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. We then consider multi-rate transmission models and show that average delay in this case typically does increase with the network size due to queues containing a small number of "residual" packets.

Index Terms—Queueing analysis

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*. While our previous queue grouping results in [7] also demonstrate that $O(1)$ delay is possible, this is the first such result for the simpler class of max-weight policies. 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 $c \log(1/(1 - \rho))/(1 - \rho)^2$, where c is a constant that does not depend on ρ or N.[1] This is in comparison to the previous delay bound of $cN/(1 - \rho)$ 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 $c \log(1/(1 - \rho))/(1 - \rho)$ 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.

Finally, we consider systems with multi-rate capabilities. We first provide a delay bound that grows linearly with N, similar to the bounds in [5] [6] but with an improved coefficient. We then provide an example multi-rate system and show that its average congestion and delay must grow at least linearly with N under *any* scheduling algorithm, due to many queues having a small number of "residual" packets. This is an important example and demonstrates that the $\Theta(N)$ behavior of the multi-rate delay bound is fundamental and cannot be avoided, highlighting a significant difference between single-rate and multi-rate systems.

It is known that order-optimal delay requires queue-based scheduling. Indeed, it is shown in [7] that average delay in an N-user 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 a delay gap between queue-aware and queue-independent algorithms 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

---

[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.

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 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 inside the capacity region. 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. Section V treats multi-rate systems. Section VI provides simulation results.

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 $S_i(t)$ is a non-negative integer that represents the current transmission rate (in units of packets/slot) available over channel i if this channel is selected for transmission on slot t. For most of this paper, we consider the simple case of ON/OFF channels, where $S_i(t) \in \{0, 1\}$ for all channels $i \in \{1, \ldots, N\}$ (multi-rate systems are treated in Section V). 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:

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

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 $\mathbb{E}\{A_i(t)^2\} < \infty$. 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:

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

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

$S(t)$ and chooses a feasible transmission rate vector $\boldsymbol{\mu}(t) \in \mathcal{F}(t)$ as a random function of $S(t)$, such that:

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

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

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

Given a rate vector $\boldsymbol{\lambda}$ interior to the capacity region Λ, a stationary, randomized, queue-independent 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:

$$\boldsymbol{\lambda} \in \rho\Lambda \qquad (4)$$

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. Because the sum of the entries of any rate vector in the capacity region Λ is no more than 1, we have by (4) that $\lambda_{tot} \leq \rho$.

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 positive integer K, 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, $r_K = K\mu_K^{sym}$. The following lemma shall be useful.

Lemma 1: For any positive integer K and any probability $p_{min} > 0$, we have $\mu_K^{sym} > \mu_{K+1}^{sym}$. That is:

$$\frac{1}{K}[1 - (1 - p_{min})^K] > \frac{1}{K+1}[1 - (1 - p_{min})^{K+1}]$$

Proof: See Appendix D. ■

Further, for $K \in \{1, \ldots, N\}$, 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 an N-dimensional vector that is 1 in all entries $i \in \mathcal{L}_K$, and zero in all other entries.

Lemma 2: For each set $\mathcal{L}_K$ of size K (for any integer K such that $1 \leq K \leq N$) we have:

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

Furthermore, for each integer k such that $1 \leq k \leq K$ and for any set $\mathcal{L}_k$ that contains k links, we have $\mu_K^{sym} \mathbf{1}_{\mathcal{L}_k} \in \Lambda$.

Proof: We first prove that $\mu_K^{sym} \mathbf{1}_{\mathcal{L}_K} \in \Lambda$. 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 by Lemma 1. Finally, the fact that $\mu_K^{sym} \mathbf{1}_{\mathcal{L}_k} \in \Lambda$ (for any integer k such that $1 \leq k \leq K$) follows because any rate vector with entries less than or equal to another rate vector in Λ is also in Λ. ■

[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.

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]

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

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. Assume that $\boldsymbol{\lambda} \in \rho\Lambda$ for some network loading ρ such that $0 < \rho < 1$. Let K be any integer such that $r_{K+1} > \lambda_{tot}$, that is:[4]

$$1 - (1 - p_{min})^{K+1} > \lambda_{tot} \quad (6)$$

Then the max-weight (LCQ) policy for the ON/OFF channel model stabilizes all queues and yields:

$$\sum_{i=1}^{N} \overline{Q}_i \le \frac{KB_\theta C}{(1-\rho)^2}$$

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

$$B_\theta \triangleq \frac{\lambda_{tot}}{2} + \frac{1}{2}\sum_{i=1}^{N} \mathbb{E}\left\{A_i(t)^2\right\} - \sum_{i=1}^{N} \lambda_i^2$$
$$+ \frac{\theta}{2}\left[\mathbb{E}\left\{A_{tot}(t)^2\right\} + \lambda_{tot} - 2\lambda_{tot}^2\right] \quad (7)$$

$$C \triangleq \begin{cases} \frac{r_{K+1}}{\frac{r_N K \lambda_{tot}}{N(1-\rho)} + \frac{r_K(r_{K+1}-\lambda_{tot})}{(1-\rho)}} & \text{if } K < N \\ 1/r_N & \text{if } K \ge N \end{cases} \quad (8)$$

$$\theta \triangleq \begin{cases} \frac{(1-\rho)(\mu_K^{sym}-\mu_N^{sym})}{r_{K+1}} & \text{if } K < N \\ 0 & \text{if } K \ge N \end{cases} \quad (9)$$
$$(10)$$

[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.

[4]Note that $\lambda_{tot} \le \rho < 1$, and hence there is always a suitably large value K such that (6) holds.

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

$$\overline{W} \le \min\left[\frac{KB_\theta C}{\lambda_{tot}(1-\rho)^2}, \frac{NB_0}{\lambda_{tot}r_N(1-\rho)}\right] \quad (11)$$

where B_0 represents the value of B_θ with $\theta = 0$, and the second expression in the above $\min[\cdot,\cdot]$ function is identical to the previous delay bound (5).

