Delay Analysis for Maximal Scheduling in Wireless Networks with Bursty Traffic

Michael J. Neely
University of Southern California
http://www-rcf.usc.edu/∼mjneely

***Abstract*—We consider the delay properties of one-hop networks with general interference constraints and multiple traffic streams with time-correlated arrivals. We first treat the case when arrivals are modulated by independent finite state Markov chains. We show that the well known maximal scheduling algorithm achieves average delay that grows at most logarithmically in the largest number of interferers at any link. Further, in the important special case when each Markov process has at most two states (such as bursty ON/OFF sources), we prove that average delay is independent of the number of nodes and links in the network, and hence is *order-optimal*. We provide tight delay bounds in terms of the individual auto-correlation parameters of the traffic sources. These are perhaps the first order-optimal delay results for controlled queueing networks that explicitly account for such statistical information.**

***Index Terms*—queueing analysis, Markov chains**

I. Introduction

This paper derives average delay bounds for one-hop wireless networks that use maximal scheduling subject to a general set of interference constraints. It is known that maximal scheduling algorithms are simple to implement and can support throughput within a constant factor of optimality. Our analysis shows that this type of scheduling also yields tight delay guarantees. In particular, when arrival processes are modulated by independent Markov processes, we show that average delay grows at most logarithmically in the number of nodes in the network. We then obtain an improved delay bound in the important special case when the individual Markov chains have at most two states (such as bursty ON/OFF sources). Average delay in this case is shown to be independent of the network size, and hence is *order-optimal*.

Specifically, we consider a network with N nodes and L links. Let $\mathcal{N}$ and $\mathcal{L}$ denote the node and link sets:

$$\mathcal{N} \triangleq \{1, 2, \ldots, N\}$$
$$\mathcal{L} \triangleq \{1, 2, \ldots, L\}$$

Each link $l \in \mathcal{L}$ represents a directed communication channel from one node to another, and we define $tran(l)$ and $rec(l)$ to be the corresponding transmitter and receiver nodes for link l (where $tran(l) \in \mathcal{N}$ and $rec(l) \in \mathcal{N}$). The network operates in slotted time with unit timeslots $t \in \{0, 1, 2, \ldots\}$. Every timeslot a decision is made about which links to activate for transmission. If a link l is activated during a particular slot, it sends exactly one packet from $tran(l)$ to $rec(l)$. However, due to scheduling and/or interference constraints, not all links can be simultaneously active during the same slot. These constraints are defined according to the general interference model of [1][2]: Each link $l \in \mathcal{L}$ is allowed to be active if and only if no other links within an *interference set* $\mathcal{S}_l$ are simultaneously active. For convenience, it is useful to define the set $\mathcal{S}_l$ to additionally include link l itself. That is, each set $\mathcal{S}_l$ consists of link l together with all possible interferers of link l. Note that these interference sets have the following pairwise symmetry property: For any two links $\omega, l \in \mathcal{L}$, we have that $\omega \in \mathcal{S}_l$ if and only if $l \in \mathcal{S}_\omega$.

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

The link sets $\mathcal{S}_l$ can be chosen to impose a variety of constraint models. For example, setting $\mathcal{S}_l$ to include all links adjacent to either the transmitter or receiver of link l imposes *matching constraints*. Matching constraint models arise naturally in scheduling problems for packet switches. They are also important for wireless networks where individual nodes can transmit or receive over at most one adjacent link, and cannot simultaneously transmit and receive (called the *node exclusive spectrum sharing* model in [3]). More general sets $\mathcal{S}_l$ can be used to model topology-dependent interference constraints for wireless networks, such as the constraint that no additional transmitting nodes can be activated within a specified distance of a node that is actively receiving.

Every timeslot, new data randomly arrives to the network. Let $A_l(t)$ represent the (integer) number of packets that arrive during slot t that are intended for transmission over link l. Packets are stored in separate queues according to their corresponding link. This is a one-hop network, so that packets exit the network once they are transmitted over their intended link. A network controller observes the current queue backlog and makes link activation decisions every slot subject to the transmission constraints.

It is well known that generalized max-weight scheduling can be used to achieve maximum throughput in such networks [4] [5]. However, this type of scheduling is difficult to implement in wireless networks with general interference constraints. In this paper we consider a simpler class of *maximal scheduling algorithms*. Maximal scheduling is of recent interest due to its low complexity and ease of distributed implementation. For $N \times N$ packet switches, maximal scheduling is known to support throughput that is within a factor of 2 of optimality, and to also have nice delay properties for i.i.d. inputs [6] [7]. Related constant factor throughput results have also been shown for wireless networks, including factor of 2 results for networks under matching constraints [3] and constant factor results for more general interference models [1] [2]. However, the work on wireless scheduling in [3][1][2] considers only throughput results and does not provide a delay analysis.

Further, while the work in [7] considers delay analysis for i.i.d. arrivals and a $N \times N$ packet switch, no existing work provides explicitly computable and order-optimal delay bounds for time-correlated arrivals.

Our work addresses the issues of general interference constraints and time-correlated "bursty" traffic simultaneously. We treat the general interference model of [1][2], but use the concept of *queue grouping* to derive the order-optimal delay results. Queue grouping techniques have been used in [8][9][6][2] to reduce scheduling complexity in switches and wireless networks. Queue groupings are also used in [10] to provide order-optimal delay results for opportunistic scheduling in a wireless downlink with a single server. The analysis in this paper particularly treats delay analysis in wireless networks with general constraint sets and time-correlated arrivals. Previous work in [11][5] also considers time-correlated arrivals for general multi-hop networks, but does not obtain order optimal delay results. Delay results within a logarithmic factor of optimality are developed in [12] for a $N \times N$ packet switch with i.i.d. inputs. Asymptotic delay optimality is studied in [13] for a heavy traffic regime where input rates are scaled so that they are very close to the capacity region boundary. Here, we focus on the case when input rates are a fixed fraction away from the boundary. We obtain tight delay bounds in terms of the individual auto-correlation parameters of the traffic sources. These are perhaps the first delay bounds for controlled queueing networks that explicitly incorporate such statistical information. This allows delay to be understood in terms of general models for network traffic. Our analysis includes the important special case of Markov modulated ON/OFF traffic sources. Specifically, we consider the case where all arrival streams are independent and modulated by a two-state Markov chain with ON and OFF states, and provide a closed form delay bound in terms of the rate and burst parameters of each source.

We first treat the case of general Markovian arrivals and prove a *structural result* about average delay, showing that average delay grows at most logarithmically in the worst case number of interferers of a given link, and hence is at most $O(\log(N))$. We then provide an explicit and tighter delay analysis for the special case when all Markov chains have at most two states. In this case, we prove average delay is *independent of* N. The time-correlated scenarios treated here are quite challenging to analyze, and we introduce a simple technique of *delayed Lyapunov analysis* to ensure the arrival processes couple sufficiently fast to a stationary distribution.

In the next section we present the network model. In Section III we provide the drift analysis, and in Section IV we present the logarithmic delay result for general time-correlated arrivals. The order-optimal delay analysis for 2-state chains is provided in Section V.

II. Network Model

Recall that $\mathcal{L}$ denotes the set of network links, and that transmission over each link is constrained by the general interference sets defined in the previous section.

