely channel-blind strategies are not necessarily optimal, and we must consider mixed strategies. We derive a unified scheduling algorithm that dynamically chooses to transmit data with full or no CSI based on queue backlog and channel statistics. Through Lyapunov analysis, we show that the unified algorithm can stabilize the downlink with optimal power consumption.

I. INTRODUCTION

To transmit data efficiently over wireless channels, it is important to accommodate the time variations of wireless channel states (due to changing environments, multi-path fading, and mobility, etc.), and the limited energy in wireless devices. In particular, the concept of *opportunistic scheduling* has been shown to enable the design of efficient control algorithms that boost supportable data rates to the limit. The intuition is that transmitting data only when channel states are good can increase the throughput of a wireless network with a limited energy budget.

There has been extensive work in this field. For example, [3][5][6][2][7] focus on throughput/utility maximization with energy constraints in wireless networks. In particular, these works share one assumption that, when making power/rate allocation decisions, current channel states are always known with negligible cost. However, acquiring channel states requires exchanging control packets with neighboring wireless nodes, which indeed consumes energy. Recent works [8][9] relax this assumption by finding the optimal number of channels that should be measured in order to maximize the throughput of wireless LANs. The key idea in [8][9] is to strike a balance between maximizing multi-user diversity gain and the corresponding timing overhead of channel probing. Their results imply that measuring all channels

regularly may not be optimal. Guha et al. [10] investigate partial channel probing and rate allocations in a wireless server allocation problem with infinitely backlogged users. They propose a polynomial time algorithm that comes within a constant factor of optimizing a linear utility function when the costs of probing different channels are different. Work in [13] considers a similar problem and shows that a utility that is arbitrarily close to optimal can be achieved in the special case when probing costs are equal. Chang et al. [11] generalize the result in [13], and investigate properties of optimal joint partial channel probing and rate allocation policies, as well as propose a two-step look-ahead algorithm which is optimal in some special cases. Kar et al. [12] study throughput-achieving scheduling algorithms under the constraint that channels are only measured every $T > 1$ slots.

In [3], a dynamic control algorithm with perfect channel state information (CSI) is proposed to achieve the capacity region in a wireless network with power consumption arbitrarily close to optimal. In this paper, we extend the result in [3] by relaxing the assumption of perfectly known CSI. In particular, we assume there is a nonzero power cost to acquire CSI, in which case the optimal control algorithm in [3] may no longer be optimal. Intuitively, it is natural to suspect that when incoming data rates are sufficiently small, scheduling without CSI may still be able to support the data rates, but requires no extra power for channel acquisition. This assumption relaxation in fact enlarges the decision space of network control policies, where now the new space consists of purely channel-aware and purely channel-blind scheduling (which are defined rigorously later), and the combination of these two types. The problem discussed in [3] can be viewed as a special case of our system model.

In the next section, we describe our mathematical model. Then we give a motivating example showing that it is necessary to have a unified treatment that incorporates both channel-aware and channel-blind scheduling policies. In section III we establish the minimum power for stability when scheduling allows dynamic channel acquisition. In section IV we propose a unified algorithm and show that it can achieve the capacity region with power consumption arbitrarily close to optimal. Through simulations, we demonstrate the performance of the unified algorithm in a special case where at most one user can transmit every slot. A low-complexity heuristic that has near-optimal performance is also introduced.

Chih-ping Li (email: chihpinl@usc.edu) and Michael J. Neely (web: <http://www-rcf.usc.edu/~mjneely>) are with the Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089, USA.

This work is supported in part by one or both of the following: The DARPA IT-MANET program grant W911NF-07-1-0028, the National Science Foundation grant OCE 0520324.

II. SYSTEM MODEL, CAPACITY REGIONS, AND MOTIVATING EXAMPLES

A. System Model

We consider a wireless base station serving L users through L independent time-varying channels. Time is slotted. Data is measured in integer units of packets. Assume packet arrivals $a_k(t)$ for user $k \in \{1, 2, 3, \dots, L\}$ in slot t are i.i.d. over slots and independent of channel states. Assume $a_k(t)$ takes values in $\{0, 1, 2, \dots, A_{max}\}$ where A_{max} is a finite integer, and define $\mathbb{E}[a_k(t)] = \lambda_k$ for all t . Assume channel states $s_k(t)$ for user k in slot t are i.i.d. over slots, and $s_k(t)$ takes values in $\mathcal{S} = \{0, 1, 2, \dots, \mu_{max}\}$, where μ_{max} is a finite integer. The value of $s_k(t)$ represents the maximum number of packets that can be transmitted if we decide to transmit data over channel k . In every slot, the base station chooses service rate vector $\boldsymbol{\mu}(t) = (\mu_1(t), \dots, \mu_L(t))$ in a feasible set Ω , where $\mu_k(t)$ is the service rate allocated to user k . We assume that $\mu_k(t) \in \mathcal{S}$ for all t and all k , and that the set Ω defines any additional system restrictions. For example, in systems where at most one channel can be served per slot, all vectors in Ω have at most one non-zero entry. A constant transmission power P_t is consumed for every channel k that is allocated a nonzero transmission rate.

At the beginning of each slot, the base station makes a decision whether or not to acquire current channel states. In this paper we assume that either states of *all* channels are acquired, with power expenditure P_m (each channel measurement consumes $\frac{P_m}{L}$ units of power), or no CSI is known. We do not consider partial channel acquisition. When channel states are acquired, the base station chooses feasible service rates $\mu_k(t) \leq s_k(t)$ for all k , and at most $\mu_k(t)$ packets can be successfully delivered for each k . Otherwise, the service rate vector $\boldsymbol{\mu}(t)$ is chosen blindly (without the knowledge of current channel states). In this case, for each channel k , at most $\mu_k(t)$ packets can be successfully delivered if $\mu_k(t) \leq s_k(t)$, and all packet transmissions fail otherwise. We assume that ACK/NACK feedback is received at the end of each timeslot via a reliable control channel (absence of an ACK signal is regarded as a NACK). The backlog level $U_k(t+1)$ of user k at time $t+1$ can thus be represented by the equation:

$$U_k(t+1) = \max\{U_k(t) - \mu_k(t)1_{[\mu_k(t) \leq s_k(t)]}, 0\} + a_k(t),$$

subject to the feasible rate allocation constraint $\boldsymbol{\mu}(t) \in \Omega$. The indicator function $1_{[\mu_k(t) \leq s_k(t)]}$ is required because of the possible blind scheduling mode. In this paper we assume that channel statistics are known and remain fixed. We say the wireless downlink is *stabilized* (by some scheduling policy) if the following inequality holds [3]:

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{k=1}^L \mathbb{E}[U_k(\tau)] < \infty.$$

B. Motivating Examples

We compare the performance of the class of *purely channel-aware* algorithms (that serve packets only if channels are measured) to the class of *purely channel-blind* algorithms (that never measure channels and can serve packets

without CSI)¹ in a simplified example. Consider the problem of allocating a server to L queues with independent i.i.d. Bernoulli ON/OFF channels (channel states are i.i.d. over slots) where at most one packet is served every slot. It is equivalent to setting $\mu_{max} = 1$ and that Ω consists of L -dimensional zero-one vectors in which at most one entry is 1. Define q_i for $i = 1, 2, \dots, L$ to be the probability of an ON channel i state. Define the *blind capacity region* Λ_{blind} to be the closure of the set of data rates that can be stabilized by purely channel-blind policies. Define the *capacity region* Λ to be the closure of the set of data rates that can be stabilized by purely channel-aware policies.² The following two lemmas characterize these capacity regions and the minimum power required to stabilize rate vectors within them (Proofs are given in Appendix).

