Delay Analysis for Max Weight Opportunistic Scheduling in Wireless Systems

Michael J. Neely

***Abstract*—We consider the delay properties of max-weight opportunistic scheduling in a multi-user ON/OFF wireless system, such as a multi-user downlink or uplink. It is well known that max-weight scheduling stabilizes the network (and hence yields maximum throughput) whenever input rates are inside the network capacity region. We show that when arrival and channel processes are independent, average delay of the max-weight policy is order-optimal, in the sense that it does not grow with the number of network links. While recent queue-grouping algorithms are known to also yield order-optimal delay, this is the first such result for the simpler class of max-weight policies.**

I. INTRODUCTION

We consider the delay properties of max-weight opportunistic scheduling in a multi-user wireless system. Specifically, we consider a system with N transmission links. Each link receives independent data that arrives randomly and must be queued for eventual transmission. Separate queues are maintained by each link $i \in \{1, \ldots, N\}$, so that data arriving to queue i must be transmitted over link i. The system works in slotted time with normalized slots $t \in \{0, 1, 2, \ldots\}$. The channel states of each link vary randomly from slot to slot, and every slot t the network controller observes the current queue backlogs and the current channel states, and selects a single link for wireless transmission.

This is a classic *opportunistic scheduling* scenario, where the network scheduler can exploit knowledge of the current state of the time varying channels. It is well known that max-weight scheduling policies are throughput optimal in such systems, in the sense that they provably stabilize all queues whenever the input rate vector is inside the network capacity region. This stability result was first shown by Tassiulas and Ephremides in [2] for the special case of ON/OFF channels, and was later generalized to multi-rate transmission models and systems with power allocation [3] [4] [5] [6]. However, the delay properties of max-weight scheduling are less understood. An average delay bound that is linear in N is derived in [5] [6]. While this bound is tight in the case of correlated arrival and channel processes, it is widely believed to be loose for independent arrivals and channels.

In this paper, we focus on the special case of ON/OFF channels, and show that the max-weight policy indeed yields average delay that is $O(1)$ under independence assumptions. Thus, average delay does not grow with the network size and hence is *order optimal*. Specifically, we first show that for any input rate vector that is within a ρ-scaled version of the capacity region (where ρ represents the network loading, and satisfies $0 < \rho < 1$), the max-weight rule yields average delay that is less than or equal to $\frac{c\log(1/(1-\rho))}{(1-\rho)^2}$, where c is a constant that does not depend on ρ or N.[1] This is in comparison to the previous delay bound of $\frac{cN}{1-\rho}$ developed for max-weight scheduling [5] [6]. Note that our new bound does not grow with N, but has a worse asymptotic in ρ. We next present a different analysis that improves the delay bound to $\frac{c\log(1/(1-\rho))}{1-\rho}$ for systems with "f-balanced" traffic rates (to be made precise in later sections). That is, if arrival rates are heterogeneous but are more balanced (so that the difference between the maximum arrival rate and the average arrival rate is sufficiently small), then order-optimal average delay is maintained while the delay asymptotic in ρ is improved. This order-optimal delay analysis is extended in our technical report [1] to treat systems with multi-rate capabilities using additional stochastic coupling techniques and using a modified max-weight policy.

It is known that order-optimal delay requires queue-based scheduling. Indeed, it is shown in [7] that average delay in a N-user wireless downlink with time varying channels grows at least linearly with N if queue-independent algorithms are used (such as round-robin or randomized schedulers). Related results are shown for $N \times N$ packet switches in [8], where an order-optimal switch scheduling algorithm is developed. Delay optimal control laws for multi-user wireless systems are mostly limited to systems with special symmetric structure [2] [9] [10]. Delay optimality results are developed in [11] for a heavy traffic regime in the limit as the system loading ρ approaches 1. Recent results on exponents of the tail of delay distributions are provided in [12] [13], and order-optimal delay for greedy maximal scheduling with ρ a constant factor away from 1 is considered in [14] [15].

The max-weight rule is also called the *Longest Connected Queue* (LCQ) scheduling rule in the special case of an ON/OFF downlink. This policy was developed by Tassiulas and Ephremides in [2], where it was shown to support the full network capacity region and to also be *delay optimal* in the special symmetric case when all arrival rates and ON/OFF probabilities are the same for each link. The fact that the actual average delay of LCQ in such symmetric cases is $O(1)$ was recently proven in [10] (which shows that doubling the size of a symmetric system does not increase the average delay) and [7] (which uses a queue-grouped Lyapunov function

Michael J. Neely is with the Electrical Engineering department at the University of Southern California, Los Angles, CA (web: http://www-rcf.usc.edu/∼mjneely).

This material is supported in part by one or more of the following: the DARPA IT-MANET program grant W911NF-07-0028, the NSF grant OCE 0520324, the NSF Career grant CCF-0747525.

[1]The value c is used here to easily express a delay scaling relationship, and represents a generic coefficient that does not depend on ρ or N. The value c is not necessarily the same in all places it is used.

to bound the average delay). Delay properties of variations of LCQ for symmetric Poisson systems are considered in [16] in the limit of asymptotically large N. For asymmetric systems, it is shown in [7] that a *different algorithm*, called the *Largest Connected Group* (LCG) algorithm, yields $O(1)$ average delay. However, the LCG algorithm requires some statistical knowledge to set up a queue-group structure. Hence, it is important to understand the delay properties of the simpler max-weight rule, which does not require statistical knowledge. In this paper, we combine the *queue grouping* concepts developed in [7] together with two novel Lyapunov functions to provide an order-optimal delay analysis of max-weight. The first Lyapunov function we use has a weighted sum of two different component functions, and is inspired by work in [17] where a Lyapunov function with a similar structure is used in a different context.

In the next section, we specify the network model and review basic concepts concerning the network capacity region. Section III proves our first delay result for the ON/OFF channel model with general heterogeneous traffic rates. Section IV provides our second bound (with a tighter asymptotic in ρ) for the case of heterogeneous traffic rates but under an f-balanced traffic assumption. Extensions to multi-rate transmission models are developed in our technical report [1].

II. System Model

Consider a multi-user wireless system with N transmission links. The system operates in slotted time with normalized slots $t \in \{0, 1, 2, \ldots\}$. We assume that data is measured in units of fixed size packets, and let $A_i(t)$ represent the number of packets that arrive to link $i \in \{1, \ldots, N\}$ during slot t. Each link i maintains a separate queue to store this arriving data, and we let $Q_i(t)$ represent the number of packets waiting for transmission over link i.

Let $S_i(t)$ represent the *channel state* for the ith channel during slot t. We assume that $S_i(t) \in \{0, 1\}$, where 0 represents an OFF channel state where no packet can be served and 1 represents an ON channel state under which a single packet can be served if this channel is selected for transmission. Define $\boldsymbol{S}(t) = (S_1(t), \ldots, S_N(t))$ as the channel state vector. Let $\mu_i(t)$ represent the control decision variable on slot t, given as follows:

$$\mu_i(t) = \begin{cases} S_i(t), & \text{if channel } i \text{ is selected on slot } t \\ 0, & \text{otherwise} \end{cases}$$

Define $\boldsymbol{\mu}(t) = (\mu_1(t), \ldots, \mu_N(t))$ as the vector of transmission decisions. We also call this the *transmission rate vector*, as it determines the instantaneous transmission rates over each link (in units of packets/slot), where the rate is either 0 or 1. The constraint that at most one channel is selected per slot translates into the constraint that $\boldsymbol{\mu}(t)$ has at most one non-zero entry (and any non-zero entry i is equal to $S_i(t)$). Define $\mathcal{F}(t)$ as the set of all such control vectors $\boldsymbol{\mu}(t)$ that are possible for slot t, called the *feasibility set* for slot t. The queue dynamics for each queue $i \in \{1, \ldots, N\}$ are given as follows:

$$Q_i(t+1) = \max[Q_i(t) - \mu_i(t), 0] + A_i(t) \tag{1}$$

subject to the constraint $\boldsymbol{\mu}(t) \in \mathcal{F}(t)$ for all t.

