Stability of the Max-Weight Protocol in Adversarial Wireless Networks

Sungsu Lim, *Student Member, IEEE*, Kyomin Jung, *Member, IEEE*, and Matthew Andrews, *Senior Member, IEEE*

***Abstract*—In this paper, we consider the MAX-WEIGHT protocol for routing and scheduling in wireless networks under an *adversarial* model. This protocol has received a significant amount of attention dating back to the papers of Tassiulas and Ephremides. In particular, this protocol is known to be *throughput-optimal* whenever the traffic patterns and propagation conditions are governed by a *stationary stochastic process*. However, the standard proof of throughput optimality (which is based on the negative drift of a quadratic potential function) does not hold when the traffic patterns and the edge capacity changes over time are governed by an arbitrary adversarial process. Such an environment appears frequently in many practical wireless scenarios when the assumption that channel conditions are governed by a stationary stochastic process does not readily apply. In this paper, we prove that even in the above adversarial setting, the MAX-WEIGHT protocol keeps the queues in the network stable (i.e., keeps the queue sizes bounded) whenever this is feasible by some routing and scheduling algorithm. However, the proof is somewhat more complex than the negative potential drift argument that applied in the stationary case. Our proof holds for any arbitrary interference relationships among edges. We also prove the same stability of ϵ-approximate MAX-WEIGHT under the adversarial model. We conclude the paper with a discussion of queue sizes in the adversarial model as well as a set of simulation results.**

***Index Terms*—Max-Weight, routing, scheduling, stability, wireless network.**

I. INTRODUCTION

WE CONSIDER the performance of the Max-Weight routing and scheduling algorithm in adversarial networks. Max-Weight has been one of the most studied algorithms [7], [15], [16] since it was introduced in the works of Tassiulas and Ephremides [20], [21] and Awerbuch and Leighton [8], [9]. The key property of Max-Weight is that, for a fixed set of flows, it is *throughput-optimal* in stochastic networks in a wide variety of scenarios [20], [11], [18], even though it may fail to provide maximum stability in a scenario with flow-level dynamics [22]. That is, for a fixed set of flows, the Max-Weight protocol keeps the queues in the network stable whenever this is feasible by some routing and scheduling algorithm. Moreover, we can obtain a bound on the amount of packets in the system that is polynomial in the network size.

However, the standard analyses of the Max-Weight algorithm make critical use of the fact that the channel conditions and the traffic patterns are governed by stationary stochastic processes. The stationary stochastic model deals with the case where traffic patterns do not deviate much from their time-average behavior. On the other hand, we shall consider the worst-case traffic scenario modeled by adversarial models. If an adversary chooses traffic patterns and interference conditions, and edge capacities change over time in an arbitrary way, then the question remains as to whether a system running under Max-Weight can be unstable. It is important to model the worst (adversarial) case because nonevenly distributed traffic patterns are observed over time in many queuing models. A typical adversarial scenario is a military communication network, in which there could exist adversarial jammers. Once it is jammed, the victim link will have zero capacity or very weak capacity. Ensuring stability under the worst case is crucial in many such systems. The aim of the current paper is to resolve this question.

Previous work has shed some light on this issue. In [5], it was shown that for a single transmitter sending data over one-hop edges to a set of mobile users, if the set of nonzero channel rates can approach zero arbitrarily closely, then no protocol can be stable. However, since this is a fairly unnatural condition, [5] looked at the more natural setting in which the set of all rates is finite. For this case, a stable protocol was given, but it was a somewhat unnatural protocol that relies on a lot of bookkeeping. The stability of a more natural protocol such as Max-Weight was left unresolved.

Previous work has looked at the stability of Max-Weight in adversarial settings. In particular [1] showed stability for static networks (with adversarial traffic), and [7] and [2] showed stability in dynamic networks for both single-commodity and multi-commodity demands, respectively. However, these proofs only applied to the case when each edge could be scheduled independently (in other words, the decision to transmit on an edge has no affect on the edge rates on other edges), This is obviously not a suitable model for wireless transmissions in which edges can clearly affect each other. As discussed in [5], the stability of Max-Weight was not known in the adversarial setting for the case of interfering edges, even if we only have one node that transmits.

In this paper, we resolve the question of the stability of Max-Weight in general adversarial networks. We present an adversarial model of interfering edges and show that the Max-Weight policy always maintains stability as long as we are

Manuscript received June 07, 2013; accepted September 29, 2013; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor T. Javidi. Date of publication November 08, 2013; date of current version December 15, 2014. This work was funded in part by the Basic Science Research Program under the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT & Future Planning (2012R1A1A1014965).

S. Lim is with the Department of Knowledge Service Engineering, KAIST, Daejeon 305-701, Korea (e-mail: ssungssu@kaist.ac.kr).

K. Jung is with the Department of Electrical and Computer Engineering, Seoul National University, Seoul 151-744, Korea (e-mail: kjung@snu.ac.kr).

M. Andrews is with Alcatel-Lucent Bell Labs, Murray Hill, NJ 07974 USA (e-mail: andrews@research.bell-labs.com).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2013.2288372

strictly within the network stability region, even when the stability region is allowed to change over time. We consider a very general adversarial model that can be applied to all the possible interference conditions, including k-hop interference [19], independent set constraints [12], [13], and node-exclusive constraints [3], [4], [14], [17]. Our proof gives a bound on the queue size that is exponential in the network size. (This is unlike the stochastic case.) However, we also demonstrate (using an example inspired by [6]) that such exponential queue sizes can occur. Although computing the optimal solution of Max-Weight is computationally NP hard for many scenarios, in many practical wireless networks, ε-approximate solutions can be computed in polynomial time [10], [12], [13]. In this paper, we also prove the same stability of any ε-approximate Max-Weight under the adversarial model, when $\varepsilon > 0$ is small enough. We conclude the paper with a set of simulation results showing actual stability of Max-Weight on adversarial setups.

A. Discussion

We now give a high-level description of the Max-Weight algorithm and discuss why the standard stochastic analyses are invalid in the adversarial case. Essentially, the protocol operates by maintaining at each node v a queue of data for each possible destination d. We denote the size of this queue at time t by $q^t_{v,d}$. For any set of edges in the network, the total *weight* on the edge set at time t is the sum over all edges in the set of the queue differentials multiplied by the instantaneous edge rates. (A formal definition will be given in Section I-C.) At all times, the MAX-WEIGHT protocol transmits data on edges so as to maximize the total weight that it gains. In many situations, computing the exact Max-Weight set of transmissions is a computationally hard problem. However, an ε-approximate Max-Weight algorithm can be implemented efficiently in many practical setups. In Section VI, we show the stability of any ε-approximate Max-Weight algorithm, when $\varepsilon > 0$ is small enough.

We say that we are in the stationary stochastic model if there is an underlying stationary Markov Chain whose state determines the channel conditions on the edges. We say that we are in the adversarial model if we do not make such assumptions. In order to make sure that the network is not inherently overloaded, the adversarial model assumes that there *exists* some way to route and schedule the packets so as to keep the network stable. However, these routes and schedules are *a priori* unknown to the Max-Weight algorithm.

Most previous analyses of Max-Weight have been performed in a stationary stochastic model, and they take the following form. Define a quadratic potential function $P(t) = \sum_{v,d}(q^t_{v,d})^2$ and show, using the assumption that the traffic arrivals are within the network stability region, that the potential function always has a negative drift up to an additive second order term of $P(t+1) - P(t) = \sum_{v,d}(q^t_{v,d} + \sigma^t_{v,d})^2 - \sum_{v,d}(q^t_{v,d})^2$, where $\sigma^t_{v,d} = q^{t+1}_{v,d} - q^t_{v,d}$. Moreover, when the potential function becomes sufficiently large, the negative drift in the first-order term is sufficient to overcome the positive second-order term. Therefore, the entire potential function has a negative drift. This determines an upper bound on $P(t) = \sum_{v,d}(q^t_{v,d})^2$, and hence we have an upper bound on $\sum_{v,d} q^t_{v,d}$.

The reason that this type of analysis does not apply in the adversarial model is that the channel rates associated with the large queues in the network may be very small (or even zero). In this case, we cannot necessarily say that a large queue implies a large drop in the potential. Thus, for any possible queue configuration, there is always the possibility that the potential function $P(t) = \sum_{v,d}(q^t_{v,d})^2$ can increase. Hence, we need a different approach to ensure stability. We discuss this in more detail in Section II.

B. Why Do Adversarial Models Make Sense?

We now briefly discuss why it is useful to consider the adversarial setting that includes the *worst*-case scenario; A model that is governed by a stationary stochastic process is not general enough to cover many widely occuring scenarios. For example, consider a cellular network in which a car is driving down a road between evenly spaced base stations. In this case, the channel conditions between the car and its closest base station will rise and fall in a periodic fashion. Moreover, when a car drives into an area of poor coverage (e.g., a tunnel), the channel rate could go to zero. In particular, this could happen in a haphazard manner that is not modeled by a stationary stochastic process.

The situation is even more severe in ad hoc networks. As nodes move around, many of the edges (i, j) will only be active for a finite amount of time. Hence, any stationary stochastic model that gives a nonzero channel rate to such an edge cannot accurately reflect the edge rate over a long time period. However, we still wish to ensure that the queue sizes will not blow up unnecessarily over time, and we believe that an adversarial analysis is one way to address this type of question.

In [2], the stability of Max-Weight in some adversarial model was proven. However, it was not sufficiently rich to capture many types of wireless interactions. First of all, in the model of [2], all edge rates were either zero or one. Second, when an edge had rate one, we could transmit on it regardless of what is happening on the other edges. In other words, the interference is assumed to be edgewise independent. However, this model cannot capture a situation in which edge rates are variable, nor can it capture a scenario with two interfering edges such that we can transmit on either one in isolation, but not both simultaneously.

In this paper, we will define a more general adversarial model in which any interference conditions are possible and edge rates can vary over time. This allows us to capture arbitrary types of wireless interference behavior. We next describe our model in more detail, after which we present our results.

C. Model

We assume a system in which time is divided into discrete time-slots. We consider a queueing model for packet transmissions. Let D be the set of possible destinations. Each destination in D can be a subset of the set of nodes. At each time step $t \in \mathbb{N}$, a set of feasible edge rate vectors $R(t) \subset \mathbb{R}^k$ is given by the adversary where $k = |E|$, and E is the set of all directed edges. Suppose that $r(t) \in R(t)$. It means that if we write $r(t) = (r_1(t), r_2(t), \ldots, r_k(t))$,, then it is possible to transmit packets on edge e at rate $r_e(t)$, *for all edges simultaneously*. In other words, we can transmit data of size $x_{1,d} \geq 0$ for any destination $d \in D$, where the sum of the amount of transmissions on edge 1 is x_1. Similarly, we can transmit data of size x_2 on edge 2, data of size x_3 on edge 3, etc., so long as $0 \leq x_e \leq r_e(t)$ for all e. Note that this means that the rates satisfy the *downward* *closed* property, i.e., we can always transmit on an edge at a rate that is less than the rate $r_e(t)$.

This is a very general setting for the interference model because it includes all the possible interference constraints, including k-hop interference, independent set constraints, and node-exclusive constraints. For example, for a dynamic network $G(V, E(t))$, $R(t) = \{(r_e(t))_{e\in E(t)} | r_e(t) = 0 \text{ or } 1 \text{ for all } e \in E(t),\ r_{e_1}(t) r_{e_2}(t) = 0 \text{ if } e_1 \text{ and } e_2 \text{ are incident in } E(t)\}$ represents a set of feasible edge rate vectors of independent set constraints on $E(t)$ that changes over time.

We make the following assumption about the adversary. (It was shown in [5] that if we do not have these conditions, then no online protocol can be stable.)

- All packet arrival and edge rates are bounded from above, and nonzero rates are bounded away from zero. In other words, there exist values $R_{\min} > 0$ and $R_{\max} > 0$ such that for each $r(t) = (r_1(t), \ldots, r_k(t)) \in R(t)$, $r_e(t) \leq R_{\max}$, and if $r_e(t) \neq 0$, then $r_e(t) \geq R_{\min}$.

We now define the (ω, ε)-adversary, where $\omega \in \mathbb{N}$, and $\varepsilon > 0$. At each time, it determines the packet arrivals and edge capacities. Then, the routing and scheduling algorithm decides the packet transfers in the network against the (ω, ε)-adversary. In this manner, our framework can be understood as a type of sequential game.