The proof of Theorem 2 is given in the next subsection. We note that the right term inside the $\min[\cdot,\cdot]$ operator in (11) is smaller in the case $K \ge N$. The above bound can be minimized over all positive integers K that satisfy $r_{K+1} > \lambda_{tot}$. For a simpler interpretation of the bound that illuminates the fact that this is an $O(1)$ delay result, note that because $(1+\rho)/2 > \rho \ge \lambda_{tot}$, we can ensure that (6) holds by choosing K to satisfy:

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

Choosing K as follows accomplishes this:

$$K = \max\left[1, \left\lceil \frac{\log(2/(1-\rho))}{\log(1/(1-p_{min}))} \right\rceil - 1\right] \quad (12)$$

Because $\lambda_{tot} \le \rho$, it is not difficult to show that with this choice of K, we have $r_{K+1} - \lambda_{tot} \ge (1-\rho)/2$. Thus, in the case $K < N$ we have:

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

Because $C = 1/r_N$ for the case $K \ge N$, we have that $C = O(1)$ (regardless of whether or not $K < N$). Further, we have from (12) that K is proportional to $\log(1/(1-\rho))$ but independent of N. Finally, if arrival processes are independent so that $\mathbb{E}\left\{A_{tot}(t)^2\right\} = O(1)$, we have $B_\theta/\lambda_{tot} = O(1)$. Therefore, the delay bound of (11) has the form:

$$\overline{W} \le \min\left[\frac{c_1\log(1/(1-\rho))}{(1-\rho)^2}, \frac{c_2N}{(1-\rho)}\right] \quad (13)$$

where c_1 and c_2 are constants that do not depend on ρ or N. If N itself is small, then the right expression in the above $\min[\cdot,\cdot]$ can be smaller than the left expression (i.e., the previous delay bound (5) can be the same as our new delay bound in the case when N is small). However, if the loading ρ is held fixed as N is scaled to infinity, then the left expression in the $\min[\cdot,\cdot]$ is always smaller and demonstrates $O(1)$ average delay (see also simulations in Figs. 1 and 2 of Section VI). Thus, LCQ is *order-optimal* with respect to N. However, the left delay bound has a worse asymptotic in ρ, and so it would be worse than the right bound in the opposite case when N is fixed and ρ is scaled to 1.

C. Lyapunov Drift Analysis

To prove Theorem 2, it suffices to consider only the case $K < N$, as the delay bound in the opposite case $K \ge N$ is identical to the previous delay bound (5). 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:

$$Q_{tot}(t) \triangleq \sum_{i=1}^{N} Q_i(t) \quad (14)$$

Define the following Lyapunov function:

$$L(\boldsymbol{Q}(t)) \triangleq \frac{1}{2}\sum_{i=1}^{N}Q_i(t)^2 + \frac{\theta}{2}\left(\sum_{j=1}^{N}Q_j(t)\right)^2 \quad (15)$$

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 inside the capacity region. The intuition on why this 2-part Lyapunov function allows a tight delay bound is as follows: The first term is a standard sum of squares of queue length, and ensures stability of the algorithm while creating a large negative drift when the number of non-empty queues is small. However, this term also has a relatively small negative drift when the number of non-empty queues is large, preventing $O(1)$ delay analysis from this term alone. To compensate, the second term is a square of the sum of all queues, which does not significantly affect the drift of the first term when the number of non-empty queues is small, but creates a large negative drift to *help* the first term when the number of non-empty queues is large.

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

$$Q_i(t+1) = Q_i(t) - \tilde{\mu}_i(t) + A_i(t) \quad (16)$$

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:

$$Q_{tot}(t+1) = Q_{tot}(t) - \tilde{\mu}_{tot}(t) + A_{tot}(t) \quad (17)$$

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]

$$\Delta(\boldsymbol{Q}(t)) \triangleq \mathbb{E}\{L(\boldsymbol{Q}(t+1)) - L(\boldsymbol{Q}(t)) \,|\, \boldsymbol{Q}(t)\} \quad (18)$$

Lemma 3: The Lyapunov drift $\Delta(\boldsymbol{Q}(t))$ for the ON/OFF channel model satisfies:

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

[5]Strictly speaking, correct notation should be $\Delta(\boldsymbol{Q}(t), t)$, as the drift could be from a non-stationary policy, although we use the simpler notation $\Delta(\boldsymbol{Q}(t))$ as formal notation for the right hand side of (18).

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

$$\begin{aligned}B(t) \triangleq \ & \frac{\tilde{\mu}_{tot}(t)}{2} + \frac{1}{2}\sum_{i=1}^{N}[A_i(t)^2 - 2A_i(t)\tilde{\mu}_i(t)] \\ & + \frac{\theta}{2}[A_{tot}(t)^2 + \tilde{\mu}_{tot}(t) - 2\tilde{\mu}_{tot}(t)A_{tot}(t)] \quad (19)\end{aligned}$$

Proof: (Lemma 3) See Appendix A. ∎

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:

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

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 3 thus yields:

$$\begin{aligned}\Delta(\boldsymbol{Q}(t)) \le \ & \mathbb{E}\{B(t)\,|\,\boldsymbol{Q}(t)\} \\ & - \sum_{i=1}^{N}Q_i(t)\mathbb{E}\{\mu_i^*(t) - \lambda_i\,|\,\boldsymbol{Q}(t)\} \\ & -\theta Q_{tot}(t)\mathbb{E}\{\tilde{\mu}_{tot}(t) - \lambda_{tot}\,|\,\boldsymbol{Q}(t)\} \quad (20)\end{aligned}$$

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 $0 \le L(t) \le N$.

- *Case 1* ($L(t) \le K$): Suppose $L(t) \le K$, and let $\mathcal{L}(t)$ represent the set of non-empty queue indices. Recall that $\mu_K^{sym}\mathbf{1}_{\mathcal{L}(t)} \in \Lambda$ (by Lemma 2) and that $\boldsymbol{\lambda}/\rho \in \Lambda$ (by assumption that $\boldsymbol{\lambda} \in \rho\Lambda$). By taking a convex combination of these two vectors and using convexity of the set Λ, it follows that:

$$\boldsymbol{\lambda} + (1-\rho)\mu_K^{sym}\mathbf{1}_{\mathcal{L}(t)} \in \Lambda \quad (21)$$

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

$$\mathbb{E}\{\boldsymbol{\mu}^*(t)\} = \boldsymbol{\lambda} + (1-\rho)\mu_K^{sym}\mathbf{1}_{\mathcal{L}(t)}$$

  Such a policy exists by (21) and Theorem 1. Thus, for all $i \in \mathcal{L}(t)$ we have:

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

  Using (22) in the drift inequality (20) and noting that $Q_i(t) = 0$ if $i \notin \mathcal{L}(t)$ yields:

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

Define ϵ as follows:

$$\epsilon \triangleq [(1-\rho)\mu_K^{sym} - \theta\lambda_{tot}] \qquad (23)$$

It follows that:

$$\Delta(\boldsymbol{Q}(t)) \leq \mathbb{E}\{B(t) \mid \boldsymbol{Q}(t)\} - \epsilon Q_{tot}(t) \qquad (24)$$

- *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 $i \in \{1, \ldots, N\}$:

$$\mathbb{E}\{\mu_i^*(t)\} = \lambda_i + (1-\rho)\mu_N^{sym} \qquad (25)$$

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:

$$\mathbb{E}\{\tilde{\mu}_{tot} \mid \boldsymbol{Q}(t)\} \geq 1 - \Pi_{i \in \hat{\mathcal{L}}_{K+1}}(1-p_i) \geq r_{K+1} \qquad (26)$$

where $\hat{\mathcal{L}}_{K+1}$ represents the set of $K+1$ links with the smallest success probabilities. Plugging (25) and (26) into the drift inequality (20) yields:

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

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 (23)):

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

Recall that we have assumed $K < N$ (as Theorem 2 is trivially true if $K \geq N$, as described at the beginning of this subsection). Thus, we have $\mu_K^{sym} > \mu_N^{sym}$ (by Lemma 1), and so we indeed have $\theta > 0$. Further, because $r_{K+1} > \lambda_{tot}$, we have that $\epsilon > 0$. Therefore, the drift inequality (24) 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 4: (Lyapunov Drift [6]) If the drift $\Delta(\boldsymbol{Q}(t))$ of a non-negative Lyapunov function satisfies the following for all t and all $\boldsymbol{Q}(t)$:

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

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

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

where

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

Using this Lyapunov drift lemma in (24) (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 equal the value B_θ 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:

$$\lambda_i \leq \lambda_{av} + f \quad \text{for all } i \in \{1, \ldots, N\} \qquad (27)$$

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:

$$\hat{N} = \lceil N/K \rceil K \qquad (28)$$

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*:

$$L(\boldsymbol{Q}(t)) \triangleq \frac{1}{2} \sum_{g \in \mathcal{G}_K} \sum_{k=1}^{K} (Q_k^{(g)}(t))^2 \qquad (29)$$

This is similar to the Lyapunov function of [7], with the exception that it sums over all possible partitions into K disjoint groups. For intuition, we note that the f-balanced traffic assumption allows the "Largest Connected Group" (LCG) argument of [7] to proceed on *any set of K disjoint groups*. However, once we fix a particular group, minimizing the drift gives rise to the LCG algorithm rather than the "max-weight" LCQ algorithm. Changing the Lyapunov function by summing over all possible K disjoint groups yields a similar negative drift as in LCG, but the "symmetry" induced by summing over all groups remarkably makes the drift minimizing algorithm the LCQ algorithm (rather than LCG).

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:

$$Q_k^{(g)}(t+1) = Q_k^{(g)}(t) - \tilde{\mu}_k^{(g)}(t) + A_k^{(g)}(t) \qquad (30)$$

Define the Lyapunov drift $\Delta(\boldsymbol{Q}(t))$ as before (given in (18)).

Lemma 5: For a general scheduling policy, the Lyapunov drift satisfies:

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

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

$$C(t) \triangleq \frac{1}{2} \sum_{g \in \mathcal{G}_K} \sum_{k=1}^{K} \left[\tilde{\mu}_k^{(g)}(t) + A_k^{(g)}(t)^2 - 2\tilde{\mu}_k^{(g)}(t) A_k^{(g)}(t)\right]$$
$$(31)$$

Proof: The proof is similar to the proof of Lemma 3. Specifically, note from (30) that:

$$Q_k^{(g)}(t+1)^2 - Q_k^{(g)}(t)^2 = \tilde{\mu}_k^{(g)}(t) + A_k^{(g)}(t)^2$$
$$- 2\tilde{\mu}_k^{(g)}(t) A_k^{(g)}(t) - 2Q_k^{(g)}(t)[\tilde{\mu}_k^{(g)}(t) - A_k^{(g)}(t)]$$

where we have used the fact that $\tilde{\mu}_k^{(g)}(t)^2 = \tilde{\mu}_k^{(g)}(t)$. The result follows by summing over all k and all groups, and taking conditional expectations. ∎

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 6: (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 5, 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:

$$\Delta(\boldsymbol{Q}(t)) \leq \mathbb{E}\{C(t) \mid \boldsymbol{Q}(t)\}$$
$$- \sum_{g \in \mathcal{G}_K} \sum_{k=1}^{K} Q_k^{(g)}(t) \mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t) - \lambda_k^{(g)} \mid \boldsymbol{Q}(t)\right\} (32)$$

The drift inequality (32) 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 (32) to yield a useful drift bound.

Lemma 7: 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:

$$\lambda_i \leq \frac{\lambda_{tot}}{\hat{N}} + \frac{\beta(1 - \rho)}{K} \qquad (33)$$

then:

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

where $|\mathcal{G}_K|$ is the cardinality of $\mathcal{G}_K$, $Q_{tot}(t)$ is the total sum backlog (defined in (14)), and z is defined:

$$z \triangleq (1 - 1/K)/(1 - 1/\hat{N}) \qquad (34)$$

Proof: The proof of Lemma 7 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 (33) imply that $\boldsymbol{\lambda}$ is f-balanced with $f = \beta(1 - \rho)/K$.

