

Controller Placement for Maximum Throughput Under Delayed CSI

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

Abstract—The performance of wireless scheduling algorithms directly depends on the availability and accuracy of network state information at the scheduler. As channel state updates must propagate across the network, they are delayed as they arrive at the controller. The location of the controller directly affects the attainable throughput, as it dictates the delays with which information is obtained to make scheduling decisions. In this paper, we analyze the optimal controller placement over a network in which CSI delays are proportional to distance. We propose a dynamic controller placement framework, in which the controller is relocated using delayed queue length information at each node, and scheduling is done using delayed QLI and CSI. We characterize the throughput region under such policies, and find a policy which stabilizes the system for all arrival rates within the throughput region.

I. INTRODUCTION

In order to schedule transmissions to achieve maximum throughput, a centralized scheduler must opportunistically make decisions based on the current state of the time-varying channels [1]. The channel state of a link can be measured by its adjacent nodes, who forward this channel state information (CSI) across the network to the scheduler. Due to the transmission and propagation delays over wireless links, it may take several time-slots for the scheduler to collect CSI throughout the network, and in that time the network state may change.

There has been extensive work on wireless scheduling [1]–[3], in which centralized approaches are used to control the network. Centralized scheduling, in which a central entity makes a scheduling decision for the entire network, yields high theoretical performance, since the central entity uses current CSI throughout the network to compute a globally optimal schedule. However, maintaining current CSI is impractical, due to the latency in acquiring CSI throughout the network. In practice, the available CSI for centralized scheduling is a delayed view of the network state. Furthermore, the delay in CSI is proportional to the distance of each link to the controller, since CSI updates must traverse the network.

Several works have studied scheduling under delayed CSI. In [4], the authors consider a system in which CSI and QLI updates are only reported once every T time-slots, but

the transmitter makes a scheduling decision every slot, using delayed information. They show that delays in the CSI reduce the achievable throughput region, while delays in QLI do not adversely affect throughput. In [5], Ying and Shakkottai study throughput optimal scheduling and routing with delayed CSI and QLI. They show that the throughput optimal policy activates a max-weight schedule, where the weight on each link is given by the product of the delayed queue length and the conditional expected channel state given the delayed CSI. This work is extended in [6], where the authors account for the uncertainty in the state of the network topology as well. Lastly, the work in [7] characterizes the impact of delayed CSI as a function of the network topology, when delays are proportional to distance.

In a centralized scheduling scheme, a node is assigned the role of a controller, and collects CSI from the rest of the network. Then, the controller uses this CSI to select a subset of the nodes to transmit in each slot, in order to maximize throughput while avoiding interference between neighboring links [7]. However, the controller must use delayed CSI to schedule transmission, and the CSI delays are proportional to the distance of the links from the controller. Since delay directly impacts the throughput performance of the scheduling algorithm, the placement of the controller affects network performance.

This paper studies the impact of the controller placement on network performance. To begin, we analyze the static controller placement problem, in which the controller placement is computed off-line, and remains fixed over time. For any static controller placement, links near the controller achieve a high throughput, while links further away from the controller attain a lower throughput, due to the increased delay in the CSI. In order to mitigate this imbalance, we propose a dynamic controller placement framework, in which the location of the controller is changed over time. This allows for the controller to be moved to a region of the network with high backlogs to increase throughput to this region and provide stability. Since at any time, each node has a different view of the network state due to the distance-based CSI delays, the controller placement algorithm must only depend on information shared by all nodes, such that no additional communication overhead is re-

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




Figure 1: Markov Chain describing the channel state evolution of each independent channel.

quired. We propose a queue-length based controller placement algorithm, and show that this algorithm offers an increased throughput over a static placement. We prove this algorithm is throughput optimal over all controller placement policies which do not use CSI.

II. STATIC CONTROLLER PLACEMENT

In this section, we consider an off-line controller placement, such that the controller remains fixed over time. We show that the optimal controller placement depends on the network topology as well as the channel transition probabilities.