Lemma 1: The blind capacity region Λ_{blind} consists of data rates $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \dots, \lambda_L)$ satisfying $\sum_{i=1}^L \lambda_i/q_i \leq 1$. Further, among the class of purely channel-blind scheduling policies, for each rate vector $\boldsymbol{\lambda}$ interior to Λ_{blind} , the minimum power consumption to stabilize $\boldsymbol{\lambda}$ is $(\sum_{i=1}^L \frac{\lambda_i}{q_i})P_t$.

Note that [1] has shown that the capacity region Λ consists of data rates $\boldsymbol{\lambda}$ satisfying, for each nonempty subset I of $\{1, 2, \dots, L\}$, $\sum_{i \in I} \lambda_i \leq 1 - \prod_{i \in I} (1 - q_i)$. Next we quantify the minimum power required for purely channel-aware scheduling to stabilize data rates interior to Λ .

Lemma 2: Among the class of purely channel-aware policies, for each rate vector $\boldsymbol{\lambda} \neq \mathbf{0}$ interior to Λ , the minimum power consumption to stabilize $\boldsymbol{\lambda}$ is $(\sum_{i=1}^L \lambda_i)P_t + \theta^*P_m$, where $\theta^* \triangleq \inf\{\theta \mid \frac{\boldsymbol{\lambda}}{\theta} \in \Lambda, 0 < \theta < 1\}$.

Next, consider the case $L = 2$. Following from Lemma 1 and 2 we observe that there is a capacity region difference between Λ and Λ_{blind} (See Fig. 1). We observe the set Λ_{blind} can be partitioned into areas according to whether data rates prefer purely channel-aware to purely channel-blind scheduling, and different power ratios $\frac{P_m}{P_t}$ may lead to different partitions. Note that purely channel-blind scheduling is impossible in the region outside Λ_{blind} (as it cannot stabilize the system). However, regardless of whether or not rates are within Λ_{blind} , in the next section we show that the minimum power policy might be *neither* purely channel-blind nor purely channel-aware. Rather, *mixed* strategies are typically required.

III. OPTIMAL POWER CONSUMPTION FOR STABILITY

In the following we prove a theorem characterizing the minimum power for stability when dynamic channel acquisition is allowed, using similar techniques of proving Theorem 1 in [3]. We show that the minimum average power required for stability is obtained by minimizing the resulting average power expenditure over the class of stationary randomized policies that achieve a time average transmission rate exactly equal to the rate vector $\boldsymbol{\lambda}$. Each stationary randomized policy makes decisions independently of queue backlog, and

¹Note that both purely channel-aware and purely channel-blind scheduling may take advantage of queue backlog information.

²Note that $\Lambda_{blind} \subset \Lambda$ because purely channel-aware scheduling can emulate purely channel-blind scheduling.




Fig. 1. For the case $L = 2$, Λ consists of data rates that are within the outer boundary (thick solid lines), while Λ_{blind} consists of data rates within the thick dotted line $\frac{\lambda_1}{q_1} + \frac{\lambda_2}{q_2} = 1$. Data rates in the shaded areas are those that prefer purely channel-blind to purely channel-aware scheduling. The shaded areas are decided under an additional assumption that $q_1 \leq q_2$. As P_m/P_t decreases, the shaded areas shrink in the directions given in the figure.

has the following structure: Every timeslot, the controller independently decides to measure channel states with some probability γ (where $0 \leq \gamma \leq 1$). If channel states are measured and the channel state \mathbf{s} is observed, the controller randomly chooses one of $L + 2$ pre-established rate vectors $\boldsymbol{\mu}_i(\mathbf{s})$ (for $i \in \{1, \dots, L+2\}$), with some probabilities $\alpha_i(\mathbf{s})$ (for $i \in \{1, \dots, L+2\}$). Otherwise, a rate vector $\boldsymbol{\mu} \in \Omega$ is chosen for blind transmission, with a particular probability distribution $\beta(\boldsymbol{\mu})$.

Theorem 1: For i.i.d. channel state processes (states independent over slots) and i.i.d. arrival processes with rate vector $\boldsymbol{\lambda}$ which is interior to Λ , the minimum power consumption to stabilize the system is the optimal objective of the following problem $\mathfrak{P}(\boldsymbol{\lambda})$ (defined in terms of auxiliary variables γ , $\{\alpha_i(\mathbf{s})\}_{i=1}^{L+2}$ for each channel vector \mathbf{s} , $\beta(\boldsymbol{\mu})$ for each $\boldsymbol{\mu} \in \Omega$, and sets of feasible rate vectors $\{\boldsymbol{\mu}_i(\mathbf{s})\}_{i=1}^{L+2}$ for each \mathbf{s} , where $\boldsymbol{\mu}_i(\mathbf{s}) \triangleq (\mu_{i,1}(\mathbf{s}), \dots, \mu_{i,L}(\mathbf{s}))$ and $\boldsymbol{\mu}_i(\mathbf{s}) \leq \mathbf{s}$ elementwise³):

$$\begin{aligned} \min. \quad & \gamma \sum_{\mathbf{s}} \pi_{\mathbf{s}} \left[\sum_{i=1}^{L+2} \alpha_i(\mathbf{s}) \left(P_m + \sum_{j=1}^L 1_{[\mu_{i,j}(\mathbf{s}) > 0]} P_t \right) \right] \\ & + (1 - \gamma) \sum_{\boldsymbol{\mu} \in \Omega} \beta(\boldsymbol{\mu}) \left(\sum_{i=1}^L 1_{[\mu_i > 0]} P_t \right) \\ \text{s.t.} \quad & \boldsymbol{\lambda} \leq \gamma \sum_{\mathbf{s}} \pi_{\mathbf{s}} \left(\sum_{i=1}^{L+2} \alpha_i(\mathbf{s}) \boldsymbol{\mu}_i(\mathbf{s}) \right) \\ & + (1 - \gamma) \sum_{\boldsymbol{\mu} \in \Omega} \beta(\boldsymbol{\mu}) (\boldsymbol{\mu} \otimes \mathbf{P}\{\boldsymbol{\mu} \leq \mathbf{s}\}), \\ & 0 \leq \gamma \leq 1, \beta(\boldsymbol{\mu}) \geq 0 \forall \boldsymbol{\mu} \in \Omega, \sum_{\boldsymbol{\mu} \in \Omega} \beta(\boldsymbol{\mu}) = 1, \\ & \alpha_i(\mathbf{s}) \geq 0 \forall \mathbf{s}, i = 1, 2, \dots, L+2; \sum_{i=1}^{L+2} \alpha_i(\mathbf{s}) = 1 \forall \mathbf{s}, \\ & \boldsymbol{\mu}_i(\mathbf{s}) \in \Omega, \boldsymbol{\mu}_i(\mathbf{s}) \leq \mathbf{s} \forall \mathbf{s}, i = 1, 2, \dots, L+2, \end{aligned}$$

where $\pi_{\mathbf{s}}$ is the steady state probability of channel states \mathbf{s} . For $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, we define $\mathbf{a} \otimes \mathbf{b} \triangleq (a_1 b_1, \dots, a_n b_n)$

³In the rest of the paper, for two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, we use the notation $\mathbf{a} \leq \mathbf{b}$ to denote $a_i \leq b_i$ for all i .

and $\mathbf{P}\{\mathbf{a} \leq \mathbf{b}\} \triangleq (\mathbf{P}\{a_1 \leq b_1\}, \dots, \mathbf{P}\{a_n \leq b_n\})$, where $\mathbf{P}\{A\}$ is the probability of event A occurring.

