Network Utility Maximization in Adversarial Environments

Qingkai Liang and Eytan Modiano Laboratory for Information and Decision Systems Massachusetts Institute of Technology, Cambridge, MA

Abstract-Stochastic models have been dominant in network optimization theory for over two decades, due to their analytical tractability. However, these models fail to capture non-stationary or even adversarial network dynamics which are of increasing importance for modeling the behavior of networks under malicious attacks or characterizing short-term transient behavior. In this paper, we consider the network utility maximization problem in adversarial network settings. In particular, we focus on the tradeoffs between total queue length and utility regret which measures the difference in network utility between a causal policy and an "oracle" that knows the future within a finite time horizon. Two adversarial network models are developed to characterize the adversary's behavior. We provide lower bounds on the tradeoff between utility regret and queue length under these adversarial models, and analyze the performance of two control policies (i.e., the Drift-plus-Penalty algorithm and the Tracking Algorithm).

I. INTRODUCTION

Stochastic network models have been dominant in network optimization theory for over two decades, due to their analytical tractability. For example, it is often assumed in wireless networks that the variation of traffic patterns and the evolution of channel capacity follow some stationary stochastic process, such as the i.i.d. model and the ergodic Markov model. Many important network control policies (e.g., MaxWeight [1] and Drift-plus-Penalty policy [2]) have been derived to optimize network performance under those stochastic network dynamics.

However, non-stationary or even adversarial dynamics have been of increasing importance in recent years. For example, modern communication networks frequently suffer from Distributed Denial-of-Service (DDoS) attacks or jamming attacks [3], where traffic injections and channel conditions are controlled by some malicious entity in order to degrade network performance. As a result, it is important to develop efficient control policies that optimize network performance even in adversarial settings. However, extending the traditional stochastic network optimization framework to the adversarial setting is non-trivial because many important notions and analytical tools developed for stochastic networks cannot be applied in adversarial settings. For example, traditional stochastic network optimization focuses on long-term network performance while in an adversarial environment the network may not have any steady state or well-defined long-term time averages. Thus, typical steady-state analysis and many equilibrium-based notions such as the network throughput region cannot be used in networks with adversarial dynamics, and it is important to understand "transient" network performance within a finite time horizon in a non-stationary/adversarial environment.

In this paper, we investigate efficient network control policies that can maximize network utility within a finite time horizon while keeping the total queue length small in an adversarial environment. In particular, we focus on the following optimization problem:

$$\max_{\alpha_t \in \mathcal{D}_{\omega_t}} \sum_{t=0}^{T-1} U(\alpha_t, \omega_t)$$

s.t.
$$\sum_{t=0}^{T-1} a_i(t) \le \sum_{t=0}^{T-1} \tilde{b}_i(t), \ \forall i$$
 (1)

where $U(\alpha_t, \omega_t)$ is the network utility gained in slot t under the control action α_t (constrained to some action space \mathcal{D}_{ω_t}) and the network event ω_t (which includes information about exogenous arrivals, link capacities, etc). The sequence of network events $\{\omega_t\}_{t=0}^{T-1}$ follows an arbitrary (possibly adversarial) process. The objective is to maximize the total network utility gained within a finite time horizon T subject to the constraint that for each queue i the total arrivals $\sum_{t=0}^{T-1} a_i(t)$ do not exceed the total departures $\sum_{t=0}^{T-1} \tilde{b}_i(t)$ during the time horizon.

A. Main Results

We develop general adversarial network models and propose a new finite-time performance metric, referred to as *utility regret* (the formal definition is given in Section II-A):

$$\mathcal{R}_{T}^{\pi} = \sum_{t=0}^{T-1} U(\alpha_{t}^{*}, \omega_{t}) - \sum_{t=0}^{T-1} U(\alpha_{t}^{\pi}, \omega_{t}),$$

where $\{\alpha_t^{\pi}\}_{t=0}^{T-1}$ is the sequence of control actions taken by a policy π , and $\{\alpha_t^{*}\}_{t=0}^{T-1}$ is the optimal sequence of actions for solving (1) generated by an "oracle" that knows the future. Note that a control policy π may trivially maximize the network utility by simply ignoring the constraint in (1) (e.g., admitting all the exogenous traffic) such that the utility

This work was supported by NSF Grant CNS-1524317 and by DARPA I2O and Raytheon BBN Technologies under Contract No. HROO I I-I 5-C-0097.

regret become zero or even negative¹. However, such an action may significantly violate the constraint in (1) and lead to large queue length. Therefore, there is a tradeoff between the utility regret and the queue length achieved by a control policy, which is similar to the well-known utility-delay tradeoff in traditional stochastic network optimization [4]. In this paper, we investigate this tradeoff in an adversarial environment. The main results are as follows.

• We prove that it is impossible to simultaneously achieve both "low" utility regret and "low" queue length if the adversary is unconstrained. In particular, there exist some adversarial network dynamics such that either the utility regret or the total queue length grows at least linearly with the time horizon T under any causal control policy. This impossibility result motivates us to study constrained adversarial dynamics.

• We develop two adversarial network models where the network dynamics are constrained to some "admissible" set. In particular, we first consider the W-constrained adversary model, where under the optimal policy the total arrivals do not exceed the total services within any window of W slots. We then propose a more general adversary model called V_T -constrained adversary, where the total queue length generated by the "oracle" during its sample path is upper bounded by V_T . By varying the values of V_T , the proposed V_T -constrained model covers a wide range of adversarial settings: from a strictly constrained adversary to a fully unconstrained adversary.

• We develop lower bounds on the tradeoffs between utility regret and queue length under both the W-contrained and the V_T -constrained adversary models. It is shown that no causal policy can simultanesouly achieve both sublinear utility regret and sublinear queue length if W or V_T grows linearly with T. We also analyze the tradeoffs achieved by two control algorithms: the Drift-plus-Penalty algorithm [2] and the Tracking Algorithm [5], [6] under the two adversarial models. In particular, both algorithms simultaneously achieve sublinear utility regret and sublinear queue length whenever W or V_T grows sublinearly with T, yet the theoretical regret bound under the Tracking Algorithm is better than that under the Drift-plus-Penalty algorithm. The Tracking Algorithm also asymptotically attains the optimal tradeoffs under the W-constrained adversary model.