A. System Model

Consider a network $G(\mathcal{N}, \mathcal{L})$ consisting of a set of nodes \mathcal{N} and links \mathcal{L} . Each link is associated with an independent, time-varying channel. Let $S_l(t) \in \{\text{OFF}, \text{ON}\}$ be the channel state of the channel at link l at time t . Assume the channel state evolves over time according to the Markov chain in Figure 1. One of the nodes is assigned to be the controller off-line, and in each time slot, activates a subset of links for transmission. Assume a primary interference constraint in which a link activation is feasible if the activation is a matching, i.e. no two neighboring links are activated. If link l is activated, and $S_l(t) = \text{ON}$, then a packet is successfully transmitted at that time slot¹. On the other hand, if the channel at link l is OFF, then the transmission fails. The objective of the controller is to activate the set of links resulting in maximum expected sum-rate throughput.

In order to determine the correct subset of links to activate, the controller obtains CSI from each link in the network, and uses the CSI to compute a feasible link activation with maximum expected throughput. Due to the physical distance between network nodes, and the propagation delay across each link, the CSI updates received at the controller are delayed relative to the distance between each link and the controller. In particular, let $d_i(l)$ be the distance in hops between node i and link l ². At time t , each node i has delayed CSI pertaining to link l from time-slot $t - d_j(l)$. In other words, node i has state information $S_i(t - d_i(l))$ for link l .

B. Controller Placement Example

To begin, consider the example topology in Figure 2, and compare the expected throughput attainable by placing the




Figure 2: Barbell Network

controller at node A, node B, or node C. Placing the controller at node A yields the same expected throughput as placing the controller at node C, due to the symmetry of the network. Consider a generalization of the network in Figure 2, where A and C have degree $k + 1$. For simplicity, assume a symmetric Markov state in Figure 1, i.e. $p = q$. Let $\gamma = (\frac{1}{2})^k$, the probability that k links are OFF. Placing the controller at node B results in an expected throughput of

$$\text{thpt}_B = \frac{1}{4}2((1 - \gamma)p_{11}^1 + \gamma p_{01}^1) + \frac{1}{2}(1 + (1 - \gamma)p_{11}^1 + \gamma p_{01}^1) + \frac{1}{4}(1 + (1 - \gamma^2)p_{11}^1 + \gamma^2 p_{01}^1). \quad (1)$$

The above expression follows from conditioning on the state of the two adjacent links to B. The first term corresponds to the expected throughput when both links are OFF, the second corresponds to the case when one is ON and the other is OFF, and the last term corresponds to both links being ON. Similarly, the expected throughput from a controller at node A is derived by conditioning on the state of the $k + 1$ links adjacent to node A.

$$\text{thpt}_A = (1 - \gamma) \left[1 + \frac{1}{2}p_{11}^1 + \frac{1}{2}((1 - \gamma)p_{11}^2 + \gamma p_{01}^2) \right] + \gamma \left[\frac{1}{2}(1 + (1 - \gamma)p_{11}^2 + \gamma p_{01}^2) + \frac{1}{2}(\frac{1}{2}p_{11}^1 + \frac{1}{2}((1 - \gamma)p_{11}^2 + \gamma p_{01}^2)) \right] \quad (2)$$

Consider the throughput obtained from a controller at A and B in the limit as k grows to infinity, in which case $\gamma = 0$.

$$\lim_{k \rightarrow \infty} \text{thpt}_A = 1 + \frac{1}{2}(1 - p) + \frac{1}{2}p_{11}^2 \quad (3)$$

$$\lim_{k \rightarrow \infty} \text{thpt}_B = \frac{3}{4} + \frac{5}{4}(1 - p). \quad (4)$$

For $p \leq \frac{1}{4}$, it is optimal to place the controller at node B, and for $p \geq \frac{1}{4}$ it is optimal to place the controller at either node A or C. This example highlights some important properties of the controller placement problem. In particular, it is clear the optimal placement depends on the channel transition probabilities. When p is small, it is advantageous to place the controller to minimize the CSI delay throughout the network. On the other hand, when p is close to $\frac{1}{2}$ the CSI is no longer useful, and it is better to maximize the amount of local CSI at the controller.

C. Optimal Controller Placement

From the previous example, it is clear that the throughput-maximizing controller placement is a function of the channel

¹In this section, assume that each node has an infinite backlog of packets to transmit.

²By convention, node i is a distance of 0 hops from its adjacent links.




