

Stability of the Max-Weight Protocol in Adversarial Wireless Networks

Sungsu Lim
KAIST
ssungssu@kaist.ac.kr

Kyomin Jung
KAIST
kyomin@kaist.edu

Matthew Andrews
Bell Labs
andrews@research.belllabs.com

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 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.

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 [1], [2], [3] since it was introduced in the work of Tassiulas and Ephremides [4], [5] and Awerbuch and Leighton [6], [7]. The key property of Max-Weight is that for a fixed set of flows it is *throughput optimal* in stochastic networks with a wide variety of scenarios [4], [8], [9], even though it may fail to provide maximum stability in a scenario with flow-level dynamics [10]. 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 non-evenly 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 [11] it was shown that for a single transmitter sending data over one-hop edges to a set of mobile users, if the set of non-zero channel rates can approach zero arbitrarily closely, then no protocol can be stable. However, since this is a fairly unnatural condition, [11] looked at the more natural setting in which all rate sets are 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.

In some adversarial setting, the stability of Max-Weight in static networks was proven in [12], and the stability of Max-Weight was proven in dynamic networks with single-commodity demands [1] and multicommodity demands [13]. 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 [11], 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 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 [14], independent set constraint [15], [16], and node exclusive

constraints [17], [18], [19], [20]. 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 [21]) 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 [22], [15], [16]. In this paper, we also prove the 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 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 $q_{v,d}^t$. For any set of edges in the network, the total *weight* on the 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 the model section below.) 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, in Section IV we show the stability of an approximate Max-Weight algorithm which can be implemented efficiently in many practical setups.

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 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 $P(t) = \sum_{v,d} (q_{v,d}^t)^2$ 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 $P(t+1) - P(t) = \sum_{v,d} (q_{v,d}^t + \sigma_{v,d}^t)^2 - \sum_{v,d} (q_{v,d}^t)^2$ where $\sigma_{v,d}^t = q_{v,d}^{t+1} - q_{v,d}^t$. Moreover, when the potential function become 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 $P(t) = \sum_{v,d} (q_{v,d}^t)^2$ and hence we have an upper bound on $\sum_{v,d} q_{v,d}^t$.

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. Hence for any possible queue configuration there is always the possibility that the potential function $P(t) = \sum_{v,d} (q_{v,d}^t)^2$ 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 which includes the *worst* case scenario; A model that is governed by a stationary stochastic process is not general enough to cover many widely occurring scenarios. For example, consider a cellular network in which a car is driving down a road between evenly spaced basestations. In this case the channel conditions between the car and its closest basestation 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 non-zero 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 [13], 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 [13] all edge rates were either zero or one. Secondly, when a edge had rate one we could transmit on it regardless of what is happening on the other edges. 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. In the next section we describe our model in more detail, after which we present our results.

C. The Model

We assume a system in which time is divided into discrete time slots. We consider a queuing 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 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 $r(t) = (r_1(t), r_2(t), \dots, r_k(t))$, then it is possible to transmit on edge e at rate $r_e(t)$, *for all edges simultaneously*. In other words we can transmit data of size x_1 on edge 1,

data of size x_2 on edge 2, 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 constraint, 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 [11] that if we do not have these conditions then no on-line protocol can be stable.)

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

We now define the (ω, ε) -adversary. 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, $A(\omega, \varepsilon)$, 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) which keeps the system stable. Let t_p be the time when a packet p is injected. Then we can define $\Psi_p = \{(e, t') | t' \in [t_p, t_p + \omega - 1], \ell(p, e, t') > 0, \text{ where } \ell(p, e, t') \text{ is a fractional amount of } p \text{ that is transmitted by } T \text{ along } e \text{ at time } t'\}$, which corresponds to the movement of packet p from its source to one of its destinations under the algorithm T . For all packet p , $(1 - \frac{\varepsilon}{2})$ fraction¹ 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

$$\sum_{\substack{p \in I^j \cup I^{j-1}, \\ t' \in \{t' \in W | (e, t') \in \Psi_{p_i}\}}} \ell(p, e, t') \leq \sum_{t' \in W_j} (1 - \varepsilon) r_e(t'),$$

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, including k-hop interference, independent set constraint, and node exclusive constraints, in dynamic networks, and this model includes

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

adversarial models used in [1], [12] and [13]. 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 $A(\omega, \varepsilon)$ for any $\varepsilon > 0$.