Lemma 8: 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 7 and 8 in the drift inequality (32) yields:

Lemma 9: If $\boldsymbol{\lambda} \in \rho\Lambda$ (for $0 < \rho < 1$) and if (33) is satisfied for all $i \in \{1, \dots, N\}$, then:

$$\Delta(\boldsymbol{Q}(t)) \leq \mathbb{E}\{C(t) \mid \boldsymbol{Q}(t)\}$$
$$- Q_{tot}(t) |\mathcal{G}_K| \frac{[r_K - \lambda_{tot} - z\beta(1 - \rho)]}{K} \quad (35)$$

Lemma 9 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:

$$[1 - (1 - p_{min})^K] \geq (1 + \rho)/2 \quad (36)$$

Suppose that $K \leq N$, and the f-balanced traffic constraints (33) are satisfied for some value β such that $0 \leq \beta < 1/(2z)$, where $z \triangleq (1 - 1/K)/(1 - 1/\hat{N})$ (note that $z \leq 1$). If the max-weight (LCQ) policy is used on this ON/OFF channel model, then average queue occupancy satisfies:

$$\overline{Q}_{tot} \leq \frac{KD}{(1 - \rho)(\frac{1}{2} - z\beta)} \quad , \quad \overline{W} \leq \frac{K(D/\lambda_{tot})}{(1 - \rho)(\frac{1}{2} - z\beta)} \quad (37)$$

where D is defined:

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

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:[6]

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

Note that the constraint (36) is satisfied by:

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

Therefore, K is independent of N, and is proportional to $\log(\frac{1}{1 - \rho})$. Assuming that traffic streams are independent, so that $\mathbb{E}\{A_{tot}(t)^2\} = O(1)$, implies that $D = O(1)$. Thus the delay bound gives $\overline{W} \leq c \log(1/(1 - \rho))/(1 - \rho)$ (where c is a constant independent of ρ and N), 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 (36) thus implies:

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

The above value is strictly positive because $z\beta < 1/2$. Using the drift inequality (35) directly in the Lyapunov Drift Lemma (Lemma 4) yields:

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

Using the definition of $C(t)$ in (31) and the fact that the system is stable (so the long term departure rate is equal to λ_{tot}) yields:

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

$$\leq |\mathcal{G}_K| \left[\frac{\lambda_{tot}}{2} + \frac{\mathbb{E}\{A_{tot}(t)^2\}}{2}\right] = |\mathcal{G}_K| D$$

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 $\mathbb{E}\{A_k^{(g)}(t)\} = \frac{\lambda_{tot}}{K}$ and $\mathbb{E}\{A_k^{(g)}(t)^2\} = \frac{\lambda_{tot}^2}{K^2} + \frac{\lambda_{tot}}{K}$ for all g, k. ∎

V. MULTI-RATE TRANSMISSION MODELS

Now suppose that for each channel $i \in \{1, \dots, N\}$, the states $S_i(t)$ are non-negative integers bounded by a finite integer $\mu_{i,max}$, where $\mu_{i,max}$ represents the maximum transmission rate over channel i.[7] That is, we have:

$$S_i(t) \in \{0, 1, \dots, \mu_{i,max}\} \quad \text{for all } t \text{ and all } i \in \{1, \dots, N\}$$

We assume that $\mu_{i,max} > 0$ for all i. The queueing dynamics are governed by (1). The capacity region Λ is known to be equal to the set of all rate vectors that can be achieved via a stationary, randomized, queue-independent algorithm that chooses $\boldsymbol{\mu}^*(t)$ as a potentially random function of only the current $\boldsymbol{S}(t)$ vector [6].

The max-weight algorithm in this case is the algorithm that observes queue backlogs and channel states every slot and selects the link $i \in \{1, \dots, N\}$ with the largest value of $Q_i(t)S_i(t)$ (breaking ties arbitrarily). Suppose the arrival rate vector satisfies $\boldsymbol{\lambda} \in \rho\Lambda$ for some loading value ρ such that $0 < \rho < 1$. The analysis in [5] [6] uses a standard Lyapunov function, given by the sum of the squares of queue backlog, to show the max-weight algorithm for a general downlink has average delay upper bounded by $cN/(1 - \rho)$, where c is a constant that is independent of N and ρ. We first present a modified version of that prior bound, which has the same structure but uses our particular $\mu_{i,max}$ notation and improves the c coefficient:

Lemma 10: Suppose $\boldsymbol{A}(t)$ is i.i.d. over slots with $\mathbb{E}\{\boldsymbol{A}(t)\} = \boldsymbol{\lambda}$, and that the channel state vector $\boldsymbol{S}(t) = (S_1(t), \dots, S_N(t))$ is also i.i.d. over slots. Suppose that $\boldsymbol{\lambda} \in \rho\Lambda$ for some value ρ that satisfies $0 < \rho < 1$. Then the system

---

[6]These bounds for symmetric Poisson traffic are obtained from the last line of the proof of Theorem 3, which gives a slightly smaller bound than that achieved by plugging $\beta = 0$, $\mathbb{E}\{A_{tot}(t)^2\} = \lambda_{tot} + \lambda_{tot}^2$ into (37).

[7]For consistency, we continue to work in integer units of packets. The analysis does not significantly change if $S_i(t)$ values are viewed as non-negative real numbers with units of bits/slot.

is stable under the max-weight algorithm and has an average delay bound given by:

$$\overline{W} \leq \frac{N\left[\frac{1}{2\lambda_{tot}}\sum_{i=1}^{N}\mathbb{E}\{A_i^2\} - \frac{3}{2\lambda_{tot}}\sum_{i=1}^{N}\lambda_i^2\right]}{(1-\rho)\mu_{sym}}$$
$$+ \frac{N\min\left[\sum_{i=1}^{N}\frac{\lambda_i\mu_{i,max}}{\lambda_{tot}}, \frac{\hat{S}^2}{\lambda_{tot}}\right]}{(1-\rho)\mu_{sym}} \qquad (38)$$

where $\hat{S}^2$ is defined:

$$\hat{S}^2 \triangleq \mathbb{E}\left\{\max_{i\in\{1,\dots,N\}} S_i(t)^2\right\}$$

and where μ_{sym} is defined as the largest value such that $(\mu_{sym}/N, \mu_{sym}/N, \dots, \mu_{sym}/N) \in \Lambda$.