Definition 1: We say that an adversary injecting the packets and controlling the edges is an (ω, ε)-adversary, $A(\omega, \epsilon)$, for some $\varepsilon > 0$ and some integer $\omega \geq 1$ called a *window* parameter, if the following holds: The adversary defines the feasible rate vectors and packet arrivals in each time step subject to the constraint that there exists a routing and scheduling algorithm T (possibly involving fractional movement of packets) that keeps the system stable. Let t_p be the time when a packet p is injected. Then, we can define $\Psi_p = \{(e,t) | t \in [t_p, t_p + \omega - 1], \ell(p,e,t) > 0$, where $\ell(p,e,t)$ is a fractional amount of p that is transmitted by T along e at time $t\}$, which corresponds to the movement of packet p from its source to one of its destinations under the algorithm T. For all packet p, $(1 - \frac{\varepsilon}{2})$ fraction[1] of p will arrive to its destination during the window $[t_p, t_p + \omega - 1]$. For any integer j, let I^j be the set of packets injected during the window $W_j = [j\omega, (j+1)\omega - 1]$. Then, the adversary assumes that the following holds:

$$\sum_{p \in I^j \cup I^{j-1}, (e,t) \in \Psi_p, t \in W_j} \ell(p,e,t) \leq \sum_{t \in W_j} (1 - \varepsilon) r_e(t)$$

where $r(t) \in R(t)$ are edge rate vectors assigned by T.

This is a very general adversarial model because it covers all the possible interference conditions in dynamic networks, including k-hop interference, independent set constraints, and node-exclusive constraints, and this model includes adversarial models used in [1], [2], and [7]. We prove the following theorem, which shows that the MAX-WEIGHT protocol is *throughput-optimal* even against the strongest adversary.

Theorem 1: The MAX-WEIGHT protocol is stable under any $A(\omega, \epsilon)$ for any $\varepsilon > 0$, and $\omega \in \mathbb{N}$.

D. Protocol

We now define the MAX-WEIGHT protocol. We assume that each node v has $|D|$ queues that correspond to each destination, respectively. Thus, we have $n|D|$ many queues. Let $Q_{v,d}$ be the queue at node v for data having destination d. Let $q_{v,d}^t$ be the total size of data in queue $Q_{v,d}$ at time t. We define a general routing and scheduling algorithm MAX-WEIGHT(β) that is parameterized by a parameter $\beta > 0$. We use MAX-WEIGHT to denote the algorithm with $\beta = 1$. In this paper, we will use the term *scheduling algorithm* to mean a combined routing and scheduling algorithm.

Algorithm MAX-WEIGHT(β)

For each time $t \in \mathbb{N}$,

1) Choose $r(t) \in R(t)$ and $d^{(e)} \in D$ for each $e = (v,u) \in E$, such that $\sum_{e\in E} s_e(t)((q_{v,d^{(e)}}^t)^\beta - (q_{u,d^{(e)}}^t)^\beta)$ is maximized (with an arbitrary tiebreaking rule) where
$$s_e(t) := \min\left\{ r_e(t), \left| \frac{q_{v,d^{(e)}}^t - q_{u,d^{(e)}}^t}{2} \right| \right\}.$$
Send data of size $s_e(t)$ from $Q_{v,d^{(e)}}$ to $Q_{u,d^{(e)}}$ along e.
2) For each node v, accept all packets injected by the Adversary to v at time t.
3) Remove all packets that arrive at their destination.

When $\beta > 0$, $(q_{v,d^{(e)}}^t)^\beta - (q_{u,d^{(e)}}^t)^\beta \geq 0$ implies $q_{v,d^{(e)}}^t - q_{u,d^{(e)}}^t \geq 0$, so it guarantees all packet movement between queues occur from a taller queue to a smaller queue.

The algorithm can be understood to be designed so that the following *potential function* decreases as much as possible (however, as discussed earlier and unlike in the stochastic case, there is no simple argument that for sufficiently large queue sizes there always *is* a decrease in potential):

$$P(t) \triangleq \sum_{v,d} \left(q_{v,d}^t\right)^{\beta+1}.$$

II. STOCHASTIC ANALYSIS

In this section, we give more details of the typical stochastic analysis and explain why this type of analysis does not directly hold in the adversarial setting. We say that we are in the stationary stochastic model if there is an underlying stationary Markov Chain $\mathcal{M}$ with state space $\{m_r\}$ and a function $f(\cdot)$ from $\{m_r\}$ to sets of feasible edge rate vectors $R(t)$ such that the Markov Chain updates its state at each time step and if it has state $\{m_r\}$ at time t then $R(t) = f(m_r)$.

Throughout this section, we will focus on the case that $\beta = 1$ and study the potential function $P(t) = \sum_{v,d}(q_{v,d}^t)^2$. Let $a_{v,d}^t$ (resp. $b_{v,d}^t$) be the amount of data arriving into (resp. departing from) $Q_{v,d}$ at time t, according to the MAX-WEIGHT algorithm. For simplicity, we shall also discuss the most basic scenario in which the distribution over feasible service rate vectors is i.i.d. at each time step. By taking into account the i.i.d. nature of the service rate vectors and the fact that the traffic injections are within the $(1 - \varepsilon)$ multiplicative interior of the network stability region, we can assume that there exists some scheduling algorithm for which the corresponding quantities ${a'}_{v,d}^t$ and ${b'}_{v,d}^t$ satisfy $E[({a'}_{v,d}^t - {b'}_{v,d}^t)] \leq -\varepsilon$ for all v, d.

[1] In fact, for any $(1 - \delta)$ fraction of p with constant $0 < \delta < \varepsilon$, all the results in this paper hold.

The expected change in $P(\cdot)$ from time step t to time $t+1$ is given by

$$\begin{aligned}
E[P(t+1) - P(t)] \\
&= E\left[\sum_{v,d}\left(q_{v,d}^{t+1}\right)^2 - \sum_{v,d}\left(q_{v,d}^{t}\right)^2\right] \\
&= E\left[\sum_{v,d}\left(q_{v,d}^{t} + a_{v,d}^{t} - b_{v,d}^{t})\right)^2 - \sum_{v,d}\left(q_{v,d}^{t}\right)^2\right] \\
&= E\left[\sum_{v,d}\left(\left(q_{v,d}^{t}\right)^2 + 2q_{v,d}^{t}\left(a_{v,d}^{t} - b_{v,d}^{t}\right) + \left(a_{v,d}^{t} - b_{v,d}^{t}\right)^2\right)\right. \\
&\qquad \left. - \sum_{v,d}\left(q_{v,d}^{t}\right)^2\right] \\
&\le E\left[\sum_{v,d}\left(2q_{v,d}^{t}\left({a'}_{v,d}^{t} - {b'}_{v,d}^{t}\right) + \left(a_{v,d}^{t} - b_{v,d}^{t}\right)^2\right)\right]. \quad (1)
\end{aligned}$$

The final inequality is due to the definition of MAX-WEIGHT since we can think of MAX-WEIGHT as always making the decision that minimizes $\sum_{v,d} q_{v,d}^t(a_{v,d}^t - b_{v,d}^t)$. Recall that $E[({a'}_{v,d}^t - {b'}_{v,d}^t)] \le -\varepsilon$ for all v, d. Moreover, since there is an upper bound on the amount of data that can be transfered between two queues at each time step, $E[(a_{v,d}^t - b_{v,d}^t)^2]$ is bounded by some quantity C that is independent of time. Hence

$$E[P(t+1) - P(t)] \le C - 2\sum_{v,d} q_{v,d}^t \varepsilon$$

and thus if there is some $Q_{v,d}$ that satisfies $q_{v,d}^t \ge C/2\varepsilon$, then the expected drift of $P(t)$ is negative at time t. This in turn implies that $P(t)$ cannot grow indefinitely over time, and so the system is stable.

We can now demonstrate why this type of argument *does not* hold in the adversarial model. In a nonstationary, adversarial environment, it is not necessarily the case that the set $R(t)$ and exogenous packet arrival rates are independent of the $q_{v,d}^t$ values. That is, we cannot assume that a large queue will have good connectivity to the rest of the network, so there is no analogue of the statement that $E[({a'}_{v,d}^t - {b'}_{v,d}^t)] \le -\varepsilon$. In particular, it may be the case that for all large $q_{v,d}^t$ and for all $r(t) \in R(t)$, the value of $r_e(t)$ is zero for all edges e that are adjacent to node v. Indeed, the fact that we have built up a large queue in one region of the network may be precisely because that region has poor connectivity to other parts of the network. Hence, we need a different type of argument to show stability in the adversarial setting, and this is the question that we address in this paper.

III. MAIN RESULTS

At the highest level, our proof proceeds as follows. We first show a result that bears some similarity to the "negative drift" result that is used to prove stability in stationary stochastic systems. In particular, in Theorem 2 we show that whenever a packet is injected, we can assign a set of partial transmissions by the MAX-WEIGHT(β) protocol to the packet such that the resulting decrease in potential almost matches the increase in potential that arises from the packet injection itself. This allows us to bound the increase in potential whenever a packet is injected. (We note as an aside that when there are no packet injections, the MAX-WEIGHT(β) protocol ensures that the potential never increases.) Moreover, Theorem 2 also shows that whenever there is an injection to a queue that is sufficiently tall, the assigned transmissions induce a decrease in potential *more than* the increase due to the packet injection. Hence, for such injections, there will always be a decrease in potential.

However, in an adversarial system, this type of argument is not sufficient to show stability since it might be the case that most packets are injected into small queues. We therefore extend the proof of Theorem 2 to a more general result that will ensure stability. In particular, we introduce the notion of a *bad injection*. This is an injection that is extra to the injections that are allowed by our definition of adversary. This notion is convenient since we will use an inductive proof in which injections to small queues that lead to a big increase in potential are treated as "extra" packets by the inductive hypothesis. In particular, we are able to use an inductive argument to show that the number of bad injections is bounded, and hence we can obtain an upper bound of the potential over all time. This immediately implies the stability of MAX-WEIGHT(β).

We now describe these ideas in a little more detail. The procedure in our setup is as follows. At each time, an adversary chooses the packet injections and interference conditions. Then, MAX-WEIGHT(β) determines the (routing and) scheduling of packet transmissions. To show the stability of MAX-WEIGHT(β), we will define an assignment of each injected packet to a set of (partial) transmissions in the network, so that any injected packet to a *tall* queue will decrease the potential function.

Definition 2: We imagine that there are $|D|$ links on each directed edge corresponding to each possible destination, respectively. Let $L = \{\ell = (e,d) | e = (v,u) \in E, d \in D\}$ be the set of all links. Let p be a packet injected at time t, and let $W = [t, t+\omega-1]$. Let $s_e(t')$ be the vector chosen by MAX-WEIGHT(β) that maximizes $\sum_e s_e(t')((q_{v,d(e)}^{t'})^\beta - (q_{u,d(e)}^{t'})^\beta)$. For a given adversary $A(\omega,\varepsilon)$, and a scheduling algorithm Alg, we say that $\Gamma_p(A(\omega,\epsilon), Alg) = (s_{p,\ell}(t'))_{\ell\in L, t'\in W} \in \mathbb{R}^{\omega|L|}$ is a *set of (partial) transmissions assigned with* p if it satisfies that for each $e \in E$, $d \in D$, $t' \in W$: (i) $s_{p,(e,d)}(t') \ge 0$; and (ii) $\sum_{p,d} s_{p,(e,d)}(t') \le s_e(t')$. For convenience, we denote it by Γ_p. By the assumption on the adversary, the sets Γ_p exist.

We note that the word *partial* is used to reflect the fact that one transmission may correspond to multiple packets p subject to the condition (ii). Conceptually, it allows the case that an injected packet can be transmitted to its destination across multiple paths. Thus, an assignment of partial transmissions Γ_p of each packet can represent many general routing patterns. Moreover, it allows the case when Γ_p does not form a set of paths. An example of this assignment is shown in Fig. 1.

Theorem 2: Consider a given adversary $A(\omega,\epsilon)$ for any $\omega \ge 1$ and $\varepsilon > 0$, and the MAX-WEIGHT(β) protocol for some $\beta > 0$. For all injected packets p, we can assign this packet with $\Gamma_p = \Gamma_p(A(\omega,\epsilon)$, MAX-WEIGHT($\beta$) simultaneously so that the sum

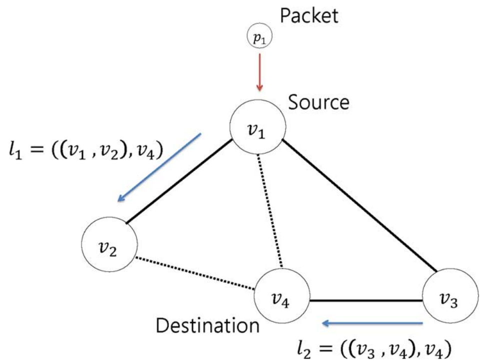

Fig. 1. Example of partial transmission assignment. Suppose that a packet p_1 is injected to a node v_1 at time t_{p_1}, and its destination is v_4. For instance, Γ_{p_1} may contain $s_{p_1,\ell_1}(t_{p_1}+2) = 0.5$ and $s_{p_1,\ell_2}(t_{p_1}-1) = 1$. Note that the assignments do not need to be at the same time, and the whole assignments do not need to form a path or multiple paths.

of total potential changes due to Γ_p is less than $-\frac{\varepsilon}{1-\varepsilon/2}\ell_p(\beta+1)q^\beta + \ell_p O(q^{\beta-1})$, where q is the height of the queue where the packet p is injected. Therefore, there is a constant q^* depending only on ω and ε, so that if $q \geq q^*$, the sum of potential changes due to the injection is less than $-\frac{\varepsilon}{2}\ell_p q^\beta$.