D. The Protocol

We now define the MAX-WEIGHT protocol. We assume that each node v has $|D|$ queues which 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)

- 1) Choose $r(t) \in R(t)$ and $d^{(e)} \in D$ for each $e = (v, u) \in E$, such that $\sum_{e \in E} s_e(t) \left((q_{v,d^{(e)}}^t)^\beta - (q_{u,d^{(e)}}^t)^\beta \right)$ is maximized (with an arbitrary tie-breaking 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 time t , and for each node v , accept all packets injected by the Adversary to v .
- 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.)

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

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. Let $a_{v,d}^t$ and $b_{v,d}^t$ be the corresponding quantities for the underlying “optimum” schedule (that keeps the system stable by assumption). The expected change in $P(\cdot)$ from time step t to time $t + 1$ is given by,

$$\begin{aligned}
E[P(t+1) - P(t)] &= E[\sum_{v,d} (q_{v,d}^{t+1})^2 - \sum_{v,d} (q_{v,d}^t)^2] \\
&= E[\sum_{v,d} (q_{v,d}^t + a_{v,d}^t - b_{v,d}^t)^2 - \sum_{v,d} (q_{v,d}^t)^2] \\
&= E[\sum_{v,d} ((q_{v,d}^t)^2 + 2q_{v,d}^t(a_{v,d}^t - b_{v,d}^t) + (a_{v,d}^t - b_{v,d}^t)^2) \\
&\quad - \sum_{v,d} (q_{v,d}^t)^2] \\
&\leq E[\sum_{v,d} (2q_{v,d}^t(a_{v,d}^t - b_{v,d}^t) + (a_{v,d}^t - b_{v,d}^t)^2)]. \quad (1)
\end{aligned}$$

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 $\sum_{v,d} q_{v,d}^t (a_{v,d}^t - b_{v,d}^t)$. By taking into account the i.i.d. nature of the service rate vectors and the fact that the traffic injections can be scheduled by the optimal algorithm, we have that $E[(a_{v,d}^t - b_{v,d}^t)] \leq -\epsilon$ for all v, d , since there is an upper bound on the amount of data that can be transferred 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,

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

and thus if there is some $Q_{v,d}$ that satisfies $q_{v,d}^t \geq C/2\epsilon$ 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 non-stationary, adversarial environment it is not necessarily the case that the set $R(t)$ and 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 work, so there is no analogue of the statement that $E[(a_{v,d}^t - b_{v,d}^t)] \leq -\epsilon$. 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 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 packet with 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 $W = [t, t + \omega - 1]$. A set of partial transmissions Γ_p assigned to p is defined as a vector of dimension $\omega|L|$. Let $s_e(t')$ be the vector chosen by MAX-WEIGHT(β) that maximizes $\sum_e s_e(t') \left((q_{v,d(e)}^{t'})^\beta - (q_{u,d(e)}^{t'})^\beta \right)$. For a given adversary $A(\omega, \epsilon)$, and a scheduling algorithm Alg , let $\Gamma_p(A(\omega, \epsilon), Alg) = (s_{p,\ell}(t'))_{\ell \in L, t' \in W}$ be a vector of size $\omega|L|$ that satisfies for each $e \in E, d \in D, t' \in W$, (i) $s_{p,(e,d)}(t') \geq 0$, and (ii) $\sum_{p,d} s_{p,(e,d)}(t') \leq s_e(t')$. We say $\Gamma_p(A(\omega, \epsilon), Alg)$ is a *set of (possibly partial) transmissions assigned with p* . For convenience, we denote it by Γ_p .

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




Fig. 1. An 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, a set of partial transmissions assigned with p_1 that contains $s_{p_1, \ell_1}(t_{p_1} + 2) = 0.5$ and $s_{p_1, \ell_2}(t_{p_1} - 1) = 1$ can be Γ_{p_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.

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 described in Fig 1.