A. Traffic Model

Suppose the arrival process $A_l(t)$ is modulated by a discrete time, stationary, ergodic Markov chain $Z_l(t)$ for each link $l \in \mathcal{L}$. Specifically, $Z_l(t)$ has finite state space $\mathcal{Z}_l = \{1, 2, \ldots, M_l\}$. For each link $l \in \mathcal{L}$ and state $m \in \mathcal{Z}_l$, arrivals $A_l(t)$ are conditionally independent and identically distributed according to mass function $p_{l,m}(a)$, where:

$$p_{l,m}(a) = Pr[A_l(t) = a \,|\, Z_l(t) = m] \quad \text{for } a \in \{0, 1, 2, \ldots\}$$

Define the conditional arrival rates $\lambda_l^{(m)}$ as follows:

$$\lambda_l^{(m)} = \mathbb{E}\left\{A_l(t) \,|\, Z_l(t) = m\right\}$$

We assume conditional second moments of arrivals are finite, so that $\mathbb{E}\left\{(A_l(t))^2 \,|\, Z_l(t) = m\right\} < \infty$ for all $l \in \mathcal{L}$ and all $m \in \mathcal{Z}_l$. Let $\pi_l^{(m)}$ represent the steady state probability that $Z_l(t) = m$. Define λ_l as the average arrival rate to link l:

$$\lambda_l = \sum_{m \in \mathcal{Z}_l} \pi_l^{(m)} \lambda_l^{(m)} \tag{1}$$

We assume that all Markov chains are in their steady state distribution at time 0, so that each $A_l(t)$ process is stationary and for all slots $t \geq 0$ and all links $l \in \mathcal{L}$ we have:

$$\mathbb{E}\left\{A_l(t)\right\} = \lambda_l$$

The Markov chains $Z_l(t)$ themselves may be correlated over different links $l \in \mathcal{L}$, although we mainly focus on the case when chains are independent. More detailed statistical information, such as the auto-correlation for individual inputs (and the spatial correlation between multiple inputs if they are not independent), is also important for delay analysis and shall be defined when needed. Note that this traffic model is quite general and includes the following important special cases:

- Case 1: $Z_l(t)$ has only one state and so arrivals $A_l(t)$ are i.i.d. over slots with some given distribution.
- Case 2: $Z_l(t)$ is a 2-state ON/OFF process where $A_l(t) = 1$ whenever $Z_l(t) = ON$ and $A_l(t) = 0$ whenever $Z_l(t) = OFF$.

B. Queueing

Define $Q_l(t)$ as the number of queued packets waiting for transmission over link l during slot t. Let $\boldsymbol{Q}(t) = (Q_l(t))_{l \in \mathcal{L}}$ be the vector of queue backlogs. Define $\mu_l(t) \in \{0, 1\}$ as the *transmission rate* offered to the link during slot t (in units of packets/slot). That is, $\mu_l(t) = 1$ if link l is scheduled for transmission on slot t, and $\mu_l(t) = 0$ otherwise. We assume the scheduler only schedules a link l that does not violate the interference constraints and that has a packet ready for transmission (so that $Q_l(t) > 0$). Let $\boldsymbol{\mu}(t) = (\mu_l(t))_{l \in \mathcal{L}}$ represent the transmission rate vector for slot t. Define $\mathcal{X}(t)$ as the set of *feasible transmission vectors* for slot t, representing all $\boldsymbol{\mu}(t)$ rate vectors that conform to the constraints defined by the interference sets $\mathcal{S}_l$ and the additional constraint that $\mu_l(t) = 1$ only if $Q_l(t) > 0$ (for each $l \in \mathcal{L}$). The queueing dynamics thus proceed as follows:

$$Q_l(t+1) = Q_l(t) - \mu_l(t) + A_l(t) \tag{2}$$

The goal is to observe the queue backlogs every slot and make scheduling decisions $\boldsymbol{\mu}(t) \in \mathcal{X}(t)$ so as to support all incoming traffic with average delay as small as possible.

C. Maximal Scheduling

Define the *network capacity region* Λ as the closure of the set of all arrival rate vectors $(\lambda_l)_{l\in\mathcal{L}}$ that can be stably supported, considering all possible scheduling algorithms that conform to the above constraints (see [5] for a discussion of capacity regions and stability). It is well known that scheduling according to a generalized *max-weight* rule every timeslot ensures stability and maximum throughput whenever arrival rates are interior to the capacity region [4] [5].[1] However, the max-weight rule involves an integer optimization that may be difficult to implement, and has delay properties that are difficult to analyze. Here, we assume scheduling is done according to a simpler *maximal scheduling* algorithm.[2] Specifically, given a queue backlog vector $\boldsymbol{Q}(t)$, a transmission vector $\boldsymbol{\mu}(t)$ is *maximal* if it satisfies the interference constraints and is such that for all links $l \in \mathcal{L}$, if $Q_l(t) > 0$ then $\mu_\omega(t) = 1$ for at least one link $\omega \in \mathcal{S}_l$. In words, this means that if link l has a packet, then either link l is selected for transmission, or some other link within the interference set $\mathcal{S}_l$ is selected. There is much recent interest in maximal scheduling because of its implementation simplicity (described briefly below) and its ability to support input rates within a constant factor of the capacity region for wireless networks [3] [1] [2] and for $N \times N$ packet switches [6].

One way to achieve a maximal scheduling is as follows: First select any non-empty link $l \in \mathcal{L}$ and label it "active." Then select any other non-empty link that does not conflict with the active link l (i.e., that is not within $\mathcal{S}_l$). Label this second link "active." Continue in the same way, selecting new non-empty links that do not conflict with any previously selected links, until no more links can be added. It is not difficult to see that this final set of links labeled "active" has the desired maximal property. Maximal link selections are not unique, and can alternatively be found in a distributed manner, where multiple nodes attempt to activate their non-conflicting, non-empty links simultaneously, and contentions are resolved locally. This distributed implementation also requires multiple iterations before the set of selected links becomes maximal.

All maximal link selections have the following important mathematical property.

Lemma 1: Under any maximal link scheduling for $\boldsymbol{\mu}(t)$, for all links $l \in \mathcal{L}$ we have:

$$Q_l(t) \sum_{\omega\in\mathcal{S}_l} \mu_\omega(t) \geq Q_l(t) \tag{3}$$

Proof: Consider any link $l \in \mathcal{L}$. If $Q_l(t) = 0$, then (3) reduces to $0 \geq 0$ and is trivially true. Else, if $Q_l(t) > 0$ then $\mu_\omega(t) = 1$ for at least one link ω within $\mathcal{S}_l$ (by definition of a maximal link selection), and so $\sum_{\omega\in\mathcal{S}_l} \mu_\omega(t) \geq 1$, which proves (3). ∎

[1]Specifically, the generalized max-weight rule schedules to maximize $\sum_{l\in\mathcal{L}} Q_l(t)\mu_l(t)$ subject to $\boldsymbol{\mu}(t) \in \mathcal{X}(t)$.

[2]Max-weight scheduling has the maximal property, and so our delay results also apply to max-weight when inputs rates are in Λ^* (of Section II-D).

In this paper, we assume transmission decisions are made every slot according to any maximal scheduling. For convenience, we further assume that the maximal scheduling has a well defined probabilistic structure given the queue backlog matrix, so that the entire queueing system can be viewed as an ergodic Markov chain with a countably infinite state space. The inequality (3) is the only additional property of maximal scheduling required in our analysis.

D. The Reduced Throughput Region

Define Λ^* as the set of all rate vectors $(\lambda_l)_{l\in\mathcal{L}}$ that satisfy the following:

$$\sum_{\omega\in\mathcal{S}_l} \lambda_\omega \leq 1 \quad \text{for all } l \in \mathcal{L}$$

The set Λ^* is overly restrictive, as it is possible to have more than one simultaneously active link within a given set $\mathcal{S}_l$ (provided that link l is idle). However, Λ^* is typically within a constant factor of the capacity region Λ. For example, in networks with matching constraints only, it is not difficult to show that $\frac{1}{2}\Lambda \subset \Lambda^*$, so that the throughput region Λ^* is within a factor of 2 of optimality. Further, in networks with general inteference sets $\mathcal{S}_l$ where each set $\mathcal{S}_l$ contains at most K links, it can be shown that $\frac{1}{K}\Lambda \subset \Lambda^*$ [1].

Throughout this paper, we assume input rates $(\lambda_l)_{l\in\mathcal{L}}$ are interior to the set Λ^*. Specifically, we assume there exists a value ρ^* such that $0 < \rho^* < 1$, where:

$$\sum_{\omega\in\mathcal{S}_l} \lambda_\omega \leq \rho^* \quad \text{for all } l \in \mathcal{L} \tag{4}$$

The value ρ^* represents the *relative network loading*, as it can be viewed as a loading factor relative to the reduced throughput region Λ^*.