A. Traffic and Channel Assumptions

We assume the arrival processes $A_i(t)$ are independent for all $i \in \{1, \ldots, N\}$. Further, each process $A_i(t)$ is i.i.d. over slots with mean $\lambda_i = \mathbb{E}\{A_i(t)\}$ and with a finite second moment $\mathbb{E}\left\{A_i(t)^2\right\} < \infty$. Similarly, we assume channel processes $S_i(t)$ are independent of each other and i.i.d. over slots with probabilities $Pr[S_i(t) = 1] = p_i$ for $i \in \{1, \ldots, N\}$.

B. The Network Capacity Region

Suppose the network control policy chooses a transmission rate vector every slot according to a well defined probability law, so that the queue states evolve according to (1).

Definition 1: A queue $Q_i(t)$ is *strongly stable* if:

$$\limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}\{Q_i(\tau)\} < \infty$$

We say that the network of queues is *strongly stable* if all individual queues are strongly stable. Throughout, we shall use the term "stability" to refer to strong stability.

Define Λ as the *network capacity region*, consisting of the closure of all arrival rate vectors $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_N)$ for which there exists a stabilizing control algorithm. In [2] it is shown that the capacity region Λ is the set of all rate vectors $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_N)$ such that for each of the $2^N - 1$ non-empty link subsets $\mathcal{L} \subset \{1, \ldots, N\}$, we have:

$$\sum_{i\in\mathcal{L}} \lambda_i \leq 1 - \Pi_{i\in\mathcal{L}}(1 - p_i) \tag{2}$$

This is an explicit description of the capacity region Λ. The following alternative implicit characterization is also useful for analysis (see [6] and references therein):

Theorem 1: (Capacity Region Λ) The capacity region Λ is equal to the set of all (non-negative) rate vectors $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_N)$ for which there exists a stationary randomized control policy that observes the current channel state vector $\boldsymbol{S}(t)$ and chooses a feasible transmission rate vector $\boldsymbol{\mu}(t) \in \mathcal{F}(t)$ as a random function of $\boldsymbol{S}(t)$, such that:

$$\lambda_i = \mathbb{E}\{\mu_i(t)\} \quad \text{for all } i \in \{1, \ldots, N\} \tag{3}$$

where the expectation is taken with respect to the random channel vector $\boldsymbol{S}(t)$ and the potentially random control action that depends on $\boldsymbol{S}(t)$. $\Box$

It is easy to see that any non-negative rate vector that is entrywise less than or equal to a vector $\boldsymbol{\lambda} \in \Lambda$ is also contained in Λ. This follows immediately from Theorem 1 by modifying the stationary randomized policy $\boldsymbol{\mu}(t)$ that yields $\mathbb{E}\{\mu_i(t)\} = \lambda_i$ to a new policy $\hat{\boldsymbol{\mu}}(t)$ by probabilistically setting each $\mu_i(t)$ value to zero with an appropriate probability q_i, yielding $\mathbb{E}\{\hat{\mu}_i(t)\} = \mathbb{E}\{\mu_i(t)\}(1 - q_i) \leq \mathbb{E}\{\mu_i(t)\}$.

It is also easy to show that the capacity region Λ is convex and compact (i.e., closed and bounded). Further, if $\mathbb{E}\{S_i(t)\} > 0$ for all $i \in \{1, \ldots, N\}$, then Λ has full dimension of size N and hence has a non-empty interior.

C. The Max-Weight Scheduling Policy

It is important to note that the capacity region is defined in terms of stationary randomized policies that make decisions based only on the channel state (and hence are independent of the queue state). Given a rate vector $\boldsymbol{\lambda}$ interior to Λ, such a stationary randomized policy could in principle be designed to stabilize the system, although this would require full knowledge of the traffic rates and channel state probabilities. However, it is well known that the following queue-aware *max-weight* policy stabilizes the system whenever the rate vector is interior to Λ, without requiring knowledge of the traffic rates or channel statistics [2]: Each slot t, observe current queue backlogs and channel states $Q_i(t)$ and $S_i(t)$ for each link i, and choose to serve the link $i^*(t) \in \{1, \ldots, N\}$ with the largest $Q_i(t)S_i(t)$ product. This is also called the *Longest Connected Queue* policy (LCQ) [2], as it serves the queue with the largest backlog among all that are currently ON.

The max-weight policy is very important because of its simplicity and its general stability properties. However, a tight delay analysis is quite challenging, and prior work provides only a loose upper bound on average delay that is $O(N)$, i.e., linear in the network size [5] [6]. It is shown in [7] that $O(1)$ average delay is possible when both channels and packet arrivals are independent across users and across timeslots, and when no traffic rate is larger than the average traffic rate by more than a specified amount. The $O(1)$ delay analysis of [7] uses an algorithm called *Largest Connected Group* that is different from the max-weight policy and that requires more statistical knowledge to implement. In the following, we use the queue grouping analysis techniques of [7] to show that the simpler max-weight policy can *also* provide $O(1)$ average delay, and does so for all traffic rates within a ρ-scaled version of the capacity region. However, the scaling in ρ is worse than that in [7]. Section IV recovers the same ρ scaling as [7] under a similar "f-balanced" traffic assumption.

III. Delay Analysis for Arbitrary Rates in Λ

Consider the ON/OFF channel model where each $S_i(t)$ is an independent i.i.d. Bernoulli process with $Pr[S_i(t) = 1] = p_i$. Assume the arrival rate vector $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_N)$ is interior to the capacity region Λ, so that there exists a value ρ such that $0 < \rho < 1$ and:

$$\boldsymbol{\lambda} \in \rho\Lambda \tag{4}$$

That is, $\boldsymbol{\lambda}$ is contained within a ρ-scaled version of the capacity region. The parameter ρ can be viewed as the *network loading*, measuring the fraction the rate vector $\boldsymbol{\lambda}$ is away from the capacity region boundary. Define $A_{tot}(t)$ as the total packet arrivals on slot t:

$$A_{tot}(t) \triangleq \sum_{i=1}^{N} A_i(t)$$

Define $\lambda_{tot} = \sum_{i=1}^{N} \lambda_i$ as the sum packet arrival rate.

A. Important Parameters of Λ

To analyze delay, it is useful to characterize the N dimensional capacity region Λ in terms of its size on subspaces of smaller dimension. To this end, define p_{min} as the smallest channel ON probability:

$$p_{min} \triangleq \min_{i \in \{1,\ldots,N\}} p_i$$

We assume that $0 < p_{min} < 1$. For each integer K such that $1 \leq K \leq N$, define parameters μ_K^{sym} and r_K as follows:

$$\mu_K^{sym} \triangleq \frac{1}{K}[1 - (1 - p_{min})^K]$$
$$r_K \triangleq 1 - (1 - p_{min})^K$$

Thus, $r_K = K\mu_K^{sym}$. It is not difficult to show that $\mu_K^{sym} > \mu_{K+1}^{sym}$ for all positive integers K. Further, let $\mathcal{L}_K$ represent a particular subset of K links within the link set $\{1, \ldots, N\}$. For each subset $\mathcal{L}_K$, define $\mathbf{1}_{\mathcal{L}_K}$ as a N-dimensional vector that is 1 in all entries $i \in \mathcal{L}_K$, and zero in all other entries.

Lemma 1: For each set $\mathcal{L}_K$ of size K, we have:

$$\mu_K^{sym}\mathbf{1}_{\mathcal{L}_K} \in \Lambda$$

Proof: By (2), it suffices to show that for any integer m such that $1 \leq m \leq K$, the sum of any m non-zero components of $\mu_K^{sym}\mathbf{1}_{\mathcal{L}_K}$ is less than or equal to r_m.[2] That is, it suffices to show that $m\mu_K^{sym} \leq r_m$. But this is equivalent to showing that $\mu_K^{sym} \leq \mu_m^{sym}$ for $m \leq K$, which is true. ∎

Thus, μ_K^{sym} can be intuitively viewed as an edge size such that any K-dimensional hypercube of this edge size (with dimensions defined along the orthogonal directions of any K axes of $\mathbb{R}^N$) can fit inside the capacity region Λ.