Theorem 2: Consider a given adversary $A(\omega, \epsilon)$ for any $\omega \geq 1$ and $\epsilon > 0$, and the MAX-WEIGHT(β) protocol for some $\beta > 0$. For any injected packet p , we can assign this packet with $\Gamma_p = \Gamma_p(A(\omega, \epsilon), \text{MAX-WEIGHT}(\beta))$ so that the sum of total potential changes is less than $-\frac{\epsilon}{1-\epsilon/2}\ell_p(\beta + 1)q^\beta + \ell_p O(q^{\beta-1})$, where q is the height of the queue where the packet is injected and ℓ_p is a size of an injected packet p . Therefore, there is a constant q^* depending on ω and ϵ , so that if $q \geq q^*$ the sum of potential changes due to the injection is less than $-\frac{\epsilon}{2}\ell_p q^\beta$.

We sketch the overall proof of Theorem 2 in Section IV-A, while the details are made available in a technical report [23].

Now, we will prove the stability of the MAX-WEIGHT protocol under any $A(\omega, \epsilon)$. 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 $\epsilon > 0$.

We define a more general adversarial model, 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 $n|D|$. An adversary allowing b many bad packets is defined as follows.

By Theorem 2, for any $A(\omega, \epsilon)$ under MAX-WEIGHT, for each injection of packet p , we can assign this packet with a set of partial transmissions Γ_p so that the sum of potential changes due to these movements are at most $-\frac{\epsilon}{1-\epsilon/2}\ell_p q + C$, where q is the height of the queue where the packet is injected and C is a constant depending on ω and ϵ (but not on n and t). In a general adversarial queueing system, we also consider the same assignment Γ_p . If the sum of potential changes due to Γ_p is at least $-\frac{\epsilon}{1-\epsilon/2}\ell_p q + C + 1$, we now say this a *bad packet*. We say 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 $\epsilon > 0$ 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 sub-system, 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 p is more than $\frac{\epsilon}{2}\ell_p q$.

The crux of our analysis will involve proving the following results (in section IV-B and a technical report).

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

Hence from Theorem 4 we obtain that Theorem 1 directly.

IV. PROOF OUTLINES

A. Proof of Theorem 2

The collection of $W_j = [j\omega, (j+1)\omega - 1]$ for all integer $j \geq 0$ is non-overlapped and the union of this collection covers all time slots t . From now on let $W = W_j$ for some j .

For each time $t' \in W_j \cup W_{j-1}$, 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 $r_p = \{d_{p,e}(t') | (e, t') \in \Psi_p\}$ of rates of directed edges used in Ψ_p as follows. Let p_1, \dots, p_m be all the packets injected in $I^j \cup I^{j-1}$. The order of p_i 's can be any ordering. Then from the definition 1, for any e used in Ψ_{p_i} ,

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

where $r^{(0)}(t) \in R(t)$ is the edge rate vector assigned by T .

We will define $d_{p_i, e}(t')$ and $r_e^{(i)}(t')$ recursively for $i = 1, \dots, m$, so that $0 \leq d_{p_i, e}(t') \leq r_e^{(i-1)}(t')$,

$$(1 - \epsilon) \sum_{t' \in W} d_{p_i, e}(t') = \sum_{t' \in W} \ell(p_i, e, t'), \quad (3)$$

and $r_e^{(i)}(t') = r_e^{(i-1)}(t') - d_{p_i, e}(t')$.

Let $(s_e(t'))$ be the vector chosen by MAX-WEIGHT(β). For any time t' , 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 T . For any p , we will define $\Gamma_p =$



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.

$(s_{p,(e,d)}(t'))_{(e,d) \in L, t' \in W}$, which satisfies the conditions (i) and (ii) of the definition 2, so that the sum of potential changes by $s_{p,(e,d)}(t')$ is equal to the sum of potential changes by $d_{p,e}(t')$ (for details, see [23]).

Then, from (3) and the above property we prove that for any $A(\omega, \epsilon)$ under MAX-WEIGHT(β), for each injection of a packet p , the sum of potential changes due to the injection of p and Γ_p is at most $-\frac{\epsilon}{1-\epsilon/2} \ell_p(\beta+1)q^\beta + \ell_p O(q^{\beta-1})$, where q is the height of the queue where the packet is injected.

