Heavy-Traffic-Optimal Scheduling with Regular Service Guarantees in Wireless Networks

Bin Li
Department of ECE
The Ohio State University
Columbus, OH 43210, USA
lib@ece.osu.edu

Ruogu Li
Department of ECE
The Ohio State University
Columbus, OH 43210, USA
lir@ece.osu.edu

Atilla Eryilmaz
Department of ECE
The Ohio State University
Columbus, OH 43210, USA
eryilmaz@ece.osu.edu

ABSTRACT

We consider the design of throughput-optimal scheduling policies in multi-hop wireless networks that also possess good mean delay performance and provide regular service for all links – critical metrics for real-time applications. To that end, we study a parametric class of maximum-weight type scheduling policies with parameter $\alpha \geq 0$, called Regular Service Guarantee (RSG) Algorithm, where each link weight consists of its own queue-length and a counter that tracks the time since the last service. This policy has been shown to be throughput-optimal and to provide more regular service as the parameter α increases, however at the cost of increasing mean delay.

This motivates us to investigate whether satisfactory service regularity and low mean-delay can be simultaneously achieved by the RSG Algorithm by carefully selecting its parameter α. To that end, we perform a novel Lyapunov-drift based analysis of the steady-state behavior of the stochastic network. Our analysis reveals that the RSG Algorithm can minimize the total mean queue-length to establish mean delay optimality under heavily-loaded conditions as long as α scales no faster than the order of $\frac{1}{\sqrt[5]{\epsilon}}$, where ϵ measures the closeness of the network load to the boundary of the capacity region. To the best of our knowledge, this is the first work that provides regular service to all links while also achieving heavy-traffic optimality in mean queue-lengths.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: Miscellaneous

General Terms

Performance

Keywords

Wireless scheduling, service regularity, throughput, mean delay, heavy-traffic analysis

MobiHoc'13, July 29–August 1, 2013, Bangalore, India.

1. INTRODUCTION

Real-time applications, such as voice over IP or live multimedia streaming, are becoming increasingly popular as smart phones proliferate in wireless networks. To support real-time applications, network algorithm design should not only efficiently manage the interference among simultaneous transmissions, but also meet the requirements of Quality-of-Service (QoS) including delay, packet delivery ratio, and jitter. Such QoS requirements, in turn, depend on the higher-order statistics of the arrival and service process, which poses significant challenges for effective network algorithm design. For example, the well-known Maximum Weight Scheduling (MWS) Algorithm (e.g., [12], [11]) that prioritizes service of links with the largest backlog levels achieves maximum throughput, but does not provide any guarantees on the regularity of service that most real-time applications demand.

In recent years, there has been an increasing understanding on the algorithm design that target various aspects of QoS, especially packet delivery ratio requirement (e.g., [5], [6], [7]). However, there has been considerably less progress associated with the regularity of service, which is clearly important for real-time applications with stringent jitter requirements. Our work is motivated by the recent advances made in [8] that provides a promising approach for managing this critical QoS metric. In particular, [8] provides a throughput-optimal algorithm that is parametrized with a design variable $\alpha \geq 0$, which improves service regularity as α increases (see Section 3 for more details). Yet, increasing α also has an averse effect on the mean delay performance, which is also vital for most applications.

With this motivation, this paper focuses on the tradeoff between the service regularity and the mean delay performance that this class of policies achieves. In particular, we are interested in identifying the range of values for α in which the mean delay performance guarantees can be provided, while the regularity characteristics are preserved. To that end, we build on the recently developed approach of using Lyapunov drifts for the steady-state analysis of queueing networks [2]. *The main result emanating from this analysis is the scaling law of α as the system gets more and more heavily loaded so that the algorithm is mean delay optimal among all feasible scheduling policies, and provides the best service regularity among this class of policies.* Specifically, we show that the heavy-traffic optimality is preserved as long as α scales[1] as $O\left(\frac{1}{\sqrt[5]{\epsilon}}\right)$, where ϵ is the heavy-traffic param-

[1]We say $a_n = O(b_n)$ if there exists a $c > 0$ such that $|a_n| \leq c|b_n|$ for two real-valued sequences $\{a_n\}$ and $\{b_n\}$.

eter characterizing the closeness of the arrival rate vector to the boundary of the capacity region.