B. The $O(1)$ Delay Bound for Arbitrary Traffic in Λ

Suppose the LCQ algorithm is used together with a stationary probabilistic tie breaking rule in cases when multiple queues have the same weight. This allows the queueing system to be viewed as a stationary Markov chain. In this case, it is well known that if the arrival rate vector is interior to the capacity region, then all queues are stable under LCQ, with a well defined steady state time average [6]. The following $O(N)$ delay bound for LCQ is given in [7]:[3]

$$\overline{W} \leq \frac{N[1 + \frac{1}{\lambda_{tot}}\sum_{i=1}^{N}\mathbb{E}\left\{A_i(t)^2\right\} - \frac{2}{\lambda_{tot}}\sum_{i=1}^{N}\lambda_i^2]}{2r_N(1-\rho)} \tag{5}$$

where $\overline{W}$ represents the average delay in the system. The bound (5) also holds for arrival vectors $\boldsymbol{A}(t)$ that are i.i.d. over slots but with possibly correlated entries $A_i(t)$ on the same slot t. The next theorem demonstrates an improved $O(1)$ bound in the case when all arrival processes $A_i(t)$ are independent.

Theorem 2: (Delay Bound for LCQ) Consider the ON/OFF channel model and assume processes $A_i(t)$ and $S_i(t)$ are independent and i.i.d. over slots. Suppose there is an integer K such that $1 \leq K < N$, and such that $r_{K+1} > \lambda_{tot}$, that is:

$$1 - (1 - p_{min})^{K+1} > \lambda_{tot}$$

[2] Note that $r_m \leq 1 - \Pi_{i \in \mathcal{L}m}(1 - p_i)$, where $\mathcal{L}_m$ is any subset of m links.

[3] The bound in [7] is of the form c/ϵ, where ϵ is any value such that $\boldsymbol{\lambda} + \boldsymbol{\epsilon} \in \Lambda$, where $\boldsymbol{\epsilon}$ is a vector with all values equal to ϵ. The bound (5) follows by observing that $\epsilon = (1 - \rho)r_N/N$ satisfies $\boldsymbol{\lambda} + \boldsymbol{\epsilon} \in \Lambda$ whenever $\boldsymbol{\lambda} \in \rho\Lambda$. A similar c/ϵ bound is given in [5] [6] for more general multi-rate systems.

Assuming that $\boldsymbol{\lambda} \in \rho\Lambda$ for some network loading ρ such that $0 < \rho < 1$, then the max-weight (LCQ) policy for the ON/OFF channel model stabilizes all queues and yields:

$$\sum_{i=1}^{N} \overline{Q}_i \leq \frac{KBC}{(1-\rho)^2} \tag{6}$$

where $\overline{Q}_i$ is the time average number of packets in queue i, and where the constants B, C, and θ are defined:

$$\begin{aligned} B &\triangleq \frac{\lambda_{tot}}{2} + \frac{1}{2}\sum_{i=1}^{N}\mathbb{E}\left\{A_i(t)^2\right\} - \sum_{i=1}^{N}\lambda_i^2 \\ &\quad + \frac{\theta}{2}\left[\mathbb{E}\left\{A_{tot}(t)^2\right\} + \lambda_{tot} - 2\lambda_{tot}^2\right] \qquad (7) \\ C &\triangleq \frac{r_{K+1}}{\frac{r_N K \lambda_{tot}}{N(1-\rho)} + \frac{r_K(r_{K+1}-\lambda_{tot})}{(1-\rho)}} \qquad (8) \\ \theta &\triangleq \frac{(1-\rho)(\mu_K^{sym} - \mu_N^{sym})}{r_{K+1}} \qquad (9) \end{aligned}$$

By Little's Theorem, average delay $\overline{W}$ thus satisfies:

$$\overline{W} \leq \frac{KBC}{\lambda_{tot}(1-\rho)^2} \;\square$$

The proof of Theorem 2 is given in the next subsection. The above bound can be minimized over all K that satisfies $r_{K+1} > \lambda_{tot}$. For a simpler interpretation of the bound that illuminates the fact that this is an $O(1)$ delay result, note that for a given $\rho < 1$ we can choose K to satisfy:

$$1 - (1 - p_{min})^{K+1} \geq (1+\rho)/2$$

Indeed, choosing K as follows accomplishes this:[4]

$$K = \max\left[1, \left\lceil \frac{\log(2/(1-\rho))}{\log(1/(1-p_{min}))} \right\rceil - 1\right] \tag{10}$$

Because $\lambda_{tot} \leq \rho$, it is not difficult to show that with this choice of K, we have $r_{K+1} - \lambda_{tot} \geq (1-\rho)/2$, and so:

$$C \leq \frac{r_{K+1}}{\frac{r_K(r_{K+1}-\lambda_{tot})}{(1-\rho)}} \leq \frac{2r_{K+1}}{r_K}$$

It follows that $C = O(1)$, $K = O(\log(\frac{1}{1-\rho}))$, and that (if arrival processes are independent so that $\mathbb{E}\left\{A_{tot}(t)^2\right\} = O(1)$) $B = O(1)$. Therefore, the above delay bound yields $\overline{W} \leq c\log(\frac{1}{1-\rho})/(1-\rho)^2$, where c is a constant that does not depend on ρ or N. Therefore, $\overline{W}$ does not grow with N and so LCQ is *order-optimal* with respect to N.

C. Lyapunov Drift Analysis

Let $\boldsymbol{Q}(t) = (Q_1(t), \ldots, Q_N(t))$ be the vector of queue backlogs. Define $Q_{tot}(t)$ as the sum queue backlog in all queues of the system:

$$Q_{tot}(t) \triangleq \sum_{i=1}^{N} Q_i(t) \tag{11}$$

[4]Note that if N is small, we may have that the value of K given in (10) is greater than or equal to N. The old delay bound (5) should be used in this case. When N is large, our new bound is tighter and demonstrates order-optimal behavior.

Define the following Lyapunov function of the queue backlog vector $\boldsymbol{Q}(t)$:

$$L(\boldsymbol{Q}(t)) \triangleq \frac{1}{2}\sum_{i=1}^{N} Q_i(t)^2 + \frac{\theta}{2}(\sum_{j=1}^{N} Q_j(t))^2 \tag{12}$$

where θ is a positive constant to be determined later. Thus, $L(\boldsymbol{Q}(t)) = \frac{1}{2}\sum_{i=1}^{N} Q_i(t)^2 + \frac{\theta}{2}Q_{tot}(t)^2$. This Lyapunov function uses the standard sum of squares of queue length, and adds a new term that is the square of the total queue backlog. This new term incorporates the *queue grouping* concept similar to [7], and will be important in obtaining tight delay bounds. The technique of composing this Lyapunov function as a sum of two different quadratic terms weighted by a θ constant shall be useful in analyzing both stability and delay in two different modes of network operation, and is inspired by a similar technique used in [17] to analyze stability in a very different context. Specifically, work in [17] considers multi-hop networks with greedy maximal scheduling and achieves stability results when input rates are a constant factor (such as a factor of 2) away from the capacity region boundary. Here, we consider a single-hop network with time-varying channels, and obtain both stability and order-optimal delay results for all input rates interior to the capacity region Λ.

