

Scheduling over Time Varying Channels with Hidden State Information

Matthew Johnston and Eytan Modiano
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA
Email: {mrj, modiano}@mit.edu

Abstract—We consider the problem of scheduling transmissions over a wireless downlink when channel state information (CSI) is not available to the transmitter. We assume channel states are time varying and evolve according to a Markov Chain. We show that using current QLI does not stabilize the system due to correlations between backlog and channel state. We show that the throughput optimal scheduling policy in this context must use delayed queue length information (QLI). We characterize the extent to which QLI must be delayed as a function of the channel state statistics.

I. INTRODUCTION

We consider the scheduling problem in a wireless downlink where channel state information (CSI) is unavailable at the base station, as in Figure 1. Packets arrive to the base station and are placed in queues to await transmission to their respective destinations. Due to wireless interference, only one transmission can be scheduled in each time slot. Furthermore, the channels to each user are independent, but evolve over time according to a Markov process. We seek a throughput optimal scheduling policy such that the queue lengths at the base station remain bounded.

Throughput optimal scheduling was pioneered by Tassioulas and Ephremides in [1], and has been studied in a variety of contexts. The optimal policies depend on the channel model and the information available to the transmitter, as summarized in Table I. If the channel state process is IID, and no CSI is available, then any work-conserving policy is throughput optimal; a commonly used throughput optimal policy in this scenario is one which schedules the longest queue. If the transmitter has current CSI and queue length information (QLI), the throughput optimal policy is one that transmits over the channel that maximizes the product of the channel rate and the queue length at the current time [1], [2]. If the CSI and QLI are delayed, Ying and Shakkottai show that the optimal policy schedules the channel with the largest product of the delayed QLI and the conditional expectation of the channel rate at the current time, given the delayed CSI [3]. If the CSI is not acquired until an acknowledgement is received from the transmission, then the throughput optimal policy

is to transmit over the channel that maximizes the product of the belief of the channel and the queue backlog [4].

While throughput optimal scheduling has been studied in a variety of contexts, to the best of our knowledge, there have been no results on throughput optimal scheduling when the controller has QLI but not CSI, and the channel process has memory. In fact, Tassioulas and Ephremides state that,

An interesting variation of the problem... is the case where the connectivity information is not available for the decision making and the server allocation can be based on queue lengths... The study of stability and optimal delay performance in [the] case of dependent connectivities are open problems for further investigation [1].

In this paper, we consider a scenario in which QLI is available to the transmitter, but not CSI. Typically, throughput optimal policies schedule the links with the largest backlogs, to ensure a fair server allocation while maintaining bounded backlogs[1], [2]. However, when the channel state process has memory, then long backlogs may be associated with poor channel qualities. Thus, giving priority to these channels is not an effective resource allocation, and the longest queue first policy does not stabilize the system.

We show that instead, the policy which schedules links based on significantly stale QLI is throughput optimal. While it has been known that the use of delayed QLI does not hurt throughput performance [5], in this scenario delayed QLI is *required* for stability. We characterize the degree by which QLI must be delayed for throughput optimality. Additionally, we provide simulation results to support the theoretical results of delayed QLI optimality, and show that using fresh QLI reduces the achievable throughput.

II. SYSTEM MODEL

Consider a system of M nodes, representing a wireless downlink, as in Figure 1. Packets arrive externally at the base station, and are destined for node i according to an i.i.d. Bernoulli arrival process $A_i(t)$ of rate λ_i . Packets are stored in a separate queue at the base station, based on the destination node, to await transmission. Let $Q_i(t)$ be the packet backlog corresponding to node i at time t . Due to

This work was supported by NSF Grant CNS-1217048 and ONR Grant N00014-12-1-0064.




Fig. 1: Wireless Downlink

Model	No CSI	Delayed CSI	Full CSI
IID Channels	$\max_i Q_i(t) \mathbb{E}[S_i(t)]$	$\max_i Q_i(t) \mathbb{E}[S_i(t)]$	$\max_i Q_i(t) S_i(t)$
Markov Channels	*This work*	$\max_i Q_i(t - \tau) \mathbb{E}[S_i(t) S_i(t - \tau)]$	$\max_i Q_i(t) S_i(t)$

TABLE I: Throughput optimal policies for different system models. Column corresponds to a different amount of information at the controller. Rows corresponds to the memory in the channel. $S_i(t)$ is the channel state and $Q_i(t)$ is the backlog.

wireless interference, the base station is able to transmit to only one node at a time, although this model can easily be extended to allow for multiple transmissions per slot.