Further, in the case when all processes $S_i(t)$ are independent and satisfy $Pr[S_i(t) = \mu_{i,max}] > 0$, we can bound μ_{sym} as follows:

$$\mu_{sym} \geq \hat{\mu}(1 - (1 - p_{min})^N)$$

where $\hat{\mu}$ and p_{min} are defined:

$$\hat{\mu} \triangleq \min_{i\in\{1,\dots,N\}} \mu_{i,max}$$
$$p_{min} \triangleq \min_{i\in\{1,\dots,N\}} Pr[S_i(t) \geq \hat{\mu}]$$

Proof: The proof uses the Lyapunov function $L(\boldsymbol{Q}(t)) = \frac{1}{2}\sum_{i=1}^{N} Q_i(t)^2$, as in [5] [6], but provides a simple modification of the argument to yield a tighter bound (given in Appendix E for completeness). ∎

We note that the above delay bound holds also in the case when $\mu_{i,max} = \infty$ for some values i, but when the second moment of transmission rates is finite (so that $\hat{S}^2$ is finite). The above delay bound has the structure $cN/(1-\rho)$, and holds even if arrival and channel vectors $\boldsymbol{A}(t)$ and $\boldsymbol{S}(t)$ have entries that are correlated over the different links $i \in \{1,\dots,N\}$. A similar argument can be used to show stability with the same structural delay bound $\hat{c}N/(1-\rho)$ for the modified max-weight policy that chooses the link $i \in \{1,\dots,N\}$ with the largest $Q_i(t)\min[Q_i(t), S_i(t)]$ value. This modified policy can sometimes provide smaller empirical average delay than the original max-weight policy, although its resulting analytical delay bound has a slightly worse coefficient $\hat{c} \geq c$ (this modified policy is equivalent to the original max-weight policy in the case of ON/OFF channels with $\mu_{i,max} = 1$ for all i). Similar to the ON/OFF case, one might suspect that for this multi-rate system, average delay that is independent of N can be achieved when arrival and channel processes are independent over each channel. However, the next subsection presents an important example that shows this is not the case.[8]

[8] We note that our original pre-print of this paper in [18] incorrectly claimed that multi-rate systems also have delay that is independent of N. The mistake in [18] arose when plugging the equation from Lemma 10 of that paper into equation (33) of that paper. Plugging one equation into the other implicitly assumed that the sum queue backlog in queues with at least μ_{max} packets is the same as the total queue backlog in the system. This is true when $\mu_{max} = 1$, but is not true in general as it neglects the "residual" packets in queues with fewer than μ_{max} packets.

A. An example showing necessity of $\Omega(N)$ delay

Here we present an example showing that the average number of queues that have at least one packet (but possibly fewer than $\mu_{i,max}$ packets) must be linear in N. This necessarily makes the average delay of *any* scheduling policy grow at least linearly with N. Consider a system with N queues with symmetric channels and traffic. Assume that $N \geq 3$ and suppose that all arrival processes $A_i(t)$ are independent and Bernoulli with $Pr[A_i(t) = 1] = 3/N$ for all $i \in \{1,\dots,N\}$ (so that $\lambda_i = 3/N$ for all i, and $\lambda_{tot} = 3$ packets/slot). Now suppose that all channels have $\mu_{i,max} = 5$, and channel state processes are i.i.d. with $Pr[S_i(t) = 5] = 1/2$, $Pr[S_i(t) = 0] = 1/2$ for all $i \in \{1,\dots,N\}$. The largest symmetric rate in the capacity region of this system is $\mu_{sym}/N = 5(1 - (1/2)^N)/N$, and hence the arrival rate vector is inside the capacity region and has ρ given by:

$$\rho = \frac{3}{5(1 - (1/2)^N)}$$

Note that $\rho < 1$ for $N \geq 3$, and ρ is approximately $3/5$ for large N. Here we show that under any scheduling policy, in steady state the average number of non-empty queues in this system must be linear in N. Specifically, consider any scheduling policy, and let $Z(t)$ represent the number of non-empty queues on slot t. For simplicity, we assume that $Z(t)$ has a well defined steady state under the scheduling policy. The intuition behind our proof is that $Z(t+1)$ is formed from $Z(t)$ by adding the number of new non-empty queues created and subtracting any non-empty queue that becomes empty. The number of non-empty queues subracted can be at most 1 (as we can serve at most one channel per slot), while the average number of new non-empty queues added is *more than one* whenever $Z(t) < N/2$.

Lemma 11: Consider any scheduling policy for which $Z(t)$ has a well defined steady state distribution. Then for the system above (with $\lambda_i = 3/N$ and $\mu_{i,max} = 5$ for all $i \in \{1,\dots,N\}$) we have that in steady state:

$$Pr[Z(t) \geq N/2] \geq 1/3$$

and hence $\mathbb{E}\{Z(t)\} \geq N/6$. That is, the average number of non-empty queues is at least $N/6$, and hence the average number of packets in the system is at least $N/6$.

Proof: Define $\Delta(t) \triangleq Z(t+1) - Z(t)$ as the change in $Z(t)$ from one slot to the next. Let t be a time at which the system is in steady state. We thus have $\mathbb{E}\{\Delta(t)\} = 0$. On the other hand, we have the following:

$$\mathbb{E}\{\Delta(t) \mid Z(t) \geq N/2\} \geq -1 \qquad (39)$$
$$\mathbb{E}\{\Delta(t) \mid Z(t) < N/2\} \geq \lambda_i N/2 - 1 = 1/2 \qquad (40)$$

where (39) follows because the drift cannot be less than -1 on any slot (as at most one non-empty queue can become an empty queue), and (40) holds because, given that $Z(t) < N/2$, the average number of *new* non-empty queues that are created on slot t is equal to the average number of new arrivals to the

empty queues, which is at least $\lambda_i(N/2)$. It follows that:

$$
\begin{aligned}
0 &= \mathbb{E}\{\Delta(t)\} \\
&= \mathbb{E}\{\Delta(t) \mid Z(t) \geq N/2\}\, Pr[Z(t) \geq N/2] \\
&\quad + \mathbb{E}\{\Delta(t) \mid Z(t) < N/2\}\,(1 - Pr[Z(t) \geq N/2]) \\
&\geq (-1)Pr[Z(t) \geq N/2] \\
&\quad + (1/2)(1 - Pr[Z(t) \geq N/2])
\end{aligned}
$$