The queue dynamics (1) can be rewritten as follows:

$$Q_i(t+1) = Q_i(t) - \tilde{\mu}_i(t) + A_i(t) \tag{13}$$

where $\tilde{\mu}_i(t) = \min[Q_i(t), \mu_i(t)]$. Define $\tilde{\mu}_{tot}(t) = \sum_{i=1}^{N}\tilde{\mu}_i(t)$, being either 0 or 1, and being 1 if and only if the system serves a packet on slot t. The dynamics for $Q_{tot}(t)$ are given by:

$$Q_{tot}(t+1) = Q_{tot}(t) - \tilde{\mu}_{tot}(t) + A_{tot}(t) \tag{14}$$

where $A_{tot}(t) = \sum_{i=1}^{N} A_i(t)$. Let $\boldsymbol{Q}(t)$ be the stochastic queue evolution process for a given control policy. Define the one-step conditional Lyapunov drift as follows:[5]

$$\Delta(\boldsymbol{Q}(t)) \triangleq \mathbb{E}\left\{L(\boldsymbol{Q}(t+1)) - L(\boldsymbol{Q}(t)) \mid \boldsymbol{Q}(t)\right\} \tag{15}$$

Lemma 2: The Lyapunov drift $\Delta(\boldsymbol{Q}(t))$ for the ON/OFF channel model satisfies:

$$\begin{aligned} \Delta(\boldsymbol{Q}(t)) &= \mathbb{E}\left\{B(t) \mid \boldsymbol{Q}(t)\right\} \\ &\quad - \sum_{i=1}^{N} Q_i(t)\mathbb{E}\left\{\mu_i(t) - \lambda_i \mid \boldsymbol{Q}(t)\right\} \\ &\quad - \theta Q_{tot}(t)\mathbb{E}\left\{\tilde{\mu}_{tot}(t) - \lambda_{tot} \mid \boldsymbol{Q}(t)\right\} \end{aligned}$$

where $\mu_i(t)$ and $\tilde{\mu}_{tot}(t)$ correspond to the LCQ policy, and where $B(t)$ is given by:

$$\begin{aligned} B(t) &\triangleq \frac{\tilde{\mu}_{tot}(t)}{2} + \frac{1}{2}\sum_{i=1}^{N}[A_i(t)^2 - 2A_i(t)\tilde{\mu}_i(t)] \\ &\quad + \frac{\theta}{2}[A_{tot}(t)^2 + \tilde{\mu}_{tot}(t) - 2\tilde{\mu}_{tot}(t)A_{tot}(t)] \qquad (16) \end{aligned}$$

Proof: (Lemma 2) See Appendix A. ■

[5]Strictly speaking, correct notation should be $\Delta(\boldsymbol{Q}(t), t)$, as the drift could be from a non-stationary policy that depends on time t, although we use the simpler notation $\Delta(\boldsymbol{Q}(t))$ as formal notation for the right hand side of (15).

Now note that the LCQ algorithm chooses $\boldsymbol{\mu}(t) \in \mathcal{F}(t)$ on each slot t to maximize $\sum_{i=1}^{N} Q_i(t)\mu_i(t)$, and hence:

$$\sum_{i=1}^{N} Q_i(t)\mu_i(t) \geq \sum_{i=1}^{N} Q_i(t)\mu_i^*(t)$$

where $\boldsymbol{\mu}^*(t) = (\mu_1^*(t), \ldots, \mu_N^*(t))$ is any other feasible transmission rate vector in $\mathcal{F}(t)$. It follows that the above inequality is preserved when taking conditional expectations given the current $\boldsymbol{Q}(t)$ value. Plugging this result into the second term on the right hand side of the drift expression in Lemma 2 thus yields:

$$\begin{aligned}\Delta(\boldsymbol{Q}(t)) \quad \leq \quad & \mathbb{E}\left\{B(t) \mid \boldsymbol{Q}(t)\right\} \\ & -\sum_{i=1}^{N} Q_i(t)\mathbb{E}\left\{\mu_i^*(t) - \lambda_i \mid \boldsymbol{Q}(t)\right\} \\ & -\theta Q_{tot}(t)\mathbb{E}\left\{\tilde{\mu}_{tot}(t) - \lambda_{tot} \mid \boldsymbol{Q}(t)\right\} \quad (17)\end{aligned}$$

where $\boldsymbol{\mu}^*(t) = (\mu_1^*(t), \ldots, \mu_N^*(t))$ is any other feasible control action on slot t. Note that $\tilde{\mu}_{tot}(t)$ in the above expression still corresponds to the LCQ policy.

Let $L(t)$ represent the number of non-empty queues on slot t, so that $0 \leq L(t) \leq N$.

- *Case 1* ($L(t) \leq K$): Suppose $L(t) \leq K$, and let $\mathcal{L}(t)$ represent the set of non-empty queue indices. Recall that $\mu_K^{sym}\mathbf{1}_{\mathcal{L}(t)} \in \Lambda$ and that $\boldsymbol{\lambda}/\rho \in \Lambda$. By taking a convex combination of these two vectors and using convexity of the set Λ, it follows that:

$$\boldsymbol{\lambda} + (1-\rho)\mu_K^{sym}\mathbf{1}_{\mathcal{L}(t)} \in \Lambda \quad (18)$$

Now let $\boldsymbol{\mu}^*(t)$ be the stationary randomized policy that makes decisions based only on the current channel state, and that yields for all $i \in \mathcal{L}(t)$:

$$\mathbb{E}\left\{\mu_i^*(t)\right\} = \lambda_i + (1-\rho)\mu_K^{sym} \quad (19)$$

Such a policy exists by (18) and Theorem 1. Using (19) in the drift inequality (17) yields:

$$\begin{aligned}\Delta(\boldsymbol{Q}(t)) \quad \leq \quad & \mathbb{E}\left\{B(t) \mid \boldsymbol{Q}(t)\right\} - \sum_{i=1}^{N} Q_i(t)(1-\rho)\mu_K^{sym} \\ & +\theta Q_{tot}(t)\lambda_{tot} \\ = \quad & \mathbb{E}\left\{B(t) \mid \boldsymbol{Q}(t)\right\} \\ & -Q_{tot}(t)[(1-\rho)\mu_K^{sym} - \theta\lambda_{tot}]\end{aligned}$$

Define ϵ as follows:

$$\epsilon \triangleq [(1-\rho)\mu_K^{sym} - \theta\lambda_{tot}] \quad (20)$$

It follows that:

$$\Delta(\boldsymbol{Q}(t)) \leq \mathbb{E}\left\{B(t) \mid \boldsymbol{Q}(t)\right\} - \epsilon Q_{tot}(t) \quad (21)$$

- *Case 2* ($L(t) > K$): Suppose $L(t) > K$, and again let $\mathcal{L}(t)$ represent the set of non-empty queue indices. Note that $\boldsymbol{\lambda}/\rho \in \Lambda$ and $\mu_N^{sym}\mathbf{1} \in \Lambda$, where $\mathbf{1}$ is the all 1 vector. By convexity of Λ, the convex combination is also in Λ:

$$\boldsymbol{\lambda} + (1-\rho)\mu_N^{sym}\mathbf{1} \in \Lambda$$

Now let $\boldsymbol{\mu}^*(t)$ be the stationary randomized policy that makes decisions independent of queue backlog, and that yields for all $i \in \{1, \ldots, N\}$:

$$\mathbb{E}\left\{\mu_i^*(t)\right\} = \lambda_i + (1-\rho)\mu_N^{sym} \quad (22)$$

Such a policy exists by Theorem 1. Note that when the number of non-empty queues is greater than K, there is a packet departure under the LCQ policy with probability at least one minus the product of the $K+1$ largest OFF probabilities:

$$\mathbb{E}\left\{\tilde{\mu}_{tot} \mid \boldsymbol{Q}(t)\right\} \geq 1 - \Pi_{i \in \hat{\mathcal{L}}_{K+1}}(1-p_i) \geq r_{K+1} \quad (23)$$

where $\hat{\mathcal{L}}_{K+1}$ represents the set of $K+1$ links with the smallest success probabilities. Plugging (22) and (23) into the drift inequality (17) yields:

$$\begin{aligned}\Delta(\boldsymbol{Q}(t)) \quad \leq \quad & \mathbb{E}\left\{B(t) \mid \boldsymbol{Q}(t)\right\} \\ & -Q_{tot}(t)[(1-\rho)\mu_N^{sym} + \theta(r_{K+1} - \lambda_{tot})]\end{aligned}$$