B. Related Work

The study of adversarial network models dates back more than two decades ago. Rene Cruz [7] provided the first concrete example of networks with adversarial dynamics, which were later generalized by Borodin *et al.* [8] under the *Adversarial Queuing Theory* (AQT) framework. In AQT, in each time slot, the adversary injects a set of packets at some of the nodes. In order to avoid trivially overloading the system, the AQT framework imposes a stringent *window* *constraints*: the maximum traffic injected in every link over any window of W time slots should not exceed the amount of traffic that the link can serve during that interval. Andrews *et al.* [9] introduced a more generalized adversary model known as the *Leaky Bucket* (LB) model that differs from AQT by allowing some traffic burst during any time interval. The AQT model and the LB model have given rise to a large number of results since their introduction, most of which are about network stability under several simple scheduling policies such as FIFO (see [10] for a review of these results).

However, the AQT and the LB models assume that only packet injections are adversarial while the underlying network topology and link states remain fixed. Such a static network model does not capture many adversarial environments, such as wireless networks under jamming attacks where the adversary can control the channel states. Andrews and Zhang [5], [6] extended the AQT model to single-hop dynamic wireless networks, where both packets injections and link states are controlled by an adversary, and prove the stability of the MaxWeight algorithm in this context. Jung *et al.* [11], [12] further extended the results of [5], [6] to multi-hop dynamic networks. Our windowbased W-constrained model is inspired by and similar to the adversarial models used in [5], [6], [11], [12].

While the above-mentioned works focused on network stability, our work is most related to the universal network utility maximization problem by Neely [13] where network utility needs to be maximized subject to stability constraints under adversarial network dynamics. Algorithm (time-average) performance is measured with respect to a so-called "W-slot look-ahead policy". Such a policy has perfect knowledge about network dynamics over the next W slots but it is required that under this policy the total arrivals to each queue should not exceed the total amount of service offered to that queue during every window of W slots. As a result, it is similar to our W-constrained model where stringent window constraints have to be enforced.

Our paper expands previous work in a number of fundamental ways. First, we develop lower bounds on the tradeoffs between utility regret and queue length under both the W-contrained and the V_T -constrained adversary models. As far as we know, none of the existing works (e.g., [5], [6], [11]-[13]) provide lower bounds in any kind of adversarial network models. Second, we provide analysis under the new V_T -constrained adversary model which generalizes the adversarial network dynamics models used by existing works. To the best of our knowledge, existing works (e.g., [5], [6], [11]–[13]) all use the W-constrained adversary model or similar windows-based variants due to its analytical tractability. In this paper, we propose a new V_T constrained adversary model which gets rid of the window constrains. Due to the lack of window-based structure, the analysis carried out in existing works cannot be applied to the V_T -constrained model. We develop new analytical results under the new V_T -constrained model by converting the V_T -constrained model to a W-constrained model with a

¹The negative utility regret may occur since any optimal solution $\{\alpha_t^*\}_{t=0}^{T-1}$ is required to satisfy the constraint in (1) while an arbitrary policy π may violate this constraint.

carefully selected window size W.

C. Organization of this Paper

The rest of this paper is organized as follows. We first introduce the system model and relevant performance metrics in Section II. We study the *W*-constrained and V_T -constrained adversary models in Sections III and IV, respectively. Finally, simulation results and conclusions are given in Sections V and VI, respectively.

II. SYSTEM MODEL

Consider a network with N queues (the set of all queues are denoted by $\mathcal{N} = \{1, \dots, N\}$). Time is slotted with a finite horizon $\mathcal{T} = \{0, \dots, T-1\}$. Let ω_t denote the *network event* that occurs in slot t, which indicates the current network parameters, such as a vector of conditions for each link, a vector of exogenous arrivals to each node, or other relevant information about the current network links and exogenous arrivals. The set of all possible network events is denoted by Ω .

At the beginning of each time slot t, the network operator observes the current network event ω_t and chooses a control action α_t from some action space \mathcal{D}_{ω_t} that can depend on ω_t . The network event ω_t and the control action α_t together produce the service vector $\mathbf{b}(\alpha_t, \omega_t) \triangleq \mathbf{b}(t) =$ $(b_1(t), \cdots, b_N(t))$ and the arrival vector $\mathbf{a}(\alpha_t, \omega_t) \triangleq \mathbf{a}(t) =$ $(a_1(t), \dots, a_N(t))$. Note that $a_i(t)$ includes both the admitted exogenous arrivals from outside the network to queue *i*, and the endogenous arrivals from other queues (i.e., routed packets from other queues to queue i). Thus, the above network model accounts for both single-hop and multi-hop networks, and the control action α_t may correspond to, for example, joint admission control, routing, rate allocation and scheduling decisions in a multi-hop network. Let $\mathbf{Q}(t) = (Q_1(t), \cdots, Q_N(t))$ be the queue length vector at the beginning of slot t (before the arrivals in that slot). The queueing dynamics are

$$Q_i(t+1) = [Q_i(t) + a_i(t) - b_i(t)]^+, \ \forall i \in \mathcal{N}, t \in \mathcal{T},$$

where $[x]^+ = \max\{x, 0\}.$

We assume that the sequence of network events $\{\omega_t\}_{t=0}^{T-1}$ are generate according to an *arbitrary* process (possibly non-stationary or even adversarial), except for the following boundedness assumption. Under any network event and any control action, the arrivals and the service rates in each slot are bounded by constants that are independent of the time horizon T: for any $\omega_t \in \Omega$ and any $\alpha_t \in \mathcal{D}_{\omega_t}$

$$0 \le a_i(\alpha_t, \omega_t) \le A, \quad 0 \le b_i(\alpha_t, \omega_t) \le B.$$

For simplicity, we assume $B \ge A$ such that both arrivals and services are upper bounded by B in each slot.