In Section IV, we will prove the stability of the MAX-WEIGHT protocol under any $A(\omega,\varepsilon)$. The same argument can be applied to prove that the MAX-WEIGHT(β) protocol with any constant $\beta > 0$ is stable under any $A(\omega,\epsilon)$ with $\varepsilon > 0$.

We define a more general adversarial model, with an adversary $A(\omega,\epsilon,b)$, which we call a *general adversarial queue system* with *bad packets*. In this model, the number of queues can be any finite number, not only of the form $n|D|$. Here, $A(\omega,\epsilon,b)$ allows b many bad packets in the system, where the notion of bad packet is defined below.

Recall that by Theorem 2, for our regular definition of $A(\omega,\epsilon)$ under MAX-WEIGHT, we can associate each injected packet p with a set of partial transmissions Γ_p so that the sum of potential changes due to these movements are at most $-\frac{\varepsilon}{1-\varepsilon/2}\ell_p q + C$, where q is the height of the queue where the packet is injected and C is a constant depending only on ω and ε (but not on n and t). For the general adversarial model, there may be a small number of injected packets for which the sum of potential changes due to Γ_p is at least $-\frac{\varepsilon}{1-\varepsilon/2}\ell_p q + C + 1$. We call such a packet a *bad packet* and say that all the other injected packets are *good packets*.

Definition 3: We say that an adversary injecting the packets and controlling edge capacities in a general adversarial queue system is an $A(\omega,\epsilon,b)$ adversary for some $\varepsilon > 0$ and some integers $\omega \geq 1$ and $b \geq 0$, if the following holds: There exists a scheduling algorithm Alg and an assignment of partial transmissions for each injected packet p (for example, the collection of Γ_p for MAX-WEIGHT protocol), such that among all the packets injected over all time, there are at most b bad packets.

In the proof of MAX-WEIGHT stability, we will use an induction on the number of queues. For a given subset of queues, we can imagine a smaller (sub)system of those queues. For an injected packet p, if *too much* of the assigned partial transmissions do not occur between the queues of the subsystem, we will consider p as a bad packet. In the analysis, we will use the following property of good packets.

Lemma 3: Consider a general adversarial queue system $A(\omega,\epsilon,b)$ with a corresponding scheduling algorithm Alg and a corresponding set of partial transmissions of packets Γ. Then, there is a constant q^* depending on ω and ε, so that for any good packet p injected to a queue of height q, if $q \geq q^*$, the sum of the *decrease of potential* due to p is more than $\frac{\varepsilon}{2}\ell_p q$.

Proof: From the definition, the sum of potential changes due to the injection of any good packet p is at most $-\frac{\varepsilon}{1-\varepsilon/2}\ell_p q + C + 1$. Let $q^* = \frac{4-2\varepsilon}{(2\varepsilon+\varepsilon^2)(\ell_p)}(C+1)$, then for any $q \geq q^*$, $\frac{\varepsilon}{1-\varepsilon/2}\ell_p q - \frac{\varepsilon}{2}\ell_p q \geq \frac{2\varepsilon+\varepsilon^2}{4-2\varepsilon}\ell_p q^* = C+1$. Thus, the decrease of potential is more than $\frac{\varepsilon}{2}\ell_p q$ for $q \geq q^*$. ■

The crux of our analysis will involve proving the following theorem (in Section IV-B).

Theorem 4: Consider any general adversarial queue system $A(\omega,\epsilon,b)$ for any constant $\varepsilon > 0$ with corresponding scheduling algorithm Alg. If Alg guarantees all packet movement between queues occurs from a taller queue to a smaller queue, then Alg makes the system stable, i.e., keeps the queue sizes bounded.

Hence from Theorem 4, we obtain Theorem 1 directly since MAX-WEIGHT(β) only transmits packets from a taller queue to a smaller one.

IV. PROOFS OF THEOREMS

A. Proof of Theorem 2

Proof: We divide time into windows of ω time steps, $[0,\omega-1],[\omega,2\omega-1],[2\omega,3\omega-1],\ldots$. Since $W_j = [j\omega,(j+1)\omega-1]$ for all integer $j \geq 0$, the collection of W_j for $j \geq 0$ is nonoverlapped, and the union of this collection covers all time-slots t. From now on, let $W = W_j$ for some integer $j \geq 0$.

For each time $t' \in W$, and for each node v, we accept all packets injected by the adversary. For each packet $p \in I^j \cup I^{j-1}$, we will associate some fraction $r_p = \{g_{p,e}(t') | (e,t') \in \Psi_p\}$ of rates of directed edges used in Ψ_p as follows. Let $p_1,\ldots p_m$ be all the packets injected in $I^j \cup I^{j-1}$; the order of p_i's can be any possible ordering. From Definition 1

$$\sum_{i=1,\ldots,m,t'\in\{t'\in W|(e,t')\in\Psi_{p_i}\}} \ell(p_i,e,t') \leq (1-\varepsilon)\sum_{t'\in W} r_e^{(0)}(t'). \tag{2}$$

where $r^{(0)}(t) \in R(t)$ are edge rate vectors assigned by T.

First, for p_1 and for each directed edge e used in Ψ_{p_1}, we can define $g_{p_1,e}(t')$ for each $t' \in W$ so that

$$0 \leq g_{p_1,e}(t') \leq r_e(t')$$
$$(1-\varepsilon)\sum_{t'\in W} g_{p_1,e}(t') = \sum_{t'\in W} \ell(p_1,e,t'). \tag{3}$$

We define $g_{p_1,e}(t') = 0$ for each directed edge e that is not used in Ψ_{p_1}. Then, from (2) and Section IV-A, for each $e \in E$ and $t' \in W$, let $r_e^{(1)}(t') = r_e(t') - g_{p_1,e}(t')$. Then, we have

$$\sum_{i=2,(e,t')\in\Psi_{p_i},t'\in W}^{m} \ell(p_i,e,t') \leq (1-\varepsilon)\sum_{t'\in W} r_e^{(1)}(t').$$

Similarly, for p_2 and for each e used in Ψ_{p_2}, we can define $g_{p_2,e}(t')$ for each $t' \in W$ so that

$$0 \le g_{p_2,e}(t') \le r_e^{(1)}(t')$$
$$(1-\varepsilon)\sum_{t'\in W} g_{p_2,e}(t') = \sum_{t'\in W} \ell(p_2,e,t').$$

By continuing this process, we can define $g_{p_i,e}(t')$ inductively for all $i \ge 2$, for each e used in Ψ_{p_i} and $t' \in W$ so that

$$0 \le g_{p_i,e}(t') \le r_e^{(i-1)}(t')$$
$$(1-\varepsilon)\sum_{t'\in W} g_{p_i,e}(t') = \sum_{t'\in W} \ell(p_i,e,t'). \tag{4}$$

At time t', think of a link (e,d) where the difference between the sizes of the queues $q_{v,d}^{t'}$ and $q_{u,d}^{t'}$ at time t' is at least a rate $r_e(t')$ of a directed edge $e=(v,u)$, i.e., $q_{v,d}^{t'} \ge q_{u,d}^{t'} + r_e(t')$. Then, the potential change $C_e(t')$ due to transmission via a link (e,d) at time t' is

$$\begin{aligned} C_e(t') &= \left(q_{u,d}^{t'} + r_e(t')\right)^{\beta+1} - \left(q_{u,d}^{t'}\right)^{\beta+1} \\ &\quad + \left(q_{v,d}^{t'} - r_e(t')\right)^{\beta+1} - \left(q_{v,d}^{t'}\right)^{\beta+1} \\ &= r_e(t')(\beta+1)\left(\left(q_{u,d}^{t'}\right)^{\beta} - \left(q_{v,d}^{t'}\right)^{\beta}\right) \\ &\quad + r_e(t')O\left(\left(q_{u,d}^{t'}\right)^{\beta-1} + \left(q_{v,d}^{t'}\right)^{\beta-1}\right). \end{aligned} \tag{5}$$

Note that this is also true when $|q_{u,d}^{t'} - q_{v,d}^{t'}| < r_e(t')$. Hence, when $g_{p,e}(t')$ amount of edge rate of e at time t' is assigned to an injected packet p, we conclude that $g_{p,e}(t')(\beta+1)((q_{u,d}^{t'})^{\beta} - (q_{v,d}^{t'})^{\beta}) + g_{p,e}(t')O((q_{u,d}^{t'})^{\beta-1} + (q_{v,d}^{t'})^{\beta-1})$ amount of potential change is induced by a packet p.

We consider the sum of potential changes at each time t' by MAX-WEIGHT(β). Let $s_e(t')$ be a vector chosen by MAX-WEIGHT(β). From (5)

$$\begin{aligned} C_e(t') &\ge s_e(t')(\beta+1)\left(\left(q_{u,d^{(e)}}^{t'}\right)^{\beta} - \left(q_{v,d^{(e)}}^{t'}\right)^{\beta}\right) \\ &\quad - R_{\max} O\left(\left(q_{u,d^{(e)}}^{t'}\right)^{\beta-1} + \left(q_{v,d^{(e)}}^{t'}\right)^{\beta-1}\right). \end{aligned} \tag{6}$$

From (6), we obtain that

$$\begin{aligned} \sum_{e\in E} C_e(t') &\ge \sum_{e\in E} s_e(t')(\beta+1)\left(\left(q_{u,d^{(e)}}^{t'}\right)^{\beta} - (q_{v,d^{(e)}}^{t'})^{\beta}\right) \\ &\quad - R_{\max} O\left(\left(q_{u,d^{(e)}}^{t'}\right)^{\beta-1} + \left(q_{v,d^{(e)}}^{t'}\right)^{\beta-1}\right). \end{aligned} \tag{7}$$

Thus, if we fix the time t', then the sum of potential changes at t' by MAX-WEIGHT(β) is less than or equal to the sum of potential changes at t' by $g_{p,e}(t')$. We then want to define Γ_p so that the sum of potential changes by $s_{p,(e,d)}(t')$ is equal to the sum of potential changes by $g_{p,e}(t')$.

First, we fix $t' \in W$. Let $p_1,\ldots,p_m$ be the packets injected in I^W. Let $E = \{e_1,\ldots,e_k\}$. The order of e_i's can be any possible ordering. For each $e_j \in E$, let

$$K_{e_j}(t') = \sum_{i=1}^{m} g_{p_i,(v_j,u_j)}(t')\left(\left(q_{v_j,d_i}^{t'}\right)^{\beta} - \left(q_{u_j,d_i}^{t'}\right)^{\beta}\right) \tag{8}$$

where d_i is the destination of p_i. Let

$$J(t') = \sum_{j=1}^{k} s_{e_j}(t')\left(\left(q_{v_j,d^{(e_j)}}^{t'}\right)^{\beta} - \left(q_{u_j,d^{(e_j)}}^{t'}\right)^{\beta}\right) \tag{9}$$

where $e_j = (v_j,u_j)$. At first, we define

$$s_{p_1,(e_1,d_1)}(t') = \min\left\{\frac{J(t')}{\left(q_{v_1,d_1}^{t'}\right)^{\beta} - \left(q_{u_1,d_1}^{t'}\right)^{\beta}},\; s_{e_1}(t'), \frac{K_{e_1}(t')}{\left(q_{v_1,d_1}^{t'}\right)^{\beta} - \left(q_{u_1,d_1}^{t'}\right)^{\beta}}\right\} \tag{10}$$

if $s_{e_1}(t') > 0$, and $s_{p_1,(e_1,d_1)}(t') = 0$ otherwise. Since $s_{e_1}(t')$ is chosen by MAX-WEIGHT(β), $s_{p_1,(e_1,d_1)}(t') \ge 0$. Next, let $s_{e_1}^{(1)}(t') = s_{e_1}(t') - s_{p_1,(e_1,d_1)}(t')$, $J^{(1)}(t') = J(t') - s_{p_1,(e_1,d_1)}(t')((q_{v_1,d_1}^{t'})^{\beta} - (q_{u_1,d_1}^{t'})^{\beta})$, and $K_{e_1}^{(1)}(t') = K_{e_1}(t') - s_{p_1,(e_1,d_1)}(t')((q_{v_1,d_1}^{t'})^{\beta} - (q_{u_1,d_1}^{t'})^{\beta})$.