Each node is connected to the base station through an independent, time-varying channel. Let $S_i(t) \in \{\text{OFF}, \text{ON}\}$ be the channel state of the channel at node i at time t . Assume the channel states evolve over time according to a Markov chain, shown in Figure 2. If a packet for node i is scheduled for transmission, and $S_i(t) = \text{ON}$, then the packet is successfully transmitted, assuming there are packets awaiting transmission. On the other hand, if the channel at node i is OFF, then the transmission fails, and the packet remains in the system. Let $P_{S_i}(1)$ and $P_{S_i}(0)$ be the steady state probability of channel i being ON or OFF respectively.

The base station has access to the history of queue lengths for each node i , but no information regarding channel states. Therefore, the base station makes a transmission decision based on QLI, but not CSI¹. Let Π be the set of transmission policies which do not use CSI.

Definition A queue with backlog $Q_i(t)$ is *stable* under policy \mathcal{P} if

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{t=0}^{n-1} \mathbb{E}[Q_i(t)] < \infty \quad (1)$$

The complete network is stable if all queues are stable.

Definition The *throughput region* Λ is the closure of the set of all rate vectors λ that can be stably supported over the network by any policy $\mathcal{P} \in \Pi$.

Definition A policy is said to be *throughput optimal* if it stabilizes the system for any arrival rate $\lambda \in \Lambda$.

In this work, we characterize the throughput region of the system above, and propose a throughput optimal scheduling policy using delayed QLI.

¹We assume packet acknowledgements occur at a separate layer, and cannot be used to predict the channel state.




Fig. 2: Markov Chain describing the channel state evolution.

III. THROUGHPUT REGION

The throughput region is characterized by the following linear program (LP).

Maximize:

$$\epsilon$$

Subject To:

$$\begin{aligned} \lambda_i + \epsilon &\leq \alpha_i P_{S_i}(1) & \forall i \in \{1, \dots, M\} \\ \sum_{i=1}^M \alpha_i &\leq 1, \alpha_i \geq 0 & \forall i \in 1, \dots, M \end{aligned} \quad (2)$$

In the above LP, α_i represents the fraction of time the base station schedules node i for transmission. To maintain stability, the arrival rate to each queue must be less than the service rate at that queue, which is a function of α_i and the statistics of the channel. Thus, the throughput region, Λ , is the set of all non-negative arrival rate vectors λ such that there exists a feasible solution to (2) for which $\epsilon \geq 0$.

Theorem 1 (Throughput Region). *For any non-negative arrival rate vector λ , the system can be stabilized by some policy $\mathcal{P} \in \Pi$ if and only if $\lambda \in \Lambda$.*

Necessity is shown in Lemma 1, and sufficiency is shown in Theorem 2 by proposing a throughput optimal scheduling policy, and proving that for all $\lambda \in \Lambda$, that policy stabilizes the system.

Lemma 1. *Suppose there exists a scheduling policy $\mathcal{P} \in \Pi$ that stabilizes the system without using CSI, then there exists an α_i such that (2) has a solution with $\epsilon \geq 0$.*

Lemma 1 shows that for all $\lambda \in \Lambda$, there exists a stationary policy $\text{STAT} \in \Pi$ that stabilizes the system, by scheduling link i with probability α_i . The value of α_i that stabilizes the network depends on the arrival rates, which are not available to the controller. In the following section, we develop a scheduling policy based on delayed QLI, that stabilizes the system without requiring knowledge of the arrival rates or channel statistics.