III. Drift Analysis

Recall that $\boldsymbol{Q}(t) = (Q_{ij}(t))$. Our technique relies on the concept of *queue grouping*. For each link $l \in \mathcal{L}$, define:

$$\hat{Q}_{\mathcal{S}_l}(t) \triangleq \sum_{\omega\in\mathcal{S}_l} Q_\omega(t) \tag{5}$$

Thus, $\hat{Q}_{\mathcal{S}_l}(t)$ is the sum of all queue backlogs of links within the interference set $\mathcal{S}_l$ of link l. Define the Lyapunov function:

$$L(\boldsymbol{Q}(t)) \triangleq \frac{1}{2} \sum_{l\in\mathcal{L}} Q_l(t)\hat{Q}_{\mathcal{S}_l}(t)$$

The queue-grouped structure of this Lyapunov function is similar the functions used in [2][8] to prove queue stability when input rates are a fixed fraction away from the capacity boundary. An alternate proof of rate-stability is given in [1]. However, the prior work in this area does not directly consider delay performance. Below we provide a more detailed drift analysis that yields explicit and tight delay bounds.

For each link $l \in \mathcal{L}$, define the *group departures* $\hat{\mu}_{\mathcal{S}_l}(t)$ and *group arrivals* $\hat{A}_{\mathcal{S}_l}(t)$ as follows:

$$\begin{aligned}\hat{\mu}_{\mathcal{S}_l}(t) &\triangleq \sum_{\omega\in\mathcal{S}_l} \mu_\omega(t) \\ \hat{A}_{\mathcal{S}_l}(t) &\triangleq \sum_{\omega\in\mathcal{S}_l} A_\omega(t)\end{aligned}$$

Thus:

$$\hat{Q}_{\mathcal{S}_l}(t+1) = \hat{Q}_{\mathcal{S}_l}(t) - \hat{\mu}_{\mathcal{S}_l}(t) + \hat{A}_{\mathcal{S}_l}(t) \quad (6)$$

Define the 1-step *unconditional Lyapunov drift* as follows:[3]

$$\Delta(t) \triangleq \mathbb{E}\left\{L(\boldsymbol{Q}(t+1)) - L(\boldsymbol{Q}(t))\right\} \quad (7)$$

where the expectation is over the randomness of $\boldsymbol{Q}(t)$ and the randomness of the system dynamics given the value of $\boldsymbol{Q}(t)$.

Lemma 2: (Drift Under Maximal Scheduling) If maximal scheduling is implemented every timeslot (using any maximal scheduling algorithm), the resulting unconditional Lyapunov drift $\Delta(t)$ satisfies the following for all slots t:

$$\Delta(t) \leq \mathbb{E}\left\{B(t)\right\} - \sum_{l\in\mathcal{L}} \mathbb{E}\left\{Q_l(t)(1 - \hat{A}_{\mathcal{S}_l}(t))\right\} \quad (8)$$

where

$$B(t) \triangleq \frac{1}{2}\sum_{l\in\mathcal{L}}\left[A_l(t)\hat{A}_{\mathcal{S}_l}(t) - 2\hat{A}_{\mathcal{S}_l}(t)\mu_l(t) + \mu_l(t)\right] \quad (9)$$

Proof: See Appendix A. ■

A. Lyapunov Drift Theorem

The drift expression (8) can be used to prove stability and delay properties of maximal matching via the following theorem:

Theorem 1: (Lyapunov drift [5]) Let $\boldsymbol{Q}(t)$ be a vector process of queue backlogs that evolve according to some probability law, and let $L(\boldsymbol{Q}(t))$ be a non-negative function of $\boldsymbol{Q}(t)$. If there exist processes $f(t)$ and $g(t)$ such that the following is satisfied for all time t:

$$\Delta(t) \leq \mathbb{E}\left\{g(t)\right\} - \mathbb{E}\left\{f(t)\right\}$$

then:

$$\limsup_{t\to\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left\{f(\tau)\right\} \leq \limsup_{t\to\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left\{g(t)\right\} \quad \square$$

A proof of this theorem is provided in Lemma 5.3 of [5].

B. Analysis for i.i.d. Arrivals

Define $\boldsymbol{A}(t) \triangleq (A_l(t))_{l\in\mathcal{L}}$ as the vector of new arrivals. Consider first the case when all arrival vectors $\boldsymbol{A}(t)$ are i.i.d. over timeslots, with rate vector $\boldsymbol{\lambda} = (\lambda_l)_{l\in\mathcal{L}}$ (the arrivals over different links in the same slot may be correlated). For each link l, define $\hat{\lambda}_{\mathcal{S}_l}$ as the sum of arrival rates over all input streams corresponding to links within the set $\mathcal{S}_l$. That is:

$$\hat{\lambda}_{\mathcal{S}_l} \triangleq \sum_{\omega\in\mathcal{S}_l}\lambda_\omega$$

Note by the loading assumption (4) that $\hat{\lambda}_{\mathcal{S}_l} \leq \rho^*$, where ρ^* is a value such that $0 < \rho^* < 1$. By independence of arrivals every slot, we have for all t:

$$\begin{aligned}\mathbb{E}\left\{Q_l(t)(1 - \hat{A}_{\mathcal{S}_l}(t))\right\} &= \mathbb{E}\left\{Q_l(t)\right\}(1 - \hat{\lambda}_{\mathcal{S}_l}) \\ &\geq \mathbb{E}\left\{Q_l(t)\right\}(1 - \rho^*)\end{aligned}$$

[3]The randomness is taken with respect to some given initial queue distribution at time 0 and the random system events that happen thereafter.

Using this inequality directly in the Lyapunov drift expression (8) yields (for all slots t):

$$\Delta(t) \leq \mathbb{E}\left\{B(t)\right\} - (1-\rho^*)\sum_{l\in\mathcal{L}}\mathbb{E}\left\{Q_l(t)\right\}$$

Plugging the above drift inequality into the Lyapunov drift theorem (Theorem 1) (using $g(t) \triangleq B(t)$ and $f(t) \triangleq (1 - \rho^*)\sum_{l\in\mathcal{L}}Q_l(t)$) yields:

$$\limsup_{t\to\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\sum_{l\in\mathcal{L}}\mathbb{E}\left\{Q_l(\tau)\right\} \leq \frac{\overline{B}}{1-\rho^*} \quad (10)$$

where:

$$\overline{B} \triangleq \limsup_{t\to\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left\{B(\tau)\right\}$$

Note from (9) that $\overline{B} < \infty$, and hence the queueing network is strongly stable with finite time average queue backlogs. Because it evolves according to an ergodic Markov chain with a countably infinite state space, it can be shown that limiting time averages exist and are equal to the steady state averages. Thus, the left hand side of (10) represents the time average total queue backlog in the system (summed over all queues). Let $\overline{Q}_l$ be the average queue backlog in link l (for each $l \in \mathcal{L}$). We thus have:

Theorem 2: (Time-Independent Arrivals) If the arrival vector process $\boldsymbol{A}(t)$ is i.i.d. over slots with a relative network loading $\rho^* < 1$ (defined in (4)), then:

(a) The average total network congestion satisfies:

$$\sum_{l\in\mathcal{L}}\overline{Q}_l \leq \frac{\sum_{l\in\mathcal{L}}\left[\mathbb{E}\left\{A_l(t)\hat{A}_{\mathcal{S}_l}(t)\right\} - 2\lambda_l\hat{\lambda}_{\mathcal{S}_l} + \lambda_l\right]}{2(1-\rho^*)}$$

(b) If arrival streams $A_l(t)$ are i.i.d. over slots and are also independent of each other, then total average network delay $\overline{W}$ satisfies:

$$\overline{W} \leq \frac{1 + \frac{1}{\lambda_{tot}}\sum_{l\in\mathcal{L}}\left[\sigma_l^2 - \lambda_l\hat{\lambda}_{\mathcal{S}_l}\right]}{2(1-\rho^*)} \quad (11)$$

where $\lambda_{tot} \triangleq \sum_{l\in\mathcal{L}}\lambda_l$, and where $\sigma_l^2 \triangleq \mathbb{E}\left\{(A_l(t))^2\right\} - \lambda_l^2$ and represents the variance of $A_l(t)$.