Figure 3: Wireless Downlink

state transition probabilities p and q , as well as the network topology. Let \mathcal{M} be the set of matchings in the network, i.e., $\forall M \in \mathcal{M}$, M is a set of links which can be scheduled simultaneously without interfering with one another. Under a throughput maximization objective, the controller schedules the matching that maximizes expected sum-rate throughput with respect to the CSI delays at that node. Consequently, the controller placement can be optimized as follows.

$$c = \arg \max_r \mathbb{E}_S \left[\max_{M \in \mathcal{M}} \sum_{l \in M} \mathbb{E}[S_l(t) | S_l(t - d_r(l)) = S_l] \right] \quad (5)$$

$$= \arg \max_r \mathbb{E}_S \left[\max_{M \in \mathcal{M}} \sum_{l \in M} p_{S_l, 1}^{d_r(l)} \right] \quad (6)$$

where $p_{i,j}^k$ is the k -step transition probability of the Markov chain in Figure 1. Equation (6) follows since the channel state satisfies $S_l(t) \in \{0, 1\}$. Computing a maximum matching requires solving an integer linear program (ILP) and it is known to be solvable in $\mathcal{O}(|\mathcal{L}|^3)$ -time [8]. However, computing the optimal controller position in (6) requires computing the expectation of the maximum matching, which involves solving the ILP for every state sequence $\mathbf{S}(t) \in \{0, 1\}^{|\mathcal{L}|}$.

III. DYNAMIC CONTROLLER PLACEMENT

For a fixed controller, the links physically close to the controller see a higher throughput than those far from the controller due to the delay in CSI. By relocating the controller, the throughput in different regions of the network can be improved. In this section, we consider policies which recompute the controller location dynamically in order to balance the throughput throughout the network.

A. System Model

Consider a system of M nodes operating under an interference constraint such that only one node can transmit at any time, as in Figure 3. Packets arrive externally to each node i according to an i.i.d. Bernoulli arrival process $A_i(t)$ of rate λ_i , and are stored in a queue at that node to await transmission. Let $Q_i(t)$ be the packet backlog of node i at time t . Each node has access to an independent time-varying ON/OFF channel as in Figure 1. If a node is scheduled for transmission, has a packet

to transmit, and has an ON channel, then a packet departs the system from node i .

The above network model applies directly to a wireless downlink or uplink; however, it can easily be extended to a network setting as follows. First, instead of the controller selecting one node to transmit, a set of non-interfering nodes is scheduled to transmit. The results of this section are extended by changing the scheduling optimization to be over all matchings in the network, rather than all individual nodes. Second, in a network, packets are required to traverse multiple hops on route to their destinations. This extension requires a modification to the throughput optimal policy of Theorem 2, analogous to the approach taken in [2].

In addition to each node i having delayed CSI pertaining to node j from the $d_j(i)$ time-slots in the past, it has delayed queue length information (QLI) as well. In other words, node i has delayed CSI $S_i(t - d_i(j))$ and delayed QLI, $Q_i(t - d_i(j))$ for each other node j . Let $\mathbf{S}(t - d_r)$ represent the vector of delayed CSI pertaining to controller r , i.e. $\mathbf{S}(t - d_r) = \{S_i(t - d_r(i))\}_i$. Let $d_{\max} = \max_{i,j} d_j(i)$, i.e. d_{\max} is the network diameter.

As described previously, one node is assigned the role of the controller. Every N time-slots, the location of the controller is recomputed. In order to do this computation, each node must be able to compute the controller at the current slot without communicating with the other nodes. Therefore, the controller selection algorithm must only depend on globally available information. In particular, we consider algorithms that are based only on delayed QLI, and do not consider CSI on deciding where to place the controller, since it is known that delayed QLI does not affect the throughput performance of the system [4]. After the controller is selected, the controller chooses a node to schedule for transmission based on the delayed CSI and QLI at the controller.

The primary objective of this work is to determine a joint controller placement and scheduling policy to stabilize the system of queues.

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

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{t=0}^{n-1} \mathbb{E}[Q_i(t)] < \infty \quad (7)$$

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 a 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 controller placement and scheduling problem above, and propose a throughput optimal joint controller placement and scheduling policy based on the information available at each node.

B. Two-Node Example




Figure 4: Example 2-node system model.

To illustrate the effect of dynamic controller relocation, consider a two-node system, as in Figure 4. Each node has instantaneous CSI pertaining to its channel at the current time, and 1-step delayed CSI of the other channel. Let Λ_1 be the throughput region when the controller is fixed at node 1, and let Λ_2 be the throughput region when the controller is fixed at node 2. The throughput regions Λ_r are characterized for each r by the following linear program (LP).