A policy π generates a sequence of control actions $(\alpha_0^{\pi}, \dots, \alpha_{T-1}^{\pi})$ within the time horizon. In each slot t, the queue length vector, the arrival vector and the service rate vector under policy π is denoted by $\mathbf{Q}^{\pi}(t)$, $\mathbf{a}^{\pi}(t)$ and $\mathbf{b}^{\pi}(t)$, respectively. A *causal* policy is one that generates the current control action α_t only based on the knowledge up

until the current slot t. In contrast, a *non-causal* policy may generate the current control action α_t based on knowledge of the future.

Let $U(\alpha_t, \omega_t)$ be the network utility gained in slot t if action α_t is taken under network event ω_t . We assume that under any control action and any network event, network utility is bounded:

$$U_{\min} \leq U(\alpha_t, \omega_t) \leq U_{\max}, \ \forall \omega_t \in \Omega, \alpha_t \in \mathcal{D}_{\omega_t}.$$

A commonly-used form of the network utility function is $U(\alpha_t, \omega_t) = \sum_i U_i(x_i(t))$ where $x_i(t)$ is the amount of admitted exogenous traffic to queue *i* in slot *t*. Typical examples include $U(\alpha_t, \omega_t) = \sum_i x_i(t)$ (total throughput), $U(\alpha_t, \omega_t) = \sum_i \log(x_i(t))$ (proportional fairness), etc. In wireless networks with power control, another widely-used network utility function is $U(\alpha_t, \omega_t) = -\sum_i P_i(t)$ where $P_i(t)$ is the power allocated to queue *i* in slot *t*. This utility function aims to minimize the total power consumption.

In this paper, we consider the following network utility maximization problem, referred to as **NUM**.

T 1

NUM:

$$\max_{\alpha_t \in \mathcal{D}_{\omega_t}} \sum_{t=0}^{T-1} U(\alpha_t, \omega_t)$$
s.t.
$$\sum_{t=0}^{T-1} a_i(t) \le \sum_{t=0}^{T-1} \tilde{b}_i(t), \ \forall i \in \mathcal{N},$$
(2)
(3)

where $\tilde{b}_i(t) = \min\{b_i(\alpha_t, \omega_t), Q_i(t)\}$ is the actual packet departures from queue *i* in slot *t*. The objective (2) is to maximize the total network utility gained in the time horizon. The constraint (3) requires that the total arrivals to each queue should not exceed the total amount of departures from that queue during the time horizon. Note that the above optimization problem is a natural analogue of the traditional stochastic network optimization problem [4], where the time-average utility is maximized subject to certain network stability constraints. Indeed, if we consider a stochastic network with an infinite time horizon, then the objective (2) is equivalent to maximizing time-average network utility, and the constraint (3) requires that the time-average arrival rate to each queue should not exceed the time-average service rate, which is equivalent to rate stability².

A. Performance Metrics

Our objective is to find a causal control policy that can maximize the network utility while keeping the total queue length small. Note that a network with adversarial dynamics may not have any steady state or well-defined time averages. Hence, it is crucial to understand the transient behavior of the network, and the traditional equilibrium-based performance metrics may not be appropriate in an adversarial setting. As a result, we introduce the notion of *utility regret* to measure the finite-time performance achieved by a control policy.

²A network is rate-stable under a control policy π if $\sum_i Q_i^{\pi}(T)/T \to 0$ as $T \to \infty$.

Definition 1 (Utility Regret). Given the time horizon T, the utility regret achieved by a policy π under a sequence of network events $\omega_0, \dots, \omega_{T-1}$ is defined to be

$$\mathcal{R}_{T}^{\pi}\Big(\{\omega_{0},\cdots,\omega_{T-1}\}\Big) = \sum_{t=0}^{T-1} U(\alpha_{t}^{*},\omega_{t}) - \sum_{t=0}^{T-1} U(\alpha_{t}^{\pi},\omega_{t}), \quad (4)$$

where $\{\alpha_t^*\}_{t=0}^{T-1}$ is an optimal solution to **NUM** generated by an "oracle" that knows the entire sequence of network events $\{\omega_0, \dots, \omega_{T-1}\}$ in advance.

In this setup, a policy π is chosen and then the adversary selects the sequence of network events $\{\omega_0, \dots, \omega_{T-1}\}$ that maximize the regret. Intuitively, the notion of utility regret captures the worst-case utility difference between a causal policy and an ideal *T*-slot lookahead non-causal policy.

Note that any optimal solution $\{\alpha_t^*\}_{t=0}^{T-1}$ to **NUM** is a utility maximizing policy subject to the constraint (3) that it clears all the backlogs within the time horizon. A causal control policy may trivially maximize the network utility by simply ignoring the stability constraint (3) (e.g., admitting all the exogenous traffic) such that the utility regret become zero or even negative. However, such an action may significantly violate the stability constraint (3) and lead to large total queue length. As a result, there is a tradeoff between the utility regret and the total queue length achieved by a causal control policy.

A desirable first order characteristic of a "good" policy π is that it simultaneously achieves sublinear utility regret and sublinear queue length w.r.t. the time horizon T, i.e., $\mathcal{R}_T^{\pi} = o(T)$ and $\sum_i Q_i^{\pi}(T) = o(T)$. Sublinear utility regret guarantees that $\mathcal{R}_T^{\pi}/T \to 0$ as the time horizon $T \to \infty$, meaning that the time-average utility gained under policy π asymptotically approaches that under the optimal non-causal policy. In other words, the long-term time-average utility is optimal. Sublinear queue length ensures $\sum_i Q_i^{\pi}(T)/T \to 0$ as $T \to \infty$, which is equivalent to rate stability. Note that simultaneously achieving sublinear utility regret and sublinear queue length is equivalent to maximizing long-term time-average utility subject rate stability, which is the goal of traditional stochastic network optimization [4].

Note that simultaneously achieving sublinear utility regret and sublinear queue length is just a coarse-grained requirement for a "good" tradeoff between utility regret and queue length. In an adversarial setting with no steady state, the fine-grained growth rates of utility regret and queue length are equally important and should also be well balanced. A better tradeoff in terms of their growth rates implies that the policy has a better learning ability and can adapt to the adversarial environment faster.

Unfortunately, the following theorem shows that in general no causal policy can simultaneously achieve both sublinear utility regret and sublinear queue length.