The average average congestion bound in part (a) of the above theorem is found by computing $\overline{B}$ using (9). Part (b) is proven via the fact that total average congestion is equal to $\lambda_{tot}\overline{W}$ (by Little's Theorem). This analysis is provided in Appendix B.

C. Delay under Poisson and Bernoulli Inputs

Note that if all arrival processes $A_l(t)$ are independent and Poisson with rate λ_l, we have that $\sigma_l^2 = \lambda_l$. The average delay bound (11) in this case reduces to:

$$\overline{W}_{Poisson} \leq \frac{1 - \frac{1}{2\lambda_{tot}}\sum_{l\in\mathcal{L}}\lambda_l\hat{\lambda}_{\mathcal{S}_l}}{(1-\rho^*)} \quad (12)$$

This demonstrates that average delay is $O(1)$, that is, it is *independent of the network size* N. Hence, maximal scheduling achieves *order optimal delay* with respect to N, provided that the arrival rates are interior to the reduced throughput region Λ^*, as described by the constraints (4). This is in contrast to the $O(N)$ average delay bounds derived for the throughput-optimal max-weight scheduling for $N \times N$ packet switches in [14] and for wireless networks in [11] [15]. The expression in the right hand side of (12) also provides an upper bound on delay in the case of independent Bernoulli arrivals, because σ_l^2 for a Bernoulli variable is less than that of a Poisson variable. Finally, we note that the term $\frac{1}{\lambda_{tot}}\sum_{l\in\mathcal{L}}\sigma_l^2$ in (11) is typically $O(1)$ for any inputs with a finite variance. For example, it can be shown to be $O(1)$ whenever there exists a constant A_{max} such that $A_l(t) \leq A_{max}$ for all $l \in \mathcal{L}$ and all t.

IV. Logarithmic Delay for Time-Correlated Arrivals

Consider the case of finite-state ergodic Markov modulated arrivals, as described in Section II-A. Assume that all traffic rates satisfy the loading constraints (4) with relative network loading ρ^*. Let $\boldsymbol{H}(t)$ represent the past history of all actual arrivals (of all processes) up to but not including time t. For a given link $l \in \mathcal{L}$, suppose there exists a non-negative function $\epsilon_l(T)$ (for $T \in \{0, 1, 2, \ldots\}$) such that, regardless of past history $\boldsymbol{H}(t)$, we have:

$$\mathbb{E}\left\{A_l(t) \mid \boldsymbol{H}(t-T)\right\} \leq \lambda_l + \epsilon_l(T)$$

and such that: $\lim_{T\to\infty}\epsilon_l(T) = 0$. That is, $\epsilon_l(T)$ characterizes the time required for the process $A_l(t)$ to converge to its stationary mean, regardless of the initial condition. It can be shown that all finite state ergodic Markov processes converge exponentially fast to their steady state (see, for example, [16]). Hence for each $l \in \mathcal{L}$ we can define $\epsilon_l(T)$ as follows:

$$\epsilon_l(T) = \nu_l \gamma_l^{T+1} \tag{13}$$

for some constant ν_l and some decay factor γ_l such that $0 < \gamma_l < 1$. The ν_l and γ_l constants can in principle be determined as parameters from the Markov chain $Z_l(t)$. Here, we prove a *structural result* concerning logarithmic delay in terms of these parameters. Define $\hat{\rho}_l(T)$ as follows:

$$\hat{\rho}(T) \triangleq \max_{l\in\mathcal{L}}\left[\hat{\lambda}_{\mathcal{S}_l} + \sum_{\omega\in\mathcal{S}_l}\epsilon_l(T)\right]$$

Note that $\hat{\lambda}_{\mathcal{S}_l} \leq \rho^* < 1$ for all $l \in \mathcal{L}$, and so there exist integers T such that $\hat{\rho}(T) < 1$.

Theorem 3: (General Time-Correlated Arrivals) If arrival processes have rates $(\lambda_L)_{l\in\mathcal{L}}$ in the interior of Λ^*, then for any integer $T \geq 0$ such that $\hat{\rho}(T) < 1$, we have:

(a) The average total network congestion satisfies:

$$\sum_{l\in\mathcal{L}}\overline{Q}_l \leq \frac{\tilde{B} + \tilde{F}_T}{1 - \hat{\rho}(T)}$$

where:

$$\tilde{B} \triangleq \frac{1}{2}\sum_{l\in\mathcal{L}}\left[\lambda_l + \mathbb{E}\left\{A_l(t)\hat{A}_{\mathcal{S}_l}(t)\right\}\right] \tag{14}$$

$$\tilde{F}_T \triangleq \sum_{l\in\mathcal{L}}\sum_{k=1}^{T}\mathbb{E}\left\{\hat{A}_{\mathcal{S}_l}(k)A_l(0)\right\} \tag{15}$$

(b) If arrival processes $A_l(t)$ for different links l are additionally independent of each other, then the constants $\tilde{B}$ and $\tilde{F}_T$ satisfy:

$$\tilde{B} = \frac{1}{2}\sum_{l\in\mathcal{L}}\left[\lambda_l + \lambda_l\hat{\lambda}_{\mathcal{S}_l} + \sigma_l^2\right] \tag{16}$$

$$\tilde{F}_T = \sum_{l\in\mathcal{L}}\sum_{k=1}^{T}[\hat{\lambda}_{\mathcal{S}_l}\lambda_l + \theta_l(k)] \tag{17}$$

where $\sigma_l^2 \triangleq \mathbb{E}\left\{(A_l(t))^2\right\} - \lambda_l^2$ is the variance of $A_l(t)$, and $\theta_l(k)$ is the auto-correlation in $A_l(t)$ and is defined:

$$\theta_l(k) \triangleq \mathbb{E}\left\{A_l(t+k)A_l(t)\right\} - \lambda_l^2$$

Proof: The proof is given in Appendix C. ■

A. Discussion of the Delay Result

By Little's Theorem, if the conditions of Theorem 3 are satisfied, then average network delay $\overline{W}$ satisfies:

$$\overline{W} \leq \frac{\frac{1}{\lambda_{tot}}(\tilde{B} + \tilde{F}_T)}{1 - \hat{\rho}(T)}$$

The parameter T only affects the delay bound and does not affect the maximal scheduling algorithm. Thus, the bound can be optimized over all integers T such that $\hat{\rho}(T) < 1$. Here we show how the resulting bound grows as a function of the network size. First note that in the case when arrival processes are independent of each other, the constant $\tilde{B}$ in (16) satisfies:

$$\tilde{B} \leq \lambda_{tot}\left[1 + \frac{1}{2\lambda_{tot}}\sum_{l\in\mathcal{L}}\sigma_l^2\right]$$

where $\lambda_{tot} \triangleq \sum_{l\in\mathcal{L}}\lambda_l$. This is because $\hat{\lambda}_{\mathcal{S}_l} \leq 1$ for all $l \in \mathcal{L}$. Likewise, the constant $\tilde{F}_T$ in (17) satisfies:

$$\tilde{F}_T \leq \lambda_{tot}T + \sum_{l\in\mathcal{L}}\sum_{k=1}^{T}\theta_l(k)$$

Therefore, when arrival processes $A_l(t)$ are independent of each other, average network delay satisfies:

$$\overline{W} \leq \frac{T + 1 + \frac{1}{2\lambda_{tot}}\sum_{l\in\mathcal{L}}\sigma_l^2 + \frac{1}{\lambda_{tot}}\sum_{l\in\mathcal{L}}\sum_{k=1}^{T}\theta_l(k)}{1 - \hat{\rho}(T)}$$

The values of $\frac{1}{2\lambda_{tot}}\sum_{l\in\mathcal{L}}\sigma_l^2$ and $\frac{1}{\lambda_{tot}}\sum_{l\in\mathcal{L}}\theta_l(k)$ are typically independent of N, and so the numerator is roughly linear in the T value. Because we have finite state ergodic Markov chains, from (13) we see the function $\hat{\rho}(T)$ has the form:

$$\hat{\rho}(T) \leq \rho^* + |\mathcal{S}|\nu\gamma^{T+1}$$

where $|\mathcal{S}|$ the cardinality of the largest interference set $\mathcal{S}_l$, and ν and γ are the largest values of ν_l and γ_l, respectively, over all links $l \in \mathcal{L}$. In this case, we have $\hat{\rho}(T) \leq (1 + \rho^*)/2$ whenever:

$$|\mathcal{S}|\nu\gamma^{T+1} \leq (1 - \rho^*)/2$$

which holds when T is chosen as the smallest integer that satisfies:

$$\frac{\log\left(2\nu\left|\mathcal{S}\right|/(1-\rho^*)\right)}{\log(1/\gamma)} - 1 \leq T \leq \frac{\log\left(2\nu\left|\mathcal{S}\right|/(1-\rho^*)\right)}{\log(1/\gamma)}$$

Thus, the above delay bound grows at most logarithmically in $|\mathcal{S}|$. A more explicit and *order-optimal* delay analysis is provided in the next section, where the special case of 2-state Markov chains is considered and average delay is shown to be independent of the network size.