To equalize the drift in both Case 1 and Case 2, we choose θ to satisfy:

$$\epsilon = (1-\rho)\mu_N^{sym} + \theta(r_{K+1} - \lambda_{tot})$$

Thus (using (20)):

$$\begin{aligned}\theta \quad & = \quad \frac{(1-\rho)(\mu_K^{sym} - \mu_N^{sym})}{r_{K+1}} \\ \epsilon \quad & = \quad \frac{(1-\rho)\left[\mu_N^{sym}\lambda_{tot} + \mu_K^{sym}(r_{K+1} - \lambda_{tot})\right]}{r_{K+1}}\end{aligned}$$

Note that indeed we have $\theta > 0$ (because $\mu_K^{sym} > \mu_N^{sym}$). Further, because $r_{K+1} > \lambda_{tot}$, we have that $\epsilon > 0$. Therefore, the drift inequality (21) holds in both Case 1 and Case 2 (and hence holds for all t and all $\boldsymbol{Q}(t)$). We now use the following well known Lyapunov drift lemma (see, for example, [6] for a proof):

Lemma 3: (Lyapunov Drift) If the drift $\Delta(\boldsymbol{Q}(t))$ of a non-negative Lyapunov function satisfies the following for all t and all $\boldsymbol{Q}(t)$:

$$\Delta(\boldsymbol{Q}(t)) \leq \mathbb{E}\left\{B(t) \mid \boldsymbol{Q}(t)\right\} - \epsilon\mathbb{E}\left\{f(t) \mid \boldsymbol{Q}(t)\right\}$$

for some stochastic processes $B(t)$, $f(t)$, and some constant $\epsilon > 0$, then:

$$\overline{f} \leq \overline{B}/\epsilon$$

where

$$\begin{aligned}\overline{f} \quad & \triangleq \quad \limsup_{t\to\infty} \frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left\{f(\tau)\right\} \\ \overline{B} \quad & \triangleq \quad \limsup_{t\to\infty} \frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left\{B(\tau)\right\} \square\end{aligned}$$

Using this Lyapunov drift lemma in (21) (using $f(t) = Q_{tot}(t)$) yields:

$$\overline{Q}_{tot} \leq \overline{B}/\epsilon$$

We note that because the system evolves according to a Markov chain with a countably infinite state space, the time averages are well defined (so that the $\limsup$ can be replaced by a regular limit). Further, using the fact that $\lim_{t\to\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\{\tilde{\mu}_i(\tau)\}=\lambda_i$, the value of $\overline{B}$ can be seen to be equal to the value B defined in (7), proving Theorem 2.

IV. A Tighter Bound for "f-Balanced Traffic"

Here we present a tighter bound on average backlog and delay of the LCQ algorithm for the ON/OFF channel model. Our bound in this section is of the form $c\log(\frac{1}{1-\rho})/(1-\rho)$, which is still $O(1)$ with respect to the network size N, but yields a better asymptotic in ρ. Unfortunately, our analysis does not hold for all rate vectors $\boldsymbol{\lambda}$ inside the capacity region Λ. Rather, we make the following assumption about a more "balanced" traffic rate vector. Let $\boldsymbol{\lambda}=(\lambda_1,\ldots,\lambda_N)$, and without loss of generality assume that $\lambda_i>0$ for all $i\in\{1,\ldots,N\}$ (else, we can redefine N to be the number of links with non-zero rates). Define $\lambda_{tot}=\sum_{i=1}^N\lambda_i$ and $\lambda_{av}=\lambda_{tot}/N$. We say that $\boldsymbol{\lambda}$ has *f-balanced rates* if there is a constant f such that:

$$\lambda_i\leq\lambda_{av}+f \quad \text{for all } i\in\{1,\ldots,N\} \tag{24}$$

That is, $\boldsymbol{\lambda}$ is f-balanced if no individual traffic rate is more than an amount f above the average rate λ_{av}. Clearly any *uniform* traffic rate vector is f-balanced for $f=0$. However, this definition of f-balanced rates also captures a large class of heterogeneous arrival rate vectors. We shall prove our delay results under the assumption that f is suitably small. A similar assumption is used in [7], and our delay analysis relies heavily on the queue-grouping techniques used there.

A. The Queue-Grouped Lyapunov Function

Fix an integer K such that $1\leq K\leq N$. Define $\hat{N}$ as the smallest multiple of K that is larger than or equal to N:

$$\hat{N}=\lceil N/K\rceil K \tag{25}$$

Now define a new rate vector $\hat{\boldsymbol{\lambda}}=(\lambda_1,\ldots,\lambda_N,0,0,\ldots,0)$, where the last $\hat{N}-N$ entries are zero. Define $\hat{N}-N$ "fictitious" queues for these last dimensions (these queues always have zero backlog, but shall be convenient to define for counting purposes). Define $\mathcal{G}_K$ as the set of all possible partitions of the link set $\{1,\ldots,\hat{N}\}$ into K disjoint sets, each with an equal size of $\hat{N}/K$ links. Let $g\in\mathcal{G}_K$ denote a particular partition, and define $\mathcal{L}_1^{(g)},\ldots,\mathcal{L}_K^{(g)}$ as the collection of sets corresponding to g (so that the union $\cup_{k=1}^K\mathcal{L}_k^{(g)}$ is equal to $\{1,\ldots,\hat{N}\}$, and the intersection $\mathcal{L}_n^{(g)}\cap\mathcal{L}_m^{(g)}$ is empty for all $m\neq n$, where $m,n\in\{1,\ldots,K\}$).

For a particular partition g, define $Q_k^{(g)}(t)$ as the sum of all queue backlogs in the kth set of g:

$$Q_k^{(g)}(t)\triangleq\sum_{i\in\mathcal{L}_k^{(g)}}Q_i(t)$$

Define the following *queue-grouped Lyapunov function*:

$$L(\boldsymbol{Q}(t))\triangleq\frac{1}{2}\sum_{g\in\mathcal{G}_K}\sum_{k=1}^K(Q_k^{(g)}(t))^2 \tag{26}$$

This is similar to the Lyapunov function of [7], with the exception that it sums over all possible partitions into K disjoint groups. Define $A_k^{(g)}(t)$ and $\tilde{\mu}_k^{(g)}(t)$ as the sum arrivals and departures from the kth group of the partition g:

$$A_k^{(g)}(t) \triangleq \sum_{i\in\mathcal{L}_k^{(g)}}A_i(t)$$

$$\tilde{\mu}_k^{(g)}(t) \triangleq \sum_{i\in\mathcal{L}_k^{(g)}}\tilde{\mu}_i(t)$$

The dynamics for the kth group of partition g thus satisfy:

$$Q_k^{(g)}(t+1)=Q_k^{(g)}(t)-\tilde{\mu}_k^{(g)}(t)+A_k^{(g)}(t) \tag{27}$$

Define the Lyapunov drift $\Delta(\boldsymbol{Q}(t))$ as before (given in (15)).

Lemma 4: For a general scheduling policy, the Lyapunov drift satisfies:

$$\Delta(\boldsymbol{Q}(t)) = \mathbb{E}\{C(t)\,|\,\boldsymbol{Q}(t)\} - \sum_{g\in\mathcal{G}_K}\sum_{k=1}^K Q_k^{(g)}(t)\mathbb{E}\left\{\tilde{\mu}_k^{(g)}(t)-\lambda_k^{(g)}\mid\boldsymbol{Q}(t)\right\}$$

where $\lambda_k^{(g)}\triangleq\sum_{i\in\mathcal{L}_k^{(g)}}\lambda_i$, and where $C(t)$ is defined:

$$C(t)\triangleq\frac{1}{2}\sum_{g\in\mathcal{G}_K}\sum_{k=1}^K\left[\tilde{\mu}_k^{(g)}(t)+A_k^{(g)}(t)^2-2\tilde{\mu}_k^{(g)}(t)A_k^{(g)}(t)\right] \tag{28}$$

Proof: The proof is similar to the proof of Lemma 2, and is omitted for brevity. ■

Remarkably, we next show that the "max-weight" LCQ algorithm for this ON/OFF channel model minimizes the final term in the right hand side of the above drift expression.