Proof: Given in Appendix.

The following corollary of Theorem 1 will be used later in the performance analysis of our proposed control algorithm:

Corollary 1: For i.i.d. arrival and channel state processes, and an interior point $\boldsymbol{\lambda}$ of Λ , the minimum average power consumption to stabilize $\boldsymbol{\lambda}$, denoted by $P_{opt}(\boldsymbol{\lambda})$, can be achieved by minimizing the power over the class of stationary randomized policies that support $\boldsymbol{\lambda}$. The optimal stationary randomized policy allocates power $(\hat{P}_1(t), \dots, \hat{P}_L(t))$ (where $\hat{P}_i(t)$ is the sum of measurement and transmission power allocated to user i in slot t for $i = 1, 2, \dots, L$) and service rates $\hat{\boldsymbol{\mu}}(t)$ that yield for every slot t ,

$$\sum_{i=1}^L \mathbb{E} \left[\hat{P}_i(t) \right] = P_{opt}(\boldsymbol{\lambda}), \mathbb{E} [\hat{\boldsymbol{\mu}}(t)] \geq \boldsymbol{\lambda}.$$

IV. THE UNIFIED ALGORITHM AND PERFORMANCE ANALYSIS

A. Dynamic Channel Acquisition Algorithm

In the previous section we established the minimum average power consumption required for stability. Here we develop a unified Dynamic Channel Acquisition (DCA) algorithm that provides stability with average power that is arbitrarily close to the minimum, with a corresponding tradeoff in average delay. The algorithm is stated below in terms of a positive control parameter V , chosen as desired to affect the tradeoff. At the beginning of each timeslot, we observe the current queue backlog $(U_1(t), \dots, U_L(t))$. We then decide to measure current channel states or not, and allocate transmission rates based on this measurement decision. The associated decision variables $m(t)$, $\boldsymbol{\mu}^{(m)}(t)$, $\boldsymbol{\mu}^{(u)}(t)$ are defined as follows: $m(t) \in \{0, 1\}$ where $m(t) = 1$ if channels are measured in slot t , and $m(t) = 0$ otherwise. Variables $\boldsymbol{\mu}^{(m)}(t)$ represent feasible transmission rate allocations in the case when channels are measured, and $\boldsymbol{\mu}^{(u)}(t)$ represent feasible transmission rates allocated when no channels are measured. Define $\boldsymbol{\chi}(t) \triangleq [m(t), \boldsymbol{\mu}^{(m)}(t), \boldsymbol{\mu}^{(u)}(t)]$ as the collection of control decision variables on slot t . The DCA algorithm observes the current queue backlog $\mathbf{U}(t)$ and chooses $\boldsymbol{\chi}(t)$ every slot to maximize the function $f(\mathbf{U}(t), \boldsymbol{\chi}(t))$, define as follows

$$\begin{aligned} f(\mathbf{U}(t), \boldsymbol{\chi}(t)) \triangleq & -m(t)VP_m + \\ & m(t)\mathbb{E}_{\mathbf{s}} \left[\sum_{i=1}^L \left(2U_i(t)\mu_i^{(m)}(t) - VP_t 1_{[\mu_i^{(m)}(t) > 0]} \right) \mid \mathbf{U}(t) \right] \\ & + \bar{m}(t) \sum_{i=1}^L \left(2U_i(t)\mu_i^{(u)}(t) \mathbf{P}\{\mu_i^{(u)}(t) \leq s_i\} - VP_t 1_{[\mu_i^{(u)}(t) > 0]} \right), \end{aligned} \quad (1)$$

where $\bar{m}(t) \triangleq 1 - m(t)$. The maximization can be achieved as follows: First, we compare the maximized multiplicands of $m(t)$ and $\bar{m}(t)$ in (1). We choose $m(t) = 1$ if its optimal multiplicand is greater than the other, and $m(t) = 0$ otherwise. If $m(t) = 1$, we measure the channel states

$s(t)$ and allocate feasible rates $\boldsymbol{\mu}^{(m)}(t)$ as the maximizer of the sum $\sum_{i=1}^L (2U_i(t)\mu_i^{(m)}(t) - VP_t 1_{[\mu_i^{(m)}(t) > 0]})$ subject to $\boldsymbol{\mu}^{(m)}(t) \leq \mathbf{s}(t)$. Otherwise, we have $m(t) = 0$, and we allocate feasible rates, without knowledge of current channel states, to maximize the multiplicand of $\bar{m}(t)$. Note that channel statistics is required in these decisions, and that the 0/1 value of $m(t)$ every slot is the most complicated part of the algorithm. In particular, the multiplicand of $m(t)$ is a conditional expectation given $\mathbf{U}(t)$, and the maximized expectation is taken over the steady state distribution π_s , under the assumption that the optimal $\boldsymbol{\mu}^{(m)}(t)$ vector is allocated for each potential observed channel state \mathbf{s} . In Section IV-B, we show that this computation can be done in real time for the special case when at most one packet can be transmitted per slot.

For proving performance of the DCA algorithm, it is important to note that the function $f(\mathbf{U}(t), \boldsymbol{\chi}(t))$ can be written as

$$f(\mathbf{U}(t), \boldsymbol{\chi}(t)) = \left(\sum_{i=1}^L 2U_i(t) \mathbb{E}_s [\hat{\mu}_i(t) | \mathbf{U}(t)] \right) - V \mathbb{E}_s \left[\sum_{i=1}^L \hat{P}_i(t) | \mathbf{U}(t) \right],$$

where

$$\hat{\mu}_i(t) \triangleq m(t)\mu_i^{(m)}(t) + \bar{m}(t)\mu_i^{(u)}(t)1_{[\mu_i^{(u)}(t) \leq s_i(t)]}, \quad (2)$$

$$\hat{P}_i(t) \triangleq m(t) \left(\frac{P_m}{L} + P_t 1_{[\mu_i^{(m)}(t) > 0]} \right) + \bar{m}(t) P_t 1_{[\mu_i^{(u)}(t) > 0]}, \quad (3)$$

where $\hat{\mu}_i(t)$ and $\hat{P}_i(t)$ are respectively the maximum number of packets that can be served for user i in slot t and the power consumption of user i in slot t .

Theorem 2: For each arrival rate $\boldsymbol{\lambda}$ interior to the capacity region Λ , the DCA algorithm implemented with any control parameter $V > 0$ stabilizes the system with time average queue backlog and power expenditure given as follows:

$$\limsup_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau-1} \sum_{i=1}^L \mathbb{E}[U_i(t)] \leq \frac{B + V(P_m + LP_t)}{2\epsilon_{max}},$$

$$\limsup_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau-1} \sum_{i=1}^L \mathbb{E}[\hat{P}_i(t)] \leq \frac{B}{V} + P_{opt}(\boldsymbol{\lambda}),$$

where $B \triangleq (\mu_{max}^2 + A_{max}^2)L$, $\epsilon_{max} > 0$ is the largest value such that $(\boldsymbol{\lambda} + \boldsymbol{\epsilon}_{max}) \in \Lambda$, where $\boldsymbol{\epsilon}_{max}$ is an all- ϵ_{max} vector. $\hat{P}_i(t)$ is defined in (3), and $P_{opt}(\boldsymbol{\lambda})$ is the minimum power consumption to stabilize $\boldsymbol{\lambda}$ from Theorem 1.