IV. DYNAMIC QLI-BASED SCHEDULING POLICY

Due to the ergodicity of the finite-state Markov chain in Figure 2, for any $\epsilon > 0$, there exists a τ_Q such that the probability of the channel state conditioned on the channel state τ_Q slots in the past is within $\frac{\epsilon}{2}$ of the steady state probability of the Markov chain.

$$\left| \mathbf{P}(\mathbf{S}(t) = \mathbf{s} | \mathbf{S}(t - \tau_Q(\epsilon))) - \mathbf{P}(\mathbf{S}(t) = \mathbf{s}) \right| \leq \frac{\epsilon}{2} \quad (3)$$




Fig. 3: Symmetric arrival rate vs. average queue backlog for a 4-queue system. Transition probabilities satisfy $p = q = 0.01$

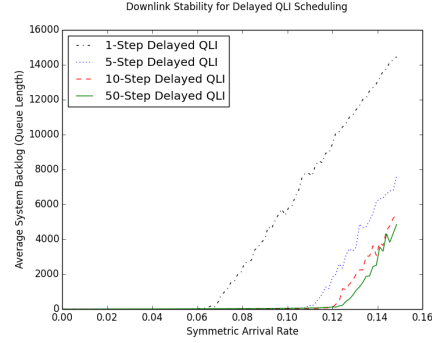


Fig. 4: Symmetric arrival rate vs. average queue backlog for a 4-queue system. Transition probabilities satisfy $p = q = 0.1$

Note that $\tau_Q(\epsilon)$ is related to the mixing time of the Markov chain. In general, the Markov chain approaches steady state exponentially fast, at a rate of $p + q$ [6].

Theorem 2 proposes the Delayed Longest Queue (DLQ) scheduling policy, which stabilizes the network whenever the input rate vector is interior to the capacity region Λ . Note, this proves sufficiency in Theorem 1.

Theorem 2. Consider the Delayed Longest Queue (DLQ) scheduling policy, which at time t schedules the following channel for transmission:

$$i^* = \underset{i}{\operatorname{argmax}} Q_i(t - \tau_Q(\epsilon)), \quad (4)$$

where $\tau_Q(\epsilon)$ is defined in (3). For any arrival rate λ , and $\epsilon > 0$ satisfying $\lambda + \epsilon \mathbf{1} \in \Lambda$, this policy stabilizes the system.

The proof of Theorem 2 is in the Appendix. The DLQ policy transmits a packet from the longest queue using *delayed QLI*. If fresher QLI is available, it cannot be used by the DLQ policy to stabilize the system. This is because at time t , the queue with the largest backlog $Q_i(t)$ is also likely to have an OFF channel. On the other hand, if sufficiently delayed QLI is used in the DLQ policy, then the QLI is independent of the current channel state, because the state process reaches its steady-state distribution over the τ_Q slots that the QLI is delayed. Under DLQ, the base station schedules queues for which the backlog is long, without favoring OFF channels.

The required delay on the QLI depends on the mixing time of the channel state process. As $p + q$ approaches 1, the Markov process approaches an IID process, and current QLI can be used. However, using further delayed QLI doesn't affect the overall throughput region. The drawback of using delayed QLI is increased packet delays. Therefore, if no CSI is available to the base station, the optimal policy must trade off between throughput and delay.

V. SIMULATION RESULTS AND CONCLUSIONS

We simulate a system of four queues, and apply the DLQ policy for different delays to QLI (τ_Q). We plot the average queue backlog over 100,000 time-slots for

different symmetric arrival rates². For small arrival rates, the average queue length remains small. As the arrival rate increases, the backlog slowly increases until a certain point, after which the backlog greatly increases. This point represents the boundary of the throughput region, and for arrival rates outside of this region, the system of queues cannot be stabilized.

For a system of four queues with symmetric channel transition probabilities $p = q$, the boundary of the stability region on the symmetric arrival rate line is given by $\frac{1}{2} \cdot \frac{1}{4} = 0.125$, since each node transmits equally often, and each channel is ON with probability $\frac{1}{2}$. Therefore, under the throughput optimal policy, the queue lengths should remain bounded for arrival rates $\lambda < 0.125$.