Similarly, for all $2 \le i \le m$, we can define

$$s_{p_i,(e_1,d_i)}(t') = \min\left\{\frac{J^{(i-1)}(t')}{\left(q_{v_1,d_i}^{t'}\right)^{\beta} - \left(q_{u_1,d_i}^{t'}\right)^{\beta}},\; s_{e_1}^{(i-1)}(t'), \frac{K_{e_1}^{(i-1)}(t')}{\left(q_{v_1,d_i}^{t'}\right)^{\beta} - \left(q_{u_1,d_i}^{t'}\right)^{\beta}}\right\} \tag{11}$$

if $s_{e_1}^{(i-1)}(t') > 0$, and $s_{p_i,(e_1,d_i)}(t') = 0$ otherwise. Let $s_{e_1}^{(i)}(t') = s_{e_1}^{(i-1)}(t') - s_{p_i,(e_1,d_i)}(t')$, $J^{(i)}(t') = J^{(i-1)}(t') - s_{p_i,(e_1,d_i)}(t')((q_{v_1,d_i}^{t'})^{\beta} - (q_{u_1,d_i}^{t'})^{\beta})$, and $K_{e_1}^{(i)}(t') = K_{e_1}^{(i-1)}(t') - s_{p_i,(e_1,d_i)}(t')((q_{v_1,d_i}^{t'})^{\beta} - (q_{u_1,d_i}^{t'})^{\beta})$.

Now, from (7), we can define inductively $s_{p_i,(e_j,d_i)}$ for all $j = 2,\ldots,k$, and $i = 2,\ldots,m$, so that

$$s_{p_i,(e_j,d_i)}(t') = \min\left\{\frac{J^{((j-1)m+(i-1))}(t')}{\left(q_{v_j,d_i}^{t'}\right)^{\beta} - \left(q_{u_j,d_i}^{t'}\right)^{\beta}},\; s_{e_j}^{(i-1)}(t'), \frac{K_{e_j}^{(i-1)}(t')}{\left(q_{v_j,d_i}^{t'}\right)^{\beta} - \left(q_{u_j,d_i}^{t'}\right)^{\beta}}\right\} \tag{12}$$

if $s_{e_j}^{(i-1)}(t') > 0$, and $s_{p_i,(e_j,d_i)}(t') = 0$ otherwise, where $e_j = (v_j,u_j)$, and d_i is the destination of p_i.

Let $\Gamma_p = (s_{p_i,(e_j,d_i)}(t'))_{e_j\in E, t'\in W}$, where d_i is the destination of p_i. for each $p_i \in I^W$. We obtain that $s_{p,(e,d)}(t') \geq 0$, $\sum_{p,d} s_{p,(e,d)}(t') \leq s_e(t')$

$$\sum_{j=1}^{k}\sum_{i=1}^{m} s_{p_i,(e_j,d_i)}(t')\left(\left(q^{t'}_{v_j,d_i}\right)^\beta - \left(q^{t'}_{u_j,d_i}\right)^\beta\right) \leq J(t') \tag{13}$$

and also the following holds:

$$\sum_{j=1}^{k}\sum_{i=1}^{m} s_{p_i,(e_j,d_i)}(t')\left(\left(q^{t'}_{v_j,d_i}\right)^\beta - \left(q^{t'}_{u_j,d_i}\right)^\beta\right) = \sum_{j=1}^{k} K_{e_j}(t'). \tag{14}$$

In the previous assignment, we first defined $s_{p_1,(e_j,d_1)}$ for $j = 1,\ldots,k$. From (13) and (14), MAX-WEIGHT algorithm guarantees that the following inequalities hold:

$$\begin{aligned}
&\sum_{j=1}^{k} s_{p_i,(e_j,d_i)}(t')\left(\left(q^{t'}_{v_j,d_i}\right)^\beta - \left(q^{t'}_{u_j,d_i}\right)^\beta\right)\\
&\quad\leq \sum_{j=1}^{k}\sum_{i=1}^{m} s_{p_i,(e_j,d_i)}(t')\left(\left(q^{t'}_{v_j,d_i}\right)^\beta - \left(q^{t'}_{u_j,d_i}\right)^\beta\right)\\
&\quad= \sum_{j=1}^{k}\sum_{i=1}^{m} g_{p_i,e_j}(t')\left(\left(q^{t'}_{v_j,d_i}\right)^\beta - \left(q^{t'}_{u_j,d_i}\right)^\beta\right)\\
&\quad\leq \sum_{j=1}^{k} s_{e_j}(t')\left(\left(q^{t'}_{v_j,d^{(e_j)}}\right)^\beta - \left(q^{t'}_{u_j,d^{(e_j)}}\right)^\beta\right).
\end{aligned} \tag{15}$$

Thus, we can assign $s_{p_i,(e_1,d_i)(t')}$ for all $i = 1,\ldots,m$, so that $\sum_{i=1}^{m} s_{p_i,(e_1,d_i)(t')}((q^{t'}_{v_1,d_i})^\beta - (q^{t'}_{u_1,d_i})^\beta) = K_{e_1}(t')$. Similarly, for all $j \geq 2$, we can assign $s_{p_i,(e_j,d_i)(t')}$ for $i = 1,\ldots,m$, so that $\sum_{i=1}^{m} s_{p_i,(e_j,d_i)(t')}((q^{t'}_{v_j,d_i})^\beta - (q^{t'}_{u_j,d_i})^\beta) = K_{e_j}(t')$. Then, by taking the sum of the above inequalities, we derive that strict equality holds in (14). Hence, Γ_p is well defined by the $s_{p,(e,d)}$ values. Thus, we assigned all packet $p \in I^W$ with Γ_p so that the assigned amount of partial packet transmissions in each link at time t' is less than or equal to the amount of packet transmissions of MAX-WEIGHT(β) in each link at time t'.

For $t, t' \in W$, $|q^t_{u,d} - q^{t'}_{u,d}| \leq nR_{\max}\omega$ since at each time-slot, at most $R_{\max}$ amount of data can move along a link from u. Hence by considering ω and n and $R_{\max}$ as constants, we obtain that for any $t, t' \in W$, $(q^{t'}_{u,d})^\beta = (q^t_{u,d})^\beta + O((q^t_{u,d})^{\beta-1})$.

Suppose that a packet p with size ℓ_p is injected at a node v_0 at time $t_0 \in W$. Let d be the destination of p. Then, the potential change due to the injection of p is

$$\begin{aligned}
&\sum_{(x,y)\in E}\sum_{t'\in W} s_{p,((x,y),d)}(t')(\beta+1)\\
&\times\left|\left(q^{t_0}_{x,d}\right)^\beta - \left(q^{t_0}_{y,d}\right)^\beta + O\left(\left(q^{t_0}_{x,d}\right)^{\beta-1} + \left(q^{t_0}_{y,d}\right)^{\beta-1}\right)\right|\\
&\quad= \sum_{e=(v,u)\in\Psi_p}\sum_{t'\in W} g_{p,e}(t')(\beta+1)\\
&\quad\times\left|\left(q^{t_0}_{v,d}\right)^\beta - \left(q^{t_0}_{u,d}\right)^\beta + O\left(\left(q^{t_0}_{v,d}\right)^{\beta-1} + \left(q^{t_0}_{u,d}\right)^{\beta-1}\right)\right|\\
&\quad\geq \sum_{e=(v,u)\in\Psi_p}\frac{1}{1-\varepsilon}\left\{\sum_{t'\in W}\ell(p,e,t')\right\}(\beta+1)\\
&\quad\times\left|\left(q^{t_0}_{v,d}\right)^\beta - \left(q^{t_0}_{u,d}\right)^\beta + O\left(\left(q^{t_0}_{v,d}\right)^{\beta-1} + \left(q^{t_0}_{u,d}\right)^{\beta-1}\right)\right|\\
&\quad\geq \sum_{e=(v,u)\in\Psi_p}\frac{1-\varepsilon/2}{1-\varepsilon}\ell_p(\beta+1)\\
&\quad\times\left|\left(q^{t_0}_{v,d}\right)^\beta - \left(q^{t_0}_{u,d}\right)^\beta + O\left(\left(q^{t_0}_{v,d}\right)^{\beta-1} + \left(q^{t_0}_{u,d}\right)^{\beta-1}\right)\right|\\
&\quad\geq \sum_{e=(v,u)\in\Psi_p}\frac{\ell_p}{1-\varepsilon/2}(\beta+1)\\
&\quad\times\left|\left(q^{t_0}_{v,d}\right)^\beta - \left(q^{t_0}_{u,d}\right)^\beta + O\left(\left(q^{t_0}_{v,d}\right)^{\beta-1} + \left(q^{t_0}_{u,d}\right)^{\beta-1}\right)\right|\\
&\quad\geq \frac{1}{1-\varepsilon/2}\ell_p(\beta+1)\left(q^{t_0}_{v_0,d}\right)^\beta + \ell_p O\left(\left(q^{t_0}_{v_0,d}\right)^{\beta-1}\right)
\end{aligned}$$

because $\{\sum_{t'\in W}\ell(p,e,t')\} \geq (1-\frac{\varepsilon}{2})\ell_p$ holds for all packet p and edge e.

The increase of potential due to the direct injection of p is $\ell_p(\beta+1)(q^{t_0}_{v_0,d})^\beta + \ell_p O((q^{t_0}_{v_0,d})^{\beta-1})$. Hence, the total change of potential induced by this injection of a packet p is

$$\begin{aligned}
&\ell_p(\beta+1)\left(q^{t_0}_{v_0,d}\right)^\beta + \ell_p O\left(\left(q^{t_0}_{v_0,d}\right)^{\beta-1}\right)\\
&\quad-\sum_{(x,y)\in E}\sum_{t'\in W} s_{p,((x,y),d)}(t')(\beta+1)\left|\left(q^{t_0}_{x,d}\right)^\beta - \left(q^{t_0}_{y,d}\right)^\beta\right.\\
&\quad\left.+O\left(\left(q^{t_0}_{x,d}\right)^{\beta-1} + \left(q^{t_0}_{y,d}\right)^{\beta-1}\right)\right|\\
&\quad\leq -\frac{\varepsilon/2}{1-\varepsilon/2}\ell_p(\beta+1)\left(q^{t_0}_{v_0,d}\right)^\beta + \ell_p O\left(\left(q^{t_0}_{v_0,d}\right)^{\beta-1}\right).
\end{aligned}$$

Hence, there is a constant q^*, depending on n, ω, and ε, so that if $q \geq q^*$, the sum of potential changes due to the injection is less than $-\frac{\varepsilon}{2}\ell_p q^\beta$. ■

B. Proof of Theorem 4

Proof: Let $\varepsilon > 0$ and let $\omega \geq 1$ be some integer. Consider a general adversarial queue system $A(\omega,\epsilon,b)$ with scheduling algorithm Alg. Let n be the number of queues in this system. We will show that there is a constant $U(n,q_0,b)$ such that for $A(\omega,\epsilon,b)$, when the size of the tallest queue at time $t = 0$ is at most q_0, the sizes of all queues over all $t \geq 0$ are bounded above by $U(n,q_0,b)$.

We induct on n to show that for any $q_0 \geq 0$ and $b \geq 0$, there exists $U(n,q_0,b)$. For the basic step, when $n = 1$, there is only one queue in the system, and thus it should be a destination queue. Hence, $U(n,q_0,b)$ exists.

For the inductive step, we assume that there is $U(m,q_0,b)$ for all $1 \leq m \leq n-1$, and for all $q_0 \geq 0$ and $b \geq 0$. Using this induction hypothesis, we will show that for any q_0, $U(n,q_0,0)$ exists. We can set $U(n,q_0,1) = U(n,U(n,q_0,0)+R_{\max},0)$, because at each time when the bad packet arrives, the size of the tallest queue is at most $U(n,q_0,0)$ and we can transmit data of size at most $R_{\max}$ on each link. Similarly, for any $i \geq 1$, we can set

$$U(n, q_0, i) = U(n, U(n, q_0, i-1) + R_{\max}, 0) \tag{16}$$

by considering the time when the ith bad packet arrives. Now we only need to prove that $U(n, q_0, 0)$ exists. Let $P(t)$ be the potential of the queues at time t. Note that each injection to a queue of size at most q^* makes the potential increase by at most $(2R_{\max}q^* + R_{\max}^2)$. By Theorem 2, the maximum possible increase of potential induced by all injections during any time window of size ω is bounded by some constant P_0. Now, for a fixed n, we define the following: Let $M_n = 0$. Given M_{k+1}, for $j = 1, 2, \ldots, k$, define

$$\begin{aligned} S_j &\triangleq U\left(n-k, M_{k+1}, \frac{(j-1)}{2}(L_1 + L_2 \ldots + L_{j-1})^2\right) \\ L_j &\triangleq \frac{2(n-k)S_j^2}{\varepsilon} \\ M_k &\triangleq \frac{L_k}{R_{\min}} + S_k + \frac{2P_0}{\varepsilon R_{\min}}. \end{aligned}$$

Then, M_k, $k = 1, 2, \ldots, n$, are decreasing over k ($M_1 \gg M_2 \gg \ldots \gg M_n = 0$). We will show that for any $A(\omega, \epsilon, 0)$ for a general adversarial queue system with n queues, for all time $t \geq 0$, $P(t)$ is bounded by some value that is independent of t. More precisely, we will show that

$$\begin{aligned} P(t) \leq (n-1)M_1^2 + \max\{&nq_0^2, nM_1^2 \\ &+ 2\sqrt{n}R_{\max}M_1 + R_{\max}^2\}. \end{aligned} \tag{17}$$

Note that the right-hand side of (17) is independent of t, so we can conclude that $U(n, q_0, 0)$ exists.