Theorem 1. For any causal policy π , there exists a sequence of network events $\omega_0, \dots, \omega_{T-1}$ such that either the utility regret $\mathcal{R}_T^{\pi}(\{\omega_0, \dots, \omega_{T-1}\}) = \Omega(T)$ or the total queue length $\sum_i Q_i^{\pi}(T) = \Omega(T)$. **Proof:** We prove this theorem by considering a specific one-hop network with 2 users and constructing a sequence of adversarial network dynamics such that either the utility regret or the total queue length grows at least linearly with the time horizon T. More specifically, the time horizon is split into two parts. In the first T/2 slots, the adversary just generates some regular network events, let the policy run and observes the queue lengths of the two users. In the remaining T/2 slots, the adversary sets the capacity to zero for the user with a longer queue and creates sufficient capacity for the other user such that the performance of the causal policy is significantly degraded while the "oracle" can still perform very well. See the technical report [14] for details.

Theorem 1 shows that is it impossible to achieve sublinear utility regret while maintaining sublinear queue length, if the adversary has unconstrained power in determining the network dynamics. As a result, in the following two sections, we develop two adversary models where the sequence of network events (i.e. network dynamics) that the adversary can select is constrained to some "admissible" set. In Section III, we consider the W-constrained adversary model that is an extension of the widely-known yet very stringent model used in Adversarial Queueing Theory. Next in Section IV, we develop a more relaxed adversary model called the V_T -constrained adversary. Lower bounds on the tradeoffs between utility regret and queue length as well as the performance of some commonly-used algorithms are analyzed under the two adversary models.

III. W-CONSTRAINED ADVERSARY MODEL

In this section, we investigate the W-constrained adversary model which is an extension of the classical Adversarial Queueing Theory (AQT) model [8]. It has stringent constraints on the set of admissible network dynamics that the adversary can set, yet is analytically tractable, which facilitates our subsequent investigation of a more relaxed adversary model in Section IV. We first give the definition of W-constrained network dynamics.

Definition 2 (W-Constrained Dynamics). Given a window size $W \in [1, T]$, a sequence of network events $\omega_0, \dots, \omega_{T-1}$ is W-constrained if

$$\sum_{\tau=t}^{t+W-1} a_i^*(\tau) \le \sum_{\tau=t}^{t+W-1} b_i^*(\tau), \ \forall i \in \mathcal{N}, t \in \mathcal{T},$$
 (5)

where $\{\mathbf{a}^*(t)\}_{t=0}^{T-1}$ and $\{\mathbf{b}^*(t)\}_{t=0}^{T-1}$ is the optimal solution to **NUM** under the above sequence of network events.

Note that if there exist multiple optimal solutions to **NUM**, then constraint (5) is only required to be satisfied by any one of them. Any network satisfying the above is called a W-constrained network. In other words, under the optimal (possibly non-causal) policy, the total amount of arrivals to each queue does not exceed the total amount of service offered to that queue during any window of W slots.

Denote by W_T the set of all sequences of network events $\{\omega_0, \dots, \omega_{T-1}\}$ that are W-constrained. Then the *W*-constrained adversary can only select the sequence of network events from the constrained set W_T .

A. Lower Bound on the Tradeoffs

The following theorem provides a lower bound on the tradeoffs between utility regret and queue length under the W-constrained adversary model.

Theorem 2. For any causal policy π , there exists a sequence of network events $\{\omega_0, \dots, \omega_{T-1}\} \in W_T$ such that

$$\mathcal{R}_T^{\pi}\Big(\{\omega_0,\cdots,\omega_{T-1}\}\Big) + c\sum_i Q_i^{\pi}(T) \ge c'W,$$

where c' > 0 is some constant.

Proof: We prove this theorem by constructing a sequence of network events such that lower bound is attained. The construction is similar to the one used in the proof of Theorem 1. The difference is that the constructed sequence of network events is *W*-constrained here. See the technical report [14] for the detailed proof.

Note that if the window size W is comparable with the time horizon T, i.e., $W = \Theta(T)$, the above theorem shows that no causal policy can simultaneously achieve sublinear utility regret and sublinear queue length under the W-constrained adversary model. On the other hand, if W = o(T), there *might* exist some causal policy that attains sublinear utility and subinear queue length simultaneously, which we investigate in the next section. In particular, we show that the above lower bound can be asymptotically attained by some causal policy.

B. Algorithm Performance in W-Constrained Networks

In this section, we analyze the tradeoffs between utility regret and queue length achieved by two network control algorithms under the W-constrained adversary model. The first is the famous Drift-plus-Penalty algorithm [2] that was proved to achieve good utility-delay tradeoffs in stochastic networks. The second is a generalized version of the Tracking Algorithm [5], [6] that was originally proposed for Adversarial Queueing Theory. In particular, we show that the Tracking Algorithm attains the tradeoff lower bound in Theorem 2.

1) Drift-plus-Penalty Algorithm: In each slot t, the Driftplus-Penalty algorithm observes the current network event ω_t and the queue length vector $\mathbf{Q}(t)$, and choose the following control action α_t^{DP} :

$$\alpha_t^{DP} = \arg \max_{\alpha_t \in \mathcal{D}_{\omega_t}} \sum_i Q_i(t) \Big(b_i(\alpha_t, \omega_t) - a_i(\alpha_t, \omega_t) \Big) + VU(\alpha_t, \omega_t),$$
(6)

where V > 0 is a parameter controlling the tradeoffs between utility regret and queue length. Note that $\sum_i Q_i(t) \left(b_i(\alpha_t, \omega_t) - a_i(\alpha_t, \omega_t) \right)$ corresponds to the drift part while $VU(\alpha_t, \omega_t)$ is the penalty part.

The control actions in the Drift-plus-Penalty algorithm can be usually decomposed into several actions. For example, in one-hop networks without routing, $\mathbf{a}(t)$ corresponds to the admitted exogenous arrival vector in slot t. Suppose that the utility function is in the form $U(\alpha_t, \omega_t) = \sum_i U_i(a_i(t))$. Then the Drift-plus-Penalty algorithm can be decomposed into the solutions of two sub-problems.