V. Analysis for Time-Correlated Arrivals with Two States

Consider the case of Markov modulated arrivals, as described in Section II-A, where all Markov chains $Z_l(t)$ have at most two states (labeled "1" and "2"). Let $\tilde{\mathcal{L}}$ be the set of all links $l \in \mathcal{L}$ that have exactly two states with different conditional rates $\lambda_l^{(1)}$ and $\lambda_l^{(2)}$. The transition probabilities are given by β_l and δ_l for each two-state chain $Z_l(t)$, as shown in Fig. 1. Note that this model includes the important special case of ON/OFF inputs, where $A_l(t)$ has a single packet arrival when in the ON state and has no arrivals when in the OFF state.[4] Assume that $0 < \delta_l < 1$ and $0 < \beta_l < 1$ for all $l \in \tilde{\mathcal{L}}$, and define $\pi_l^{(1)}$ and $\pi_l^{(2)}$ as the steady state probabilities for each two-state chain $Z_l(t)$:

$$\pi_l^{(1)} = \frac{\delta_l}{\beta_l + \delta_l} \quad , \quad \pi_l^{(2)} = \frac{\beta_l}{\beta_l + \delta_l}$$

Arrival processes $A_l(t)$ for the remaining links $l \notin \tilde{\mathcal{L}}$ are either i.i.d. over slots (effectively "one-state" chains), or have two states but with $\lambda_l^{(1)} = \lambda_l^{(2)}$ (in the latter case, the different states may correspond to different conditional second moments).

The Markov chains $Z_l(t)$ are possibly correlated over different links $l \in \mathcal{L}$, although we focus primarily on the case when all chains are independent. The chains are assumed to be *stationary*, so that for each link $l \in \mathcal{L}$ we have $\mathbb{E}\{A_l(t)\} = \lambda_l$ for all time t. The time average rates λ_l for all 2-state arrival processes are given by:

$$\lambda_l = \pi_l^{(1)}\lambda_l^{(1)} + \pi_l^{(2)}\lambda_l^{(2)}$$

The traffic rates $(\lambda_l)_{l\in\mathcal{L}}$ are assumed to satisfy the loading constraints (4) with relative network loading $\rho^* < 1$. In the previous section we showed that the system is strongly stable and ergodic with finite steady state queue backlogs. Here we provide a tighter delay analysis using a combination of Markov chain theory and Lyapunov drift theory.

[4]The conditional rates for such an ON/OFF example are given by $\lambda_l^{(1)} = 1$, $\lambda_l^{(2)} = 0$, where states 1 and 2 are associated with the ON and OFF states, respectively, as shown in Fig. 1.

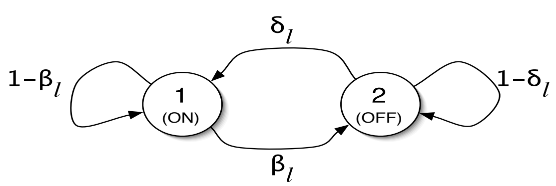

Fig. 1. The 2-state Markov chain $Z_l(t)$ for link l.

A. 2-State Drift Analysis

Recall the definition $\hat{A}_{\mathcal{S}_l}(t) \triangleq \sum_{\omega\in\mathcal{S}_l} A_\omega(t)$. Thus, from the drift inequality (8) we have that the unconditional Lyapunov drift $\Delta(t)$ satisfies the following every slot t:

$$\begin{aligned}\Delta(t) \quad \leq \quad & \mathbb{E}\{B(t)\} - \sum_{l\in\mathcal{L}} \mathbb{E}\{Q_l(t)\} \\ & + \sum_{l\in\mathcal{L}} \sum_{\omega\in\mathcal{S}_l} \mathbb{E}\{Q_l(t)A_\omega(t)\} \qquad (18)\end{aligned}$$

Now assume that the system is in *steady state* at time t, and is also in steady state at time $t-1$. Fix a link $l \in \mathcal{L}$ and an arrival process $A_\omega(t)$. We shall derive a relationship between $\mathbb{E}\{Q_l(t)A_\omega(t)\}$ and $\mathbb{E}\{Q_l(t)\}$. To this end, first note that for any link $l \in \mathcal{L}$ we have:

$$\mathbb{E}\{Q_l(t)A_\omega(t)\} = \mathbb{E}\{Q_l(t)\}\lambda_\omega \text{ if } \omega \notin \tilde{\mathcal{L}} \qquad (19)$$

That is, the expectation has a simple product form in the special case when $\omega \notin \tilde{\mathcal{L}}$, because $A_\omega(t)$ is either i.i.d. over slots or has the same conditional rate $\lambda_\omega^{(1)} = \lambda_\omega^{(2)} = \lambda_\omega$.

For the rest of this subsection we consider the opposite and more challenging case when $\omega \in \tilde{\mathcal{L}}$. In this case, we have:

$$\mathbb{E}\{Q_l(t)\} = \sum_{m=1}^{2} \pi_\omega^{(m)} \mathbb{E}\{Q_l(t) \mid Z_\omega(t) = m\} \qquad (20)$$

However, we also have:

$$\begin{aligned}&\mathbb{E}\{Q_l(t)A_\omega(t)\} \\ &\quad = \sum_{m=1}^{2} \pi_\omega^{(m)} \mathbb{E}\{Q_l(t)A_\omega(t) \mid Z_\omega(t) = m\} \\ &\quad = \sum_{m=1}^{2} \pi_\omega^{(m)} \lambda_\omega^{(m)} \mathbb{E}\{Q_l(t) \mid Z_\omega(t) = m\} \qquad (21)\end{aligned}$$

This final equality holds because the expectation of $Q_l(t)$ is conditionally independent of $A_\omega(t)$ given $Z_\omega(t)$. Because the system is in steady state, the quantities $\mathbb{E}\{Q_l(t)\}$, $\mathbb{E}\{Q_l(t)A_\omega(t)\}$, and $\mathbb{E}\{Q_l(t) \mid Z_\omega(t) = m\}$ do not depend on t, and we define:

$$\begin{aligned}\mathbb{E}\{Q_l\} &\triangleq \mathbb{E}\{Q_l(t)\} \\ \mathbb{E}\{Q_l A_\omega\} &\triangleq \mathbb{E}\{Q_l(t)A_\omega(t)\} \\ x_{l,\omega}^{(m)} &\triangleq \pi_\omega^{(m)} \mathbb{E}\{Q_l(t) \mid Z_\omega(t) = m\}\end{aligned}$$

The equalities (20) and (21) can be re-written:

$$\mathbb{E}\{Q_l\} = x_{l,\omega}^{(1)} + x_{l,\omega}^{(2)} \qquad (22)$$

$$\mathbb{E}\{Q_l A_\omega\} = \lambda_\omega^{(1)} x_{l,\omega}^{(1)} + \lambda_\omega^{(2)} x_{l,\omega}^{(2)} \qquad (23)$$