Now suppose that we are given a general adversarial queue system with n queues controlled by an $A(\omega, \epsilon, 0)$ and some given scheduling algorithm Alg and a corresponding set Γ of partial transmissions assigned with packets, such that all the initial queue sizes are at most q_0.

Suppose that for all time t, $P(t) < nM_1^2$ holds. Then, it implies that the given scheduling algorithm Alg is stable and (17) is satisfied. Now suppose that there is t_0 such that $P(t_0) \geq nM_1^2$. By choosing the smallest such t_0, we may assume that $P(t_0) \leq \max\{nq_0^2, nM_1^2 + 2\sqrt{n}R_{\max}M_1 + R_{\max}^2\}$ since if $P(t_0 - 1) < nM_1^2$, the change of potential between time $t_0 - 1$ and t_0 is at most $2\sqrt{n}R_{\max}M_1 + R_{\max}^2$. Note that if $P(t_0) \geq nM_1^2$, then there is a queue of size at least M_1, and hence the size of tallest queue at that time is at least M_1.

Let $q_1 \geq q_2 \geq \cdots \geq q_n = 0$ be the ordered sizes of the queues at time t_0. For $1 \leq j \leq n$, let Q_j be the corresponding jth tallest queue at time t_0. Then, since $q_1 \geq M_1$ and $q_n = M_n = 0$, there exists some $1 \leq k \leq (n-1)$ such that $q_k \geq M_k$ and $q_{k+1} \leq M_{k+1}$. Hence, $q_k \gg q_{k+1}$ and the sizes of the small queues stay much smaller than q_k, and so the sizes of the tall queues are much bigger than those of the small queues. We will show that, for all the time afterward, the size of the $(k+1)$th tallest queue stays much smaller than M_k. A precise description will appear later.

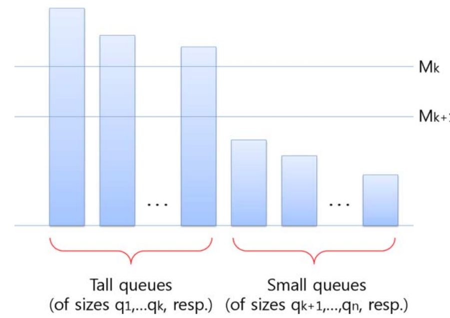

Fig. 2. All the queues having size at least M_k at time t_0 are called "*tall queues*," and all the other queues are called "*small queues*." Then, tall queues are much higher than small queues.

Now fix one such k. We will say all the queues having size at least M_k at time t_0 are "*tall queues*," and all the other queues "*small queues*" (Fig. 2). Recall that by our assumption on the MAX-WEIGHT protocol, data from a small queue will never move to a tall queue. Hence, we can consider the set of all the small queues as a separate general adversarial queue system. We will call this queue system a *system of small queues*. Afterward, we will use an inductive argument on this system of small queues to guarantee that their sizes are bounded by a constant S_j for some $1 \leq j \leq k$ during some period of time.

Let t_1 be the first time after t_0 such that there is an injection of a packet to a tall queue or a transmission of a packet from a tall queue to a small queue. Note that if there is no such $t_1 > t_0$, then for this $A(\omega, \epsilon, 0)$, the argument that will be presented in the proof of Lemma 5 shows that the sizes of all the small queues cannot be bigger than M_k for any time $t \geq t_0$. Since a packet in a small queue will never move to a tall queue, the potential of tall queues are nonincreasing over all time. Hence, we obtain that $P(t)$ is bounded by $(n-1)M_k^2 + P(t_0) \leq (n-1)M_1^2 + \max\{nq_0^2, nM_1^2 + 2\sqrt{n}R_{\max}M_1 + R_{\max}^2\}$ for all $t \geq t_0$ as required in (17).

When there is such t_1, our main argument is that during time $t_0 \leq t \leq t_1$, the system of small queues is maintained. By Lemma 3, we are able to show a net decrease in the potential in the system, as long as there are "sufficient" injections into queues that are large enough. Hence, one injection to a tall queue or one transmission of a packet from a tall queue to a small queue creates a sufficient decrease in potential. We can therefore show that the potential remains bounded as long as the increase in potential between times t_0 and t_1 is less than the decrease in potential due to the injection or transmission at time t_1. We will prove the following lemma.

Lemma 5: There is t^*, satisfying $t_0 < t^* \leq t_1 + \omega - 1$, such that $P(t^*) \leq P(t_0)$, and during $t_0 \leq t \leq t^*$ the sizes of small queues are bounded by M_k.

The proof of Lemma 5 will appear later after we conclude the proof of Theorem 4. By applying this, the potential of all the small queues, $P_S(t) \triangleq \sum_{i=k+1}^{n}(q_i^t)^2$, is bounded above by $(n-1)M_k^2 \leq (n-1)M_1^2$ since the sizes of all the small queues cannot be bigger than M_k for any time $t_0 \leq t \leq t^* - 1$. Note also that until the time $t^* - 1$, the potential of all the tall queues, $P_T(t) \triangleq \sum_{i=1}^{k}(q_i^t)^2$, is nonincreasing over time. Since at time t^* we know that the total potential $P(t^*) \leq P(t_0)$, for

$t_0 \leq t \leq t^*$, the potential $P(t)$ is bounded by $(n-1)M_1^2 + \max\{nq_0^2, nM_1^2 + 2\sqrt{n}R_{\max}M_1 + R_{\max}^2\}$. We now choose the first time $t \geq t^*$, if there exists such t, so that $P(t) \geq nM_1^2$, and set this time as a new t_0. Then, by applying the same argument, we obtain that for all time $t \geq 0$, (17) holds. Hence, $U(n, q_0, 0)$ exists. It implies (16), which in turn proves Theorem 4. ∎

C. Proof of Lemma 5

To complete the proof of Theorem 4, here we give the following proof of Lemma 5.

Proof: Note that, for all time $t_0 \leq t \leq t_1$, there may be some injection of packets to a small queue so that its corresponding set of partial transmissions includes some links between tall queues that yields the amount of potential change at least 1. We will regard these kinds of injected packets as "bad packets" for the system of small queues, and we will call these injections "bad injections." That is, each bad injection in the system of small queues makes the potential change among tall queues by at least 1. Note that by considering these packets as bad packets, the dynamics of small queues can be thought as an independent general adversarial queue system having $n-k$ queues, which means that it is a kind of subsystem of the original system. Then, essentially, we will show that the total amount of these bad injections over all time $t_0 \leq t \leq t_1$ is bounded by some number that is independent of t. Note that each bad injection in the system of small queues makes potential change among tall queues at least 1.

We will derive how we can obtain the required t^* on a case-by-case basis. We consider the following two cases:

Case I) if there is no bad injection to small queues for all time $t_0 \leq t \leq t_1$;

Case II) if there are some bad injections in that time window.

From the definition of the MAX-WEIGHT(β) algorithm, $r_e(t) \geq R_{\min}$ for each link (e, d) such that $q_{v,d}^t - q_{v,d}^t \geq 0$, so we send data along e at least $R_{\min}$ at once if we can. Without loss of generality, we can assume that $R_{\min}$ and $R_{\max}$ satisfy $R_{\min} \leq \ell_p \leq R_{\max}$ for each $p \in I^W$.

Case I: If there is no bad injection to small queues for all time $t_0 \leq t \leq t_1$, then by the induction hypothesis, for all $t_0 \leq t \leq t_1$, the sizes of small queues are bounded above by $S_1 = U(n-k, M_{k+1}, 0)$. Thus, the potential of all the small queues at time t_1 is at most $\frac{\varepsilon}{2}L_1 = (n-k)S_1^2$. By Lemma 3, the decrease of potential due to an injection to a tall queue is at least $\frac{\varepsilon}{2}R_{\min}M_k$, and the decrease of potential due to a transmission from a tall queue to a small queue at time t_1 is at least $\{(M_k)^2 - (S_1)^2\} - \{(M_k - R_{\min})^2 - (S_1 + R_{\min})^2\} = 2R_{\min}(M_k - S_1)$. Thus, the decrease of potential due to an injection to a tall queue or the decrease of potential due to a transmission from a tall queue to a small queue at time t_1 is at least $\min\{\frac{\varepsilon}{2}R_{\min}M_k, 2R_{\min}(M_k - S_1)\} \geq \frac{\varepsilon}{2}R_{\min}(M_k - S_1)$. Note that from the definition of M_k

$$\frac{\varepsilon}{2}R_{\min}(M_k - S_1) \geq \frac{\varepsilon}{2}L_1 + P_0.$$

Therefore, the decrease of potential due to an injection to a tall queue or a transmission from a tall queue to a small queue at time t_1 is at least P_0 more than the potential of all the small queues at time t_1. Note also that the maximum possible increase of the potential induced by injections during the time $[t_1, t_1 + \omega - 1]$ is bounded by P_0, and that all the packet movement associated with the injection to a tall queue at time t_1 occurs in this time window of size ω. Since there was no injection to any of the tall queues during $t_0 \leq t \leq (t_1 - 1)$, the potential of the tall queues is nonincreasing for $t_0 \leq t < t_1$. Hence, by letting $t^* = t_1 + \omega - 1$, we have $P(t^*) \leq P(t_0)$.

Case II: Suppose that there are some bad injections to small queues. Let $0 \leq r_1 \leq r_2 \leq \cdots \leq r_{k-1}$ be the ordered list of $(q_1 - q_2), (q_2 - q_3), \ldots, (q_{k-1} - q_k)$. As $\{M_1, M_2, \cdots\}$ is a set of queue thresholds, $\{L_1, L_2, \cdots\}$ defines a set of thresholds for the above list of queue differences, and $\{S_1, S_2, \cdots\}$ gives a bound on the sizes of the small queues during some period of time in the following cases. Note that these numbers are independent of t. We can divide Case II into the following three cases:

Case II-A) if $r_1 > L_1$;

Case II-B) if there is $1 \leq m < k-1$ such that for all $1 \leq j \leq m$, $r_j \leq L_j$, and $r_{m+1} > L_{m+1}$;

Case II-C) if $r_m \leq L_m$ for all $1 \leq m \leq k-1$.

Case II-A: Suppose that $r_1 > L_1$. Then, any transmission between two tall queues at some time $t_0 < t \leq t_1$ will make the decrease of potential more than L_1. Let t^* be the smallest time $t^* > t_0$ so that there is a transmission between two tall queues at time t^*. By the induction hypothesis, for all time $t_0 \leq t \leq t^*$, the sizes of the small queues are bounded by $S_1 = U(n-k, M_{k+1}, 0)$, and the potential of the small queues is bounded by $\frac{\varepsilon}{2}L_1 = (n-k)S_1^2$. Then, from the same argument as Case I, $P(t^*) \leq P(t_0)$.

Case II-B: Suppose that there is $1 \leq m < k-1$ such that for all $1 \leq j \leq m$, $r_j \leq L_j$, and $r_{m+1} > L_{m+1}$. We will show that the potential of all the small queues is bounded by $\frac{\varepsilon}{2}L_{m+1}$. We may assume that bad injections to small queues induce transmissions just between neighboring tall queues. Note also that the amount of bad injections to small queues during some period of time is bounded by the total amount of transmissions between tall queues during that period of time.

We say a link $e_j = (Q_j, Q_{j+1})$ between two neighboring tall queues is a *tall link* if $q_j - q_{j+1} > L_{m+1}$ and a *small link* otherwise. We can divide Case II-B into the following two cases:

Case II-B-1) if there is no transmission via tall links for all time $t_0 \leq t \leq t_1$;

Case II-B-2) if there is a transmission via some tall link for some time $t_0 < t \leq t_1$.

We will use the following lemma.