Our analysis relates to the vast literature on heavy-traffic analysis of queueing networks (for example, [13], [3], [1], [14], [10], [9]), and in particular extends the Lyapunov drift-based approach in [2]. A critical step in most of these results is to establish a *state-space collapse* along a single dimension, and thus relate the multi-dimensional system operation to a *resource-pooled* single dimensional system. Our construction also follows such line of argument in broad strokes. However, the new dynamics of the considered class of algorithms require new Lyapunov functions and techniques in establishing their heavy-traffic optimality.

Note on Notation: We use bold and script font of a variable to denote a vector and a set. Also, let $|\mathcal{A}|$ to denote the cardinality of the set $\mathcal{A}$. We use $\text{Int}(\mathcal{A})$ to denote the set of interior points of the set $\mathcal{A}$. We use $\langle \mathbf{x}, \mathbf{y} \rangle, \mathbf{x} \cdot \mathbf{y}$ to denote the inner product and component-wise product of the vector $\mathbf{x}$ and $\mathbf{y}$, respectively. We use $\mathbf{x}^2$ and $\sqrt{\mathbf{x}}$ to denote the component-wise square and square root of the vector $\mathbf{x}$, respectively. We also use $\preceq, \succeq, \prec, \succ$ to denote component-wise comparison of two vectors, respectively. Let $\|\mathbf{x}\|_1$ and $\|\mathbf{x}\|$ denote the l_1 and l_2 norm of the vector $\mathbf{x}$, respectively.

2. SYSTEM MODEL

We consider a wireless network represented by a graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$, where $\mathcal{N}$ is the set of nodes and $\mathcal{L}$ is the set of links. A node represents a wireless transmitter or receiver, while a link represents a pair of transmitter and receiver that are within the transmission range of each other. We use $L \triangleq |\mathcal{L}|$ for convenience. We consider the *link-based conflict model*, where links conflicting with each other cannot be active at the same time. We call a set of links that can be active simultaneously as a *feasible schedule* and denote it as $\mathbf{S}[t] = (S_l[t])_{l \in \mathcal{L}}$, where $S_l[t] = 1$ if the link l is scheduled in time slot t and $S_l[t] = 0$, otherwise.

We capture the channel fading over link l in time slot t via a *non-negative-integer-valued* random variable $C_l[t]$, with $C_l[t] \leq C_{\max}, \forall l, t$, for some $C_{\max} < \infty$, which measures the maximum amount of service available in slot t, if scheduled. We use $\mathcal{J}$ to denote the set of global channel states (with finite cardinality). Let $J[t] \in \mathcal{J}$ denote the global state of the channel states of all links in time slot t. We assume that $\{J[t] \in \mathcal{J}\}_{t \geq 0}$ is an independently and identically distributed (i.i.d.) sequence of random variables with $\psi_j \triangleq \Pr\{J[t] = j\}$. Let $\mathcal{S}^j$ denote the set of feasible schedules when the channel is in state $j \in \mathcal{J}$. Then, the *capacity region* is defined as

$$\mathcal{R} \triangleq \sum_{j \in \mathcal{J}} \psi_j \cdot \text{CH}\{\mathcal{S}^j\}, \tag{1}$$

where $\text{CH}\{\mathcal{A}\}$ denotes the convex hull of the set $\mathcal{A}$.