Equations (22) and (23) are two linear equations that express a relationship between 4 unknowns (where the unknowns are $\mathbb{E}\{Q_l\}, \mathbb{E}\{Q_l A_\omega\}$, $x_{l,\omega}^{(1)}$ and $x_{l,\omega}^{(2)}$). To express a direct linear relationship between $\mathbb{E}\{Q_l\}$ and $\mathbb{E}\{Q_l A_\omega\}$, we require an additional equation. To this end, note that:

$$Q_l(t) = Q_l(t-1) - \mu_l(t-1) + A_l(t-1)$$

Therefore:

$$\mathbb{E}\{Q_l(t)A_\omega(t)\} = \mathbb{E}\{Q_l(t-1)A_\omega(t)\} - D_{l,\omega} + C_{l,\omega} \qquad (24)$$

where $C_{l,\omega}$ and $D_{l,\omega}$ are defined as:

$$C_{l,\omega} \triangleq \mathbb{E}\left\{A_l(t-1)A_\omega(t)\right\} \tag{25}$$
$$D_{l,\omega} \triangleq \mathbb{E}\left\{\mu_l(t-1)A_\omega(t)\right\} \tag{26}$$

Now:

$$\begin{aligned}\mathbb{E}\left\{Q_l(t-1)A_\omega(t)\right\} &\\ = \sum_{m=1}^{2}\pi_\omega^{(m)}\mathbb{E}\left\{Q_l(t-1)A_\omega(t) \mid Z_\omega(t-1)=m\right\} &\\ = \sum_{m=1}^{2}\pi_\omega^{(m)}h_\omega^{(m)}\mathbb{E}\left\{Q_l(t-1) \mid Z_\omega(t-1)=m\right\} & \quad (27)\end{aligned}$$

where $h_\omega^{(m)}$ is defined:

$$h_\omega^{(m)} \triangleq \mathbb{E}\left\{A_\omega(t) \mid Z_\omega(t-1)=m\right\} \tag{28}$$

The last equality follows again because $Q_l(t-1)$ is conditionally independent of $A_\omega(t)$ given $Z_\omega(t-1)$. However, because the system is in steady state at time t and also at time $t-1$, it follows that:

$$\pi_\omega^{(m)}\mathbb{E}\left\{Q_l(t-1) \mid Z_\omega(t-1)=m\right\} = x_{l,\omega}^{(m)} \tag{29}$$

Therefore, using (29) and (27), equation (24) becomes the following:

$$\mathbb{E}\left\{Q_lA_\omega\right\} = C_{l,\omega} - D_{l,\omega} + h_\omega^{(1)}x_{l,\omega}^{(1)} + h_\omega^{(2)}x_{l,\omega}^{(2)} \tag{30}$$

The constants $h_\omega^{(m)}$, defined in (28), can be computed directly from the transition probabilities for chain $Z_\omega(t)$:

$$h_\omega^{(1)} = (1-\beta_\omega)\lambda_\omega^{(1)} + \beta_\omega\lambda_\omega^{(2)}$$
$$h_\omega^{(2)} = \delta_\omega\lambda_\omega^{(1)} + (1-\delta_\omega)\lambda_\omega^{(2)}$$

The linear equations (22), (23), and (30) involve three equations and four unknowns, and can be shown to be linearly independent whenever $\lambda_\omega^{(1)} \neq \lambda_\omega^{(2)}$ (which holds for all $\omega \in \tilde{\mathcal{L}}$). This directly leads to the following lemma that relates $\mathbb{E}\left\{Q_l\right\}$ and $\mathbb{E}\left\{Q_lA_\omega\right\}$.

Lemma 3: For all links $l \in \mathcal{L}$, we have:

$$\mathbb{E}\left\{Q_lA_\omega\right\} = \mathbb{E}\left\{Q_l\right\}\lambda_\omega + \frac{C_{l,\omega} - D_{l,\omega}}{\beta_\omega + \delta_\omega} \quad \text{if } \omega \in \tilde{\mathcal{L}} \tag{31}$$

Proof: The result follows by eliminating the $x_{l,\omega}^{(1)}$ and $x_{l,\omega}^{(2)}$ variables from (22), (23), and (30), and the computation is omitted for brevity. ■

Note from the definitions (25) and (26) that $C_{l,\omega} \geq 0$, $D_{l,\omega} \geq 0$ for all $l \in \mathcal{L}$ and all $\omega \in \tilde{\mathcal{L}}$. Define $C_{l,\omega} = 0$ for $\omega \notin \tilde{\mathcal{L}}$. Using (31) and (19) we find that for all link pairs $l, \omega \in \mathcal{L}$, we have:

$$\mathbb{E}\left\{Q_lA_\omega\right\} \leq \mathbb{E}\left\{Q_l\right\}\lambda_\omega + \frac{C_{l,\omega}}{\beta_\omega + \delta_\omega}$$

where the inequality comes because we have neglected the $D_{l,\omega}$ constant. Using this expression directly in the Lyapunov drift inequality (18) yields the following drift expression that holds at any time t at which the system is in steady state:

$$\begin{aligned}\Delta(t) \leq\ & \mathbb{E}\left\{B(t)\right\} - \sum_{l\in\mathcal{L}}\mathbb{E}\left\{Q_l(t)\right\}(1-\rho^*) \\ & + \sum_{l\in\mathcal{L}}\sum_{\omega\in\mathcal{S}_l\cap\tilde{\mathcal{L}}}\frac{C_{l,\omega}}{\beta_\omega+\delta_\omega}\end{aligned} \tag{32}$$

where we have used the fact that $\sum_{\omega\in\mathcal{S}_l}\lambda_\omega \leq \rho^*$. In the next subsection, this drift expression is used with the Lyapunov drift Theorem (Theorem 1) to prove a tight delay bound.

B. The Delay Bound for 2-State Markov Modulated Arrivals

Theorem 4: If the input rates $(\lambda_l)_{l\in\mathcal{L}}$ satisfy the loading constraints (4) for a given relative network loading $\rho^* < 1$, then:

(a) The system is stable with steady state average congestion that satisfies:

$$\sum_{l\in\mathcal{L}}\overline{Q}_l \leq \frac{\tilde{B}+\tilde{C}}{1-\rho^*}$$

where:

$$\tilde{B} \triangleq \frac{1}{2}\sum_{l\in\mathcal{L}}\left[\mathbb{E}\left\{A_l(t)\hat{A}_{\mathcal{S}_l}(t)\right\} + \lambda_l\right]$$
$$\tilde{C} \triangleq \sum_{l\in\mathcal{L}}\sum_{\omega\in\mathcal{S}_l\cap\tilde{\mathcal{L}}}\frac{\mathbb{E}\left\{A_l(t-1)A_\omega(t)\right\}}{\beta_\omega+\delta_\omega}$$

(b) If all Markov chains $Z_l(t)$ are additionally independent of each other, then:

$$\tilde{B} = \frac{1}{2}\sum_{l\in\mathcal{L}}\left[\lambda_l\hat{\lambda}_{\mathcal{S}_l} + \sigma_l^2 + \lambda_l\right]$$
$$\mathbb{E}\left\{A_l(t-1)A_\omega(t)\right\} = \lambda_l\lambda_\omega \text{ if } l \neq \omega$$

Hence, average delay $\overline{W}$ satisfies:

$$\begin{aligned}\overline{W} \leq\ & \frac{1+\frac{1}{2\lambda_{tot}}\sum_{l\in\mathcal{L}}\sigma_l^2}{1-\rho^*} \\ & + \max_{\omega\in\tilde{\mathcal{L}}}\left[\frac{1}{\beta_\omega+\delta_\omega}\right]\left(\frac{1+\frac{1}{\lambda_{tot}}\sum_{l\in\tilde{\mathcal{L}}}\theta_l[1]}{1-\rho^*}\right)\end{aligned}$$

where $\lambda_{tot} \triangleq \sum_{l\in\mathcal{L}}\lambda_l$, and $\theta_l[1] \triangleq \mathbb{E}\left\{A_l(t-1)A_l(t)\right\} - (\lambda_l)^2$ is the 1-slot auto-correlation for process $A_l(t)$.