Maximize:

ϵ

Subject To:

$$\begin{aligned} \lambda_i + \epsilon &\leq \sum_{(s_1, s_2) \in \mathcal{S}} \mathbf{P}(S(t - d_r) = (s_1, s_2)) \\ &\quad \cdot \alpha_i(s_1, s_2) \mathbb{E}[S_i(t) | S_i(t - d_r(i)) = s_i] \quad \forall i \in \{1, 2\} \\ \sum_{i=1}^M \alpha_i(s_1, s_2) &\leq 1 \quad \forall \mathbf{s} \in \mathcal{S} \\ \alpha_i(s_1, s_2) &\geq 0 \quad \forall \mathbf{s} \in \mathcal{S}, i \in 1, 2 \end{aligned} \quad (8)$$

In the above LP, $\alpha_i(s_1, s_2)$ represents the fraction of time link i is scheduled when delayed CSI at the controller is (s_1, s_2) . To maintain stable queue lengths, the arrival rate to each queue must be less than the service rate at that queue, which is determined by the fraction of time the node transmits, and the expected throughput obtained over that link. For the case when the controller is at node r , Λ_r is the set of arrival rate pairs $\lambda = (\lambda_1, \lambda_2)$ such that there exists a solution to (8) satisfying $\epsilon > 0$. The proof that Λ_r is in fact the stability region of the system is found in [5].

The throughput regions Λ_1 and Λ_2 are plotted in Figure 5 for the case when $p = q = 0.1$. The throughput region is larger in the dimension of the controller, as a higher throughput is obtained at the node for which current CSI is available. The other node cannot attain the same throughput due to the CSI delay. Now consider a time-sharing policy, alternating between placing the controller at node 1 and node 2. The resulting throughput region Λ is given by the convex hull of Λ_1 and Λ_2 , which is shown as the dotted black line in Figure 5. Time-sharing between controller placements allows for higher throughputs than if the controller is fixed at either node. For example, the point $(\lambda_1, \lambda_2) = (\frac{3}{8} - \epsilon, \frac{3}{8} - \epsilon)$, for ϵ small, is not attainable by any fixed controller placement; however, this




Figure 5: Throughput regions for different controller scenarios. Assume the channel state model satisfies $p = 0.1, q = 0.1$, and $d_1(2) = d_2(1) = 1$.

throughput point is achieved by an equal time-sharing between controller locations.

The correct time sharing between controller placements depends on the arrival rate. This information is usually unavailable to the controller, and the control policy must stabilize the system even if the arrival rates change. Thus, we propose a dynamic controller placement and scheduling policy which achieves the full throughput region Λ using only delayed QLI for controller placement, and delayed CSI and QLI for scheduling, with no information pertaining to the arrival rates.

C. Queue Length-based Dynamic Controller Placement

Next, we consider controller placement policies that depend on delayed QLI, but not on delayed CSI³. Let Π be the set of policies which make a controller-placement decision based on QLI and not CSI, and make a scheduling decision based on the delayed CSI and QLI at the controller. In this section, we characterize the throughput region under such policies, and propose the *dynamic controller placement and scheduling (DCPS)* policy, which is proven to stabilize the system for all arrival rates within the throughput region.

Theorem 1 shows that the throughput region is characterized by the following LP.

Max.

ϵ

S.t.:

$$\begin{aligned} \lambda_i + \epsilon &\leq \sum_{\mathbf{s} \in \mathcal{S}} \mathbf{P}_S(\mathbf{s}) \sum_{r=1}^M \beta_r \alpha_i^r(\mathbf{s}) \mathbb{E}[S_i(t) | S_i(t - d_r(i))] \\ &\quad \forall i \in \{1, \dots, M\} \\ \sum_{i=1}^M \alpha_i^r(\mathbf{s}) &\leq 1, \quad \sum_{r=1}^M \beta_r \leq 1 \quad \forall \mathbf{s} \in \mathcal{S} \\ \alpha_i^r(\mathbf{s}) &\geq 0, \beta_r \geq 0 \quad \forall \mathbf{s} \in \mathcal{S}, i, r \in 1, \dots, M \end{aligned} \quad (9)$$

³For networks with a large diameter, the common CSI may be too stale to be used in controller placement; thus, we restrict our attention to policies which utilize QLI to make controller placement decisions, but not CSI.