• (Admission Control) Choose

$$\mathbf{a}(t) = \arg\max_{\mathbf{a}} \sum_{i} \Big(V U_i(a_i) - Q_i(t) a_i \Big).$$

• (Resource Allocation and Scheduling) Choose

$$\mathbf{b}(t) = \arg\max_{\mathbf{b}} \sum_{i} Q_i(t) b_i.$$

The first part is usually a convex optimization problem while the second part corresponds to the MaxWeight policy [1].

The following theorem gives the performance of the Driftplus-Penalty algorithm in *W*-constrained networks.

Theorem 3. In any W-constrained network, the Drift-plus-Penalty algorithm with parameter V achieves $O\left(\frac{TW}{V}\right)$ utility regret and the total queue length is $O\left(\sqrt{T(W+V)}\right)$.

Proof: The proof is based on the Lyapunov drift analysis. However, instead of considering the one-slot drift as in the traditional stochastic analysis, we find upper bounds on the *W*-slot drift-plus-penalty term and make sample-path arguments. See the technical report [14] for details.

There are several important observations about Theorem 3. First, if parameter V is set appropriately, then sublinear utility regret and sublinear queue length can be simultaneously achieved by the Drift-plus-Penalty algorithm in W-constrained networks as long as W = o(T). For example, if $W = \Theta(T^{1/2})$, then setting $V = \Theta(T^{3/4})$ yields the utility regret of $O(T^{3/4})$ and the total queue length of $O(T^{7/8})$.

Noticing that sublinear utility regret and sublinear queue length cannot be achieved simultaneously by any causal policy if $W = \Omega(T)$ (Theorem 2), we have the following corollary.

Corollary 1. Under the W-constrained adversary model, sublinear utility regret and sublinear queue length are simultaneously achievable if and only if W = o(T).

Second, the performance of the Drift-plus-Penalty algorithm could be much worst than the lower bound in Theorem 2. For example, if $W = \Theta(T^{1/2})$, then one of the tradeoffs implied by the lower bound is that the utility regret is $\Theta(T^{1/2})$ and the total queue length is also $\Theta(T^{1/2})$, which is not achievable by the Drift-plus-Penalty algorithm. In the next section, we develop an algorithm that has a better performance and attains the lower bound.

2) Tracking Algorithm: The tradeoff bounds achieved by the Drift-plus-Penalty algorithm is relatively loose as compared to the lower bound in Theorem 2. In this section, we develop the Tracking Algorithm that has a better performance and attains the lower bound in Theorem 2.

The original Tracking Algorithm was proposed in [5], [6] to solve a scheduling problem under the Adversarial Queueing Theory model. However, it only works for a very specific network model: (i) the network has to be singlehop where the arrival vector is independent of the control action, and (ii) the control action has to satisfy the primary interference constraints, i.e., only one link incident on the same node can be activated in each slot. Next, we extend the original Tracking Algorithm to accommodate the general network model considered in this paper.

Let Ω be the set of all possible network events that could happen in each slot. In order for the Tracking Algorithm to work, the cardinality of Ω has to be finite (otherwise it could be discretized into a finite set as in [5]). For example, in a single-hop network, suppose each network event ω_t corresponds to a couple $(\mathbf{A}(t), \mathbf{S}(t))$ where $\mathbf{A}(t)$ is a vector of exogenous packet arrivals in slot t and $\mathbf{S}(t)$ a vector of link states in slot t. For any link i and time t, assume that $0 \le A_i(t) \le B$ and $A_i(t)$ is an integer, and each link only has a finite number of S states. Then $|\Omega| = (SB)^N$.

The Tracking Algorithm is given in Algorithm 1. It maintains an action queue Q_{ω} for each type of network events $\omega \in \Omega$. The action queue \mathcal{Q}_{ω} stores the optimal actions that the Tracking Algorithm should have taken when network event ω occurred. Note that the sequence of optimal control actions cannot be calculated online but can be calculated every W slots due to the window structure (5). In the Tracking Algorihtm, the sequence of optimal actions during each window are added to the action queues in batch at the end of this window (steps 8-9). Here, the optimal actions during a window [t - W + 1, t] corresponds to any optimal solution to (7) (which is also a part of the optimal solution to NUM). In each slot t, the Tracking Algorithm first observes the current network event $\omega_t = \omega$. If the corresponding action queue \mathcal{Q}_{ω} is not empty (i.e., there are some actions we should have taken but have not taken yet), the algorithm just sets the control action as the first action in the action queue Q_{ω} , and the action is removed from the action queue \mathcal{Q}_{ω} (steps 3-5). If the action queue is empty, the algorithm may take any feasible action. In our analysis, we assume that no action is taken when the action queue is empty.

$$\max_{\alpha_{\tau} \in \mathcal{D}_{\omega_{\tau}}} \sum_{\tau=t-W+1}^{t} U(\alpha_{\tau}, \omega_{\tau})$$
s.t.
$$\sum_{\tau=t-W+1}^{t} a_{i}(\alpha_{\tau}, \omega_{\tau}) \leq \sum_{\tau=t-W+1}^{t} b_{i}(\alpha_{\tau}, \omega_{\tau}), \quad \forall i$$
(7)

The following theorem gives the tradeoff between utility regret and queue length achieved by the Tracking Algorithm under the *W*-constrained adversary model.

Theorem 4. In any W-constrained network, the Tracking Algorithm achieves O(W) utility regret and the total queue length is O(W).

Proof: Since the Tracking Algorithm updates the optimal actions every W slots and replays these actions whenever possible, the number of unfulfilled actions in any action queue is at most W. Thus, the performance gap between the

Algorithm 1 Tracking Algorithm (TA)

- 1: Initialize $Q_{\omega} = \emptyset$ for each $\omega \in \Omega$.
- 2: for $t = 0, \cdots, T 1$ do
- 3: Observe the current network event $\omega_t = \omega$.
- 4: **if** action queue Q_{ω} is not empty **then**
- 5: Choose the control action α_t^{TA} as the first action in Q_{ω} and remove this action from Q_{ω} .
- 6: **end if**
- 7: **if** mod(t, W) = W 1 **then**
- 8: Compute the sequence of optimal control actions $\{\alpha_{\tau}^*\}_{\tau=t-W+1}^t$ in the past window [t-W+1, t], which is any optimal solution to (7).
- 9: For each slot τ in the past window [t W + 1, t], enqueue the computed optimal action α_{τ}^* into the action queue $Q_{\omega_{\tau}}$, where ω_{τ} is the network event occurring in slot τ .