Lemma 5: (Max Weight Matching) Every slot t, the LCQ algorithm chooses a transmission rate vector $\boldsymbol{\mu}(t)\in\mathcal{F}(t)$ that maximizes the following expression over all alternative feasible transmission rate vectors:

$$\sum_{g\in\mathcal{G}_K}\sum_{k=1}^K Q_k^{(g)}(t)\tilde{\mu}_k^{(g)}(t)$$

Proof: See Appendix B. ■

It follows that we can replace the variables $\tilde{\mu}_k^{(g)}(t)$ in the final term of the drift expression in Lemma 4, which correspond to the LCQ policy, with variables $\tilde{\mu}_k^{(g)*}(t)$ that correspond to any other feasible rate vector $\boldsymbol{\mu}^*(t)\in\mathcal{F}(t)$, while creating an inequality relationship:

$$\Delta(\boldsymbol{Q}(t)) \leq \mathbb{E}\{C(t)\,|\,\boldsymbol{Q}(t)\} - \sum_{g\in\mathcal{G}_K}\sum_{k=1}^K Q_k^{(g)}(t)\mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t)-\lambda_k^{(g)}\mid\boldsymbol{Q}(t)\right\} \tag{29}$$

The drift inequality (29) is quite subtle: It is defined in terms of any other *single* feasible rate vector $\boldsymbol{\mu}^*(t)$ (where this vector does not depend on the partition g). Note that the variables $\tilde{\mu}_k^{(g)*}(t)$ are defined for different partitions $g\in\mathcal{G}_K$, but for each particular g these variables are still derived from the *same* vector $\boldsymbol{\mu}^*(t)$. They are derived from $\boldsymbol{\mu}^*(t)$ by summing the components of this rate vector that have non-empty queues over the dimensions that correspond to the groups within the particular partition g.

B. Optimizing the Drift Bound

Here we manipulate the sum in the right-hand side of (29) to yield a useful drift bound.

Lemma 6: For any vector $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_{\hat{N}})$, if there is a value β such that $0 < \beta < 1$ such that for all $i \in \{1, \ldots, N\}$ we have:

$$\lambda_i \leq \frac{\lambda_{tot}}{\hat{N}} + \frac{\beta(1-\rho)}{K} \tag{30}$$

then:

$$\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\lambda_k^{(g)} \leq Q_{tot}(t)\,|\mathcal{G}_K|\,\frac{[\lambda_{tot}+\beta(1-\rho)]}{K}$$

where $|\mathcal{G}_K|$ is the cardinality of $\mathcal{G}_K$ and $Q_{tot}(t)$ is the total sum backlog (defined in (11)).

Proof: The proof of Lemma 6 follows from simple counting arguments, and is given in Appendix C. ■

Note that $\lambda_{tot}/\hat{N} \leq \lambda_{tot}/N$ with approximate equality when N is large (so that $\hat{N}/N \approx 1$). The constraints (30) imply that the rate vector $\boldsymbol{\lambda}$ is f-balanced with $f = \beta(1-\rho)/K$.

Lemma 7: There exists a single randomized strategy that observes queue backlogs and channel states for slot t and chooses $\boldsymbol{\mu}^*(t) \in \mathcal{F}(t)$ such that:

$$\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t)\mid \boldsymbol{Q}(t)\right\} \geq Q_{tot}(t)\,|\mathcal{G}_K|\,\frac{r_K}{K}$$

where $r_K = 1-(1-p_{min})^K$.

Proof: See Appendix C. ■

Using Lemmas 6 and 7 in the drift inequality (29) yields:

Lemma 8: If $\boldsymbol{\lambda} \in \rho\Lambda$ (for $0<\rho<1$) and if (30) is satisfied for all $i \in \{1, \ldots, N\}$, then:

$$\begin{aligned}\Delta(\boldsymbol{Q}(t)) \leq\ & \mathbb{E}\{C(t)\mid \boldsymbol{Q}(t)\} \\ & -Q_{tot}(t)\,|\mathcal{G}_K|\,\frac{[r_K-\lambda_{tot}-\beta(1-\rho)]}{K}\end{aligned} \tag{31}$$

Lemma 8 leads immediately to the delay theorem stated in the next subsection.

C. An Improved Delay Bound for f-Balanced Traffic

Theorem 3: (Delay Bound for ON/OFF Channels with f-Balanced Traffic) Suppose $\boldsymbol{\lambda} \in \rho\Lambda$ for $0<\rho<1$. Let K be the smallest integer that satisfies $r_K \geq (1+\rho)/2$, that is:

$$[1-(1-p_{min})^K] \geq (1+\rho)/2 \tag{32}$$

Suppose that the f-balanced traffic constraints (30) are satisfied for some value β such that $0 \leq \beta < 1/2$. If the max-weight (LCQ) policy is used on this ON/OFF channel model, then average queue occupancy satisfies:

$$\overline{Q}_{tot} \leq \frac{KD}{(1-\rho)(\frac{1}{2}-\beta)} \quad , \quad \overline{W} \leq \frac{K(D/\lambda_{tot})}{(1-\rho)(\frac{1}{2}-\beta)}$$

where D is defined:

$$D \triangleq \frac{1}{2}\left[\lambda_{tot}+\mathbb{E}\left\{A_{tot}(t)^2\right\}\right]$$

Further, in the special case when N is a multiple of K, and when traffic is uniform and Poisson with $\lambda_i = \lambda_{tot}/N$ for all i, we have $\beta = 0$ and:

$$\overline{Q}_{tot} \leq \frac{[2K\lambda_{tot}-\lambda_{tot}^2]}{1-\rho} \quad , \quad \overline{W} \leq \frac{2K-\lambda_{tot}}{1-\rho}$$

Note that the constraint (32) is satisfied by:

$$K = \left\lceil \frac{\log(\frac{2}{1-\rho})}{\log(1/(1-p_{min}))} \right\rceil$$

Therefore, $K = O(\log(\frac{1}{1-\rho}))$. Assuming that traffic streams are independent, so that $\mathbb{E}\left\{A_{tot}(t)^2\right\} = O(1)$, implies that $D = O(1)$. Thus the delay bound gives $\overline{W} \leq O(\frac{\log(1/(1-\rho))}{1-\rho})$, being independent of the network size N and having an asymptotic in ρ that is better than that of Theorem 2.

Proof: (Theorem 3) Because $\boldsymbol{\lambda} \in \rho\Lambda$, we have $\lambda_{tot} \leq \rho$ (as the maximum sum rate is at most $r_N \leq 1$). The assumption on r_K in (32) thus implies:

$$[r_K - \lambda_{tot} - \beta(1-\rho)] \geq (1-\rho)(\frac{1}{2}-\beta)$$

The above value is strictly positive because $\beta < 1/2$. Using the drift inequality (31) directly in the Lyapunov Drift Lemma (Lemma 3) yields:

$$\overline{Q}_{tot} \leq \frac{K\overline{C}}{|\mathcal{G}_K|\,(1-\rho)(\frac{1}{2}-\beta)}$$

Using the definition of $C(t)$ in (28) and the fact that the system is stable (so the long term departure rate is equal to λ_{tot}) yields:

$$\begin{aligned}\overline{C} &= \frac{|\mathcal{G}_K|\,\lambda_{tot}}{2} + \frac{1}{2}\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K}\left[\mathbb{E}\left\{A_k^{(g)}(t)^2\right\} - 2(\lambda_k^{(g)})^2\right] \\ &\leq |\mathcal{G}_K|\left[\frac{\lambda_{tot}}{2} + \frac{\mathbb{E}\left\{A_{tot}(t)^2\right\}}{2}\right] = |\mathcal{G}_K|\,D\end{aligned}$$

The above bound on $\overline{C}$ proves the first part of the theorem. The second part, for uniform Poisson traffic, follows by the above equality for $\overline{C}$ (without the bound), using $\mathbb{E}\left\{A_k^{(g)}(t)\right\} = \frac{\lambda_{tot}}{K}$ and $\mathbb{E}\left\{A_k^{(g)}(t)^2\right\} = \frac{\lambda_{tot}^2}{K^2} + \frac{\lambda_{tot}}{K}$ for all g, k. ■

V. Conclusions

We have presented an improved delay analysis for the max-weight scheduling algorithm. For ON/OFF channels, max-weight is equivalent to Longest Connected Queue (LCQ), and yields average delay that is order-optimal, being independent of the network size N. If an f-balanced traffic assumption holds, average delay was shown to maintain independence of N while allowing an improved asymptotic in ρ. Our technical report [1] extends this analysis to a modified max-weight policy for the case of multi-rate channels, using similar techniques and using additional stochastic coupling arguments. Our delay analysis makes use of the technique of queue grouping. The particular Lyapunov functions introduced for this delay analysis are powerful and may be useful in other contexts.