Figures 3 and 4 show the results for transition probabilities $p = q = 0.01$ and $p = q = 0.1$ respectively. As shown in Figure 3, when the QLI is insufficiently delayed, the system becomes unstable before the boundary of the stability region (0.125). For $\tau_Q = 1$, the system becomes unstable at approximately $\lambda = 0.03$, representing a 75% reduction in the stability region. As τ_Q increases, the maximum arrival rate supportable by the DLQ policy increases. At $\tau_Q = 150$, it appears that the system becomes stable for all arrival rates within the stability region.

Similar results are shown in Figure 4 for a channel with less memory. In this case, the attainable throughput of the DLQ policy is less sensitive to the magnitude of the delays in QLI. The simulation results suggest that $\tau_Q = 100$ is sufficient to achieve the full throughput region in this case.

In summary, using current QLI does not stabilize the system when the channel state process has memory, and significantly stale QLI, based on the amount of memory in the channel, must be used for throughput optimality.

VI. APPENDIX

Proof of Theorem 2: Let $\tau_Q = \tau_Q(\epsilon)$, where the dependence on ϵ is clear. Let $\mathbf{Y}(t)$ be the history of queue-lengths in the system up to time t , i.e. $\mathbf{Y}(t) = \{\mathbf{Q}(0), \dots, \mathbf{Q}(t)\}$. The vector $\mathbf{Y}(t)$ forms a Markov pro-

²A symmetric arrival rate implies that each node sees the same arrival rate.

cess. Define the following quadratic Lyapunov function:

$$L(\mathbf{Q}(t)) = \frac{1}{2} \sum_{i=1}^M Q_i^2(t). \quad (5)$$

The T -step Lyapunov drift is computed as

$$\Delta_T(\mathbf{Y}(t)) = \mathbb{E} \left[L(\mathbf{Q}(t+T)) - L(\mathbf{Q}(t)) \middle| \mathbf{Y}(t) \right] \quad (6)$$

We show that under the DLQ policy, the T -step Lyapunov drift is negative for large backlogs, implying the stability of the system under the DLQ policy for all arrival rates within Λ , which follows from the Foster-Lyapunov criteria [7]. We bound the Lyapunov drift by combining (7), (5) and (6), and showing for large queue lengths, this upper bound is negative.

Consider the DLQ scheduling policy. Let $D_i(t)$ be the departure process of queue i , such that $D_i(t) = 1$ if there is a departure from queue i at time t under policy DLQ. Consider the evolution of the queues over T time slots.

$$Q_i(t+T) \leq \left(Q_i(t) - \sum_{k=0}^{T-1} D_i(t+k) \right)^+ + \sum_{k=0}^{T-1} A_i(t+k) \quad (7)$$

Equation (7) is an inequality rather than an equality due to the assumption that the departures are taken from the backlog at the beginning of the T -slot period, and the arrivals occur at the end of the T slots. The Lyapunov drift in (6) is bounded as follows:

$$\begin{aligned} \Delta_T(\mathbf{Y}(t)) &\leq B + \mathbb{E} \left[\sum_{i=1}^M Q_i(t) \left(\sum_{k=0}^{T-1} A_i(t+k) - \sum_{k=0}^{T-1} D_i(t+k) \right) \middle| \mathbf{Y}(t) \right] \end{aligned} \quad (8)$$

where B is a finite constant, which exists due to the boundedness of the second moment of the arrival process.

The difference between queue lengths at any two times t and s is bounded using the following inequality:

$$Q_i(t) - Q_i(s) \leq |t - s|, \quad (9)$$

which holds assuming that an arrival occurs in each slot, and no departures occur, or vice versa. This inequality establishes a relationship between current queue lengths and delayed queue lengths.

$$Q_i(t) \leq Q_i(t+k-\tau_Q) + |k-\tau_Q| \quad (10)$$

$$Q_i(t) \geq Q_i(t+k-\tau_Q) - |k-\tau_Q| \quad (11)$$

The inequalities in (10) and (11) are used in (8) to upper bound the Lyapunov drift in terms of the delayed QLI.