10: **end if**

11: end for

Tracking Algorithm and optimal policy is also O(W). See the technical report [14] for details.

There are several important observations about Theorem 4. First, under the W-constrained adversary model, sublinear utility regret and sublinear queue length can be simultaneously achieved by the Tracking Algorithm as long as W = o(T). Moreover, the tradeoff achieved by the Tracking Algorithm is better than that of the Drift-plus-Penlaty algorithm, in terms of their dependence on W and T. For example, if $W = \Theta(T^{1/2})$, the Tracking Algorithm can achieve $O(T^{1/2})$ utility regret and $O(T^{1/2})$ total queue length, while such a tradeoff is not attainable by the Driftplus-Penalty algorithm.

Second, the Tracking Algorithm asymptotically achieves the lower bound in Theorem 2 in the sense that it ensures that $\mathcal{R}_T(\{\omega_0, \cdots, \omega_{T-1}\}) + \sum_i Q_i(T) = O(W)$ for any $\{\omega_0, \cdots, \omega_{T-1}\} \in \mathcal{W}_T$. As a result, the Tracking Algorithm asymptotically achieves the optimal tradeoff between utility regret and queue length w.r.t. W and T.

Third, the Tracking Algorithm needs to maintain a virtual queue for each type of network events while the size of the network event space Ω may be exponential in the number of users N. As a result, the Tracking Algorithm may not be a practical algorithm. The purpose of presenting the Tracking Algorithm is to demonstrate that the lower bound in Theorem 2 could be asymptotically achieved by a causal policy. Note that Andrews and Zhang [5] proposed a method to get rid of the exponential dependence on N, at the expense of much more involved algorithm.

Finally, the Tracking Algorithm described in Algorithm 1 only achieves one point in the tradeoff curve since it only tracks the optimal solution to **NUM**. We leave it for future work to enable tunable tradeoffs for the Tracking Algorithm.

IV. V_T -Constrained Adversary Model

The aforementioned W-constrained model is relatively restrictive, where the stringent constraints (5) have to be satisfied for every window of W slots. In this section, we consider a general adversary model where the window constraints (5) are relaxed.

The new adversary model is parameterized by the inherent variation in the sequence of network events, which is measured as follows. Given a sequence of network events $\omega_0, \dots, \omega_{T-1}$ and a (possibly non-causal) policy, we define

$$V^{\pi}\Big(\{\omega_0,\cdots,\omega_{T-1}\}\Big) = \max_{t\leq T}\sum_i Q_i^{\pi}(t).$$

The above function measures the peak queue length achieved by policy π during its sample path. We further define $V^*(\{\omega_0, \dots, \omega_{T-1}\})$ to be the peak queue length during the sample path of the optimal solution to **NUM** under the sequence of network events $\omega_0, \dots, \omega_{T-1}$. If there are multiple optimal solutions to **NUM**, then the one with the smallest value of $V^*(\cdot)$ is considered. Note that $V^*(\cdot)$ only depends on $\{\omega_0, \dots, \omega_{T-1}\}$ and measures the inherent variations in the sequence of network events.

Now we define the notion of V_T -constrained network dynamics where the value of $V^*(\cdot)$ is constrained by some budget V_T .

Definition 3 (V_T -Constrained Dynamics). Given some $V_T \in [0, NTB]$, a sequence of network events $\omega_0, \dots, \omega_{T-1}$ is V_T -constrained if

$$V^*\Big(\{\omega_0,\cdots,\omega_{T-1}\}\Big) \le V_T$$

Any network satisfying the above is called a V_T -constrained network. Denote by \mathcal{V}_T the set of all possible sequences of network events that are V_T -constrained. A V_T -constrained adversary can only select the sequence of network events from the set \mathcal{V}_T .

Note that we restrict the range of V_T to [0, NTB] since the peak queue length within T slots is at most NTB. Any larger value of V_T has the same effect as $V_T = NTB$. Note also that the larger V_T is, the more variations the network could have. By varying the value of V_T from 0 to NTB, the above V_T -constrained adversary model covers a wide range of adversarial settings: from a strictly constrained adversary ($V_T = 0$, i.e., the arrivals should not exceed the services for each queue in every slot) to a completely unconstrained adversary ($V_T = NTB$).

A. Lower Bound on the Tradeoffs

The following theorem provides a lower bound on the tradeoffs between utility regret and queue length under the V_T -constrained adversary model.

Theorem 5. For any causal policy π , there exists a sequence of network events $\{\omega_0, \dots, \omega_{T-1}\} \in \mathcal{V}_T$ such that

$$\mathcal{R}_T^{\pi}\Big(\{\omega_0,\cdots,\omega_{T-1}\}\Big) + c\sum_i Q_i^{\pi}(T) \ge c' V_T$$

where c' > 0 is some constant.

Proof: The proof is the same as that for Theorem 2 except that we replace W with V_T , thus omitted for brevity.

Theorem 5 shows that if $V_T = \Omega(T)$, then no causal policy can simultaneously achieve sublinear utility regret and sublinear queue length under the V_T -constrained adversary model. On the other hand, if $V_T = o(T)$, there *might* exist some causal policy that attains sublinear utility regret and sublinear queue length simultaneously, which we investigate in Section IV-B.

B. Algorithm Performance in V_T-Constrained Networks

In this section, we analyze the tradeoffs between utility regret and queue length achieved by two algorithms in V_T constrained networks: the Drift-plus-Penalty algorithm and the Tracking Algorithm. In particular, we show that both algorithms simultaneously achieve sublinear utility regret and sublinear queue length if $V_T = o(T)$.

1) Drift-plus-Penalty Algorithm: The Drift-plus-Penalty algorithm discussed in Section III-B can be directly applied to the V_T -constrained setting. The following theorem gives the tradeoffs between utility regret and queue length achieved by the Drift-plus-Penalty algorithm under the V_T -constrained adversary model.