Proof: For $i = 1, 2, \dots, L$, the queueing dynamics of user i can be written as

$$U_i(t+1) = \max(U_i(t) - \hat{\mu}_i(t), 0) + a_i(t), \quad (4)$$

where $U_i(t)$ is the backlog of user i data at time t , $a_i(t)$ is the number of user i arrivals in slot t , and $\hat{\mu}_i(t)$ is define in (2). By squaring (4) for each i and the facts that

$$(\max(U_i(t) - \hat{\mu}_i(t), 0))^2 \leq (U_i(t) - \hat{\mu}_i(t))^2,$$

$$U_i(t) - \hat{\mu}_i(t) \leq U_i(t), \hat{\mu}_i(t) \leq \mu_{max}, a_i(t) \leq A_{max},$$

we have

$$\sum_{i=1}^L (U_i^2(t+1) - U_i^2(t)) \leq B - 2 \sum_{i=1}^L U_i(t)(\hat{\mu}_i(t) - a_i(t)), \quad (5)$$

where $B \triangleq (\mu_{max}^2 + A_{max}^2)L$. Next we define the Lyapunov function $L(t) \triangleq \sum_{i=1}^L U_i^2(t)$ and the one-step Lyapunov drift $\Delta(\mathbf{U}(t)) \triangleq \mathbb{E}[L(t+1) - L(t) | \mathbf{U}(t)]$. By taking expectation of (5) conditioning on current backlog $\mathbf{U}(t)$ and noting that arrival processes are i.i.d. over slots, it is easy to show that

$$\Delta(\mathbf{U}(t)) \leq B + 2 \sum_{i=1}^L U_i(t)\lambda_i - \sum_{i=1}^L 2U_i(t)\mathbb{E}[\hat{\mu}_i(t) | \mathbf{U}(t)] \quad (6)$$

Motivated by the performance optimal Lyapunov optimization technique developed in [3] [4], we add the cost metric $V\mathbb{E}[\sum_{i=1}^L \hat{P}_i(t) | \mathbf{U}(t)]$ which is weighted by V to both sides of (6), yielding

$$\Delta(\mathbf{U}(t)) + V\mathbb{E} \left[\sum_{i=1}^L \hat{P}_i(t) | \mathbf{U}(t) \right] \leq B + 2 \sum_{i=1}^L U_i(t)\lambda_i - \left(\sum_{i=1}^L 2U_i(t)\mathbb{E}[\hat{\mu}_i(t) | \mathbf{U}(t)] - V\mathbb{E} \left[\sum_{i=1}^L \hat{P}_i(t) | \mathbf{U}(t) \right] \right). \quad (7)$$

The DCA algorithm is designed to minimize the right hand side of (7) over all possible control decisions $\boldsymbol{\chi}(t)$. In other words, the DCA algorithm minimizes the last term of (7). This can be achieved by maximizing $f(\mathbf{U}(t), \boldsymbol{\chi}(t))$ over all feasible control decisions $\boldsymbol{\chi}(t)$, which follows the procedures described in Section IV-A.

For performance analysis, note that the resulting right hand side of (7) under the DCA algorithm is less than or equal to corresponding right hand side if *some other policy* is used. In particular, we choose the *some other policy* to be the optimal stationary randomized policy, denoted by ω_r , associated with the optimal solution of the problem $\mathfrak{P}(\boldsymbol{\lambda} + \boldsymbol{\epsilon})$ in Theorem 1, where the vector $\boldsymbol{\epsilon} > \mathbf{0}$ is some all- ϵ vector such that $\boldsymbol{\lambda} + \boldsymbol{\epsilon}$ is interior to Λ . Let $\hat{\boldsymbol{\mu}}^r(t)$ and $(\hat{P}_1^r(t), \dots, \hat{P}_L^r(t))$ be the service rates and power consumption associated with policy ω_r in slot t . Then from (7) we have

$$\Delta(\mathbf{U}(t)) + V\mathbb{E} \left[\sum_{i=1}^L \hat{P}_i(t) | \mathbf{U}(t) \right] \leq B + 2 \sum_{i=1}^L U_i(t)\lambda_i - \left(\sum_{i=1}^L 2U_i(t)\mathbb{E}[\hat{\mu}_i^r(t) | \mathbf{U}(t)] - V\mathbb{E} \left[\sum_{i=1}^L \hat{P}_i^r(t) | \mathbf{U}(t) \right] \right). \quad (8)$$

By Corollary 1, the policy ω_r makes control decisions $\boldsymbol{\chi}_r(t)$, independent of current queue backlog $\mathbf{U}(t)$, resulting in service rates $\hat{\boldsymbol{\mu}}_r(t)$ and power consumption $(\hat{P}_1^r(t), \dots, \hat{P}_L^r(t))$ that satisfy

$$\mathbb{E}[\hat{\boldsymbol{\mu}}_r(t) | \mathbf{U}(t)] \geq \boldsymbol{\lambda} + \boldsymbol{\epsilon}, \quad (9)$$

$$\mathbb{E} \left[\sum_{i=1}^L \hat{P}_i^r(t) | \mathbf{U}(t) \right] = P_{opt}(\boldsymbol{\lambda} + \boldsymbol{\epsilon}) \quad (10)$$

in every slot t . Plugging (9) and (10) into (8) results in

$$\begin{aligned} \Delta(\mathbf{U}(t)) + V\mathbb{E} \left[\sum_{i=1}^L \hat{P}_i(t) \mid \mathbf{U}(t) \right] \\ \leq B - 2\epsilon \sum_{i=1}^L U_i(t) + VP_{opt}(\boldsymbol{\lambda} + \epsilon). \end{aligned} \quad (11)$$

Taking expectation of (11) over $\mathbf{U}(t)$, summing it from $t = 0$ to $\tau - 1$, and dividing the sum by τ yield

$$\begin{aligned} \frac{2\epsilon}{\tau} \sum_{t=0}^{\tau-1} \sum_{i=1}^L \mathbb{E}[U_i(t)] \leq B + \frac{\mathbb{E}[L(\mathbf{U}(0))] - \mathbb{E}[L(\mathbf{U}(\tau))]}{\tau} \\ + VP_{opt}(\boldsymbol{\lambda} + \epsilon) - \frac{V}{\tau} \mathbb{E} \left[\sum_{t=0}^{\tau-1} \sum_{i=1}^L \hat{P}_i(t) \right]. \end{aligned} \quad (12)$$

By taking limsup of (12) as $\tau \rightarrow \infty$ and the fact that $P_{opt}(\boldsymbol{\lambda} + \epsilon) \leq P_m + LP_t$,⁴ we have

$$\limsup_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau-1} \sum_{i=1}^L \mathbb{E}[U_i(t)] \leq \frac{B + V(P_m + LP_t)}{2\epsilon}, \quad (13)$$

$$\limsup_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau-1} \sum_{i=1}^L \mathbb{E}[\hat{P}_i(t)] \leq \frac{B}{V} + P_{opt}(\boldsymbol{\lambda} + \epsilon). \quad (14)$$