B. Proof of Theorem 4

Let $\epsilon > 0$, 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 $t \geq 0$ is 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 any $i \geq 1$, we can set

$$U(n, q_0, i) = U(n, U(n, q_0, i-1) + R_{\max}, 0) \quad (4)$$

by considering the time when the i th bad packet arrives. By applying this, we only need to prove that $U(n, q_0, 0)$ exists.

We define recursively a set of queue thresholds $M_1 > M_2 > \dots > M_n = 0$, which correspond to some potential of the queues. Let $q_1 \geq q_2 \geq \dots \geq q_n = 0$ be the ordered sizes of the queues at time t_0 . Then we can take some $1 \leq k \leq n-1$ that satisfies $q_k \geq M_k$ and $q_{k+1} \leq M_{k+1}$.

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”. 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. 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. We prove the following Lemma in the technical report [23].

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 .

By applying Lemma 5, for $t_0 \leq t \leq t^*$, the potential $P(t)$ is bounded by some value that is independent of t . Since $P(t^*) \leq P(t_0)$, by the same argument, we obtain that for all time $t \geq 0$, $P(t)$ is bounded by the same value. Hence $U(n, q_0, 0)$ exists, which in turn proves Theorem 4.

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 $(k^2 \epsilon^{-1} (R_{\max})^2)^2$. More precisely, if the potential ever does get larger than that amount then some queue size must be larger than $k^2 \epsilon^{-1} (R_{\max})^2$. At that point the expression for the change in network potential implies the expected drift in potential is non-positive. One consequence of this is that whenever an individual queue size becomes larger than $(k^2 \epsilon^{-1} (R_{\max})^2)^2$ the expected drift in potential is non-positive.

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. 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 [21] in which it was shown that we can get exponential queue sizes in a critically loaded scenario (i.e. where $\epsilon = 0$). We now show that this is actually possible in a subcritically loaded example (with $\epsilon > 0$).

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 $i' = \min\{i : q_i(t) < (1-\epsilon)2^i\}$. If $i' = 0$ then the adversary sets $r_0(t) = 1$ and $a_0(t) = 1 - \epsilon$. If $i' > 0$ then it sets $r_{i'-1}(t) = 1 - \epsilon$, $r_{i'}(t) = \frac{1-\epsilon}{2}$ and $a_{i'}(t) = \frac{(1-\epsilon)^2}{2}$. 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 $A(1, \epsilon)$ adversary.

Lemma 6: 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 $q_i(t') \geq (1-\epsilon)2^i$.

Proof: We prove the above statement by induction on i . Suppose that $q_0(t) < 1 - \epsilon$. Then for this time step i' is set to 0 and so $a_0(t) = 1 - \epsilon$. Once data has been served for edge 0 and the arriving data has been added to the edge’s queue we have $q_0(t+1) \geq 1 - \epsilon$. (Note that this assumes that data arrives in a queue after data has 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 $q_i(t) < (1-\epsilon)2^i$ 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 7: 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 protocol adversarial models do not necessarily lead to exponentially large queues. In [21] 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 feel 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 $A(\omega, \varepsilon)$ assures that there is a set of fractional movement of packets Ψ_p for each $p \in I^W$ and there is a edge rate vector $r_e \in R(t)$ for each $t \in W$, so that each edge is used at most $(1 - \varepsilon)$ times of 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 $\frac{1}{1-\varepsilon}$ 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 ε -approximate solution for $r(t)$ can be computed in polynomial time. For example, [22] presented an ε -approximate solution to find the MWIS (maximum weight independent set) on planar graphs, and this was extended by several authors to more general classes of graphs. In [15], [16], an ε -approximate MWIS for a large class of wireless networks in the Euclidean space 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 $A(\omega, \varepsilon)$ for $\varepsilon > 0$, if $0 < \hat{\varepsilon} < \varepsilon$.

Theorem 8: For $0 < \hat{\varepsilon} < \varepsilon$, any $\hat{\varepsilon}$ -approximate MAX-WEIGHT is stable under $A(\omega, \varepsilon)$.

Proof: Let $\tilde{\varepsilon} = \frac{\varepsilon - \hat{\varepsilon}}{1 - \hat{\varepsilon}}$, then $A(\omega, \varepsilon)$ is $A(\omega, \tilde{\varepsilon})$. 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 $\frac{\tilde{\varepsilon}}{2} \ell_p q^\beta$, then all the other arguments in the proof of Theorem 4 holds when we replace ε with $\tilde{\varepsilon}$.