Therefore $Pr[Z(t) \geq N/2] \geq 1/3$, completing the proof. ∎

VI. SIMULATIONS

Here we present simulations for the ON/OFF system with independent channel and arrival processes. We assume that $Pr[S_i(t) = ON] = 1/2$ for all $i \in \{1, \ldots, N\}$. We first consider symmetric Bernoulli arrivals, so that $\lambda_i = \lambda$ for all i, where λ is chosen so that $\boldsymbol{\lambda} \in \rho\Lambda$ with $\rho = 0.8$. We simulated the system over 10^6 slots for values of N between 3 and 300. The resulting simulated queue averages are shown in Fig. 1, together with the two $O(1)$ bounds and the previous $O(N)$ bound. Note that the previous $O(N)$ bound is a considerable overestimate of queue backlog. Our new bounds do not grow with N, and our second $O(1)$ bound (derived for f-balanced traffic rates) is indeed tighter than the first $O(1)$ bound (for $N \geq 9$), although it applies only to f-balanced traffic while the first bound applies to any traffic rates in $\rho\Lambda$. However, there is still a significant gap (roughly a factor of 10 in this example) between our tightest bound and the simulated value. We next consider heterogeneous traffic rates implemented on the same ON/OFF system. We assume that N is odd, and choose rates λ_i given as follows:

$$
\lambda_i = \begin{cases} \lambda & \text{for } i \in \{1, \ldots, (N-1)/2\} \\ 2\lambda & \text{for } i \in \{(N-1)/2+1, \ldots, N-1\} \\ 4\lambda & \text{for } i = N \end{cases}
$$

where λ is chosen so that $\boldsymbol{\lambda} \in \rho\Lambda$ for $\rho = 0.8$. The results are shown in Fig. 2. Note that we plot only the first $O(1)$ bound (for heterogeneous traffic) in this case, although the f-balanced traffic assumption also applies in this case when N is sufficiently large.

Simulations of the multi-rate system example in Section V-A were also conducted, and it was verified that average backlog indeed grows linearly with N due to the "residual" packets in queues i that have fewer than $\mu_{i,max}$ packets (simulation plots omitted for brevity). However, it was observed in the simulations that the total backlog due to queues with *at least* $\mu_{i,max}$ packets is $O(1)$. This suggests that, although the total average backlog in multi-rate systems may have a fundamental $O(N)$ term due to residual packets, the average backlog may be $O(1)$ after a term of at most $\sum_{i=1}^{N}(\mu_{i,max}-1)$ is subtracted out.

VII. 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
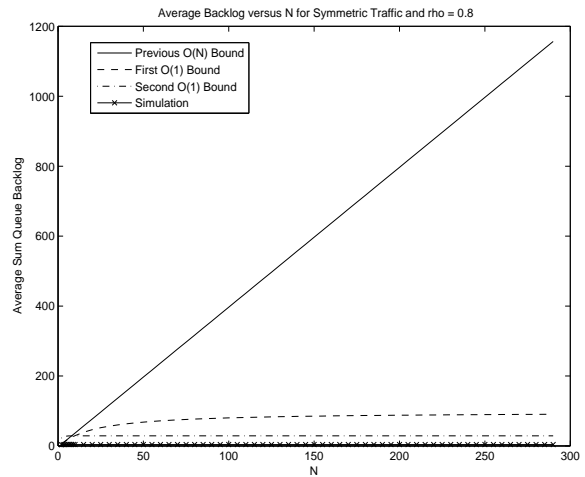

Fig. 1.   Simulation and Bounds for the ON/OFF system with symmetric traffic and $\rho = 0.8$.
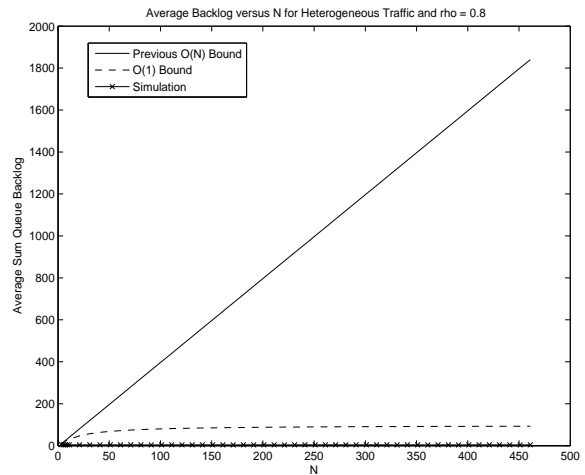

Fig. 2.   Simulation and Bounds for the ON/OFF system with heterogeneous traffic and $\rho = 0.8$.

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. For multi-rate channels, a delay bound of $O(N)$ applies. Conversely, it is shown for a simple multi-rate example that, unlike ON/OFF channels, average backlog must be *at least* linear in N due to "residual" packets. 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 3

Here we prove Lemma 3. 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 (16) 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:

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

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:

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

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 (19) that $B(t) = B_1(t) + \theta B_2(t)$ yields the result of Lemma 3.

APPENDIX B — PROOF OF LEMMA 6

Define the integer $m = \hat{N}/K$. Here we prove Lemma 6. 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 6. 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:

$$
\begin{aligned}
f(t) &= \sum_{i=1}^{\hat{N}} \tilde{\mu}_i(t)Q_i(t)\,|\mathcal{G}_K|\,(1 - \frac{m-1}{\hat{N}-1}) \\
&\quad + \sum_{i=1}^{\hat{N}} \tilde{\mu}_i(t)\left(\sum_{j=1}^{\hat{N}} Q_j(t)\right)|\mathcal{G}_K|\frac{m-1}{\hat{N}-1} \quad (41)
\end{aligned}
$$

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 a non-negative constant. 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 6.

APPENDIX C — PROOF OF LEMMAS 7 AND 8

Proof: (Lemma 7) Define the integer $m = \hat{N}/K$. Using a counting argument similar to that of Appendix B (compare with (41)), we have that for any vector $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_{\hat{N}})$:

$$
\begin{aligned}
\sum_{g \in \mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\lambda_k^{(g)} &= \sum_{i=1}^{\hat{N}} Q_i(t)\lambda_i\,|\mathcal{G}_K|\,(1 - \frac{m-1}{\hat{N}-1}) \\
&\quad + Q_{tot}(t)\lambda_{tot}\,|\mathcal{G}_K|\frac{m-1}{\hat{N}-1} \quad (42)
\end{aligned}
$$

Using the bound on λ_i given in (33) yields:

$$
\sum_{g \in \mathcal{G}_K}\sum_{k=1}^{K} Q_k^{(g)}(t)\lambda_k^{(g)} \le
$$

$$
\begin{aligned}
&Q_{tot}(t)\lambda_{tot}\,|\mathcal{G}_K|\left[\frac{1}{\hat{N}} - \frac{m-1}{\hat{N}(\hat{N}-1)} + \frac{m-1}{\hat{N}-1}\right] \\
&\quad + Q_{tot}(t)\,|\mathcal{G}_K|\,(1 - \frac{m-1}{\hat{N}-1})\beta(1-\rho)/K
\end{aligned}
$$