Lemma 6: Let $r_1, r_2, \ldots, r_m$ be the sizes of the small links at time t_0 and assume that $r_j \leq L_j$ for all $1 \leq j \leq m$. If there is no transmission via *tall links* for $t_0 \leq t < t'$ and all the transmissions occur via small links, then the total amount of packet transmissions via *small links* during that period of time is bounded by

$$\frac{m}{2}(r_1 + r_2 + \cdots + r_m)^2 \leq \frac{m}{2}(L_1 + L_2 \ldots + L_m)^2.$$

Proof: Let $e_{j_1}, e_{j_2}, \ldots e_{j_m}$ be the set of small links, where $j_1 < j_2 < \cdots < j_m$. For $1 \leq i \leq m$, let s_i be $(q_{j_i} - q_{j_i+1})$. Hence, $\{s_i\}_{1 \leq i \leq m}$ is a permutation of $\{r_i\}_{1 \leq i \leq m}$.

Recall that the sizes of the queues at time t_0 are nonincreasing with respect to their indices. Moreover, note that if $j_{i+1}-j_i \geq 2$ for some i, then any packet p that was originally located at Q_m, with $m \leq j_i+1$ cannot move to Q_{j_i+2} for all time $t_0 \leq t \leq t'$. Hence, we can consider each subset of consecutive small links separately. For example, if $j_1,\dots,j_m$ are 2, 3, 5, 6, 7, then we will consider 2, 3 and 5, 6, 7 separately. Suppose that $j_1, j_2 \dots, j_s$ are consecutive integers. Since $q^{t'}_{j_1}+\cdots+q^{t'}_{j_{s+1}} = q^{t_0}_{j_1}+\cdots+q^{t_0}_{j_{s+1}}$, we obtain that

$$\left(q^{t'}_{j_1}\right)^2+\cdots+\left(q^{t'}_{j_{s+1}}\right)^2 \geq \sum_{i=1}^{s+1}\left(\frac{q^{t_0}_{j_1}+\cdots+q^{t_0}_{j_{s+1}}}{s+1}\right)^2.$$

Thus, the amount of packet transmission via $e_{j_1},\dots,e_{j_s}$ is

$$\begin{aligned}
&\left\{\left(q^{t_0}_{j_1}\right)^2+\cdots+\left(q^{t_0}_{j_{s+1}}\right)^2\right\}-\left\{\left(q^{t'}_{j_1}\right)^2+\cdots+\left(q^{t'}_{j_{s+1}}\right)^2\right\}\\
&\leq \sum_{i=1}^{s+1}(q^{t_0}_{j_i})^2-(s+1)\left(\frac{q^{t_0}_{j_1}+\cdots+q^{t_0}_{j_{s+1}}}{s+1}\right)^2\\
&=\frac{1}{s+1}\left(s\sum_{i=1}^{s+1}\left(q^{t_0}_{j_i}\right)^2-2\sum_{1\leq i<k\leq s+1}q^{t_0}_{j_i}q^{t_0}_{j_k}\right)\\
&=\frac{1}{s+1}\left(\sum_{1\leq i<k\leq s+1}\left(q^{t_0}_{j_i}-q^{t_0}_{j_k}\right)^2\right)\\
&\leq \frac{1}{s+1}\binom{s+1}{2}(r_1+\cdots+r_{s+1})^2\\
&=\frac{s}{2}(r_1+\cdots+r_{s+1})^2\\
&\leq \frac{s}{2}(L_1+\cdots+L_{s+1})^2.
\end{aligned}$$

A similar argument holds for other consecutive indices, separately. Hence, the sum of total amount of transmissions via small links during time $t_0 \leq t \leq t'$ is bounded by $\frac{m}{2}(L_1+L_2+\cdots+L_m)^2$. ■

Case II-B-1: If there is no transmission via tall links for all time $t_0 \leq t \leq t_1$, then by Lemma 6, the total amount of bad injections to the small queues during $t_0 \leq t \leq t_1$ is bounded by $\frac{m}{2}(L_1+L_2+\cdots+L_m)^2$. Since each bad injection in the system of small queues makes the potential change among tall queues by at least 1, we conclude that the number of bad packets to the system of small queues is also at most $\frac{m}{2}(L_1+L_2+\cdots+L_m)^2$. Therefore, for all time $t_0 \leq t \leq t_1$, the sizes of the small queues are bounded by

$$S_{m+1}=U\left(n-k, M_{k+1}, \frac{m}{2}(L_1+L_2+\cdots+L_m)^2\right)$$

by the induction hypothesis. Hence, the potential of all the small queues at time t_1 is at most $\frac{\varepsilon}{2}L_{m+1}=(n-k)S^2_{m+1}$.

Note that the potential for the tall queues is nonincreasing for $t_0 \leq t \leq t_1$. By Lemma 3, the decrease in potential due to an injection to a tall queue is at least $\frac{\varepsilon}{2}R_{\min}M_k$, and the decrease in potential due to a transmission from a tall queue to a small queue at time t_1 is at least $2R_{\min}(M_k-S_{m+1})$. Thus, the decrease of potential due to an injection to a tall queue or the decrease of potential due to a transmission from a tall queue to a small queue at time t_1 is at least $\min\{\frac{\varepsilon}{2}R_{\min}M_k, 2R_{\min}(M_k-S_{m+1})\} \geq \frac{\varepsilon}{2}R_{\min}(M_k-S_{m+1})$. Note that from the definition of M_k

$$\frac{\varepsilon}{2}R_{\min}(M_k-S_{m+1}) \geq \frac{\varepsilon}{2}L_{m+1}+P_0.$$

Therefore, the decrease of the potential at time t_1 is at least P_0 more than the potential of all the small queues at t_1. By letting $t^*=t_1+\omega-1$, we have $P(t^*) \leq P(t_0)$.

Case II-B-2: If there is a transmission via some tall link for some time $t_0 < t \leq t_1$, let t^* be the smallest such t. Then, similarly, by Lemma 6, the total amount of bad injections to the small queues during $t_0 \leq t \leq t^*$ is bounded by $\frac{m}{2}(L_1+L_2+\cdots+L_m)^2$. Hence, the sizes of the small queues during this time interval are bounded by S_{m+1} by the induction hypothesis and from the definition of t_1, so the potential of all the small queues at time t^* is at most $\frac{\varepsilon}{2}L_{m+1}=(n-k)S^2_{m+1}$. Moreover, during $t_0 \leq t \leq t^*$, for any tall link $e_j=(Q_j,Q_{j+1})$, q_j is nondecreasing and q_{j+1} is nonincreasing because any transmission via small links can make q_j bigger (when e_{j-1} is a small link), or q_{j+1} smaller (when e_{j+1} is a small link), but it cannot increase q_j-q_{j+1}. Thus, $q_j-q_{j+1} \geq L_{m+1}$ at $t=t^*$. Hence, a transmission via a tall link at time t^* will make the potential decrease by at least $\frac{\varepsilon}{2}L_{m+1}$, which is more than the potential of all the small queues at time t^*. Note also that the potential for the tall queues is nonincreasing for $t_0 \leq t < t^*$. Hence, we have $P(t^*) \leq P(t_0)$.

Case II-C: Finally, consider the case when $r_m \leq L_m$ for all $1 \leq m \leq k-1$. Then, by Lemma 6 and the induction hypothesis, for all time $t_0 \leq t \leq t_1$, the sizes of small queues are bounded by

$$S_k=U\left(n-k, M_{k+1}, \frac{(k-1)}{2}(L_1+L_2+\cdots+L_{k-1})^2\right).$$

Hence, the potential of all the small queues at time t_1 is at most $\frac{\varepsilon}{2}L_k=(n-k)S_k^2$. By Lemma 3, the decrease of the potential due to an injection to a tall queue is at least $\frac{\varepsilon}{2}R_{\min}M_k$, and the decrease of the potential due to a transmission from a tall queue to a small queue at time t_1 is at least $2R_{\min}(M_k-S_k)$. Thus, the decrease of the potential due to an injection to a tall queue or the decrease of the potential due to a transmission from a tall queue to a small queue at time t_1 is at least $\min\{\frac{\varepsilon}{2}R_{\min}M_k, 2R_{\min}(M_k-S_k)\} \geq \frac{\varepsilon}{2}R_{\min}(M_k-S_k)$. Then, from the definition of M_k, $\frac{\varepsilon}{2}R_{\min}(M_k-S_k)=\frac{\varepsilon}{2}L_k+P_0$, which is at least P_0 more than the potential of all the small queues at time t_1. Note also that the potential of the tall queues is nonincreasing for $t_0 \leq t < t_1$. Hence, by letting $t^*=t_1+\omega-1$, we have $P(t^*) \leq P(t_0)$.

Hence, in all the cases, we have $P(t^*) \leq P(t_0)$, and for $t_0 \leq t \leq t^*$, the sizes of small queues are bounded by $S_j+\omega nR_{\max}$ for some $1 \leq j \leq k$, so they are bounded by M_k. ■

V. Characterization of the Queue Sizes

We now consider the behavior of the queue sizes under the adversarial model. In the case of a stationary stochastic network, the typical "negative drift" argument that we described earlier essentially shows that the potential in the system cannot grow much larger than $(n^2\varepsilon^{-1}(R_{\max})^2)^2$. More precisely, if the potential ever does get larger than that amount, then some queue size must be larger than $n^2\varepsilon^{-1}(R_{\max})^2$. At that point, the expression for the change in network potential implies the expected drift in potential is nonpositive. One consequence of this is that whenever an individual queue size becomes larger than $(n^2\varepsilon^{-1}(R_{\max})^2)^2$, the expected drift in potential is nonpositive.

In contrast, for the MAX-WEIGHT protocol in the adversarial model, the bound on queue size implied by the analysis of Section IV is actually exponential in the number of users n. We now briefly show that this is necessary. In particular, we present an example where the MAX-WEIGHT protocol does indeed give rise to exponentially sized queues. Our example is close to an example given in [6] in which it was shown that we can get exponential queue sizes in a critically loaded scenario (i.e., where $\varepsilon = 0$). We now show that this is actually possible in a subcritically loaded example (with $\varepsilon > 0$).

We consider a set of N single-hop edges (numbered $0, \dots, N-1$) that are all mutually interfering, i.e., only one edge can transmit data at a time. Let $a_i(t)$ be the amount of data injected for edge i at time t, and let $r_i(t)$ be the edge rate. The adversary defines these quantities in the following simple manner. At any given time t, let $i' = \min\{i : q_i(t) < (1-\varepsilon)2^i\}$. If $i' = 0$, then the adversary sets $r_0(t) = 1$ and $a_0(t) = 1-\varepsilon$. If $i' > 0$, then it sets $r_{i'-1}(t) = 1-\varepsilon$, $r_{i'}(t) = \frac{1-\varepsilon}{2}$, and $a_{i'}(t) = \frac{(1-\varepsilon)^2}{2}$. In both cases, all other $r_i(t)$ and $a_i(t)$ values are set to 0. It is clear that these definitions are consistent with an $A(1,\varepsilon)$ adversary.

Lemma 7: With the above patterns of data arrivals and edge rates, for each t and for each i, there exists a $t' \geq t$ such that $q_i(t') \geq (1-\varepsilon)2^i$.

Proof: We prove the above statement by induction on i. Suppose that $q_0(t) < 1-\varepsilon$. Then, for this time step, i' is set to 0, and so $a_0(t) = 1-\varepsilon$. Once data have been served for edge 0 and the arriving data have been added to the edge's queue, we have $q_0(t+1) \geq 1-\varepsilon$. (Note that this assumes that data arrive in a queue after data have been served. This is a reasonable assumption, but if it does not hold, then we can simply set $\omega \geq 2$ and have all the arrivals in a window of length ω arrive at the beginning of the window.) This completes the base case.

For the inductive step, suppose that $q_i(t) < (1-\varepsilon)2^i$ for an $i > 0$. The inductive hypothesis implies that there exists some time $t' \geq t$ at which $i' = i$. Suppose that t' is the first such time step. Between t and t', note that we must have $i' < i$, and so the value of q_i does not change. When we reach time step t', it must be the case that $q_i(t') < 2q_{i-1}(t')$. Moreover, by the definition of the edge rates, $r_{i-1}(t') = 1-\varepsilon$ and $r_i(t') = \frac{1-\varepsilon}{2}$. Hence, the MAX-WEIGHT protocol serves queue $i-1$, but the arrivals are for queue i. Hence, $q_i(t')$ is strictly greater than $q_i(t)$. By repeating this process, we eventually reach a time t'' at which $q_i(t'') \geq (1-\varepsilon)2^i$.

By the inductive hypothesis, there must be a time $t''' \geq t''$ for which $q_j(t''') \geq (1-\varepsilon)2^j$ for all $j \leq i-1$. Between times t'' and t''', the value of q_i cannot decrease. Hence, at time t''', we have $q_j(t''') \geq (1-\varepsilon)2^j$ for all $j \leq i$. The inductive step is complete. ■

Corollary 8: There exists a network configuration with N edges and an $A(1,\varepsilon)$ adversary such that some queue grows to size $(1-\varepsilon)2^{N-1}$.