Note that (13) and (14) hold for any $\epsilon > 0$ satisfying that $\boldsymbol{\lambda} + \epsilon$ is interior to the capacity region Λ . Thus we can tighten the bounds by setting $\epsilon = \epsilon_{max}$ in (13) where $\epsilon_{max} > 0$ is the largest real number satisfying $\boldsymbol{\lambda} + \epsilon_{max} \in \Lambda$, and by setting $\epsilon = 0$ in (14). Thus we have

$$\limsup_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau-1} \sum_{i=1}^L \mathbb{E}[U_i(t)] \leq \frac{B + V(P_m + LP_t)}{2\epsilon_{max}}, \quad (15)$$

$$\limsup_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau-1} \sum_{i=1}^L \mathbb{E}[\hat{P}_i(t)] \leq \frac{B}{V} + P_{opt}(\boldsymbol{\lambda}). \quad (16)$$

□

Note that the two bounds in (15) and (16) are parameterized by the positive scalar V , where a larger V pushes the average power consumption asymptotically to optimal, at the expense of the linearly increasing average congestion bound (which, by Little's Theorem, yields an average delay bound).

B. Server Allocation Problem and Algorithm Implementation

Consider the simplified example of the L -queue downlink given in Section II-B. We show that the DCA algorithm, following the explanation in Section IV-A, can be simply implemented as follows: In every slot if channel states are acquired, we allocate the server to the user, among all users with ON channel state, with the largest positive $f_i^{(m)}(t) \triangleq 2U_i(t) - VP_t$. If $f_i^{(m)}(t)$ is non-positive for all users with ON state, we idle the server. Further, if channel acquisition is not performed, we allocate the server to the user with the largest positive $f_i^{(u)}(t) \triangleq 2U_i(t)\mathbf{P}\{s_i = ON\} - VP_t$. If $f_i^{(u)}(t)$ is non-positive for all users, we idle the server.

⁴ $(P_m + LP_t)$ is the maximum power that can be consumed every slot.

Next, to decide whether or not to acquire channel states, we compare the optimal multiplicands of $m(t)$ and $\bar{m}(t)$ in (1). First note that choosing the largest positive $f_i^{(u)}(t)$ gives us the optimal multiplicand of $\bar{m}(t)$. The optimal multiplicand of $m(t)$ can be computed as

$$\begin{aligned} -VP_m + \sum_{i=1}^L (2U_i(t) - VP_t) 1_{[2U_i(t) > VP_t]} \mathbf{P}\{s_i(t) = ON\} \times \\ \mathbf{P}\{s_j(t) < \frac{U_i(t)}{U_j(t)}, \forall j < i, s_k(t) \leq \frac{U_i(t)}{U_k(t)}, \forall k > i\}. \end{aligned} \quad (17)$$

It is because we assign the server to user i when $s_i(t) = 1$, $f_i^{(m)}(t) > 0$ and $U_i(t) \geq s_j(t)U_j(t)$ for all $j \neq i$ (we break ties by choosing the smallest index), which occurs with probability

$$\mathbf{P}\{s_i(t) = ON, s_j(t) < \frac{U_i(t)}{U_j(t)}, \forall j < i, s_k(t) \leq \frac{U_i(t)}{U_k(t)}, \forall k > i\}$$

Then we acquire channel states if the optimal multiplicand of $m(t)$ is no less than that of $\bar{m}(t)$.

Unlike the EECA algorithm in [3] which does not require channel statistics, we observe that the DCA algorithm indeed requires channel statistics in the decisions of channel acquisitions. Knowing the joint probability distribution is easier for uncorrelated channels, which may not be the case for correlated channels. As a result, for the case of a server allocation problem with independent i.i.d. Bernoulli ON/OFF channels, i.e., no spacial correlation among channels, (17) can be computed in polynomial time with polynomial memory spaces. In other words, the DCA algorithm is a polynomial time algorithm in this case. Further, we observe that choosing service rates is trivial after channel acquisition decisions are made, and the implementation complexity of the DCA algorithm completely lies in whether or not to acquire channel states.

C. Simulations and Low Complexity Alternatives

We simulate the DCA algorithm for the simplified downlink with $L = 2$ in Section II-B. Assume every slot channel 1 and 2 are ON with probability 0.5 and 0.8, respectively. We let $A_{max} = 1$, meaning that at most one packet arrives for each user in every slot. The arrival rate vector (λ_1, λ_2) is set to be $\rho(0.2, 0.8)$, where ρ is a nonnegative controlled loading factor. The value of ρ is at most 0.89 in our simulation, corresponding to an operating point very close to the boundary of the capacity region Λ ($\rho = 0.9$ is right on the boundary). We set the transmission power $P_t = 10$ units, channel acquisition power $P_m = 4.5$ units, control variable $V = 10$, and the total simulation time to be 10^7 slots. We note that these chosen parameters yield that, compared to purely channel-aware scheduling, purely channel-blind scheduling is preferred for all data rates pushed toward the boundary of the blind capacity region Λ_{blind} along the direction $(0.2, 0.8)$. The simulation results are given in Fig. 2 and 3.

In theory, the DCA algorithm stabilizes the capacity region Λ . We show in Fig. 2 as an example that the DCA algorithm

stabilizes the system for all data rates pushed toward the boundary of Λ along the direction $(0.2, 0.8)$. Further, Fig. 3 shows that the DCA algorithm yields average power consumption strictly less than the theoretical minimum power that purely channel-aware and purely channel-blind scheduling can achieve.

Regarding the complexity of making channel acquisition decisions, we also examine the performance of a simple heuristic called *Simplified* algorithm. The *Simplified* algorithm, instead of computing an optimal expectation, uses the maximum of

$$\begin{aligned} \max. \quad & \sum_{i=1}^L (2U_i(t)\mu_i(t) - VP_t 1_{[\mu_i(t)>0]}) - VP_m \\ \text{s.t.} \quad & 0 \leq \mu_i(t) \leq \mu_{max}, \mu_i(t) > 0 \text{ for at most one } i. \end{aligned}$$

as the optimal multiplicand of $m(t)$ in (1). In other words, we choose $m(t)$ without requiring the joint probability distribution of channel states, based on the assumption that all channels are ON. This assumption greatly reduces the complexity of making channel acquisition decisions. Note that the new multiplicand of $m(t)$ is no less than the one in the DCA algorithm. Therefore it potentially measures the channel states more frequently and remains a stabilizing policy (because we can emulate channel-blind transmissions when channel states are known). Surprisingly, Fig. 2 and 3 show that the Simplified algorithm seems to approximate the DCA algorithm very well.




Fig. 2. The average backlog of the DCA and the Simplified algorithm, compared to the theoretical bound.




Fig. 3. The average power consumption of the DCA algorithm, compared to purely channel-aware, purely channel-blind, and the Simplified algorithm. Note that curves for the DCA and the Simplified algorithm are nearly overlapped. The curve of purely channel-blind is drawn only to $\rho = 0.71$, which corresponds to data rates very close to the blind capacity region boundary, because purely channel-blind cannot support data rates outside the blind capacity region.

V. CONCLUSION

In this paper, under the assumption that channel acquisition is no longer for free, we propose an algorithm that

dynamically acquires channel states to stabilize a wireless downlink. This proposed Dynamic Channel Acquisition (DCA) algorithm is a unified treatment of incorporating both channel-aware and channel-blind scheduling policies to achieve energy optimality. Through Lyapunov analysis, we prove that the DCA algorithm can stabilize the downlink with average power consumption arbitrarily close to optimal, at the expense of increasing average network delays.