This LP is an extension of the LP given in (8) to M nodes, with the addition of a time-sharing between controller locations. The optimization variables β_r and $\alpha_i^r(\mathbf{s})$ correspond to controller placement and link scheduling policies respectively. The variables β_r represent the fraction of the time that node r is elected to be a controller, and $\alpha_i^r(\mathbf{s})$ is the fraction of time that controller r schedules node i when the controller observes a delayed CSI of $\mathbf{S}(t - d_r) = \mathbf{s}$. Note that $\mathbf{P}_S(\mathbf{s})$ is the stationary probability of the Markov chain in Figure 1. The throughput region Λ , is the set of all non-negative arrival rate vectors λ such that there exists a feasible solution to (9) for which $\epsilon \geq 0$. This implies that there exists a stationary policy such that the effective service rate at each queue is greater than the arrival rate to that queue.

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 joint scheduling and controller placement algorithm, and proving that for all $\lambda \in \Lambda$, that policy stabilizes the system.

Lemma 1. *Suppose there exists a policy $\mathcal{P} \in \Pi$ that stabilizes the network for all $\lambda \in \Lambda$. Then, there exists a β_r and $\alpha_i^r(\mathbf{s})$ such that (9) 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, which places the controller at r with probability β_r , and schedules i to transmit when the delayed CSI at controller r is \mathbf{s} with probability $\alpha_i^r(\mathbf{s})$.

Due to the ergodicity of the finite state Markov chain controlling the channel state process, for any $\delta > 0$, there exists an N such that the probability of the channel state conditioned on the channel state N slots in the past is within δ of the steady state probability of the Markov chain.

$$\left| \mathbf{P}(\mathbf{S}(t) = \mathbf{s} | \mathbf{S}(t - N)) - \mathbf{P}(\mathbf{S}(t) = \mathbf{s}) \right| \leq \delta \quad (10)$$

Define $T_{SS}(\epsilon)$ is to be a large constant such that when $N = T_{SS}(\epsilon)$, (10) is satisfied for $\delta = \frac{\epsilon}{2|\mathcal{S}|}$, where $|\mathcal{S}| = 2^M$. In other words,

$$\left| \mathbf{P}(\mathbf{S}(t) = \mathbf{s} | \mathbf{S}(t - T_{SS}(\epsilon))) - \mathbf{P}(\mathbf{S}(t) = \mathbf{s}) \right| \leq \frac{\epsilon}{2|\mathcal{S}|} \quad (11)$$

T_{SS} is related to the time it takes the Markov chain to approach its steady state distribution.

Next, we propose the dynamic controller placement and scheduling (DCPS) policy, and show that this policy stabilizes the network whenever the arrival rate vector is interior to the capacity region Λ . Additionally, this proves the sufficient condition of Theorem 1. While the problem formulation is such that the controller is repositioned every N time-slots. In this section we prove throughput optimality for $N = 1$. The extension to general N is straightforward, and is discussed in Section IV-A.

Theorem 2. *Consider the dynamic controller placement and scheduling (DCPS) policy, which operates in two steps. First, choose a controller by solving the following optimization as a function of the delayed queue backlogs $Q_i(t - \tau_Q)$.*

$$r^* = \arg \max_r \left(\sum_{\mathbf{s} \in \mathcal{S}} \mathbf{P}_S(\mathbf{S}(t - d_r) = \mathbf{s}) \max_i Q_i(t - \tau_Q) p_{s_i,1}^{d_r(i)} \right) \quad (12)$$

where $P_S(\mathbf{s})$ is the steady state probability of the channel-state process. Then the controller uses its observed CSI $S(t - d_{r^*}(i)) = \mathbf{s}$, and schedules the following queue to transmit.

$$i^* = \arg \max_i Q_i(t - \tau_Q) p_{s_i,1}^{d_{r^*}(i)} \quad (13)$$

For any arrival rate λ , and $\epsilon > 0$ satisfying $\lambda + \epsilon \mathbf{1} \in \Lambda$, the DCPS policy stabilizes the system if $\tau_Q \geq d_{\max} + T_{SS}(\epsilon)$ for $T_{SS}(\epsilon)$ defined in (11).