We remark in conclusion that with a different routing and scheduling protocol, adversarial models do not necessarily lead to exponentially large queues. In [6] another protocol was presented (which directly keeps track of the past history of edge rates and arrivals) which ensures a maximum queue size of $O(\omega k|\mathcal{R}|^2 R_{\max})$, where $\mathcal{R}$ is the set of feasible rate values. However, we still believe that it is of interest to study the performance and stability of the MAX-WEIGHT protocol in adversarial networks since it is extremely simple to implement and it has been proposed so many times in the literature as a solution to the scheduling problem in wireless networks.

VI. STABILITY OF APPROXIMATE MAX-WEIGHT

As remarked in the Introduction, computing the exact Max-Weight set of feasible transmissions is in general an NP-hard problem. Hence, a natural question to ask is what can be achieved if at each time step we only find an approximate Max-Weight set of feasible transmissions. In this section, we address this question.

Recall that $A(\omega,\epsilon)$ assures that there is a set of fractional movement of packets Ψ_p for each $p \in I^W$ and there is an edge rate vector $r_e \in R(t)$ for each $t \in W$, so that each edge is used at most $(1-\varepsilon)$ times the sum of rates associated at e during the time window W. Thus, it guarantees that each edge e can transmit more data than is actually required by a $\frac{1}{1-\varepsilon}$ factor. Hence, the actual packet movement by MAX-WEIGHT induces potential changes that are *about* $\frac{1}{1-\varepsilon}$ times greater than necessary.

For an optimization problem, an ε-approximation algorithm is an algorithm that provides an approximate solution within $(1 \pm \varepsilon)$ factor of the optimal solution. Although computing the optimal solution of MAX-WEIGHT is computationally very hard, in many practical wireless networks, an ε-approximate solution can be computed in polynomial time. For example, [10] presented an ε-approximate solution to find the maximum weight independent set (MWIS) on planar graphs, and this was extended by several authors to more general classes of graphs. In [12] and [13], an ε-approximate MWIS for a large class of wireless networks in Euclidean space and polynomially growing graphs[2] [13] is provided. In our model, we assume an ε-approximate MAX-WEIGHT computes an ε-approximate solution $r'(t)$ for each time t so that the potential decrease is at least $(1-\varepsilon)$ times the maximum possible potential decrease at time t. We will prove the stability of any $\hat{\varepsilon}$-approximate MAX-WEIGHT protocol under any $A(\omega,\epsilon)$ for $\varepsilon > 0$, if $0 < \hat{\varepsilon} < \varepsilon$.

Theorem 9: For $0 < \hat{\varepsilon} < \varepsilon$, any $\hat{\varepsilon}$-approximate MAX-WEIGHT(β) is stable under any $A(\omega,\epsilon)$.

Proof: Let $\tilde{\varepsilon} = \frac{\varepsilon-\hat{\varepsilon}}{1-\hat{\varepsilon}}$, then $A(\omega,\epsilon)$ is a $A(\omega,\tilde{\varepsilon})$ since $0 < \hat{\varepsilon} < \varepsilon$. As in the statement of Theorem 2, if we can associate the injection of $p \in I^j \cup I^{j-1}$ with a set of partial transmissions Γ'_p, so that the sum of potential changes due to this injection to a queue of height $q \geq q^*$ is less than $-\frac{\tilde{\varepsilon}}{2}(\beta+1)\ell_p q^\beta$, then all the other arguments in the proof of Theorem 4 hold when we replace ε with $\tilde{\varepsilon}$.

[2]A sequence of graphs is said to be polynomially growing if there are constants $C, \rho > 0$ such that for all vertex v and $r > 0$, the number of vertices whose shortest path distance from v is within r is at most $C \cdot r^\rho$.

As in the proof of Theorem 2, we define $d'_{p,e}(t') = \frac{1-\varepsilon}{1-\tilde{\varepsilon}} g_{p,e}(t')$ for all $e \in E$, $t' \in W$, and $p \in I^j \cup I^{j-1}$. Then, we show that $d'_{p,e}(t')$ satisfy (IV-A) if we substitute ε by $\tilde{\varepsilon}$. As in the proof of Theorem 2, we define

$$\begin{aligned} K'_{e_j}(t') &= \sum_{i=1}^{m} d'_{p_i,(v_j,u_j)}(t') \left(\left(q^{t'}_{v_j,d_i}\right)^{\beta} - \left(q^{t'}_{u_j,d_i}\right)^{\beta} \right) \\ &= (1-\hat{\varepsilon}) K_{e_j}(t'). \end{aligned} \tag{18}$$

By the definition of $\hat{\varepsilon}$-approximate MAX-WEIGHT, we can take $\hat{s}_{e_j}(t')$ for each $e_j = (v_j, u_j) \in E$, $t' \in W$ such that

$$\begin{aligned} J'(t') &:= \sum_{j=1}^{k} \hat{s}_{e_j}(t) \left\{ \left(q^t_{v_j,d_a^{(e_j)}}\right)^{\beta} - \left(q^t_{u_j,d_a^{(e_j)}}\right)^{\beta} \right\}. \\ &\geq \sum_{j=1}^{k} (1-\hat{\varepsilon}) s_{e_j}(t) \left\{ \left(q^t_{v_j,d^{(e_j)}}\right)^{\beta} - \left(q^t_{u_j,d^{(e_j)}}\right)^{\beta} \right\} \end{aligned} \tag{19}$$

for some destinations $d_a^{(e_j)}$ for each e_j. By (18) and (19), we can recursively assign

$$\hat{s}_{p_i,(e_j,d_i)}(t') = \min \left\{ \frac{J'^{((j-1)m+(i-1))}(t')}{\left(q^{t'}_{v_j,d_i}\right)^{\beta} - \left(q^{t'}_{u_j,d_i}\right)^{\beta}}, \hat{s}^{(i-1)}_{e_j}(t'), \frac{K'^{(i-1)}_{e_j}(t')}{\left(q^{t'}_{v_j,d_i}\right)^{\beta} - \left(q^{t'}_{u_j,d_i}\right)^{\beta}} \right\}$$

where $e_j = (v_j, u_j)$ and d_i is the destination of p_i, in the same manner as in the proof of Theorem 2. Then, for all $e_j \in E$, $t' \in W$, and $p_i \in I^j \cup I^{j-1}$

$$\begin{aligned} &\sum_{e \in E} \sum_{p,d} \hat{s}_{p,(e,d)}(t') \left(\left(q^{t'}_{v,d}\right)^{\beta} - \left(q^{t'}_{u,d}\right)^{\beta} \right) \\ &\qquad \leq \sum_{e \in E} \hat{s}_e(t') \left(\left(q^{t'}_{v,d^{(e)}}\right)^{\beta} - \left(q^{t'}_{u,d^{(e)}}\right)^{\beta} \right). \end{aligned}$$

Let $\Gamma'_{p_i} = (\hat{s}_{p_i,e_j}(t'))_{e_j \in E, t' \in W}$ for each $p_i \in I^j \cup I^{j-1}$.

We obtain that the sum of potential changes due to this injection is less than $-\frac{\tilde{\varepsilon}}{2} \ell_{p_i} (\beta+1) q^{\beta}$ by using the same argument as in Section IV-A. This in turn implies that $\hat{\varepsilon}$-MAX-WEIGHT(β) algorithm is stable under $A(\omega, \epsilon)$. ■

VII. EXPERIMENTS

A. Simulation Setup

We now describe a numerical experiment that aims to understand the queue size dynamics of the MAX-WEIGHT protocol under the adversarial model. Consider an $n_1 \times n_2$ simple grid graph G, and let $n = n_1 n_2$. Then, there are $4n - 2n_1 - 2n_2$ directed edges in the graph. We assume that all single nodes can be a destination. We let $n_1 = 3, n_2 = 4$, so $n = 12$, and $4n - 2n_1 - 2n_2 = 34$.

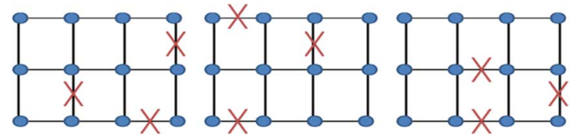

Fig. 3. Underlying graphs express which edges are not available under $r^{(1)}, r^{(2)}, r^{(3)}$.

In our simulation, we used three different edge rate vectors $r^{(1)}, r^{(2)}, r^{(3)} \in \mathbb{R}^{34}$ for G. For each $r^{(i)}$, $1 \leq i \leq 3$, we select three edges among 17 possible undirected edges and remove them. The underlying graphs of $r^{(1)}, r^{(2)}, r^{(3)}$ are described in Fig. 3. Other directed edges have edge rates chosen independently and uniformly at random from [0.5, 2]. We used the node-exclusive interference model [14], i.e., matching constraint model.[3]

Among $n(n-1)$ many distinct source–destination pairs (S-D pairs), we randomly chose K many S-D pairs $(s_1, d_1), \ldots, (s_K, d_K)$ for $K = 10$. We fix the set of feasible edge rate vectors R for all time $t \geq 0$, and we define the feasible *arrival rate* as follows. The collection of all the feasible arrival rate vectors S is called the network stability region.

Definition 4: The *arrival rate vector* $\gamma = (\gamma_1, \ldots, \gamma_K) \in [0,1]^K$ corresponding to the S-D pairs $(s_1, d_1), \ldots, (s_K, d_K)$ is said to be *feasible* if there exist flows, $(f^1, \ldots, f^K)$ such that the following applies.

1) For each $1 \leq j \leq K$, f^j routes a flow of at least γ_j from s_j to d_j.
2) The induced net flow on the directed edges, $\hat{f} = \sum_{i=1}^{K} f^j$ the interior of $co(R)$, where $co(R)$ is the convex hull of R.

If an arrival rate vector is in S, i.e., feasible, and the arrivals are identical for all time, then MAX-WEIGHT is stable [12]. Moreover, if an arrival rate vector is not contained in $\bar{S}$, where $\bar{S}$ is the closure of S, then MAX-WEIGHT is unstable. We chose K many source–destination pairs at random. For each $r^{(i)}$,, $1 \leq i \leq 3$, we computed three different feasible arrival rate vectors that are close to the boundary of the network stability region. To do so, we fixed random arrival rate vectors $\gamma^{(1)}, \gamma^{(2)}, \gamma^{(3)}$ such that each entry has a value from [0.5, 2]. We computed constants c_{ij}, by binary search, for edge rate vector $r^{(i)}$, and arrival rate vector $\gamma^{(j)}$ so that $c_{ij}\gamma^{(j)}$ is stable under MAX-WEIGHT, and $(c_{ij}+0.001)\gamma^{(j)}$ is not stable under MAX-WEIGHT, as described in Fig. 4. Each c_{ij} varied from 0.098 to 0.178 in our simulation. We used a sufficiently large time window of size 10^6 so that we could check the stability.

We did three experiments. In all three experiments, we divided the time $t \geq 0$ into nonoverlapped subwindows of ordered phases. The first phase is $t \in [1, \lceil 1.5 \rceil]$, the second phase is $t \in [\lceil 1.5 \rceil + 1, \lceil 1.5 + (1.5)^2 \rceil]$, and for each $i \geq 1$, the ith phase is $t \in [\lceil \sum_{j=1}^{i} (1.5)^{j-1} \rceil + 1, \lceil \sum_{j=1}^{i} (1.5)^j \rceil]$.

In the first experiment, we fixed the edge rate vector $r^{(i)}$ for some $i \in \{1, 2, 3\}$. Over time, the adversary injects packets as follows. For $t \geq 0$, if t is in the jth phase, then inject

[3]In the matching constraint model, the only way in which two edges could interfere is by sharing a common node.

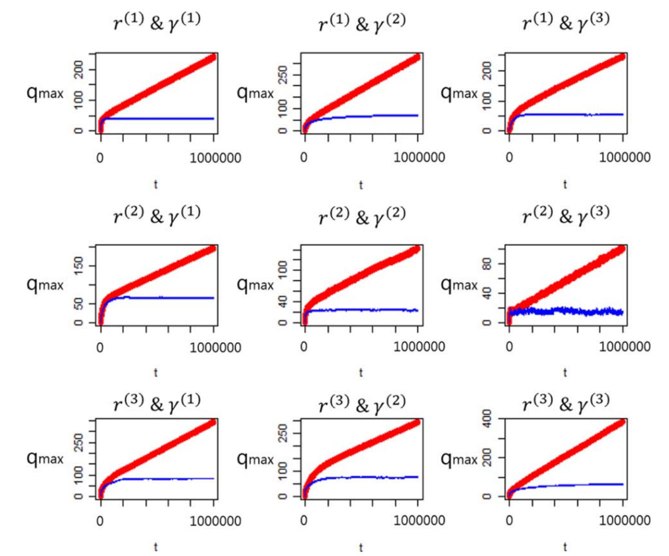

Fig. 4. For each pair of edge rate and arrival rate vector, the plot represents the change of the maximum size of queues for $c_{ij}\gamma^{(j)}$ and $(c_{ij}+0.01)\gamma^{(j)}$ in the time window $[0, 10^6]$.

packets with an arrival rate $c_{i\bar{j}}\gamma^{(\bar{j})}$ where $\bar{j} \in \{1,2,3\}$ and $\bar{j} \equiv j \pmod 3$. We call the above arrival rate vector a cyclic arrival rate vector.