$$\begin{aligned} \Delta_T(\mathbf{Y}(t)) &\leq B + \mathbb{E} \left[\sum_{k=0}^{T-1} \sum_{i=1}^M (Q_i(t+k-\tau_Q) + |k-\tau_Q|) A_i(t+k) \right. \\ &\quad \left. - \sum_{k=0}^{T-1} \sum_{i=1}^M (Q_i(t+k-\tau_Q) - |k-\tau_Q|) D_i(t+k) \middle| \mathbf{Y}(t) \right] \end{aligned} \quad (12)$$

$$\leq B' + \mathbb{E} \left[\sum_{k=0}^{T-1} \sum_{i=1}^M Q_i(t+k-\tau_Q) (\lambda_i - D_i(t+k)) \middle| \mathbf{Y}(t) \right] \quad (13)$$

Equation (13) follows from upper bounding the per-slot arrival and departure rate each by 1, defining $B' = B + 2MT^2$, and using $\mathbb{E}[A_i(t+k)] = \lambda_i$. To bound (13), we require the channel state at slot $t+k$ to be independent from $\mathbf{Y}(t)$, which only holds in slots where k is sufficiently large. Thus, we break the summation in (13) into two parts: a smaller number of slots for which k is small, and a larger number of slots for which k is large. A trivially conservative bound is used for $k < \tau_Q$, but the frame size is chosen to ensure the first τ_Q slots is a small fraction of the overall T slots.

$$\begin{aligned} \Delta_T(\mathbf{Y}(t)) &\leq B' + \sum_{k=0}^{T-1} \mathbb{E} \left[\sum_{i=1}^M Q_i(t+k-\tau_Q) (\lambda_i - D_i(t+k)) \middle| \mathbf{Y}(t) \right] \\ &\leq B' + \sum_{k=0}^{\tau_Q-1} \mathbb{E} \left[\sum_{i=1}^M Q_i(t+k-\tau_Q) (\lambda_i - D_i(t+k)) \middle| \mathbf{Y}(t) \right] \\ &\quad + \sum_{k=\tau_Q}^{T-1} \mathbb{E} \left[\sum_{i=1}^M Q_i(t+k-\tau_Q) (\lambda_i - D_i(t+k)) \middle| \mathbf{Y}(t) \right] \end{aligned} \quad (14)$$

For values of $k < \tau_Q$, the upper bound follows by trivially upper bounding the arrival rate by 1 and lower bounding the departures by 0 in each slot.

$$\begin{aligned} \sum_{k=0}^{\tau_Q-1} \mathbb{E} \left[\sum_{i=1}^M Q_i(t+k-\tau_Q) (\lambda_i - D_i(t+k)) \middle| \mathbf{Y}(t) \right] \\ \leq \sum_{k=0}^{\tau_Q-1} \sum_{i=1}^M \mathbb{E} [Q_i(t+k-\tau_Q) | \mathbf{Y}(t)] \end{aligned} \quad (15)$$

$$\leq \sum_{k=0}^{\tau_Q-1} \sum_{i=1}^M Q_i(t-\tau_Q) + \sum_{k=0}^{\tau_Q-1} \sum_{i=1}^M k \quad (16)$$

$$\leq \tau_Q \sum_{i=1}^M Q_i(t-\tau_Q) + \frac{1}{2} (\tau_Q)^2 M \quad (17)$$

where (16) follows from (9).

Now consider slots for which $k \geq \tau_Q$. The last term on the right hand side of (14) can be rewritten by conditioning on the delayed QLI at the current slot $t+k$ and using the law of iterated expectations. For exposition, define $\hat{Q}_i^{k-\tau_Q} = Q_i(t+k-\tau_Q)$.

$$\begin{aligned} \mathbb{E} \left[\sum_{k=\tau_Q}^{T-1} \sum_{i=1}^M \hat{Q}_i^{k-\tau_Q} (\lambda_i - D_i(t+k)) \middle| \mathbf{Y}(t) \right] \\ = \mathbb{E} \left[\sum_{k=\tau_Q}^{T-1} \mathbb{E} \left[\sum_{i=1}^M \hat{Q}_i^{k-\tau_Q} (\lambda_i - D_i(t+k)) \middle| \hat{\mathbf{Q}}^{k-\tau_Q} \right] \middle| \mathbf{Y}(t) \right] \end{aligned} \quad (18)$$