Theorem 6. In any V_T -constrained network, the Driftplus-Penalty algorithm with parameter V > 0 achieves $O\left(\frac{V_T^{2/3}T^{4/3}}{V} + \frac{V_T^{1/3}T^{7/6}}{V^{1/2}}\right)$ utility regret and the total queue length is $O\left(V_T^{1/3}T^{2/3} + T^{1/2}V^{1/2}\right)$.

Proof: We first divide the time horizon into frames of W slots. Then we apply the analysis used in the W-constrained adversary model and derive bounds on the W-slot drift-plus-penalty term, which further leads to upper bounds on utility regret and queue length. The value of W is carefully chosen to optimize these bounds. See the technical report [14] for details.

There are several observations about Theorem 6. First, the Drift-plus-Penalty algorithm achieves sublinear utility regret and sublinear queue length under the V_T -constrained adversary model whenever $V_T = o(T)$. For example, if $V_T = \Theta(T^{1/2})$ and we set $V = \Theta(T^{4/5})$, then the utility regret and the total queue length are both $O(T^{11/12})$. Notice that sublinear utility regret and sublinear queue length cannot be simultaneously achieved by any causal policy if $V_T = \Omega(T)$ (Theorem 5). We have the following corollary.

Corollary 2. Under the V_T -constrained adversary model, sublinear utility regret and sublinear queue length are simultaneously achievable if and only if $V_T = o(T)$.

Second, the Drift-plus-Penalty algorithm does not attain the lower bound in Theorem 5. For example, if $V_T = \Theta(T^{1/2})$, one of the tradeoffs implied by the lower bound is that the utility regret is $\Theta(T^{1/2})$ and the total queue length is also

 $\Theta(T^{1/2})$, which is not achievable by the Drift-plus-Penalty algorithm. In fact, although the Drift-plus-Penalty algorithm can achieve sublinear utility regret and sublinear queue length, the tradeoff bound in Theorem 6 is relatively loose. In the next section, we show that the Tracking Algorithm can achieve a better tradeoff bound than the Drift-plus-Penalty algorithm.

2) Tracking Algorithm: The Tracking Algorithm introduced under the W-constrained adversary model requires that the window constraints (5) be satisfied for some window size W. However, there might be no window structure under the V_T -constrained adversary model and thus the Tracking Algorithm cannot be directly applied in V_T -constrained networks. We slightly modify the Tracking Algorithm in two aspects. First, the window size W is set to be $W = \sqrt{TV_T}$ under the V_T -constrained adversary model. Second, in step 8 of the original Tracking Algorithm, the optimization problem (7) is modified to be

$$\max \sum_{\tau=t-W+1}^{t} U(\alpha_{\tau}, \omega_{\tau})$$

s.t.
$$\sum_{\tau=t-W+1}^{t} a_{i}(\alpha_{\tau}, \omega_{\tau}) \leq \sum_{\tau=t-W+1}^{t} b_{i}(\alpha_{\tau}, \omega_{\tau}) + V_{T}, \ \forall i$$
$$\alpha_{\tau} \in \mathcal{D}_{\omega_{\tau}}, \ \forall \tau.$$
(8)

In particular, the first constraint in (7) is relaxed by allowing some bursts up to V_T . Note that by the definition of V_T constrained networks, the optimal solution to **NUM** is also a feasible solution to (8). Under the above setting, the utility regret and the total queue length achieved by the Tracking Algorithm in V_T -constrained networks is given by the following theorem³.

Theorem 7. Under the V_T -constrained adversary model, the Tracking Algorithm achieves $O(\sqrt{TV_T})$ utility regret and the total queue length is $O(\sqrt{TV_T})$.

Proof: The proof is similar to the analysis under the W-constrained adversary model, except that an additional V_T terms is added in the first constraint of (8). See the technical report [14] for details.

There are several important observations about Theorem 7. First, the Tracking Algorithm can simultaneously achieve sublinear utility regret and sublinear queue length whenever $V_T = o(T)$. Second, the performance of the Tracking Algorithm is better than that under the Drift-plus-Penalty algorithm in V_T -constrained networks. For example, if we set $W = \Theta(\sqrt{V_T T})$ and $V_T = \Theta(\sqrt{T})$, then the Tracking Algorithm achieves $O(T^{3/4})$ utility regret and $O(T^{3/4})$ queue length, which is not achievable by the Drift-plus-Penalty algorithm. Finally, the Tracking Algorithm does not attain the tradeoff lower bound in Theorem 5. Thus, finding a causal policy that can close the gap remains an open problem.

C. Relationship between Adversary Models

The V_T -constrained adversary model generalizes the Wconstrained adversary model: any sequence of network events that are W-constrained must also be V_T -constrained with $V_T = O(W)$ due to the window structure (note that the peak queue length under the optimal policy is at most NWB). The analysis in the V_T -constrained adversary model also gives a more general condition for sublinear utility regret and sublinear queue length.

V. SIMULATIONS

In this section, we empirically validate the theoretical bounds derived in this paper and compare the performance of the Drift-plus-Penalty and the Tracking Algorithm.

In our simulations, we consider a single-hop network with N = 2 users. In each slot t, the central controller observes the current network event $\omega_t = (\mathbf{A}(t), \mathbf{S}(t))$, where $\mathbf{A}(t)$ is the exogenous arrival vector and $\mathbf{S}(t)$ is the channel rate vector for each link in slot t. Then the controller makes an admission control and a scheduling decision. The constraint on the admission control action is $0 \le a_i(t) \le A_i(t)$ for each link *i*, and the constraint on the scheduling decision is that at most one of the links can be served in each slot. The network utility is $U(\alpha_t, \omega_t) = \sum_i \log (1 + a_i(t))$ (proportional fairness). We consider a scenario where the channel rate vector in each slot is controlled by an adaptive adversary. Time is divided into frames of W slots. In the first [W/2] slots of each frame, the exogenous arrivals to each user are 10 packets/slot and the channel rate for each user is also 10 packets/slot. In the remaining slots of each frame, there are no exogenous arrivals to both users while the channel rate is zero for the user with a longer queue and 10 packets/slot for the other user. If the two users have the same queue length, ties are broken randomly. Such a scenario is similar to the one that we use to prove the tradeoff lower bound under the W-constrained adversary model (see the proof of Theorems 2), and it has been shown that this is a W-constrained adversary (and also a V_T -constrained adversary with $V_T = 5W$).