In the second experiment, over time the adversary determines edge rate vectors and packet arrivals as follows. For $t \geq 0$, if t is in the ith phase, we assign an edge rate vector $r^{(\bar{i})}$ where $\bar{i} \in \{1,2,3\}$ and $\bar{i} \equiv i \pmod 3$, and we assign an arrival rate vector $c_{\bar{i}\tilde{i}}\gamma^{(\tilde{i})}$ where $\tilde{i} \in \{1,2,3\}$ and $\tilde{i} \equiv \lfloor (i - \lceil i/9 \rceil)/3 \rfloor + 1 \pmod 9$. We call the above rate vector a cyclic edge and arrival rate vector.

In the third experiment, we follow the same setup of the second experiment with an $\hat{\varepsilon}$-approximate MAX-WEIGHT, where the condition $0 < \hat{\varepsilon} < \varepsilon$ of Theorem 9 is satisfied. Specifically, we set $\varepsilon = 0.1$ and $\hat{\varepsilon} = 0.05$. Then, we use an $\hat{\varepsilon}$-approximate MAX-WEIGHT protocol that uses the same algorithm with the MAX-WEIGHT protocol, but sends data of size $(1-\hat{\varepsilon})s_e(t)$ along e at each time t.

Notice that, in all three experiments, the average of the arrival rate vectors until time T does not converge as T goes to infinity. Also in the second and third experiments, the same holds for the edge rate vectors. Hence, all setups cannot be expressed by any stationary stochastic processes. However, the above injections satisfy the definition of $A(\omega, \varepsilon)$ for some $\omega > 0$ and a small $\varepsilon > 0$. In all setups, we observed the dynamics of the maximum queue sizes over time.

B. Simulation Results

For the first experiment, as Fig. 5 shows, for each edge rate vector, MAX-WEIGHT is stable with a cyclic arrival rate vector. Interestingly, the maximum queue size may increase in some subwindow, but it decreases rapidly when the new subwindow starts. This is because the congested edges are different for each arrival rate vector, and the traffic congestions are resolved when the arrival rate is changed. Notice that the maximum queue sizes for the cyclic arrival case are bounded above and bounded below by some fixed arrival rate vector cases, respectively.

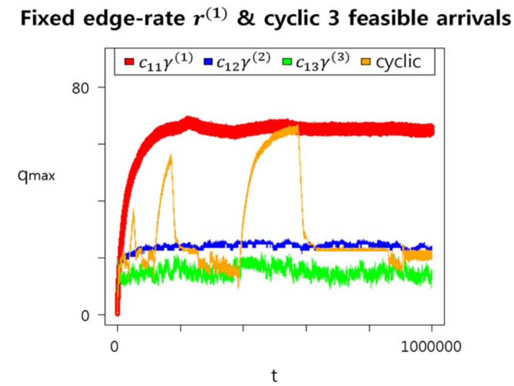

Fig. 5. For the edge rate vector $r^{(1)}$, we plot the maximum queue size when we use fixed arrival rate vectors $c_{11}\gamma^{(1)}$, $c_{12}\gamma^{(2)}$, $c_{13}\gamma^{(3)}$, and a cyclic arrival rate vector.

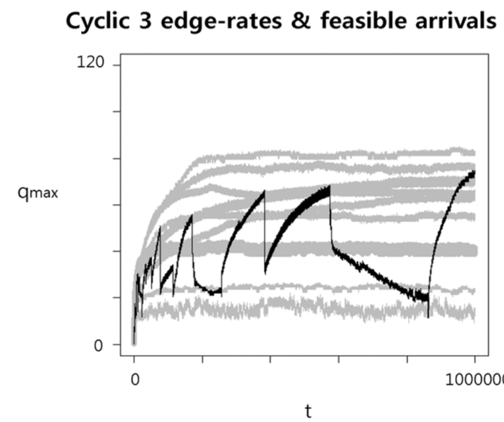

Fig. 6. We use a cyclic edge and arrival rate in the second experiment. It shows the stability of MAX-WEIGHT.

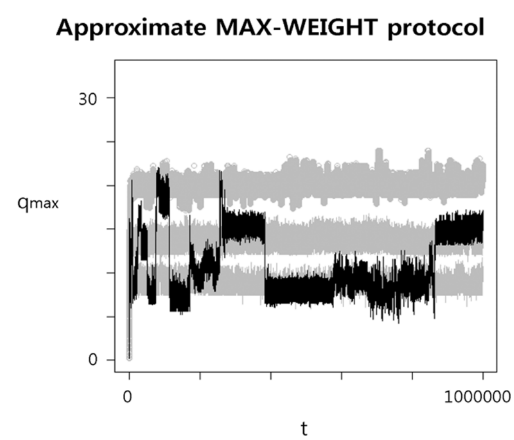

Fig. 7. The result shows that maximum queue sizes are bounded over time for approximate MAX-WEIGHT protocols as stated in Theorem 9.

The queue dynamics for the second experiment are described in Fig. 6. The gray lines describe queue sizes for nine possible fixed edge and arrival rate vectors. The black line describes the queue size for the cyclic edge and arrival rate vector case. Again, the maximum queue sizes for this case are bounded above and bounded below by some fixed edge and arrival rate vector cases, respectively. From our two experiments, we observe that MAX-

WEIGHT makes the system stable under $A(\omega, \varepsilon)$ even when the edge and arrival rate vectors do not converge over time.

In the third experiment, the queue dynamics are described in Fig. 7. The gray lines describe queue sizes for three pairs of fixed edge and arrival rate vectors among nine possible pairs. The result verifies that the stability results of the MAX-WEIGHT protocol can be extended to approximate MAX-WEIGHT as stated in Theorem 9.

VIII. CONCLUSION

In this paper, we have shown that the MAX-WEIGHT protocol remains stable even when the traffic arrivals and edge rates are determined in an adversarial manner.

In our opinion, the most natural open question concerns the bound on queue size. Our analysis gives a bound that is exponential in the network size, and we have shown in Section V that such a bound is unavoidable in the general case. However, achieving these large queue sizes involves choosing the achievable edge rate vectors $R(t)$ in a very specific manner. We are interested in whether there are any simple sufficient conditions on the sets $R(t)$, which would ensure that such large queues do not occur.

REFERENCES

[1] W. Aiello, E. Kushilevitz, R. Ostrovsky, and A. Rosen, "Adaptive packet routing for bursty adversarial traffic," in *Proc. 30th Annu. ACM Symp. Theory Comput.*, Dallas, TX, USA, May 1998, pp. 359–368.

[2] M. Andrews, K. Jung, and A. Stolyar, "Stability of the max-weight routing and scheduling protocol in dynamic networks and at critical loads," in *Proc. 39th Annu. ACM Symp. Theory of Comput.*, Jun. 2007, pp. 145–154.

[3] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting, "CDMA data QoS scheduling on the forward link with variable channel conditions," Bell Labs Tech. Memo., Apr. 2000.

[4] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting, "Providing quality of service over a shared wireless link," *IEEE Commun. Mag.*, vol. 39, no. 2, pp. 150–154, Feb. 2001.

[5] M. Andrews and L. Zhang, "Scheduling over a time-varying user dependent channel with applications to high speed wireless data," in *Proc. 43rd Symp. Found. Comput. Sci.*, Nov. 2002, pp. 293–302.

[6] M. Andrews and L. Zhang, "Scheduling over non-stationary wireless channels with finite rate sets," in *Proc. IEEE INFOCOM*, Mar. 2004, pp. 1694–1704.

[7] E. Anshelevich, D. Kempe, and J. Kleinberg, "Stability of load balancing algorithms in dynamic adversarial systems," in *Proc. 34th Annu. ACM Symp. Theory Comput.*, Montreal, QC, Canada, May 2002, pp. 399–406.

[8] B. Awerbuch and T. Leighton, "A simple local-control approximation algorithm for multicommodity flow," in *Proc. 34th Annu. Symp. Found. Comput. Sci.*, Nov. 1993, pp. 459–468.

[9] B. Awerbuch and T. Leighton, "Improved approximation algorithms for the multi-commodity flow problem and local competitive routing in dynamic networks," in *Proc. 26th Annu. ACM Symp. Theory Comput.*, 1994, pp. 487–496.

[10] B. S. Baker, "Approximation algorithms for np-complete problems on planar graphs," *J. ACM*, vol. 41, no. 1, pp. 153–180, Jan. 1994.

[11] L. Georgiadis, M. J. Neely, and L. Tassiulas, "Esource allocation and cross-layer control in wireless networks," *Found. Trends Netw.*, vol. 1, no. 1, pp. 1–149, 2006.

[12] R. Gummadi, K. Jung, D. Shah, and R. Sreenivas, "Feasible rate allocation in wireless networks," in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 995–1003.

[13] R. Gummadi, K. Jung, D. Shah, and R. Sreenivas, "Computing capacity region of a wireless network," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 1341–1349.

[14] B. Hajek and G. Sasaki, "Link scheduling in polynomial time," *IEEE Trans. Inf. Theory*, vol. 34, no. 5, pp. 910–917, Sep. 1988.

[15] N. W. McKeown, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switch," in *Proc. IEEE INFOCOM*, San Francisco, CA, USA, Mar. 1996, pp. 296–302.

[16] S. Muthukrishnan and R. Rajaraman, "An adversarial model for distributed dynamic load balancing," in *Proc. 10th ACM SPAA*, 1998, pp. 47–54.

[17] M. Neely, E. Modiano, and C. Rohrs, "Power and server allocation in a multi-beam satellite with time varying channels," in *Proc. IEEE INFOCOM*, New York, NY, USA, Jun. 2002, pp. 1451–1460.

[18] D. Shah and D. Wischik, "Optimal scheduling algorithms for input queued switches," in *Proc. IEEE INFOCOM*, Barcelona, Spain, Apr. 2006, pp. 1–11.

[19] G. Sharma, R. R. Mazumdar, and N. B. Shroff, "On the complexity of scheduling in wireless networks," in *Proc. 12th Annu. MobiCom*, Sep. 2006, pp. 227–238.

[20] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.

[21] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. Inf. Theory*, vol. 39, no. 2, pp. 466–478, Mar. 1993.

[22] P. van de Ven, S. Borst, and S. Shneer, "Instability of maxweight scheduling algorithm," in *Proc. IEEE INFOCOM*, Rio de Janeiro, Brazil, Apr. 2009, pp. 1701–1709.

Sungsu Lim (S'12) received the B.S. and M.S. degree in mathematical sciences from KAIST, Daejeon, Korea, in 2009 and 2011, respectively, and is currently pursuing the Ph.D. degree in knowledge service engineering at KAIST.

His research interests include mining and modeling large-scale social and information networks and their theoretical aspects.

Kyomin Jung (M'12) received the B.Sc. degree from Seoul National University, Seoul, Korea, in 2003, and the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2009, both in mathematics.

He has been an Assistant Professor with the Electrical and Computer Engineering Department, Seoul National University, since 2013. He worked with the Computer Science Department, KAIST, Daejeon, Korea, since 2009, where he had joint appointments with the Department of Electrical Engineering and the Department of Mathematics. During his Ph.D. studies, he worked with Microsoft Research Cambridge, Cambridge, U.K., in 2008; IBM T. J. Watson Research, Yorktown Heights, NY, USA, in 2007; and Bell Labs, Murray Hill, NJ, USA, in 2006 as summer research internships.

Dr. Jung was a recipient of the Samsung Lee Kun Hee Scholarship for graduate study, and he was a gold medalist at the International Mathematical Olympiad (IMO) in 1995. He is a recipient of the excellent new faculty funding from NRF Korea in 2012.

Matthew Andrews (SM'11) received the B.A. degree in mathematics with first-class honors from the University of Oxford, Oxford, U.K., in 1993, and the Ph.D. degree in the theory of computing from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1997.

He is a Distinguished Member of Technical Staff with Alcatel-Lucent Bell Labs, Murray Hill, NJ, USA. His research interests include combinatorial optimization, operations research, and data analytics, with applications to resource allocation and revenue management in wireless networks, energy-efficient telecommunications, and the analysis of large graphs.

Dr. Andrews is serving as Program Chair for WiOpt 2014 and was a recipient of the Best Paper Award at the 2010 IEEE Symposium on Foundations of Computer Science (FOCS).