We assume a per-link traffic model[2], where $A_l[t]$ denotes the number of packets arriving at link l in slot t that are independently distributed over links and i.i.d. over time with finite mean λ_l, and $A_l[t] \leq A_{\max}, \forall l, t$, for some $A_{\max} < \infty$. Accordingly, a queue is maintained for each link l with $Q_l[t]$ denoting its queue length at the beginning of time slot t. Let $U_l[t] = \max\{0, C_l[t]S_l[t] - Q_l[t] - A_l[t]\}$ be the unused service for queue l in slot t. Then, the evolution of queue l is described as follows:

$$Q_l[t+1] = Q_l[t] + A_l[t] - C_l[t]S_l[t] + U_l[t], \quad \forall l. \tag{2}$$

We say that the queue l is *strongly stable* if it satisfies

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}[Q_l[t]] < \infty.$$

We call an algorithm *throughput-optimal* if it makes all queues strongly stable for any arrival rate vector $\boldsymbol{\lambda} = (\lambda_l)_l$ that lies strictly within the capacity region.

Our goal is to design a *throughput-optimal* scheduling algorithm that also possesses the following desirable properties for satisfying the QoS requirements: (i) provides *regular services* in the sense that the *second-moment of the inter-service times* of the links is small; and (ii) achieves *low mean delay* in the sense that the total mean queue-lengths is small, especially in the regime where the system is *heavily-loaded* – when delay effects are most pronounced.

Next, we provide a *regular service scheduler* that possesses throughput-optimality and regular service guarantees, and then investigate its mean-delay performance under the heavy-traffic regime.

3. REGULAR SERVICE SCHEDULER

One of our goals is to provide regular services for each link, which is related to the second moment of the inter-service times. To characterize the inter-service time, we introduce a counter T_l for each link l, namely Time-Since-Last-Service (TSLS), to keep track of the time since link l was last served. In particular, each T_l increases by 1 in each time slot when link l has zero transmission rate, either because it is not scheduled, or because its channel is unavailable, i.e., $C_l[t] = 0$, and drops to 0, otherwise. More precisely, the evolution of T_l is described as follows:

$$T_l[t+1] = \begin{cases} 0 & \text{if } S_l[t]C_l[t] > 0; \\ T_l[t] + 1 & \text{if } S_l[t]C_l[t] = 0. \end{cases} \tag{3}$$

Thus, the TSLS records the link "age" since the last time it received service, and is closely related to the inter-service time. Indeed, the authors in [8] showed that the normalized second moment of the inter-service times of each link is proportional to the mean value of its TSLS for any stabilizing policy. Thus, the TSLS has a direct impact on service regularity: the smaller the mean TSLS value, the more regular the service.

This connection motivates the following maximum-weight type algorithm that uses a combination of queue-lengths and TSLS values as its weights, extending the algorithm in [8] to multi-hop fading networks:

Regular Service Guarantee (RSG) Algorithm:

In each time slot t, select the schedule as

$$\mathbf{S}^*[t] \in \arg\max_{\mathbf{S} \in \mathcal{S}} \sum_{l=1}^{L} (Q_l[t] + \alpha T_l[t]) C_l[t] S_l, \tag{4}$$

where $\alpha > 0$ is a design parameter.

Note that the RSG Algorithm coincides with the MWS Algorithm when $\alpha = 0$. Yet, the true significance of the RSG algorithm is observed for large α, since as α increases,

[2]We note that our algorithm can be extended to serve multi-hop traffic, but the notion of service regularity is clearer in the per-link context.

the RSG Algorithm prioritizes the schedule with the larger TSLS, hence providing more regular services for each link. We can show that the RSG Algorithm in the multi-hop setup not only achieves throughput optimality but also provides regular service guarantees, which extends the results in [8].

Yet, large values of α may also deteriorate the mean delay performance. We demonstrate this tradeoff in a single-hop non-fading network with 4 links, where the number of packets arriving at each link follows a Bernoulli distribution with the arrival rate of 0.225. Figure 1 shows the mean delay and service regularity performance of the RSG Algorithm with varying α.