APPENDIX

Proof of Lemma 1: First assume rate vector $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_L)$ can be stabilized by a scheduling policy. Since successfully transmitting a packet through channel $i = 1, 2, \dots, L$ takes on average $\frac{1}{q_i}$ attempts, the fraction of time the system is busy is equal to $\sum_{i=1}^L \frac{\lambda_i}{q_i}$, which must be less than or equal to 1 for stability. The associated necessary average power consumption is equal to $(\sum_{i=1}^L \frac{\lambda_i}{q_i})P_t$. Conversely, for each rate vector $\lambda \neq \mathbf{0}$ satisfying $\sum_{i=1}^L \frac{\lambda_i}{q_i} < 1$, we define $\rho \triangleq \sum_{i=1}^L \frac{\lambda_i}{q_i}$, and there exists some $\epsilon > 0$ such that $\rho + \epsilon < 1$. Consider the policy of every slot assigning the server to queue i with probability $(\rho + \epsilon)\alpha_i$, where $\alpha_i \triangleq \frac{\lambda_i}{q_i \rho}$, and being idle with probability $1 - \rho - \epsilon$. The associated average power consumption is $(\rho + \epsilon)P_t$. Then it is easy to see that this policy yields average transmission rates being strictly greater than λ elementwise, and thus stabilizes the system. By passing $\epsilon \rightarrow 0$, the data rate vector λ is stabilized with average power consumption arbitrarily close to $\rho P_t = (\sum_{i=1}^L \frac{\lambda_i}{q_i})P_t$. \square

Proof of Lemma 2: Suppose the rate vector λ can be stabilized by a purely channel-aware policy ω . For simplicity, we assume the policy ω is ergodic with well-defined time averages (the general case can be proven similarly, as in [3]). Define θ as the fraction of time that policy ω acquires channel states to stabilize λ . Then the average power consumption to stabilize λ is equal to $(\sum_{i=1}^L \lambda_i)P_t + \theta P_m$. Suppose that θ satisfies $0 < \theta < \theta^*$. Using the same proving technique in Appendix A of [3], we can show that there exists a rate vector $\tilde{\lambda} \in \theta \Lambda$ such that $\lambda \leq \tilde{\lambda}$. In other words, $\lambda \in \theta \Lambda$. It contradicts the definition of θ^* , and finishes the proof of necessity. Conversely, since $\lambda \in \theta^* \Lambda$ (equivalent to $\frac{\lambda}{\theta^*} \in \Lambda$), λ is an interior point of the set $(\theta^* + \epsilon)\Lambda$ for some $\epsilon > 0$, $\theta^* + \epsilon < 1$. Then measuring channels every slot with probability $(\theta^* + \epsilon)$ and serving the longest ON queue whenever channels are measured is a λ -stabilizing policy [1] with power consumption $(\sum_{i=1}^L \lambda_i)P_t + (\theta^* + \epsilon)P_m$. Letting $\epsilon \rightarrow 0$ finishes the proof of sufficiency. \square

Proof of Theorem 1: Assume rate vector λ is stabilizable. Then there exists a λ -stabilizing policy ω which decides in which slots channel measurement is performed, and allocates transmission rates $\mu(t) = (\mu_1(t), \dots, \mu_L(t))$ accordingly in every slot t . Using the stability definition and Lemma 1 in [2], the following necessary condition holds with probability 1:

$$\lambda \leq \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \hat{\mu}(\tau), \quad (18)$$

where $\hat{\boldsymbol{\mu}}(\tau)$ denotes the effective transmission rates allocated in slot τ . When channels are measured in slot t , we have $\hat{\boldsymbol{\mu}}(t) = \boldsymbol{\mu}(t)$.⁵ Otherwise, $\hat{\boldsymbol{\mu}}(t) = \boldsymbol{\mu}(t) \otimes \mathbf{1}_{[\boldsymbol{\mu}(t) \leq \mathbf{s}(t)]}$, where $\mathbf{1}_{[\boldsymbol{\mu}(t) \leq \mathbf{s}(t)]} \triangleq (\mathbf{1}_{[\mu_1(t) \leq s_1(t)]}, \dots, \mathbf{1}_{[\mu_L(t) \leq s_L(t)]})$. Define $\mathbf{P}(t) = (P_1(t), \dots, P_L(t))$ to be the power vector where $P_i(t)$ is the power consumed by channel i in slot t , including both measurement and transmission power. When channels are measured, we have $P_i(t) = \frac{P_m}{L} + \mathbf{1}_{[\mu_i(t, \mathbf{s}(t)) > 0]} P_t$, and $P_i(t) = \mathbf{1}_{[\mu_i(t) > 0]} P_t$ otherwise.

For some time horizon $M > 0$, let $T_M^{(m)}$ and $T_M^{(u)}$ be the sets of slots in $[0, M]$ in which channel states are acquired and unknown, respectively. Without loss of generality assume $T_M^{(m)}$ and $T_M^{(u)}$ are nonempty. Define

$$\hat{\boldsymbol{\mu}}_{av}^{(m)}(M) \triangleq \frac{1}{M} \sum_{\tau \in T_M^{(m)}} \hat{\boldsymbol{\mu}}(\tau) = \frac{1}{M} \sum_{\tau \in T_M^{(m)}} \boldsymbol{\mu}(\tau, \mathbf{s}(\tau)),$$

$$\hat{\boldsymbol{\mu}}_{av}^{(u)}(M) \triangleq \frac{1}{M} \sum_{\tau \in T_M^{(u)}} \hat{\boldsymbol{\mu}}(\tau) = \frac{1}{M} \sum_{\tau \in T_M^{(u)}} \boldsymbol{\mu}(\tau) \otimes \mathbf{1}_{[\boldsymbol{\mu}(\tau) \leq \mathbf{s}(\tau)]}.$$

The empirical service rate $\hat{\boldsymbol{\mu}}_{av}(M) \triangleq \frac{1}{M} \sum_{\tau=0}^{M-1} \hat{\boldsymbol{\mu}}(\tau)$ over $[0, M]$ is equal to $\hat{\boldsymbol{\mu}}_{av}^{(m)}(M) + \hat{\boldsymbol{\mu}}_{av}^{(u)}(M)$. Define

$$P_{av}^{(m)}(M) \triangleq \frac{1}{M} \sum_{\tau \in T_M^{(m)}} \sum_{i=1}^L P_i(\tau),$$

$$P_{av}^{(u)}(M) \triangleq \frac{1}{M} \sum_{\tau \in T_M^{(u)}} \sum_{i=1}^L P_i(\tau).$$