Let ϕ_i be a binary variable denoting whether queue i is scheduled under the DLQ policy as a function of the delayed QLI. For these time-slots, we evaluate the expected departure rate, and compare it to the departure rate of the STAT policy in Lemma 1, which we know stabilizes the system. The expected departure rate is expanded as

$$\mathbb{E} \left[\sum_{i=1}^M \hat{Q}_i^{k-\tau_Q} D_i(t+k) \middle| \hat{\mathbf{Q}}^{k-\tau_Q}, \mathbf{Y}(t) \right] \quad (19)$$

$$= \sum_{i=1}^M \hat{Q}_i^{k-\tau_Q} \mathbb{E} [\phi_i(\hat{\mathbf{Q}}^{k-\tau_Q}) S_i(t+k) | \hat{\mathbf{Q}}^{k-\tau_Q}, \mathbf{Y}(t)] \quad (20)$$

$$= \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \phi_i(\widehat{\mathbf{Q}}_i^{k-\tau_Q}) \mathbb{E}[S_i(t+k) | \widehat{\mathbf{Q}}_i^{k-\tau_Q}, \mathbf{Y}(t)]. \quad (21)$$

Equation (21) follows since the scheduling under DLQ is completely determined by the delayed QLI.

Note that the throughput optimal policy maximizes the expression in (21); however, the expectation cannot be computed because it requires knowledge of the conditional distribution of the channel state sequence given QLI, which requires knowledge of the arrival rates to compute. However, when QLI is sufficiently delayed, the bound in (3) can be used to remove the conditioning on QLI.

$$\begin{aligned} & \mathbf{P}(S_i(t+k) = s | \widehat{\mathbf{Q}}_i^{k-\tau_Q}, \mathbf{Y}(t)) \\ &= \sum_{s' \in \mathcal{S}} \mathbf{P}(S_i(t+k-\tau_Q) = s' | \widehat{\mathbf{Q}}_i^{k-\tau_Q}, \mathbf{Y}(t)) \\ & \quad \cdot \mathbf{P}(S_i(t+k) = s | S_i(t+k-\tau_Q) = s') \quad (22) \\ & \geq \mathbf{P}(S_i(t+k) = s) - \frac{\epsilon}{2} \quad (23) \end{aligned}$$

Equation (22) follows from the law of total probability, and the Markov property of the channel state. Equation (23) holds from the definition of τ_Q in (3), which implies the conditional state distribution is within $\frac{\epsilon}{2}$ of the stationary distribution. The expression in (21) can now be bounded in terms of an unconditional expectation.

$$\begin{aligned} & \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \phi_i(\widehat{\mathbf{Q}}_i^{k-\tau_Q}) \mathbb{E}[S_i(t+k) | \widehat{\mathbf{Q}}_i^{k-\tau_Q}, \mathbf{Y}(t)] \\ &= \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \phi_i(\widehat{\mathbf{Q}}_i^{k-\tau_Q}) \mathbf{P}(S_i(t+k) = 1 | \widehat{\mathbf{Q}}_i^{k-\tau_Q}, \mathbf{Y}(t)) \quad (24) \end{aligned}$$

$$\geq \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \phi_i(\widehat{\mathbf{Q}}_i^{k-\tau_Q}) \mathbf{P}_{S_i(t+k)}(1) - \frac{\epsilon}{2} \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \quad (25)$$

$$= P_S(1) \max_i \widehat{Q}_i^{k-\tau_Q} - \frac{\epsilon}{2} \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q}. \quad (26)$$

Equation (24) follows from the distribution of channel state. The inequality in (25) follows from applying (23) and upper bounding $\phi_i(\mathbf{Q}) \leq 1$. Equation (26) follows from applying the DLQ policy. Combining equation (26) with equation (18) yields

$$\sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \lambda_i - P_S(1) \max_i \widehat{Q}_i^{k-\tau_Q} + \frac{\epsilon}{2} \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \quad (27)$$

Now, we reintroduce the stationary policy of Lemma 1 to complete the bound. Recall that for any $\lambda \in \Lambda$, there exists a stationary policy which schedules node i for transmission with probability α_i , and satisfies

$$\lambda_i + \epsilon \leq \alpha_i P_S(1) \quad \forall i \in \{1, \dots, M\}. \quad (28)$$