APPENDIX A — PROOF OF LEMMA 2

Here we prove Lemma 2. Define $\Delta_1(\boldsymbol{Q}(t))$ and $\Delta_2(\boldsymbol{Q}(t))$ as the conditional drift for the sum of squares term and the square of queue backlog term, respectively, so that $\Delta(\boldsymbol{Q}(t)) = \Delta_1(\boldsymbol{Q}(t)) + \theta\Delta_2(\boldsymbol{Q}(t))$. Squaring (13) and using the fact that $\tilde{\mu}_i(t)^2 = \tilde{\mu}_i(t)$ and $Q_i(t)\tilde{\mu}_i(t) = Q_i(t)\mu_i(t)$ (because $\tilde{\mu}_i(t) \in \{0,1\}$, and $\tilde{\mu}_i(t) = \mu_i(t)$ if $Q_i(t) > 0$): yields:

$$\begin{aligned}\frac{1}{2}Q_i(t+1)^2 &= \frac{1}{2}Q_i(t)^2 + \frac{(A_i(t) - \tilde{\mu}_i(t))^2}{2} \\ &\quad -Q_i(t)(\mu_i(t) - A_i(t)) \\ &= \frac{1}{2}\left[Q_i(t)^2 + A_i(t)^2 + \tilde{\mu}_i(t)\right] - A_i(t)\tilde{\mu}_i(t) \\ &\quad -Q_i(t)(\mu_i(t) - A_i(t))\end{aligned}$$

Therefore:

$$\begin{aligned}\Delta_1(\boldsymbol{Q}(t)) &= \mathbb{E}\left\{B_1(t) \,|\, \boldsymbol{Q}(t)\right\} \\ &\quad -\sum_{i=1}^{N} Q_i(t)\mathbb{E}\left\{\mu_i(t) - \lambda_i \,|\, \boldsymbol{Q}(t)\right\}\end{aligned}$$

where

$$B_1(t) \triangleq \sum_{i=1}^{N}\left[\frac{1}{2}\left[A_i(t)^2 + \tilde{\mu}_i(t)\right] - A_i(t)\tilde{\mu}_i(t)\right]$$

Similarly:

$$\begin{aligned}\frac{1}{2}Q_{tot}(t+1)^2 &= \frac{1}{2}\left[Q_{tot}(t)^2 + \tilde{\mu}_{tot}(t)^2 + A_{tot}(t)^2\right] \\ &\quad -\tilde{\mu}_{tot}(t)A_{tot}(t) \\ &\quad -Q_{tot}(t)(\tilde{\mu}_{tot}(t) - A_{tot}(t))\end{aligned}$$

Therefore:

$$\begin{aligned}\Delta_2(\boldsymbol{Q}(t)) &= \mathbb{E}\left\{B_2(t) \,|\, \boldsymbol{Q}(t)\right\} \\ &\quad -Q_{tot}(t)\mathbb{E}\left\{\tilde{\mu}_{tot}(t) - \lambda_{tot} \,|\, \boldsymbol{Q}(t)\right\}\end{aligned}$$

where $\tilde{\mu}_{tot}(t)^2 = \tilde{\mu}_{tot}(t)$ (because it is either 0 or 1) , and

$$B_2(t) \triangleq \frac{1}{2}\left[\tilde{\mu}_{tot}(t) + A_{tot}(t)^2\right] - \tilde{\mu}_{tot}(t)A_{tot}(t)$$

Summing $\Delta_1(\boldsymbol{Q}(t))$ and $\theta\Delta_2(\boldsymbol{Q}(t))$ and noting from (16) that $B(t) = B_1(t) + \theta B_2(t)$ yields the result of Lemma 2.

APPENDIX B — PROOF OF LEMMA 5

Define the integer $m = \hat{N}/K$. Here we prove Lemma 5. Given a particular queue backlog vector $\boldsymbol{Q}(t)$, the LCQ algorithm maximizes the expression $\sum_{i=1}^{\hat{N}} Q_i(t)\mu_i(t)$ over all $\boldsymbol{\mu}(t) \in \mathcal{F}(t)$. We now show that this *also* maximizes the expression given in Lemma 5. To this end, we have:

$$\begin{aligned}\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\tilde{\mu}_k^{(g)}(t) &= \sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\sum_{i\in\mathcal{L}_k^{(g)}}\tilde{\mu}_i(t) \\ &= \sum_{i=1}^{\hat{N}}\tilde{\mu}_i(t)Q_i(t)\,|\mathcal{G}_K| \\ &\quad + \sum_{i=1}^{\hat{N}}\tilde{\mu}_i(t)\sum_{j\neq i}Q_j(t)\frac{|\mathcal{G}_K|\,(m-1)}{\hat{N}-1}\end{aligned}$$

where the final equality holds because link i is in every group that multiplies the $\tilde{\mu}_i(t)$ term, and all other links multiply this term the same number of times (by group symmetry). The above also uses the fact that (by symmetry) the number of group partitions for which a particular link j is in the same group as link i is equal to the total number of partitions multiplied by the probability that a randomly chosen partition includes i and j in the same group. Define the above expression as $f(t)$ for simplicity. Therefore:

$$\begin{aligned}f(t) &= \sum_{i=1}^{\hat{N}}\tilde{\mu}_i(t)Q_i(t)\,|\mathcal{G}_K|\,(1 - \frac{m-1}{\hat{N}-1}) \\ &\quad + \sum_{i=1}^{\hat{N}}\tilde{\mu}_i(t)\left(\sum_{j=1}^{\hat{N}}Q_j(t)\right)|\mathcal{G}_K|\,\frac{m-1}{\hat{N}-1}\end{aligned} \quad (33)$$

The $\tilde{\mu}_i(t)$ values in the expression for $f(t)$ are the only ones affected by the control action on slot t. The final term on the right hand side is given by $\sum_i \tilde{\mu}_i(t)$ (the total departures on slot t) multiplied by the total sum backlog on slot t. This final term is maximized by any work conserving policy that always transmits a packet when there is a non-empty connected queue. The first term on the right hand side is a non-negative constant multiplied by the term $\sum_i \tilde{\mu}_i(t)Q_i(t)$. But note that $\tilde{\mu}_i(t)Q_i(t) = \mu_i(t)Q_i(t)$, and thus the LCQ policy maximizes this first term. As LCQ is work conserving, it also maximizes the second term, and thus maximizes $f(t)$, proving Lemma 5.