As in the proof of Theorem 2, we define $d'_{p,e}(t') = \frac{1-\tilde{\varepsilon}}{1-\varepsilon} d_{p,e}(t')$ 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 (3) if we substitute ε by $\tilde{\varepsilon}$. As in the proof of Theorem 2 in [23], we define

$$\begin{aligned} K'_{e_j}(t') &= \sum_{i=1}^m d'_{p_i, (v_j, u_j)}(t') \left((q'_{v_j, d_i})^\beta - (q'_{u_j, d_i})^\beta \right) \\ &= (1 - \hat{\varepsilon}) K_{e_j}(t'). \end{aligned} \quad (5)$$

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

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

for some destinations $d_a^{(e_j)}$ for each e_j . By (5) and (6), we can inductively assign

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

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 in [23]. Then, for all $e_j \in E$, $t' \in W$, and $p_i \in I^j \cup I^{j-1}$,

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

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 $-\tilde{\varepsilon} \ell_{p_i} q^\beta$ by using the same argument in section IV-A. This in turn implies that $\hat{\varepsilon}$ -MAX-WEIGHT algorithm is stable under $A(\omega, \varepsilon)$. ■

VII. EXPERIMENTS

A. Simulation Setup

We now describe a numerical experiment that aims to understand the queue size dynamics of the MAX-WEIGHT




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

protocol under the adversarial model. Consider a $n_1 \times n_2$ 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 $4n - 2n_1 - 2n_2 = 34$.

In our simulation, we used 3 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 3 edges among 17 possible 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 constraint model, i.e., matching constraint model.

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

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

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

If an arrival rate vector is in the interior of $co(R)$, and the arrivals are identical for all time, then MAX-WEIGHT is stable [15]. Moreover, if an arrival rate vector is in the interior of $co(R)^c$, then MAX-WEIGHT is unstable. We chose K many source-destination pairs at random. For each $r^{(i)}$, $1 \leq i \leq 3$, we compute 3 different feasible arrival rate vectors that are closed 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 two set of experiments. In both of those experiments, we divided the time $t \geq 0$ into non-overlapped sub-windows 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 i th phase is: $t \in [\lceil \sum_{j=1}^i (1.5)^{j-1} \rceil + 1, \lceil \sum_{j=1}^{i+1} (1.5)^{j-1} \rceil]$.

In the first experiment, we fixed the edge rate vector $r^{(i)}$




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 $[0, 10^6]$.

Fixed edge-rate $r^{(1)}$ & cyclic 3 feasible arrivals




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.

for some $i \in \{1, 2, 3\}$. Over time the adversary injects packets as follows. For $t \geq 0$, if t is in the j -th phase, then inject packets with an arrival rate $c_{i\bar{j}}\gamma^{(\bar{j})}$ where $\bar{j} \in \{1, 2, 3\}$ and $\bar{j} \equiv j \pmod{3}$.

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 i -th 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 $c_{i\bar{j}}\gamma^{(\bar{j})}$ at random.

Notice that, in both experiments, the average of the arrival rate vectors until time T does not converge as T goes to infinity. Also in the second experiment, the same holds for the edge rate vectors. However the above injections satisfy the definition of $A(\omega, \varepsilon)$ for some $\omega > 0$ and a small $\varepsilon > 0$. In both setups, we observed the dynamics of the maximum queue sizes over time.

Cyclic 3 edge-rates & feasible arrivals




Fig. 6. We use the randomly cyclic edge and arrival rate pairs. It shows the stability of MAX-WEIGHT.

B. Simulation Results

For the first experiment, as Fig 5 shows, for each edge rate vector, MAX-WEIGHT is stable with the above cyclic rate vectors. Interestingly, the maximum queue size may increase in some sub-window, but it decreases rapidly when the new sub-window 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 rate vector case are bounded above and bounded below by some fixed arrival rate vector cases respectively.

The queue dynamics for the second experiment are described in Fig 6. The gray lines describe queue sizes for fixed edge and arrival rate vectors. The black line describes the queue size for the cyclic rate vector case. Again, the maximum queue sizes for the cyclic rate vector 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 make the system stable under $A(\omega, \varepsilon)$ even when the edge and arrival rate vectors do not converge over time.