Under policy DCPS, the controller is placed at the node maximizing the expected max weight schedule, over all possible states. Then, the controller observes the delayed CSI and schedules the max-weight schedule for transmission as in [2] and [5]. Moving the controller to nodes with high backlog increases the throughput at those nodes, keeping the system stable.

The proof of Theorem 2 is outlined as follows. Consider a scheduling and controller placement policy $\mathcal{P} \in \Pi$. Let $D_i^{\mathcal{P}}(t)$ be the departure process of queue i , such that $D_i^{\mathcal{P}}(t) = 1$ if there is a departure from queue i at time t under policy \mathcal{P} . Consider the evolution of the queues over T time slots, subject to a scheduling policy \mathcal{P} .

$$Q_i(t+T) \leq \left(Q_i(t) - \sum_{k=0}^{T-1} D_i^{\mathcal{P}}(t+k) \right)^+ + \sum_{k=0}^{T-1} A_i(t+k) \quad (14)$$

Equation (14) 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. Under this assumption, the packets that arrive within the T -slot period cannot depart within this period.

Since the CSI is delayed by different amounts of time depending on the location of the controller, and the controller changes locations over time, the relevant system state $\mathbf{Y}(t)$ is defined to include all possible combinations of delayed CSI, as well as the complete history of QLI.

$$\mathbf{Y}(t) = \left\{ \mathbf{S}(t - d_{\max}) \dots \mathbf{S}(t), \mathbf{Q}(0) \dots \mathbf{Q}(t) \right\} \quad (15)$$

This definition insures the system state is Markovian. Note that the system state $\mathbf{Y}(t)$ is not completely available to the controller, since each node has delayed CSI. Because d_{\max} is the largest delay to CSI in the network, values of $S(\tau)$ for $\tau < t - d_{\max}$ do not affect the evolution of the system. Define




Figure 6: Example star network topology where each node measures its own channel state instantaneously, and has d -step delayed CSI of each other node.

the following quadratic Lyapunov function:

$$L(\mathbf{Q}(t)) = \frac{1}{2} \sum_{i=1}^M Q_i^2(t). \quad (16)$$

The T -step Lyapunov drift is computed as

$$\Delta_T(\mathbf{Y}(t)) = \mathbb{E} \left[L(\mathbf{Q}(t+T)) - L(\mathbf{Q}(t)) \middle| \mathbf{Y}(t) \right]. \quad (17)$$

Theorem 2 follows by showing that as the system backlogs grow large under DCPS, the drift in (17) becomes negative, implying system stability [2]. We consider the Lyapunov drift over a T -slot window, where T is large enough that the system reaches its steady state distribution. The complete proof can be found in [7].

The throughput optimal controller placement uses delayed QLI $\mathbf{Q}(t - \tau_Q)$. The delay τ_Q must be sufficiently large such that $\mathbf{Q}(t - \tau_Q)$ is available at every node, i.e. $\tau_Q \geq d_{\max}$. Furthermore, we require that $\tau_Q \geq d_{\max} + T_{SS}(\epsilon)$, where $T_{SS}(\epsilon)$ is the time required for the channel state process to approach its steady state distribution (i.e. the mixing time of the Markov process). Even though QLI is available at much less delay, the controller must use an older version of the QLI for throughput optimality. The reasoning behind this is related to the fact that long queues are typically located at nodes with OFF channels, and relocating the controller to those channels will not increase throughput; however, if the QLI is sufficiently delayed, the dependency between backlog and channel state is removed.

While the throughput optimal controller placement is given in Theorem 2, for specific topologies, the throughput optimal controller placement in (12) takes on a simpler form. In particular, consider a topology for which each node is equidistant from all other nodes, as in Figure 6.

Corollary 1. Consider a system of M nodes, where only one can transmit at each time. Assume the controller has full knowledge of its own channel state and d -slot delayed CSI for each other channel, as in Figure 6. At time t , the DCPS policy places the controller at the node with the largest backlog at time $t - \tau_Q$.


$$r^* = \arg \max_r Q_r(t - \tau_Q) \quad (18)$$

Corollary 1 follows due to the symmetry of the system. Note


the queue lengths in the above theorem must still be delayed according to Theorem 2.