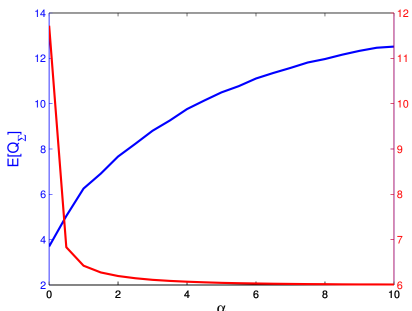

Figure 1: Performance of the RSG Algorithm

Figure 1 reveals that the improved regularity of the RSG Algorithm with increasing α comes at the cost of larger mean delays. We can show that the mean of the total TSLS value is minimized as α goes to ∞ (see [8]). On the other hand, it is known (e.g. [10, 2]) that the mean queue-lengths are minimized under heavily-loaded conditions (cf. Section 4 for more detail) when $\alpha = 0$. In view of the tradeoff observed in the above figure, our objective is to understand whether *both the regularity and the mean-delay optimality characteristics of the RSG Algorithm can be preserved, especially under heavily-loaded conditions, by carefully selecting* α.

In the next section, we answer this question in the affirmative by explicitly characterizing how α should scale with respect to the traffic load in order to achieve the heavy-traffic optimality while also preserving the regularity performance of the RSG Algorithm.

4. HEAVY-TRAFFIC OPTIMALITY RESULT

In this section, we present our main result for the RSG Algorithm in terms of its mean delay optimality under the heavy-traffic limit, where the arrival rate vector approaches the boundary of the capacity region.

We first note that the capacity region $\mathcal{R}$ is a polyhedron due to the discreteness and finiteness of the service rate choices, and thus has a finite number of faces. We consider the exogenous arrival vector process $\{\mathbf{A}^{(\epsilon)}[t]\}_{t\geq 0}$ with mean vector $\boldsymbol{\lambda}^{(\epsilon)} \in \text{Int}(\mathcal{R})$, where ϵ measures the Euclidean distance of $\boldsymbol{\lambda}^{(\epsilon)}$ to the boundary of $\mathcal{R}$ (see Figure 2).

In heavy-traffic analysis, we study the system performance as ϵ decreases to zero, i.e., as the arrival rate vector approaches $\boldsymbol{\lambda}^{(0)}$ belonging to the *relative interior* of a face, referred to as the *dominant hyperplane* $\mathcal{H}^{(\mathbf{c})}$. We denote $\mathcal{H}^{(\mathbf{c})} \triangleq \{\mathbf{r} \in \mathbb{R}^L : \langle \mathbf{r}, \mathbf{c}\rangle = b\}$, where $b \in \mathbb{R}$, and $\mathbf{c} \in \mathbb{R}^L$ is the normal vector of the hyperplane $\mathcal{H}^{(\mathbf{c})}$ satisfying $\|\mathbf{c}\| = 1$ and $\mathbf{c} \succeq \mathbf{0}$.

We are interested in understanding the steady-state queue-length values with vanishing ϵ. To that end, we first provide

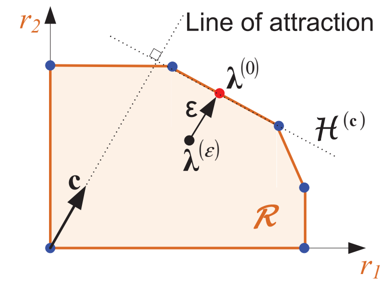

Figure 2: Geometric structure of capacity region

a generic lower bound for all feasible schedulers by constructing a hypothetical single-server queue with the arrival process $\langle \mathbf{c}, \mathbf{A}^{(\epsilon)}[t]\rangle$, and the i.i.d service process $\beta[t]$ with the probability distribution

$$\Pr\{\beta[t] = b_j\} = \psi_j, \qquad \text{for each } j \in \mathcal{J}, \tag{5}$$