The empirical average power consumption $P_{av}(M) \triangleq \frac{1}{M} \sum_{\tau=0}^{M-1} \sum_{i=1}^L P_i(\tau)$ is equal to $P_{av}^{(m)}(M) + P_{av}^{(u)}(M)$. It is easy to show the $(L+1)$ -dim vector $(\hat{\boldsymbol{\mu}}_{av}^{(m)}(M); P_{av}^{(m)}(M))$ satisfies:

$$(\hat{\boldsymbol{\mu}}_{av}^{(m)}(M); P_{av}^{(m)}(M)) = \frac{1}{M} \sum_{\tau \in T_M^{(m)}} (\boldsymbol{\mu}(\tau, \mathbf{s}(\tau)); \sum_{i=1}^L P_i(\tau))$$

$$= \gamma_M \sum_{\mathbf{s}} \sigma_M(\mathbf{s}) \mathbf{x}_M(\mathbf{s}),$$

where $\gamma_M \triangleq \frac{|T_M^{(m)}|}{M}$, $\sigma_M(\mathbf{s}) \triangleq \frac{|T_M^{(m)}(\mathbf{s})|}{|T_M^{(m)}|}$, where $T_M^{(m)}(\mathbf{s}) \subset T_M^{(m)}$ consists of slots in which channel states are \mathbf{s} , and $\mathbf{x}_M(\mathbf{s}) \triangleq \frac{1}{|T_M^{(m)}(\mathbf{s})|} \sum_{\tau \in T_M^{(m)}(\mathbf{s})} (\boldsymbol{\mu}(\tau, \mathbf{s}); \sum_{i=1}^L P_i(\tau))$. Since $\mathbf{x}_M(\mathbf{s})$ is a convex combination of vectors of the form $(\boldsymbol{\mu}(\tau, \mathbf{s}); \sum_{i=1}^L P_i(\tau))$, by Caratheodory Theorem, there exist $L+2$ feasible rate vectors $\boldsymbol{\mu}_i^M(\mathbf{s}) = (\mu_{i,1}^M(\mathbf{s}), \dots, \mu_{i,L}^M(\mathbf{s}))$ satisfying $\boldsymbol{\mu}_i^M(\mathbf{s}) \leq \mathbf{s}$ for $i = 1, 2, \dots, L+2$, and an associated probability distribution $\{\alpha_i^M(\mathbf{s})\}_{i=1}^{L+2}$, such that $\mathbf{x}_M(\mathbf{s}) = \sum_{i=1}^{L+2} \alpha_i^M(\mathbf{s}) (\boldsymbol{\mu}_i^M(\mathbf{s}); \sum_{j=1}^L P_i^j(\mathbf{s}))$, where $P_i^j(\mathbf{s}) \triangleq \frac{P_m}{L} +$

⁵Without loss of generality, we only consider the class of policies that, when channels are measured in slot t , always allocate transmission rates $\boldsymbol{\mu}(t)$ satisfying $\boldsymbol{\mu}(t) \leq \mathbf{s}(t)$. It is because allocating transmission rates that cannot be supported by current known channel states is equivalent to assigning zero rates. In the rest of the paper we use the notation $\boldsymbol{\mu}(t, \mathbf{s}(t))$ to emphasize that the allocated rates are supported by current known channel states.

$\mathbf{1}_{[\mu_{i,j}^M(\mathbf{s}) > 0]} P_t$. Also, it is easy to show that the vector $(\hat{\boldsymbol{\mu}}_{av}^{(u)}(M); P_{av}^{(u)}(M))$ satisfies:

$$(\hat{\boldsymbol{\mu}}_{av}^{(u)}(M); P_{av}^{(u)}(M))$$

$$= \frac{1}{M} \sum_{\tau \in T_M^{(u)}} (\boldsymbol{\mu}(\tau) \otimes \mathbf{1}_{[\boldsymbol{\mu}(\tau) \leq \mathbf{s}(\tau)]}; \sum_{i=1}^L P_i(\tau))$$

$$= (1 - \gamma_M) \sum_{\boldsymbol{\mu} \in \Omega} \beta_M(\boldsymbol{\mu}) \mathbf{y}_M(\boldsymbol{\mu}),$$

where $1 - \gamma_M = \frac{|T_M^{(u)}|}{M}$ because $|T_M^{(m)}| + |T_M^{(u)}| = M$, $\beta_M(\boldsymbol{\mu}) \triangleq \frac{|T_M^{(u)}(\boldsymbol{\mu})|}{|T_M^{(u)}|}$, where $T_M^{(u)}(\boldsymbol{\mu}) \subset T_M^{(u)}$ consists of slots in which transmission rate vector $\boldsymbol{\mu}$ is channel-blindly allocated, and $\mathbf{y}_M(\boldsymbol{\mu}) \triangleq \frac{1}{|T_M^{(u)}(\boldsymbol{\mu})|} \sum_{\tau \in T_M^{(u)}(\boldsymbol{\mu})} (\boldsymbol{\mu} \otimes \mathbf{1}_{[\boldsymbol{\mu} \leq \mathbf{s}(\tau)]}; \sum_{i=1}^L P_i(\tau))$. Note that $\sum_{\boldsymbol{\mu}} \beta_M(\boldsymbol{\mu}) = 1$. Then we observe that the sequence $\{(\hat{\boldsymbol{\mu}}_{av}^{(m)}(M); P_{av}^{(m)}(M))\}$ indexed by M is a bounded sequence. By Weierstrass Theorem, there exists a converging subsequence $\{M_n\}$. Since subsequences $\{\gamma_{M_n}\}$, $\{\sigma_{M_n}(\mathbf{s})\}$, $\{\alpha_i^{M_n}(\mathbf{s})\}$, $\{\boldsymbol{\mu}_i^{M_n}(\mathbf{s})\}$, $\{\beta_{M_n}(\boldsymbol{\mu})\}$, and $\{\mathbf{y}_{M_n}(\boldsymbol{\mu})\}$ are all bounded, by iteratively applying Weierstrass Theorem, there exists a subsequence $\{M_k\}$ of $\{M_n\}$ such that as $k \rightarrow \infty$, there exist some γ satisfying $0 \leq \gamma \leq 1$, probability distributions $\{\alpha_i(\mathbf{s})\}_{i=1}^{L+2}$ for each \mathbf{s} , feasible transmission rates $\boldsymbol{\mu}_i(\mathbf{s}) \leq \mathbf{s}$ for each i and \mathbf{s} , and a probability distribution $\{\beta(\boldsymbol{\mu})\}_{\boldsymbol{\mu}}$, such that $\gamma_{M_k} \rightarrow \gamma$, $\alpha_i^{M_k}(\mathbf{s}) \rightarrow \alpha_i(\mathbf{s})$, $\boldsymbol{\mu}_i^{M_k}(\mathbf{s}) \rightarrow \boldsymbol{\mu}_i(\mathbf{s})$, and $\beta_{M_k}(\boldsymbol{\mu}) \rightarrow \beta(\boldsymbol{\mu})$. Define $\pi_{\mathbf{s}}$ to be the steady state probability of channel state \mathbf{s} . Then $\sigma_{M_k}(\mathbf{s})$ converges to $\pi_{\mathbf{s}}$ as $k \rightarrow \infty$ by Law of Large Numbers. It is because the channel acquisition decision in a slot is independent of the channel states in that slot. Further, we observe that, for each $\boldsymbol{\mu}$, the vectors $\boldsymbol{\mu} \otimes \mathbf{1}_{[\boldsymbol{\mu} \leq \mathbf{s}(\tau)]}$ are i.i.d. for $\tau \in T_{M_k}^{(u)}(\boldsymbol{\mu})$. Thus by Law of Large Numbers $\mathbf{y}_{M_k}(\boldsymbol{\mu})$ converges to $(\boldsymbol{\mu} \otimes \mathbf{P}\{\boldsymbol{\mu} \leq \mathbf{s}\}; \sum_{i=1}^L \mathbf{1}_{[\mu_i > 0]} P_t)$ as $k \rightarrow \infty$. In sum, as $k \rightarrow \infty$, $\{(\hat{\boldsymbol{\mu}}_{av}(M_k); P_{av}(M_k))\}$ converges to

$$\gamma \sum_{\mathbf{s}} \pi_{\mathbf{s}} \left[\sum_{i=1}^{L+2} \alpha_i(\mathbf{s}) \left(\boldsymbol{\mu}_i(\mathbf{s}); P_m + \sum_{j=1}^L \mathbf{1}_{[\mu_{i,j}(\mathbf{s}) > 0]} P_t \right) \right]$$

$$+ (1 - \gamma) \sum_{\boldsymbol{\mu} \in \Omega} \beta(\boldsymbol{\mu}) \left(\boldsymbol{\mu} \otimes \mathbf{P}\{\boldsymbol{\mu} \leq \mathbf{s}\}; \sum_{i=1}^L \mathbf{1}_{[\mu_i > 0]} P_t \right).$$