Proof: The result follows directly from (32) via the Lyapunov drift theorem (Theorem 1), and is omitted for brevity. ■

Note that $\theta_l[1] \leq \lambda_l\lambda^{(max)}$, where $\lambda_l^{(max)}$ is the largest conditional rate over all links and states. Thus, the numerator in the final term in the above delay bound satisfies:

$$\frac{1}{\lambda_{tot}}\sum_{l\in\tilde{\mathcal{L}}}\theta_l[1] \leq \lambda^{max}$$

Therefore, the above bound is $O(1)$ (independent of the network size N). In the special case of ON/OFF sources, where a single packet arrives from stream l when $Z_l(t) = ON$ and no packet arrives when $Z_l(t) = OFF$, we have $\lambda_l^{(1)} = 1$ and $\lambda_l^{(2)} = 0$, and $\lambda_l = \pi_l^{(1)}$. Further:

$$\sigma_l^2 = \lambda_l(1-\lambda_l) \leq \lambda_l$$
$$\theta_l[1] = \lambda_l(1-\lambda_l)(1-(\beta_l+\delta_l)) \leq \lambda_l$$

Thus, average delay in this ON/OFF example satisfies:

$$\overline{W}_{ON/OFF} \leq \frac{3/2 + \max_{\omega\in\tilde{\mathcal{L}}}\left[2/(\beta_\omega+\delta_\omega)\right]}{1-\rho^*}$$

Note that $1/\beta_l$ is the average burst size (i.e., the average time spent in the ON state), and so the numerator roughly grows linearly in the largest average burst size over any input.

APPENDIX A — PROOF OF LEMMA 2

To compute $\Delta(t)$, note that using (6) and (2) yields:

$$\begin{aligned} Q_l(t+1)\hat{Q}_{\mathcal{S}_l}(t+1) &= Q_l(t)\hat{Q}_{\mathcal{S}_l}(t) \\ &\quad +(A_l(t)-\mu_l(t))(\hat{A}_{\mathcal{S}_l}(t)-\hat{\mu}_{\mathcal{S}_l}(t)) \\ &\quad -Q_l(t)(\hat{\mu}_{\mathcal{S}_l}(t)-\hat{A}_{\mathcal{S}_l}(t)) \\ &\quad -\hat{Q}_{\mathcal{S}_l}(t)(\mu_l(t)-A_l(t)) \end{aligned}$$

Thus, the 1-step unconditional Lyapunov drift is given by:

$$\begin{aligned} \Delta(t) &= \mathbb{E}\{B(t)\} \\ &\quad -\frac{1}{2}\sum_{l\in\mathcal{L}}\mathbb{E}\left\{Q_l(t)(\hat{\mu}_{\mathcal{S}_l}(t)-\hat{A}_{\mathcal{S}_l}(t))\right\} \\ &\quad -\frac{1}{2}\sum_{l\in\mathcal{L}}\mathbb{E}\left\{\hat{Q}_{\mathcal{S}_l}(t)(\mu_l(t)-A_l(t))\right\} \end{aligned} \tag{33}$$

where

$$B(t)\triangleq\frac{1}{2}\sum_{l\in\mathcal{L}}\left[(A_l(t)-\mu_l(t))(\hat{A}_{\mathcal{S}_l}(t)-\hat{\mu}_{\mathcal{S}_l}(t))\right] \tag{34}$$

We now use the following important structural property of the interference sets.

Lemma 4: (Sum Switching) For any function $f(l,\omega)$ (where $l\in\mathcal{L}$, $\omega\in\mathcal{L}$), we have:

$$\sum_{l\in\mathcal{L}}\sum_{\omega\in\mathcal{S}_l} f(l,\omega) = \sum_{\omega\in\mathcal{L}}\sum_{l\in\mathcal{S}_\omega} f(l,\omega) \ \Box$$

The above lemma follows directly from the *pairwise symmetry property* of the interference sets: For any two links $l,\omega\in\mathcal{L}$, we have that $\omega\in\mathcal{S}_l$ if and only if $l\in\mathcal{S}_\omega$, and hence $\{l\in\mathcal{L}\,|\,\omega\in\mathcal{S}_l\}=\{l\in\mathcal{L}\,|\,l\in\mathcal{S}_\omega\}$. Using this lemma we can re-write the final term in (33):

$$\begin{aligned} &\sum_{l\in\mathcal{L}}\hat{Q}_{\mathcal{S}_l}(t)(\mu_l(t)-A_l(t)) \\ &= \sum_{l\in\mathcal{L}}\sum_{\omega\in\mathcal{S}_l} Q_\omega(t)(\mu_l(t)-A_l(t)) && (35) \\ &= \sum_{\omega\in\mathcal{L}}\sum_{l\in\mathcal{S}_\omega} Q_\omega(t)(\mu_l(t)-A_l(t)) && (36) \\ &= \sum_{\omega\in\mathcal{L}} Q_\omega(t)(\hat{\mu}_{\mathcal{S}_\omega}(t)-\hat{A}_{\mathcal{S}_\omega}(t)) && (37) \\ &= \sum_{l\in\mathcal{L}} Q_l(t)(\hat{\mu}_{\mathcal{S}_l}(t)-\hat{A}_{\mathcal{S}_l}(t)) && (38) \end{aligned}$$

where (35) follows by the definition of $\hat{Q}_{\mathcal{S}_l}(t)$ given in (5), (36) follows by the Sum Switching Lemma (Lemma 4), and (38) follows by re-labeling the indices. Plugging the equality (38) directly into the drift expression (33) yields:

$$\Delta(t) = \mathbb{E}\{B(t)\} - \sum_{l\in\mathcal{L}}\mathbb{E}\left\{Q_l(t)(\hat{\mu}_{\mathcal{S}_l}(t)-\hat{A}_{\mathcal{S}_l}(t))\right\} \tag{39}$$

Further, we note that the expression for $B(t)$ in (34) is equivalent to that given in (9). This can be seen by using a sum-switching argument similar to (35)-(38) on the summation $\sum_{l\in\mathcal{L}}\hat{\mu}_{\mathcal{S}_l}(t)A_l(t)$, and by noting that $\mu_l(t)\hat{\mu}_{\mathcal{S}_l}(t)=\mu_l(t)$. The latter equality holds because if $\mu_l(t)=0$ we have $0=0$ which is trivially true, while if $\mu_l(t)=1$ then no other links within $\mathcal{S}_l$ can be active, and so $\hat{\mu}_{\mathcal{S}_l}(t)=1$.

We now use the fact that maximal scheduling is performed every timeslot. Specifically, we recall that any maximal scheduling algorithm satisfies (3) every timeslot. Using the definition of $\hat{\mu}_{\mathcal{S}_l}(t)$, we note that (3) is equivalent to:

$$Q_l(t)\hat{\mu}_{\mathcal{S}_l}(t)\geq Q_l(t) \quad \text{for all } l\in\mathcal{L} \text{ and all } t$$

Plugging the above inequality directly into (39) yields the expression (8) for $\Delta(t)$ under maximal scheduling.

APPENDIX B — PROOF OF THEOREM 2

Recall from (10) that:

$$\sum_{l\in\mathcal{L}}\overline{Q}_l \leq \frac{\overline{B}}{1-\rho^*} \tag{40}$$

where $B(t)$ is defined in (9). Note that we are using the fact that time averages exist. Using the fact that arrivals are i.i.d. over slots yields:

$$\mathbb{E}\{B(t)\} = \frac{1}{2}\sum_{l\in\mathcal{L}}\mathbb{E}\left\{A_l(t)\hat{A}_{\mathcal{S}_l}(t) - 2\mu_l(t)\hat{\lambda}_{\mathcal{S}_l} + \mu_l(t)\right\}$$

Because $\mu_l(t)\leq 1$, it follows that $\mathbb{E}\{B(t)\}\leq L/2+\frac{1}{2}\sum_{l\in\mathcal{L}}\mathbb{E}\left\{A_l(t)\hat{A}_{\mathcal{S}_l}(t)\right\}$, and hence from (40) we know that the network is strongly stable with a finite time average backlog (recall that second moments of arrivals are assumed to be finite). It follows that the time average input rate λ_l is equal to the time average transmission rate $\lim_{t\to\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\{\mu_l(\tau)\}$ for all links $l\in\mathcal{L}$ [17]. Therefore, we have:

$$\overline{B} = \frac{1}{2}\sum_{l\in\mathcal{L}}\left[\mathbb{E}\left\{A_l(t)\hat{A}_{\mathcal{S}_l}(t)\right\} - 2\lambda_l\hat{\lambda}_{\mathcal{S}_l} + \lambda_l\right]$$

This proves Theorem 2 part (a).