where $b_j \triangleq \max_{\mathbf{s}\in\mathcal{S}^{(j)}} \langle \mathbf{c}, \mathbf{s}\rangle$ is the maximum $\mathbf{c}$-weighted service rate achievable in channel state $j \in \mathcal{J}$. By the construction of capacity region $\mathcal{R}$, we have $\mathbb{E}[\beta[t]] = b$. Also, it is easy to show that the constructed single-server queue-length $\{\Phi[t]\}_{t\geq 0}$ is stochastically smaller than the queue-length process $\{\langle \mathbf{c}, \mathbf{Q}^{(\epsilon)}[t]\rangle\}_{t\geq 0}$ under any feasible scheduling policy. Hence, by using Lemma 4 in [2], we have the following lower bound on the expected limiting queue-length vector under any feasible scheduling policy.

Proposition 1. *Let* $\overline{\mathbf{Q}}^{(\epsilon)}$ *be a random vector with the same distribution as the steady-state distribution of the queue length processes under any feasible scheduling policy. Consider the heavy-traffic limit* $\epsilon \downarrow 0$*, suppose that the variance vector* $\left(\boldsymbol{\sigma}^{(\epsilon)}\right)^2$ *of the arrival process* $\{\mathbf{A}^{(\epsilon)}[t]\}_{t\geq 0}$ *converges to a constant vector* $\boldsymbol{\sigma}^2$*. Then,*

$$\lim_{\epsilon\downarrow 0} \epsilon\mathbb{E}\left[\langle \mathbf{c}, \overline{\mathbf{Q}}^{(\epsilon)}\rangle\right] \geq \frac{\zeta}{2}, \tag{6}$$

where $\zeta \triangleq \langle \mathbf{c}^2, \boldsymbol{\sigma}^2\rangle + Var(\beta)$*.*

This fundamental lower bound of all feasible scheduling policies motivates the following definition of *heavy-traffic optimality* of a scheduler.

Definition 1. *(Heavy-Traffic Optimality) A scheduler is called heavy-traffic optimal, if its limiting queue length vector* $\overline{\mathbf{Q}}^{(\epsilon)}$ *satisfies*

$$\lim_{\epsilon\downarrow 0} \epsilon\mathbb{E}\left[\langle \mathbf{c}, \overline{\mathbf{Q}}^{(\epsilon)}\rangle\right] \leq \frac{\zeta}{2}, \tag{7}$$

where ζ *is defined in Proposition 1.*

It is well-known that the MWS Algorithm, which corresponds to the RSG Algorithm with $\alpha = 0$, is heavy-traffic optimal (e.g., [10, 2]). This is shown by first establishing a *state-space collapse*, i.e., the deviations of queue lengths from the direction $\mathbf{c}$ are bounded, independent of heavy-traffic parameter ϵ. Since the lower bound of mean queue length is of order of $\frac{1}{\epsilon}$, the deviations from the direction $\mathbf{c}$ are negligible compared to the large queue length for a sufficiently small ϵ, and thus the queue lengths concentrate along the normal vector $\mathbf{c}$. Because of this, we also call the normal vector $\mathbf{c}$ the *line of attraction*.

However, as discussed in Section 3, we are interested in large values of α to provide satisfactory service regularity. Yet, it is unknown whether the RSG Algorithm can remain heavy-traffic optimal when α is non-zero, since larger values of α leads to higher mean queue-lengths (cf. Figure 1). Also, the state-space collapse result is not applicable since the deviations from the line of attraction depend on α. This raises the question of how $\alpha(\epsilon)$ should scale with ϵ in order to achieve heavy-traffic optimality while allowing $\alpha(\epsilon)$ to take large values (providing more regular services). We answer this interesting and challenging question by providing the following main result, proved in Section 6.