IV. SIMULATION RESULTS

To begin, we simulate a 6-Queue system with Bernoulli arrival processes of different rates. Assume the controller has instantaneous CSI for its channel, and homogeneously delayed (2 slots) CSI of each other channel, as in Figure 6. For each symmetric arrival rate vector λ , we simulate the evolution of the system over 100,000 time slots, and compute the average system backlog over those time slots. The results are plotted in Figure 7. Clearly, for small arrival rates, the average queue length remains very small. As the arrival rates increase towards the boundary of the stability region, the average system backlog starts to slightly increase. When the arrival rate grows beyond the stability region, the average queue length increases greatly, since packets arrive faster than they can be served in the system, implying that the system is unstable in this region.



(a) 2-Step Delayed QLI



(b) 100-Step Delayed QLI

Figure 7: Simulation results for different controller placement policies, with channel model parameters $p = 0.1, q = 0.1$.

Figure 7 compares the average queue backlogs under several controller placement policies. The black curve corresponds to a fixed controller, as in Section II. The blue curve corresponds

to a policy that chooses a controller at each time uniformly at random. Note that this random policy is optimal when the arrival rate is the same to each node, as it represents the correct stationary policy to stabilize the system. The red curve corresponds to the DCPS policy in Theorem 2.

In Figure 7a, the DCPS policy uses 2-step delayed QLI to place the controller. In this case, the DCPS policy fails to stabilize the system for the same set of arrival rates as the time-sharing policy, implying that the DCPS policy is not throughput optimal. However, in Figure 7b, the delay on the QLI is increased to 100 time-slots. In this scenario, the DCPS policy does stabilize the system for all symmetric arrival rates in the stability region. Thus, using further delayed information is required for throughput optimality. In general, dynamically changing the controller location provides a 7% increase in capacity region over the static controller placement.




Figure 8: Effect of QLI-delay on system stability, for $p = q = 0.1$. Each curve corresponds to a different value of τ_Q .

Figure 8 illustrates the effect of the delay in QLI on the stability of the system. This figure presents four different values for τ_Q , the delay to QLI used by the controller placement policy. For larger QLI delays, the system remains stable for more arrival rates. As τ_Q increases, the improvements to the stability region become smaller, as the stability region of the policy approaches the actual throughput region. In this example, using sufficiently delayed QLI yields a 16% increase in the stability region of the system over the policy which uses current QLI.




Figure 9: Two-level binary tree topology.

Additionally, we simulate the controller placement problem over the network in Figure 9, to compare the dynamic controller placement with the static controller placement in Section




Figure 10: Average queue length versus symmetric arrival rate for tree network in Figure 9 for different controller placement policies.

II. Figure 10 analyzes the stability of the system over different controller placement policies. The black solid line represents the DCPS policy, with QLI delay $\tau_Q = 150$. This policy is compared with the policy that randomly selects the controller and the policy that places the controller at node 3. These results show that relocating the controller according to the DCPS policy shows improvements over both the static placement, and an equal time-sharing between controller placements.

Figure 11 shows the fraction of time each node is selected as the controller under the DCPS policy for the binary-tree topology of Figure 9. For small transition probabilities (e.g. $p = q = 0.1$), the central node 3 is chosen as the controller most frequently. When the transition probabilities increase (e.g. $p = q = 0.3$), then more time is spent with nodes 2 and 4 as controllers. This corresponds to the analysis for static controller placement shown in Section II. Moreover, as the arrival rate increases toward the boundary of the stability region, the results resemble the static results even more closely. Note that the DCPS policy can be applied to any network topology, but we only consider smaller topologies in this work due to the computational complexity of computing the optimal controller location.

A. Infrequent Controller Relocation

Throughout this paper, we assume that a new controller placement occurs at every time slot. This is justified by ensuring that the controller placement algorithm depends only on information that is available to each node in the network. Thus, there is no additional communication overhead required to compute the controller placement. However, there may be an additional cost associated with relocating the controller due to the computation required. Therefore, in this section, we consider the case in which the controller placement occurs infrequently.

Consider the controller placement problem, in which the controller is relocated every N time slots. As discussed in Section III, the throughput region is not affected by infrequent controller placement. Lemma 1 shows that any arrival rate




 (a) Symmetric arrival rate $\lambda = 0.2$.

 (b) Symmetric arrival rate $\lambda = 0.25$.