Note that the ϵ in the theorem statement and in (28) are designed to be equal. The expression in (27) is bounded by adding and subtracting identical terms corresponding to the stationary policy.

$$\sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} (\lambda_i - \alpha_i P_S(1)) + \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \alpha_i P_S(1)$$

$$- P_S(1) \max_i \widehat{Q}_i^{k-\tau_Q} + \frac{\epsilon}{2} \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \quad (29)$$

$$\begin{aligned} & \leq -\epsilon \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} + \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \alpha_i P_S(1) \\ & - P_S(1) \max_i \widehat{Q}_i^{k-\tau_Q} + \frac{\epsilon}{2} \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \quad (30) \end{aligned}$$

$$\leq -\frac{\epsilon}{2} \sum_{i=1}^M \widehat{Q}_i^{k-\tau_Q} \quad (31)$$

$$\leq -\frac{\epsilon}{2} \sum_{i=1}^M Q_i(t-\tau_Q) + \frac{\epsilon}{2} M k \quad (32)$$

Equation (30) follows from (28), and equation (31) follows from the fact that since $\sum_i \alpha_i \leq 1$, the weighted sum of queue lengths is maximized by placing all the weight at the largest queue length. Equation (32) follows from (9).

The upper bound in (32) for slots $k \geq \tau_Q$ is combined with (17) for $k < \tau_Q$ to bound the drift in (14).

$$\begin{aligned} \Delta_T(\mathbf{Y}(t)) & \leq B' + \tau_Q \sum_{i=1}^M Q_i(t-\tau_Q) + \frac{1}{2} (\tau_Q)^2 M \\ & \quad + \sum_{k=\tau_Q}^{T-1} \mathbb{E} \left[-\frac{\epsilon}{2} \sum_{i=1}^M Q_i(t-\tau_Q) + \frac{\epsilon}{2} M k \mid \mathbf{Y}(t) \right] \quad (33) \end{aligned}$$

$$\begin{aligned} & \leq B' + \frac{1}{2} \tau_Q^2 M (1 - \frac{\epsilon}{2}) + \frac{\epsilon}{4} M T^2 \\ & \quad + \tau_Q \sum_{i=1}^M Q_i(t-\tau_Q) - (T-\tau_Q) \frac{\epsilon}{2} \sum_{i=1}^M Q_i(t-\tau_Q) \quad (34) \end{aligned}$$

Thus, for any $\xi > 0$, and T satisfying

$$T \geq \frac{2(1 + \frac{\epsilon}{2})\tau_Q + 2\xi}{\epsilon}, \quad (35)$$

there exists a positive constant K such that

$$\Delta_T(\mathbf{Y}(t)) \leq K - \xi \sum_{i=1}^M Q_i(t-\tau_Q). \quad (36)$$

Thus, for large enough queue backlogs, the T -slot Lyapunov drift is negative, and from [2] it follows that the overall system is stable. \blacksquare

REFERENCES

- [1] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *Information Theory, IEEE Transactions on*, vol. 39, no. 2, pp. 466–478, 1993.
- [2] M. Neely, E. Modiano, and C. 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, 2005.
- [3] L. Ying and S. Shakkottai, "On throughput optimality with delayed network-state information," *Information Theory, IEEE Transactions on*, 2011.
- [4] K. Jagannathan, S. Mannor, I. Menache, and E. Modiano, "A state action frequency approach to throughput maximization over uncertain wireless channels," *Internet Mathematics*, vol. 9, no. 2-3, pp. 136–160, 2013.
- [5] K. Kar, X. Luo, and S. Sarkar, "Throughput-optimal scheduling in multichannel access point networks under infrequent channel measurements," *Wireless Communications, IEEE Transactions on*, 2008.
- [6] R. G. Gallager, *Stochastic processes: theory for applications*. Cambridge University Press, 2013.
- [7] M. J. Neely, "Dynamic power allocation and routing for satellite and wireless networks with time varying channels," Ph.D. dissertation, Massachusetts Institute of Technology, 2003.