Proposition 2. *Let $\overline{\mathbf{Q}}^{(\epsilon)}$ be a random vector with the same distribution as the steady-state distribution of the queue length processes under the RSG Algorithm. Consider the heavy-traffic limit $\epsilon \downarrow 0$, suppose that the variance vector $\left(\boldsymbol{\sigma}^{(\epsilon)}\right)^2$ of the arrival process $\{\mathbf{A}^{(\epsilon)}[t]\}_{t\geq 0}$ converges to a constant vector $\boldsymbol{\sigma}^2$. Suppose the channel fading satisfies the mild assumption*[3] $\Pr\{C_l[t]=0\}>0$, *for all $l \in \mathcal{L}$. Then,*

$$\epsilon\mathbb{E}\left[\langle \mathbf{c}, \overline{\mathbf{Q}}^{(\epsilon)}\rangle\right] \leq \frac{\zeta^{(\epsilon)}}{2} + \overline{B}^{(\epsilon)}, \tag{8}$$

where $\zeta^{(\epsilon)} \triangleq \langle \mathbf{c}^2, \left(\boldsymbol{\sigma}^{(\epsilon)}\right)^2\rangle + Var(\beta) + \epsilon^2$ and $\overline{B}^{(\epsilon)}$ is defined in (22).

Further, if $\alpha(\epsilon) = O(\frac{1}{\sqrt[5]{\epsilon}})$, then $\lim_{\epsilon\downarrow 0}\overline{B}^{(\epsilon)} = 0$ and thus the RSG Algorithm is heavy-traffic optimal.

This result is interesting in that it provides an explicit scaling regime in which the design parameter $\alpha(\epsilon)$ can be increased to utilize the service regulating nature of the RSG Algorithm without sacrificing the heavy-traffic optimality. Intuitively, if $\alpha(\epsilon)$ scales slowly as ϵ vanishes, each link weight is dominated by its own queue length in the heavy-traffic regime and thus the heavy-traffic optimality may be maintained; otherwise, the heavy-traffic optimality result may not hold, as will be demonstrated in the next section.

5. SIMULATION RESULTS

In this section, we provide simulation results to compare the mean delay and service regularity performance of the RSG Algorithm with the MWS Algorithm. In the simulation, we consider a single-hop non-fading network with 4 links. Its capacity region is $\mathcal{R} = \{\boldsymbol{\lambda} = (\lambda_l)_{l=1}^4 \succeq \mathbf{0} : \sum_{l=1}^4 \lambda_l < 1\}$. We use arrival process where the number of arrivals in each slot follows a Bernoulli distribution. We consider the symmetric case $\boldsymbol{\lambda}^{(\epsilon)} = (1-\frac{\epsilon}{2}) \times [\frac{1}{4},\frac{1}{4},\frac{1}{4},\frac{1}{4}]$, and the asymmetric case $\boldsymbol{\lambda}^{(\epsilon)} = [\frac{1}{2},\frac{1}{4},\frac{1}{8},\frac{1}{16}] + (1-\frac{\epsilon}{32}) \times [\frac{1}{64},\frac{1}{64},\frac{1}{64},\frac{1}{64}]$.

From Figure 3a and 4a, we can observe that the RSG Algorithm with both $\alpha = 1$ and $\alpha = \frac{1}{\sqrt[5]{\epsilon}}$, and the MWS Algorithm converge to the theoretical lower bound and thus is heavy-traffic optimal, which confirms our theoretical results. Yet, the RSG Algorithm with $\alpha = \frac{1}{\epsilon}$ has large mean

[3]We note that our result holds in single-hop network topologies without this assumption, and its extension to more general settings is part of our future work.