The result of Lemma 7 follows by using the identities:

$$
\left[\frac{1}{\hat{N}} - \frac{m-1}{\hat{N}(\hat{N}-1)} + \frac{m-1}{\hat{N}-1}\right] = \frac{1}{K} \quad (43)
$$

$$
\left[1 - \frac{m-1}{\hat{N}-1}\right] = z
$$

∎

Proof: (Lemma 8) 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 $\hat{S}(t)$ 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:

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

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

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

Using this equality gives:

$$\sum_{g \in \mathcal{G}_K} \sum_{k=1}^{K} Q_k^{(g)}(t) \mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t) \mid \boldsymbol{Q}(t)\right\} =$$
$$\frac{r_l}{l} \sum_{g \in \mathcal{G}_K} \sum_{k=1}^{K} Q_k^{(g)}(t) l_k^{(g)} \quad (44)$$

Now note that the $l_k^{(g)}$ values are structurally similar to the $\lambda_k^{(g)}$ values, and hence (similar to (42)) we have (using $Q_i(t) l_i = Q_i(t)$ and $l_{tot} = l$):

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

Using this in (44) yields:

$$\sum_{g \in \mathcal{G}_K} \sum_{k=1}^{K} Q_k^{(g)}(t) \mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t) \mid \boldsymbol{Q}(t)\right\}$$
$$= \frac{r_l}{l} Q_{tot}(t) |\mathcal{G}_K| \left[1 - \frac{m-1}{\hat{N}-1} + \frac{l(m-1)}{\hat{N}-1}\right] \quad (45)$$
$$\geq r_l Q_{tot}(t) |\mathcal{G}_K| \left[\frac{1}{\hat{N}} - \frac{(m-1)}{\hat{N}(\hat{N}-1)} + \frac{m-1}{\hat{N}-1}\right]$$
$$= r_l Q_{tot}(t) |\mathcal{G}_K| / K \quad (46)$$

where the last equality holds by (43). The above holds for $L(t) = l \in \{1, \dots, N\}$. Suppose now that $l \geq K$. In this case we have $r_l \geq r_K$, proving the result of Lemma 8 for $l \geq K$.

Consider now the final case where $l \in \{1, \dots, K-1\}$. Then from (45) we have:

$$\sum_{g \in \mathcal{G}_K} \sum_{k=1}^{K} Q_k^{(g)}(t) \mathbb{E}\left\{\tilde{\mu}_k^{(g)*}(t) \mid \boldsymbol{Q}(t)\right\}$$
$$\geq \frac{r_l}{l} Q_{tot}(t) |\mathcal{G}_K| \quad (47)$$

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

APPENDIX D — PROOF OF LEMMA 1

Here we prove that $\mu_K^{sym} > \mu_{K+1}^{sym}$. Specifically, we show that if p is a value such that $0 < p \leq 1$, then for any positive integer K we have:

$$\frac{1}{K}[1 - (1-p)^K] > \frac{1}{K+1}[1 - (1-p)^{K+1}] \quad (48)$$

To show this, note that it is trivially true for the case $p = 1$. In the opposite case where $0 < p < 1$, we can multiply (48) by $K(K+1)$ and rearrange terms to see that the inequality is equivalent to the following:

$$(1-p)^K + Kp(1-p)^K < 1 \quad (49)$$

Thus, it suffices to prove that (49) is true. To this end, we have:

$$(1-p)^K + Kp(1-p)^K < (1-p)^K + Kp(1-p)^{K-1}$$
$$\leq \sum_{i=0}^{K} \binom{K}{i} p^i (1-p)^{K-i}$$
$$= ((1-p) + p)^K = 1$$

where the first (strict) inequality holds because $0 < p < 1$ and hence $(1-p)^K < (1-p)^{K-1}$. This establishes (49) and completes the proof of Lemma 1.

APPENDIX E — PROOF OF LEMMA 10

The queueing dynamics are given by $Q_i(t+1) = Q_i(t) - \tilde{\mu}_i(t) + A_i(t)$, where $\tilde{\mu}_i(t) = \min[\mu_i(t), Q_i(t)]$. Using the Lyapunov function $L(\boldsymbol{Q}(t)) \triangleq \frac{1}{2} \sum_{i=1}^{N} Q_i(t)^2$ and performing a standard quadratic drift computation (see, for example, [6]), it is not difficult to show the drift satisfies:

$$\Delta(\boldsymbol{Q}(t)) = \frac{1}{2} \sum_{i=1}^{N} \mathbb{E}\{A_i(t)^2\}$$
$$- \sum_{i=1}^{N} \mathbb{E}\left\{\lambda_i \tilde{\mu}_i(t) - \frac{\tilde{\mu}_i(t)^2}{2} \mid \boldsymbol{Q}(t)\right\}$$
$$+ \sum_{i=1}^{N} \lambda_i Q_i(t) - \sum_{i=1}^{N} Q_i(t) \mathbb{E}\{\mu_i(t) \mid \boldsymbol{Q}(t)\}$$
$$+ \sum_{i=1}^{N} Q_i(t) \mathbb{E}\{\mu_i(t) - \tilde{\mu}_i(t) \mid \boldsymbol{Q}(t)\}$$

By definition of $\tilde{\mu}_i(t)$, we have:

$$Q_i(t)(\mu_i(t) - \tilde{\mu}_i(t)) = \tilde{\mu}_i(t)\mu_i(t) - \tilde{\mu}_i(t)^2$$
$$\leq \min[\mu_{i,max}\tilde{\mu}_i(t), \mu_i(t)^2] - \tilde{\mu}_i(t)^2$$

Hence:

$$\Delta(\boldsymbol{Q}(t)) \leq \frac{1}{2} \sum_{i=1}^{N} \mathbb{E}\{A_i(t)^2\}$$
$$- \sum_{i=1}^{N} \mathbb{E}\left\{\lambda_i \tilde{\mu}_i(t) + \frac{\tilde{\mu}_i(t)^2}{2} \mid \boldsymbol{Q}(t)\right\}$$
$$+ \sum_{i=1}^{N} \lambda_i Q_i(t) - \sum_{i=1}^{N} Q_i(t) \mathbb{E}\{\mu_i(t) \mid \boldsymbol{Q}(t)\}$$
$$+ \sum_{i=1}^{N} \mathbb{E}\left\{\min\left[\mu_{i,max}\tilde{\mu}_i(t), \mu_i(t)^2\right] \mid \boldsymbol{Q}(t)\right\} (50)$$