APPENDIX C — PROOF OF LEMMAS 6 AND 7

Proof: (Lemma 6) Define the integer $m = \hat{N}/K$. Using a counting argument similar to that of Appendix B (compare with (33)), we have that for any vector $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_{\hat{N}})$:

$$\begin{aligned}\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\lambda_k^{(g)} &= \sum_{i=1}^{\hat{N}} Q_i(t)\lambda_i\,|\mathcal{G}_K|\,(1 - \frac{m-1}{\hat{N}-1}) \\ &\quad + Q_{tot}(t)\lambda_{tot}\,|\mathcal{G}_K|\,\frac{m-1}{\hat{N}-1}\end{aligned} \quad (34)$$

Using the bound on λ_i given in (30) yields:

$$\begin{aligned}&\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\lambda_k^{(g)} \leq \\ &\quad Q_{tot}(t)\lambda_{tot}\,|\mathcal{G}_K|\left[\frac{1}{\hat{N}} - \frac{m-1}{\hat{N}(\hat{N}-1)} + \frac{m-1}{\hat{N}-1}\right] \\ &\quad + Q_{tot}(t)\,|\mathcal{G}_K|\,(1 - \frac{m-1}{\hat{N}-1})\beta(1-\rho)/K\end{aligned}$$

The result of Lemma 6 follows by using the identity:

$$\left[\frac{1}{\hat{N}} - \frac{m-1}{\hat{N}(\hat{N}-1)} + \frac{m-1}{\hat{N}-1}\right] = \frac{1}{K} \quad (35)$$

■

Proof: (Lemma 7) Let $L(t)$ represent the number of non-empty queues on slot t. If $L(t) = 0$, then $Q_{tot}(t) = 0$ and the result is trivial. Now suppose that $L(t) = l$, where $l \in \{1, 2, \ldots, N\}$. Define $(l_1, \ldots, l_{\hat{N}})$ to be a $0/1$ vector with

$l_i = 1$ if and only if $Q_i(t) > 0$. Define $l_k^{(g)}$ to be the number of non-empty queues in the kth group of partition g. Consider the following randomized policy for $\boldsymbol{\mu}^*(t) \in \mathcal{F}(t)$: First observe all channel states $S_i(t)$ for non-empty queues i, and define *new* channel states $\hat{S}_i(t)$ as follows: If $S_i(t) = 0$ (OFF), assign $\hat{S}_i(t) = 0$. If $S_i(t) = 1$ (ON), independently assign $\hat{S}_i(t) = 1$ with probability p_{min}/p_i (this is a valid probability because $p_{min} \leq p_i$). It follows that the new channel state vector $\hat{\boldsymbol{S}}(t)$ has independent and symmetric ON probabilities p_{min}. Now independently, randomly, and uniformly choose a queue to serve over all non-empty queues i with $\hat{S}_i(t) = 1$. It follows that for all non-empty queues i we have:

$$\mathbb{E}\left\{\tilde{\mu}_i^*(t) \mid \boldsymbol{Q}(t)\right\} = \frac{1-(1-p_{min})^l}{l} = \frac{r_l}{l}$$

Further, for any $g \in \mathcal{G}_K$ and any $k \in \{1, \ldots, K\}$ we have:

$$\mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t) \mid \boldsymbol{Q}(t)\right\} = \sum_{i \in \mathcal{L}_k^{(g)}} \mathbb{E}\left\{\tilde{\mu}_i^*(t) \mid \boldsymbol{Q}(t)\right\} = l_k^{(g)}\frac{r_l}{l}$$

Using this equality gives:

$$\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t) \mid \boldsymbol{Q}(t)\right\} = \frac{r_l}{l}\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)l_k^{(g)} \quad (36)$$

Now note that the $l_k^{(g)}$ values are structurally similar to the $\lambda_k^{(g)}$ values, and hence (similar to (34)) we have (using $Q_i(t)l_i = Q_i(t)$ and $l_{tot} = l$):

$$\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)l_k^{(g)} = \sum_{i=1}^{\hat{N}} Q_i(t)\left|\mathcal{G}_K\right|(1-\frac{m-1}{\hat{N}-1}) + Q_{tot}(t)l\left|\mathcal{G}_K\right|\frac{m-1}{\hat{N}-1}$$

Using this in (36) yields:

$$\begin{aligned}&\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t) \mid \boldsymbol{Q}(t)\right\} \\ &= \frac{r_l}{l}Q_{tot}(t)\left|\mathcal{G}_K\right|\left[1-\frac{m-1}{\hat{N}-1}+\frac{l(m-1)}{\hat{N}-1}\right] \quad (37)\\ &\geq r_lQ_{tot}(t)\left|\mathcal{G}_K\right|\left[\frac{1}{\hat{N}}-\frac{(m-1)}{\hat{N}(\hat{N}-1)}+\frac{m-1}{\hat{N}-1}\right] \\ &= r_lQ_{tot}(t)\left|\mathcal{G}_K\right|/K \quad (38)\end{aligned}$$

where the last equality holds by (35). The above holds for $L(t) = l \in \{1, \ldots, N\}$. Suppose now that $l \geq K$. In this case we have $r_l \geq r_K$, proving the result of Lemma 7 for $l \geq K$.

Consider now the final case where $l \in \{1, \ldots, K-1\}$. Then from (37) we have:

$$\sum_{g\in\mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t) \mid \boldsymbol{Q}(t)\right\} \geq \frac{r_l}{l}Q_{tot}(t)\left|\mathcal{G}_K\right| \quad (39)$$

Using the fact that $\frac{r_l}{l} \geq \frac{r_{K-1}}{K-1} \geq \frac{r_K}{K}$ yields the result. ■

REFERENCES

[1] M. J. Neely. Delay analysis for max weight opportunistic scheduling in wireless systems. *arXiv:0806.2345v1*, June 2008.

[2] L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory*, vol. 39, pp. 466-478, March 1993.

[3] N. Kahale and P. E. Wright. Dynamic global packet routing in wireless networks. *Proc. IEEE INFOCOM*, 1997.

[4] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, and P. Whiting. Providing quality of service over a shared wireless link. *IEEE Communications Magazine*, vol. 39, no.2, pp.150-154, 2001.

[5] 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.

[6] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1-149, 2006.

[7] M. J. Neely. Order optimal delay for opportunistic scheduling in multi-user wireless uplinks and downlinks. *Proc. of Allerton Conf. on Communication, Control, and Computing (invited paper)*, Sept. 2006.

[8] M. J. Neely, E. Modiano, and Y.-S. Cheng. Logarithmic delay for $n \times n$ packet switches under the crossbar constraint. *IEEE Transactions on Networking*, vol. 15, no. 3, pp. 657-668, June 2007.

[9] E. M. Yeh. *Multiaccess and Fading in Communication Networks*. PhD thesis, Massachusetts Institute of Technology, Laboratory for Information and Decision Systems (LIDS), 2001.

[10] A. Ganti, E. Modiano, and J. N. Tsitsiklis. Optimal transmission scheduling in symmetric communication models with intermittent connectivity. *IEEE Transactions on Information Theory*, vol. 53, no. 3, March 2007.

[11] S. Shakkottai, R. Srikant, and A. Stolyar. Pathwise optimality of the exponential scheduling rule for wireless channels. *Advances in Applied Probability*, vol. 36, no. 4, pp. 1021-1045, Dec. 2004.

[12] V. J. Venkataramanan and X. Lin. Structural properties of ldp for queue-length based wireless scheduling algorithms. *Proc. of 45th Annual Allerton Conference on Communication, Control, and Computing, Monticello, Illinois*, September 2007.

[13] A. Stolyar. Large deviations of queues under qos scheduling algorithms. *Proc. of 44th Annual Allerton Conference on Communication, Control, and Computing, Monticello, Illinois*, September 2006.

[14] S. Deb, D. Shah, and S. Shakkottai. Fast matching algorithms for repetitive optimization: An application to switch scheduling. *Proc. of 40th Annual Conference on Information Sciences and Systems (CISS), Princeton, NJ*, March 2006.

[15] M. J. Neely. Delay analysis for maximal scheduling in wireless networks with bursty traffic. *Proc. IEEE INFOCOM*, April 2008.

[16] M. Alanyali and M. Dashouk. On power-of-choice in downlink transmission scheduling. *Information Theory and Applications Workshop (ITA), University of California, San Diego*, January 2008.

[17] X. Wu, R. Srikant, and J. R. Perkins. Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. *IEEE Transactions on Mobile Computing*, June 2007.