Figure 1 illustrates the growth of the total queue length and the utility regret with the time horizon T under the Drift-plus-Penalty algorithm (with different values of V) and the Tracking algorithms. First, when $W = \Theta(\sqrt{T})$, the Drift-plus-Penalty can simultaneously achieve sublinear utility regret and sublinear queue length, if the parameter V is set appropriately (for example, $V = \Theta(T^{3/4})$). Note that setting V to some very large value (e.g., $V = \Theta(T^2)$) still achieves sublinear utility regret and sublinear queue length, though the theoretical bound on queue length (see Theorem 3) is at least linear in T when $V = \Omega(T)$, which shows that the performance upper bound is not tight in this scenario. The Tracking Algorithm also simultaneously achieves sublinear utility regret and sublinear queue length when $W = \Theta(\sqrt{T})$. However, when $W = \Theta(T)$, both algorithms fail to achieve desirable performance: either the utility regret or the queue length grows linearly with T.

³As is discussed in Section III-B2, the set of possible network events should be finite in order for the Tracking Algorithm to work.



Fig. 1. Growth of total queue length and utility regret with the time horizon T under an adaptive W-adversary.

In fact, the lower bound in Theorem 2 shows that no causal policy can achieve both sublinear utility regret and sublinear queue length if $W = \Theta(T)$.

Figure 2 shows the tradeoffs between utility regret and queue length under the Drift-plus-Penalty algorithm and the Tracking Algorithm, where we fix the time horizon to be $T = 10^4$ slots and the window size $W = \Theta(\sqrt{T})$. Note that for the Drift-plus-Penalty algorithm, we plot a tradeoff curve (since it achieves different tradeoffs by tuning the parameter V), while only a single tradeoff point is plotted for the Tracking Algorithm. It is observed that the Tracking Algorithm achieves a better tradeoff point that is not achievable by the Drift-plus-Penalty algorithm. In addition, the theoretical lower bound for any causal policy (Theorem 2) and the theoretical performance upper bounds for both algorithms (Theorems 3 and 4) are also validated in the figure.



Fig. 2. Tradeoffs between utility and total queue length (double log scale). The time horizon is fixed to be $T = 10^4$ slots and $W = \Theta(\sqrt{T})$.

VI. CONCLUSIONS

In this paper, we focus on optimizing network utility within a finite time horizon under adversarial network models. We show that no causal policy can simultaneously achieve both sublinear utility regret and sublinear queue length if the network dynamics are unconstrained, and investigate two constrained adversary models. We first consider the restrictive W-constrained adversary model and then propose a more relaxed V_T -constrained adversary model. Lower bounds on the tradeoffs between utility regret and queue length are derived under the two adversary models, and the performance of two control policies is analyzed, i.e., the Drift-plus-Penalty algorithm and the Tracking Algorithm. It is shown that the Tracking Algorithm asymptotically attains the optimal tradeoffs under the *W*-constrained adversary model and that the Tracking Algorithm has a better tradeoff bound than that of the Drift-plus-Penalty policy.

REFERENCES

- 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, 1992.
- [2] M. J. Neely, E. Modiano, and C.-P. Li, "Fairness and optimal stochastic control for heterogeneous networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 16, no. 2, pp. 396–409, 2008.
- [3] Y. Zou, J. Zhu, X. Wang, and L. Hanzo, "A survey on wireless security: Technical challenges, recent advances, and future trends," *Proceedings of the IEEE*, vol. 104, no. 9, pp. 1727–1765, 2016.
- [4] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synthesis Lectures on Communication Networks*, vol. 3, no. 1, pp. 1–211, 2010.
- [5] M. Andrews and L. Zhang, "Scheduling over a time-varying userdependent channel with applications to high speed wireless data," in *The 43rd Annual IEEE Symposium on Foundations of Computer Science*, 2002. Proceedings., 2002, pp. 293–302.
- [6] —, "Scheduling over nonstationary wireless channels with finite rate sets," *IEEE/ACM Transactions on Networking*, vol. 14, no. 5, pp. 1067–1077, Oct 2006.
- [7] R. L. Cruz, "A calculus for network delay. i. network elements in isolation," *IEEE Transactions on information theory*, vol. 37, no. 1, pp. 114–131, 1991.
- [8] A. Borodin, J. Kleinberg, P. Raghavan, M. Sudan, and D. P. Williamson, "Adversarial queuing theory," *Journal of the ACM* (*JACM*), vol. 48, no. 1, pp. 13–38, 2001.
- [9] M. Andrews, B. Awerbuch, A. Fernández, T. Leighton, Z. Liu, and J. Kleinberg, "Universal-stability results and performance bounds for greedy contention-resolution protocols," *Journal of the ACM (JACM)*, vol. 48, no. 1, pp. 39–69, 2001.
- [10] V. Cholvi and J. Echagüe, "Stability of fifo networks under adversarial models: State of the art," *Computer Networks*, vol. 51, no. 15, pp. 4460–4474, 2007.
- [11] 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 Thirty-ninth Annual ACM Symposium* on Theory of Computing, ser. STOC '07. ACM, 2007, pp. 145–154.
- [12] S. Lim, K. Jung, and M. Andrews, "Stability of the max-weight protocol in adversarial wireless networks," *IEEE/ACM Trans. Netw.*, vol. 22, no. 6, pp. 1859–1872, Dec. 2014.
- [13] M. J. Neely, "Universal scheduling for networks with arbitrary traffic, channels, and mobility," in *Decision and Control (CDC), 2010 49th IEEE Conference on*. IEEE, 2010, pp. 1822–1829.
- [14] Q. Liang and E. Modiano, "Network utility maximization in adversarial environments," *Technical Report*, https://arxiv.org/abs/1712.08672, 2016.