Figure 11: Fraction of time each node is selected as the controller under DCPS for the topology in Figure 9. Blue bars correspond to system with $p = q = 0.1$, and red bars correspond to system with $p = q = 0.3$.

$\lambda \in \Lambda$ corresponds to a stationary policy which stabilizes the system. The throughput region Λ is formed by a time-sharing between controller placements. Consequently, the frequency of changing the controller placement does not affect throughput, but rather the overall fraction of time spent in each controller state.

The DCPS policy of Section III extends directly to the case of infrequent controller placement as follows.

Theorem 3. *Consider the dynamic controller placement and scheduling policy (DCPS), which operates in two steps. First, at each time $t = k * N$, choose a controller by solving the following optimization as a function of the delayed queue backlogs $Q_i(kN - \tau_Q)$.*

$$r^* = \arg \max_r \left(\sum_{\mathbf{s} \in \mathcal{S}} \mathbf{P}_{\mathbf{S}}(\mathbf{s}) \max_i Q_i(kN - \tau_Q) p_{s_{i,1}}^{d_r(i)} \right) \quad (19)$$

where $P_{\mathbf{S}}(\mathbf{s})$ is the steady state probability of the channel-state process. At the subsequent time slots $t = kN + j$, the controller uses its observed CSI $S(kN + j - d_r^*(i)) = \mathbf{s}$, and schedules the following queue to transmit.

$$i^* = \arg \max_i Q_i(kN - \tau_Q) p_{s_{i,1}}^{d_r^*(i)} \quad (20)$$

For any arrival rate λ , and $\epsilon > 0$ satisfying $\lambda + \epsilon \mathbf{1} \in \Lambda$, the DCPS policy stabilizes the system if $\tau_Q \geq d_{\max} + T_{SS}(\epsilon)$ for $T_{SS}(\epsilon)$ defined in (11).

The DPCS policy of Theorem 3 differs from that of Theorem 2 in that controller placement decisions are only made in time slots which are multiples of N , but the controller placement calculation is the same as in Theorem 2. The scheduling portion of Theorem 3 uses the delayed QLI with respect to the time at which the controller was placed, rather than the current time slot. This additional delay in QLI, does not affect the throughput optimality of the policy. The proof of Theorem 3 follows similarly to the proof of Theorem 2, except using a T -slot drift argument at every time slot $t = kN$ rather than every time slot.

V. CONCLUSION

This paper studies the optimal controller placement in wireless networks. First, we formulated the location of the optimal static controller placement. Then, we consider dynamically placing controllers, using QLI to move the controller to the heavily backlogged areas of the network. We characterize the throughput region under dynamic controller placement, and propose a throughput optimal joint controller placement and scheduling policy. This policy uses significantly delayed QLI to place the controllers, and the CSI available at the controller to schedule links.

An interesting result in this paper is that because the controller placement depends only on delayed QLI, the throughput optimal policy uses a very delayed version of the QLI, even if better QLI is available. This is due to the fact that QLI is related to CSI, particularly if there is a high degree of memory in the system.

REFERENCES

- [1] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *Automatic Control, IEEE Transactions on*, vol. 37, no. 12, pp. 1936–1948, 1992.
- [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] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synthesis Lectures on Communication Networks*, vol. 3, no. 1, pp. 1–211, 2010.
- [4] K. Kar, X. Luo, and S. Sarkar, "Throughput-optimal scheduling in multi-channel access point networks under infrequent channel measurements," *Wireless Communications, IEEE Transactions on*, 2008.
- [5] L. Ying and S. Shakkottai, "On throughput optimality with delayed network-state information," *Information Theory, IEEE Transactions on*, 2011.
- [6] —, "Scheduling in mobile ad hoc networks with topology and channel-state uncertainty," *Automatic Control, IEEE Transactions on*, vol. 57, no. 10, pp. 2504–2517, 2012.
- [7] M. Johnston, "The role of control information in wireless link scheduling," Ph.D. dissertation, MIT, 2015.
- [8] C. H. Papadimitriou and K. Steiglitz, *Combinatorial optimization: algorithms and complexity*. Courier Dover Publications, 1998.
- [9] D. P. Bertsekas, *Introduction to Probability: Dimitri P. Bertsekas and John N. Tsitsiklis*. Athena Scientific, 2002.
- [10] 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.