To prove part (b), note that if arrival processes are independent over different streams, then:

$$\mathbb{E}\left\{A_l(t)\hat{A}_{\mathcal{S}_l}(t)\right\} = \sigma_l^2 + \lambda_l\sum_{\omega\in\mathcal{S}_l}\lambda_\omega$$

where $\sigma_l^2\triangleq\mathbb{E}\left\{(A_l(t))^2\right\}-\lambda_l^2$ is the variance of the random variable $A_l(t)$. The result then follows by plugging into part (a) and using Little's Theorem.

APPENDIX C – PROOF OF THEOREM 3

To prove Theorem 3, we introduce an artificial delay in the final term of the drift expression (8) to decouple correlations between queue state and arrivals. This is similar to the T-slot technique of [11] [5], although, unlike [11][5], it allows a tight logarithmic delay result. To begin, fix an integer $T\geq 0$, and note that for $t\in\{0,1,2,\ldots,\}$ we have:

$$Q_l(t)\leq Q_l(t-T)+\sum_{v=0}^{T-1}A_l(t-T+v)$$

Using the above inequality in (8) yields

$$\begin{aligned} \Delta(t) &\leq \mathbb{E}\{B(t)+F_T(t)\} - \sum_{l\in\mathcal{L}}\mathbb{E}\{Q_l(t)\} \\ &\quad + \sum_{l\in\mathcal{L}}\mathbb{E}\left\{Q_l(t-T)\hat{A}_{\mathcal{S}_l}(t)\right\} \end{aligned} \tag{41}$$

where $F_T(t)$ is defined:

$$F_T(t) \triangleq \sum_{l\in\mathcal{L}} \mathbb{E}\left\{ \hat{A}_{\mathcal{S}_l}(t) \sum_{v=0}^{T-1} A_l(t-T+v) \right\}$$

We now use the $\hat{\rho}(T)$ function to modify the final term on the right hand side of (41):

$$\begin{aligned} &\mathbb{E}\left\{Q_l(t-T)\hat{A}_{\mathcal{S}_l}(t)\right\} \\ &= \mathbb{E}\left\{Q_l(t-T)\mathbb{E}\left\{\hat{A}_{\mathcal{S}_l}(t) \mid \boldsymbol{Q}(t-T)\right\}\right\} \\ &\leq \mathbb{E}\left\{Q_l(t-T)\left(\hat{\lambda}_{\mathcal{S}_l} + \sum_{\omega\in\mathcal{S}_l} \epsilon_\omega(T)\right)\right\} \\ &\leq \mathbb{E}\left\{Q_l(t-T)\right\}\hat{\rho}(T) \end{aligned} \quad (42)$$

Using (42) in (41), it follows that unconditional Lyapunov drift satisfies:

$$\begin{aligned} \Delta(t) \leq\; & \mathbb{E}\left\{B(t) + F_T(t)\right\} - \sum_{l\in\mathcal{L}} \mathbb{E}\left\{Q_l(t)\right\} \\ & + \hat{\rho}(T) \sum_{l\in\mathcal{L}} \mathbb{E}\left\{Q_l(t-T)\right\} \end{aligned} \quad (43)$$

Fixing the integer T and using the Lyapunov drift theorem (Theorem 1) in the drift expression (43) yields:

$$\begin{aligned} &\limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{l\in\mathcal{L}} \mathbb{E}\left\{Q_l(\tau)\right\} \leq \left(\overline{B} + \overline{F_T}\right) \\ &+ \hat{\rho}(T) \limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{l\in\mathcal{L}} \mathbb{E}\left\{Q_l(\tau-T)\right\} \end{aligned} \quad (44)$$

where

$$\overline{B} \triangleq \limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}\left\{B(\tau)\right\}, \; \overline{F_T} \triangleq \limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}\left\{F_T(\tau)\right\}$$

However, it is not difficult to see that:

$$\begin{aligned} &\limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{l\in\mathcal{L}} \mathbb{E}\left\{Q_l(\tau-T)\right\} \\ &= \limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{l\in\mathcal{L}} \mathbb{E}\left\{Q_l(\tau)\right\} \end{aligned} \quad (45)$$

Indeed, the equality (45) follows by noting the time-delayed version of the limit on the left hand side does not affect the overall time average, as any contribution to the sum over the extra T slots is finite and becomes negligible as $t \to \infty$. Therefore, because it is assumed that $\hat{\rho}(T) < 1$, the inequality (44) simplifies to:

$$\limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{l\in\mathcal{L}} \mathbb{E}\left\{Q_l(\tau)\right\} \leq \frac{\overline{B} + \overline{F_T}}{1 - \hat{\rho}(T)} \quad (46)$$

It is easy to show that:

$$\overline{B} \leq \frac{1}{2} \sum_{l\in\mathcal{L}} \left[\lambda_l + \mathbb{E}\left\{A_l(t)\hat{A}_{\mathcal{S}_l}(t)\right\}\right]$$

$$\overline{F_T} = \sum_{l\in\mathcal{L}} \mathbb{E}\left\{ \hat{A}_{\mathcal{S}_l}(t) \sum_{v=0}^{T-1} A_l(t-T+v) \right\}$$

It follows that $\overline{B} \leq \tilde{B}$ and $\overline{F_T} \leq \tilde{F}_T$, where $\tilde{B}$ and $\tilde{F}_T$ are defined in (14) and (15). Finally, we note that because the queueing dynamics are described by a countable state space, irreducible Markov chain, the $\limsup$ in the left hand side of (46) can be replaced by a regular limit, which proves part (a) of Theorem 3. Part (b) of Theorem 3 follows by computing $\tilde{B}$ and $\tilde{F}_T$ for the case when arrival processes $A_l(t)$ are independent of each other.

REFERENCES

[1] P. Chaporkar, K. Kar, and S. Sarkar. Throughput guarantees through maximal scheduling in wireless networks. *Proc. of 43rd Annual Allerton Conf. on Communication Control and Computing*, September 2005.

[2] 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.

[3] X. Lin and N. B. Shroff. The impact of imperfect scheduling on cross-layer rate control in wireless networks. *Proc. IEEE INFOCOM*, 2005.

[4] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936-1949, Dec. 1992.

[5] 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.

[6] D. Shah. Maximal matching scheduling is good enough. *Proc. IEEE Globecom*, Dec. 2003.

[7] 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.

[8] J. G. Dai and B. Prabhakar. The throughput of data switches with and without speedup. *Proc. IEEE INFOCOM*, 2000.

[9] A. Mekkittikul and N. McKeown. A practical scheduling algorithm to achieve 100% throughput in input-queued switches. *Proc. IEEE INFOCOM*, 1998.

[10] 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.

[11] M. J. Neely, E. Modiano, and C. E Rohrs. Dynamic power allocation and routing for time varying wireless networks. *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, pp. 89-103, January 2005.

[12] 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.

[13] 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.

[14] E. Leonardi, M. Mellia, F. Neri, and M. Ajmone Marsan. Bounds on average delays and queue size averages and variances in input-queued cell-based switches. *Proc. IEEE INFOCOM*, 2001.

[15] 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.

[16] S. Ross. *Stochastic Processes*. John Wiley & Sons, Inc., New York, 1996.

[17] M. J. Neely. Energy optimal control for time varying wireless networks. *IEEE Transactions on Information Theory*, vol. 52, no. 7, July 2006.