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 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] E. Anshelevich, D. Kempe, and J. Kleinberg, "Stability of load balancing algorithms in dynamic adversarial systems," in *Proceedings of the 34th*

Annual ACM Symposium on Theory of Computing, Montreal, Canada, May 2002.

[2] N. W. McKeown, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switch," in *Proceedings of IEEE INFOCOM '96*, San Francisco, CA, March 1996, pp. 296–302.

[3] S. Muthukrishnan and R. Rajaraman, "An adversarial model for distributed dynamic load balancing," in *10th ACM Symposium on Parallel Algorithms and Architectures, SPAA 1998*, 1998, pp. 47–54.

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

[5] —, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Transaction on Information Theory*, vol. 30, pp. 466–478, 1993.

[6] B. Awerbuch and T. Leighton, "A simple local-control approximation algorithm for multicommodity flow," in *Proceedings of the 34th annual symposium on Foundations of Computer Science*, November 1993, pp. 459–468.

[7] —, "Improved approximation algorithms for the multi-commodity flow problem and local competitive routing in dynamic networks," in *Proceedings of the 26th annual ACM symposium on Theory of computing*, 1994.

[8] L. Georgiadis, M. J. Neely, and L. Tassiulas, "Resource allocation and cross-layer control in wireless networks," *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1–149, 2006.

[9] D. Shah and D. Wischik, "Optimal scheduling algorithms for input-queued switches," in *Proceedings of IEEE INFOCOM '06*, Barcelona, Spain, April 2006, pp. 1–11.

[10] P. van de Ven, S. Borst, and S. Shneer, "Instability of maxweight scheduling algorithm," in *Proceedings of IEEE INFOCOM '09*, Rio de Janeiro, Brazil, April 2009, pp. 1701–1709.

[11] M. Andrews and L. Zhang, "Scheduling over a time-varying user-dependent channel with applications to high speed wireless data," in *Proceedings of the 43rd Symposium on Foundations of Computer Science*, November 2002, pp. 293–302.

[12] W. Aiello, E. Kushilevitz, R. Ostrovsky, and A. Rosen, "Adaptive packet routing for bursty adversarial traffic," in *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, Dallas, TX, May 1998, pp. 359–368.

[13] M. Andrews, K. Jung, and A. Stolyar, "Stability of the max-weight routing and scheduling protocol in dynamic networks and at critical loads," in *Proceedings of the 39th annual ACM symposium on Theory of computing*, June 2007, pp. 145–154.

[14] G. Sharma, R. R. Mazumdar, and N. B. Shroff, "On the complexity of scheduling in wireless networks," in *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, September 2006, pp. 227–238.

[15] R. Gummadi, K. Jung, D. Shah, and R. Sreenivas, "Feasible rate allocation in wireless networks," in *Proceedings of IEEE INFOCOM '08*, April 2008, pp. 995–1003.

[16] —, "Computing capacity region of a wireless network," in *Proceedings of IEEE INFOCOM '09*, April 2009, pp. 1341–1349.

[17] 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 Technical Memorandum*, April 2000.

[18] —, "Providing quality of service over a shared wireless link," *IEEE Communications Magazine*, February 2001.

[19] B. Hajek and G. Sasaki, "Link scheduling in polynomial time," *IEEE Transaction on Information Theory*, vol. 34, pp. 910–917, September 1988.

[20] M. Neely, E. Modiano, and C. Rohrs, "Power and server allocation in a multi-beam satellite with time varying channels," in *Proceedings of IEEE INFOCOM '02*, New York, NY, June 2002.

[21] M. Andrews and L. Zhang, "Scheduling over non-stationary wireless channels with finite rate sets," in *Proceedings of IEEE INFOCOM '04*, March 2004, pp. 1694–1704.

[22] B. S. Baker, "Approximation algorithms for NP-complete problems on planar graphs," *Journal of the Association for Computing Machinery*, vol. 41, no. 1, pp. 153–180, January 1994.

[23] S. Lim, K. Jung, and M. Andrews, "Stability of the max-weight protocol in adversarial wireless networks," in *arXiv:cs.NI*, January 2012.