Equation (18) then yields

$$\lambda \leq \lim_{k \rightarrow \infty} \frac{1}{M_k} \sum_{\tau=0}^{M_k-1} \hat{\boldsymbol{\mu}}(\tau)$$

$$= \lim_{k \rightarrow \infty} (\hat{\boldsymbol{\mu}}_{av}^{(m)}(M_k) + \hat{\boldsymbol{\mu}}_{av}^{(u)}(M_k))$$

$$= \gamma \sum_{\mathbf{s}} \pi_{\mathbf{s}} \left(\sum_{i=1}^{L+2} \alpha_i(\mathbf{s}) \boldsymbol{\mu}_i(\mathbf{s}) \right)$$

$$+ (1 - \gamma) \sum_{\boldsymbol{\mu} \in \Omega} \beta(\boldsymbol{\mu}) (\boldsymbol{\mu} \otimes \mathbf{P}\{\boldsymbol{\mu} \leq \mathbf{s}\}).$$
(19)

The corresponding average power consumption is

$$\begin{aligned} & \gamma \sum_{\mathbf{s}} \pi_{\mathbf{s}} \left[\sum_{i=1}^{L+2} \alpha_i(\mathbf{s}) \left(P_m + \sum_{j=1}^L 1_{[\mu_{i,j}(\mathbf{s}) > 0]} P_t \right) \right] \\ & + (1 - \gamma) \sum_{\boldsymbol{\mu} \in \Omega} \beta(\boldsymbol{\mu}) \left(\sum_{i=1}^L 1_{[\mu_i > 0]} P_t \right). \end{aligned} \quad (20)$$

The first inequality of (19) is due to the fact that \liminf of a sequence is a lower bound of a limit point of the sequence.

Equation (19) (20) show that, for any $\boldsymbol{\lambda}$ -stabilizing policy ω , there exists a stationary randomized policy $\hat{\omega}$ which yields average service rates greater than or equal to $\boldsymbol{\lambda}$. The parameter γ , probability distributions $\{\beta(\boldsymbol{\mu})\}_{\boldsymbol{\mu}}$ and $\{\alpha_i(\mathbf{s})\}_i$ for all \mathbf{s} , and transmission rates $\{\mu_i(\mathbf{s})\}_i$ for all \mathbf{s} associated with $\hat{\omega}$ constitute a feasible solution to $\mathfrak{P}(\boldsymbol{\lambda})$. Thus the power consumption (20) of $\hat{\omega}$, or the necessary power to stabilize $\boldsymbol{\lambda}$, is no less than the optimal objective of $\mathfrak{P}(\boldsymbol{\lambda})$.

Conversely, for each rate vector $\boldsymbol{\lambda}$ interior to Λ , there exists a positive scalar ϵ such that $\boldsymbol{\lambda} + \epsilon$ is interior to Λ , where ϵ is the vector of which each entry is ϵ . By definition of Λ , rate vector $\boldsymbol{\lambda} + \epsilon$ is supportable. Then the optimal solution to $\mathfrak{P}(\boldsymbol{\lambda} + \epsilon)$ yields a stationary randomized policy $\underline{\omega}$ whose average service rates are greater than or equal to $\boldsymbol{\lambda} + \epsilon$. Thus policy $\underline{\omega}$ stabilizes $\boldsymbol{\lambda}$ by Lemma 3.6 of [4], with average power consumption being equal to the optimal objective of $\mathfrak{P}(\boldsymbol{\lambda} + \epsilon)$. By pushing ϵ to zero, there exists a stationary randomized policy which stabilizes $\boldsymbol{\lambda}$ with average power consumption arbitrarily close to the optimal objective of $\mathfrak{P}(\boldsymbol{\lambda})$. \square

REFERENCES

- [1] Leandros Tassiulas and A. Ephremides, Dynamic Server Allocation to Parallel Queues with Randomly Varying Connectivity, *IEEE Transactions on Information Theory*, vol.39, no.2, pp. 466-478, March 1993.
- [2] M. J. Neely, E. Modiano, and C. E. Rohrs, Power Allocation and Routing in Multi-Beam Satellites with Time Varying Channels, *IEEE Transactions on Networking*, vol. 11, no. 1, pp. 138-152, Feb. 2003.
- [3] M. J. Neely, Energy Optimal Control for Time Varying Wireless Networks, *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 2915-2934, July 2006.
- [4] L. Georgiadis, M. J. Neely, L. Tassiulas, Resource Allocation and Cross-Layer Control in Wireless Networks, *Foundations and Trends in Networking*, Vol. 1, no. 1, pp. 1-144, 2006.
- [5] P. Giaccone, B. Prabhakar, and D. Shah. Energy constrained wireless switching. *Proc. of Allerton Conference on Communications, Control, and Computing*, Oct. 2003.
- [6] E. M. Yeh and A. S. Cohen. Throughput optimal power and rate control for multiaccess and broadcast communications. *Proc. of the International Symposium on Information Theory*, June 2004.
- [7] J. W. Lee, R. R. Mazumdar, and N. B. Shroff. Opportunistic power scheduling for dynamic multiserver wireless systems. *IEEE Transactions on Wireless Communications*, vol. 5, no.5, May 2006.
- [8] Z. Ji, Y. Yang, J. Zhou, M. Takai, and R. Bagrodia, Exploiting medium access diversity in rate adaptive wireless lans, *ACM MOBICOM*, 2004.
- [9] A. Sabharwal, A. Khoshnevis, and K. Knightly, Opportunistic spectral usage: Bounds and a multi-band csma/ca protocol, *IEEE/ACM Trans. on Networking*, August 2007.
- [10] Sudipto Guha, Kamesh Munagala, and Saswati Sarkar, Performance guarantees through partial information based control in multichannel wireless networks, *Technical Report*, September 2006.
- [11] Nicholas B. Chang, and Mingyan Liu, Optimal Channel Probing and Transmission Scheduling in a Multichannel System, *ITA workshop*, San Diego, 2007.
- [12] Koushik Kar, Xiang Luo, Saswati Sarkar, Throughput-optimal Scheduling in Multichannel Access Point Networks under Infrequent Channel Measurement, *IEEE INFOCOM*, 2007.

- [13] Sudipto Guha, Kamesh Munagala, Saswati Sarkar, Approximation Schemes for Information Acquisition and Exploitation in Multichannel Wireless Networks, *44th Allerton Conference*, 2007.