Using the fact that $\mathbb{E}\{\min[\cdot,\cdot]\} \leq \min[\mathbb{E}\{\cdot\},\mathbb{E}\{\cdot\}]$ (by Jensen's inequality and concavity of the $\min[\cdot,\cdot]$ operator), that the sum of a $\min$ is less than or equal to the $\min$ of a sum, and that $\sum_{i=1}^{N} \mathbb{E}\{\mu_i(t)^2 \mid \boldsymbol{Q}(t)\} \leq \hat{S}^2$, the final term on the right hand side of (50) can be bounded by:

$$\min\left[\sum_{i=1}^{N} \mu_{i,max}\mathbb{E}\{\tilde{\mu}_i(t) \mid \boldsymbol{Q}(t)\}, \hat{S}^2\right]$$

Because the max-weight policy maximizes $\sum_{i=1}^{N} Q_i(t)\mu_i(t)$ (given queue backlogs $\boldsymbol{Q}(t)$) we have:

$$\sum_{i=1}^{N} Q_i(t)\mathbb{E}\{\mu_i(t) \mid \boldsymbol{Q}(t)\} \geq \sum_{i=1}^{N} Q_i(t)\mathbb{E}\{\mu_i^*(t) \mid \boldsymbol{Q}(t)\} \quad (51)$$

where $\mu_i^*(t)$ represents any alternative scheduling decision. Noting that $\boldsymbol{\lambda}/\rho \in \Lambda$ and $\boldsymbol{\mu}_{sym}/N \in \Lambda$, we have by convexity of Λ:

$$\boldsymbol{\lambda} + (1-\rho)\boldsymbol{\mu}_{sym}/N \in \Lambda$$

Thus, there exists a stationary randomized policy that chooses $\boldsymbol{\mu}^*(t)$ independent of queue backlog to yield:

$$\mathbb{E}\{\boldsymbol{\mu}^*(t) \mid \boldsymbol{Q}(t)\} = \mathbb{E}\{\boldsymbol{\mu}^*(t)\} = \boldsymbol{\lambda} + (1-\rho)\boldsymbol{\mu}_{sym}/N$$

Plugging this into (51) and then into (50) yields:

$$\begin{aligned}
\Delta(\boldsymbol{Q}(t)) \leq \; & \frac{1}{2}\sum_{i=1}^{N}\mathbb{E}\{A_i(t)^2\} \\
& - \sum_{i=1}^{N}\mathbb{E}\left\{\lambda_i\tilde{\mu}_i(t) + \frac{\tilde{\mu}_i(t)^2}{2} \mid \boldsymbol{Q}(t)\right\} \\
& - \frac{(1-\rho)\mu_{sym}}{N}\sum_{i=1}^{N} Q_i(t) \\
& + \min\left[\sum_{i=1}^{N}\mu_{i,max}\mathbb{E}\{\tilde{\mu}_i(t) \mid \boldsymbol{Q}(t)\}, \hat{S}^2\right] \quad (52)
\end{aligned}$$

Using the Lyapunov drift lemma (Lemma 4) on the above drift and noting that the system is stable with well defined time average limits yields:

$$\sum_{i=1}^{N}\overline{Q}_i \leq \frac{N\left[\frac{1}{2}\sum_{i=1}^{N}\mathbb{E}\{A_i^2\} - \frac{3}{2}\sum_{i=1}^{N}\lambda_i^2\right]}{(1-\rho)\mu_{sym}} + \frac{N\min\left[\sum_{i=1}^{N}\lambda_i\mu_{i,max}, \hat{S}^2\right]}{(1-\rho)\mu_{sym}}$$

where we have used the fact that $\lim_{t\to\infty}\mathbb{E}\{\tilde{\mu}_i(t)\} = \lambda_i$ and $\lim_{t\to\infty}\mathbb{E}\{\tilde{\mu}_i^2(t)\} \geq \lim_{t\to\infty}\mathbb{E}\{\tilde{\mu}_i(t)\}^2 = \lambda_i^2$. Using Little's theorem on this congestion bound proves (38).

Now suppose that all $S_i(t)$ processes are independent and $Pr[S_i(t) = \mu_{i,max}] \geq p_{min}$ for all i. We derive the bound on μ_{sym} given in Lemma 10. Define $\hat{\mu} \triangleq \min_{i\in\{1,\ldots,N\}}\mu_{i,max}$. Consider the stationary and randomized algorithm $\boldsymbol{\mu}^*(t)$ that observes channel states $\boldsymbol{S}(t)$ and probabilistically places each link $i \in \{1,\ldots,N\}$ in a set $\chi(t)$ with probability 0 if $S_i(t) < \hat{\mu}$, and with probability $p_{min}/Pr[S_i(t) \geq \hat{\mu}]$ if $S_i(t) \geq \hat{\mu}$. Then $\chi(t)$ contains a random number of links, and each link appears in $\chi(t)$ independently with probability p_{min}. Select

a link to serve on slot t uniformly and randomly with equal probability over all links in $\chi(t)$ (remaining idle if $\chi(t)$ is empty). It follows that under this policy, a particular link i is selected for transmission with probability exactly $(1 - (1 - p_{min})^N)/N$, and is selected only if $S_i(t) \geq \hat{\mu}$. Hence:

$$\mathbb{E}\{\mu_i^*(t)\} \geq \frac{\hat{\mu}(1 - (1 - p_{min})^N)}{N} \quad \text{for all } i \in \{1,\ldots,N\}$$

It follows that the symmetric rate vector with all N entries equal to the right hand side in the above expression is in the capacity region Λ, so that μ_{sym}/N is greater than or equal to this value.

REFERENCES

[1] M. J. Neely. Delay analysis for max weight opportunistic scheduling in wireless systems. *Proc. 46th Annual Allerton Conf. on Communication, Control, and Computing*, Sept. 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. *IEEE/ACM Transactions on Networking*, vol. 16, no. 5, pp. 1188-1199, October 2008.

[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.

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