Minimizing Queue Length Regret Under Adversarial Network Models

Qingkai Liang, Eytan Modiano
LIDS, Massachusetts Institute of Technology, USA

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 focus on minimizing queue length regret under adversarial network models, which measures the finite-time queue length difference between a causal policy and an “oracle” that knows the future. Two adversarial network models are developed to characterize the adversary’s behavior. We provide lower bounds on queue length regret under these adversary models and analyze the performance of two control policies (i.e., the MaxWeight policy and the Tracking Algorithm). We further characterize the stability region under adversarial network models, and show that both the MaxWeight policy and the Tracking Algorithm are throughput-optimal even in adversarial settings.

ACM Reference Format:

Qingkai Liang, Eytan Modiano. 2018. Minimizing Queue Length Regret Under Adversarial Network Models. In *SIGMETRICS '18 Abstracts: ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems Abstracts, June 18–22, 2018, Irvine, CA, USA*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3219617.3219630>

1 BACKGROUND AND MOTIVATION

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 policy [7]) have been derived to optimize network performance (e.g., throughput) 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 [8], where traffic

This work was supported by NSF Grant CNS-1524317 and by DARPA I2O and Raytheon BBN Technologies under Contract No. HROO11-1 5-C-0097. The full version of this paper can be found in [4].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS '18 Abstracts, June 18–22, 2018, Irvine, CA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5846-0/18/06.

<https://doi.org/10.1145/3219617.3219630>

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, the 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 optimize network performance (i.e., queue length) within a finite time horizon in an adversarial environment.

2 MAIN RESULTS

Consider a network with N queues. 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.

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 $\mathcal{D}_{\omega_t}$ that can depend on ω_t . The network event ω_t and the control action α_t together produce the service vector $\mathbf{b}(\alpha_t, \omega_t) \triangleq \mathbf{b}(t) = (b_1(t), \dots, b_N(t))$ 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 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). Let $\mathbf{Q}(t) = (Q_1(t), \dots, Q_N(t))$ be the queue length vector at the beginning of slot t (before the arrivals in that slot). The queuing dynamics are

$$Q_i(t+1) = [Q_i(t) + a_i(t) - b_i(t)]^+,$$

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

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 controlled 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.

Our objective is to find a causal control policy that keeps the total queue length as small as possible. We propose a new finite-time performance metric, referred to as *queue length regret*:

$$\mathcal{R}_T^\pi = \sum_i Q_i^\pi(T) - \sum_i Q_i^*(T),$$

where $\sum_i Q_i^\pi(T)$ is the total queue length achieved by control policy π after a finite time horizon T , and $\sum_i Q_i^*(T)$ is the minimum queue length achieved by some “oracle” that has perfect knowledge about the future. A desirable first order characteristic of a “good” policy π is that it achieves sublinear regret $\mathcal{R}_T^\pi = o(T)$ such that $\mathcal{R}_T^\pi/T \rightarrow 0$ as the time horizon $T \rightarrow \infty$. In other words, the time-average queue growth rate asymptotically approaches the one achieved by the optimal non-causal policy. Under mild conditions, sublinear queue length regret is equivalent to *rate stability*, i.e., the long-term arrival rate to each queue equals the long-term departure rate from that queue.

We first prove that it is impossible to achieve “low” queue length regret if the adversary is unconstrained. In particular, there exist some adversarial network dynamics such that the queue length regret 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 then study two adversarial network models where the network dynamics are constrained to some “admissible” set. In particular, we first consider the (W, ϵ) -constrained adversary model, where the total arrivals are less than $(1 - \epsilon)$ times of the total services within any window of W slots. Although this window-based model is relatively limited, it is widely used by existing works (e.g., [1–3, 5, 6]) due to its analytical tractability and serves as a foundation for understanding more generalized adversary models.

Observing the limitation of (W, ϵ) -constrained model, we then propose a more generalized V_T -constrained model, where the total queue length generated by the “oracle” during its sample path is upper bounded by V_T . By varying the values V_T , the proposed V_T -constrained adversary model can cover a wide range of adversarial settings: from a strictly constrained adversary to a fully unconstrained adversary.

Under the above two adversary models, we develop lower bounds on queue length regret. It is shown that no causal policy can achieve sublinear queue length regret if W or V_T grows linearly with T . We also analyze the queue length regret of two control algorithms: the MaxWeight policy [7] and the Tracking Algorithm [2, 3] under the two adversarial models. In particular, both the MaxWeight policy and the Tracking Algorithm achieve sublinear queue length regret whenever W or V_T grows sublinearly with T , yet the theoretical regret bound under the Tracking Algorithm is better than that under the MaxWeight policy. The Tracking Algorithm is also asymptotically regret-optimal under the (W, ϵ) -constrained adversary model. We summarize these results in Table 1.

Table 1: Queue Length Regret Bounds

	(W, ϵ) -Constrained Adversary	V_T -Constrained Adversary
Lower Bound	$\Omega(W)$	$\Omega(V_T)$
MaxWeight	$O(\sqrt{TW})$ if $\epsilon = 0$ $O(W/\epsilon^3)$ if $\epsilon > 0$	$O(V_T^{1/3}T^{2/3})$
Tracking Alg.	$O(W)$	$O(\sqrt{TV_T})$

Finally, based on the above analytical results and the observation that sublinear queue length regret is equivalent rate stability, we characterize the stability region under adversarial network models, and show that both the MaxWeight policy and the Tracking Algorithm are throughput-optimal even in adversarial settings.

REFERENCES

- [1] Matthew Andrews, Kyomin Jung, and Alexander Stolyar. 2007. 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 (STOC '07)*. ACM, 145–154.
- [2] M. Andrews and L. Zhang. 2002. Scheduling over a time-varying user-dependent channel with applications to high speed wireless data. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.* 293–302. <https://doi.org/10.1109/SFCS.2002.1181952>
- [3] M. Andrews and L. Zhang. 2006. Scheduling Over Nonstationary Wireless Channels With Finite Rate Sets. *IEEE/ACM Transactions on Networking* 14, 5 (Oct 2006), 1067–1077. <https://doi.org/10.1109/TNET.2006.882835>
- [4] Qingkai Liang and Eytan Modiano. 2018. Minimizing Queue Length Regret Under Adversarial Network Models. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2, 1 (2018), 11.
- [5] Sungsu Lim, Kyomin Jung, and Matthew Andrews. 2014. Stability of the Max-weight Protocol in Adversarial Wireless Networks. *IEEE/ACM Trans. Netw.* 22, 6 (Dec. 2014), 1859–1872.
- [6] Michael J Neely. 2010. Universal scheduling for networks with arbitrary traffic, channels, and mobility. In *Decision and Control (CDC), 2010 49th IEEE Conference on.* IEEE, 1822–1829.
- [7] Leandros Tassioulas and Anthony Ephremides. 1992. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE transactions on automatic control* 37, 12 (1992), 1936–1948.
- [8] Yulong Zou, Jia Zhu, Xianbin Wang, and Lajos Hanzo. 2016. A survey on wireless security: Technical challenges, recent advances, and future trends. *Proc. IEEE* 104, 9 (2016), 1727–1765.