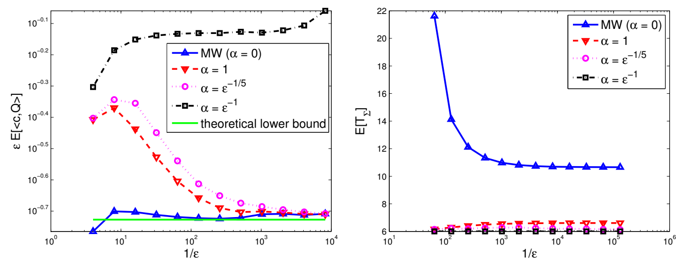

(a) Mean queue length (b) Service regularity

Figure 3: Symmetric arrivals in a single-hop network

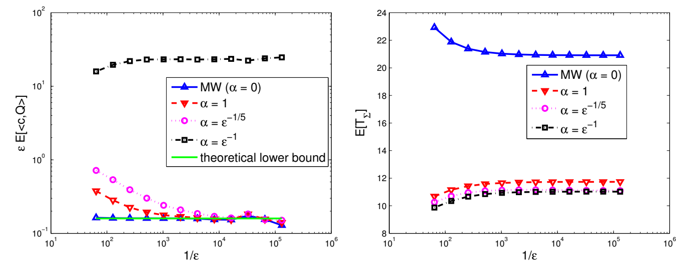

(a) Mean queue length (b) Service regularity

Figure 4: Asymmetric arrivals in a single-hop network

queue length, which does not match with the theoretical lower bound and thus is not heavy-traffic optimal. Hence, α should scale as slowly as $O(\frac{1}{\sqrt[5]{\epsilon}})$ to preserve heavy-traffic optimality.

From Figure 3b and 4b, we can see that the RSG Algorithm with even $\alpha = 1$ significantly outperforms the MWS Algorithm in terms of service regularity. More remarkably, the RSG Algorithm with $\alpha = \frac{1}{\sqrt[5]{\epsilon}}$ can achieve the lower bound (see [8]) achieved by the round robin policy under symmetric arrivals.

6. DETAILED HEAVY-TRAFFIC ANALYSIS

In this section, we prove Proposition 2 by using the analytical approach in [2], which includes two parts: (i) showing state-space collapse; (ii) using the state-space collapse result to obtain an upper bound on the mean queue lengths. Yet, it is worth noting that the strong coupling between queue length processes and TSLS counters in the RSG Algorithm poses significant challenges in heavy-traffic analysis. In particular, it requires new Lyapunov functions and a novel technique to establish heavy-traffic optimality of the RSG Algorithm.

6.1 State-Space Collapse

We have mentioned in Section 3 that the RSG Algorithm is throughput-optimal, i.e., it stabilizes all queues for any arrival rate vector that are strictly within the capacity region. Let $\{\mathbf{Q}^{(\epsilon)}\}_{t\geq 0}$ and $\{\mathbf{T}^{(\epsilon)}\}_{t\geq 0}$ be queue-length processes and TSLS counters under the RSG Algorithm, respectively. Also, we use $\overline{\mathbf{Q}}^{(\epsilon)}$ and $\overline{\mathbf{T}}^{(\epsilon)}$ to denote their limiting queue-length random vector and limiting TSLS random vector, respectively. Then, by the continuous mapping theorem, we have

$$\mathbf{Q}_{\|}^{(\epsilon)} \Rightarrow \overline{\mathbf{Q}}_{\|}^{(\epsilon)}, \qquad \mathbf{Q}_{\perp}^{(\epsilon)} \Rightarrow \overline{\mathbf{Q}}_{\perp}^{(\epsilon)}; \tag{9}$$

$$\mathbf{T}_{\|}^{(\epsilon)} \Rightarrow \overline{\mathbf{T}}_{\|}^{(\epsilon)}, \qquad \mathbf{T}_{\perp}^{(\epsilon)} \Rightarrow \overline{\mathbf{T}}_{\perp}^{(\epsilon)}, \tag{10}$$

where $\Rightarrow$ denotes convergence in distribution, and we define the projection and the perpendicular vector of any given L-dimensional vector $\mathbf{I}$ with respect to the normal vector $\mathbf{c}$ as:

$$\mathbf{I}_{\|} \triangleq \langle \mathbf{c}, \mathbf{I}\rangle \mathbf{c}, \quad \mathbf{I}_{\perp} \triangleq \mathbf{I} - \mathbf{I}_{\|}.$$

Next, we will show that under the RSG Algorithm, the second moment of $\|\overline{\mathbf{Q}}_{\perp}^{(\epsilon)}\|$ is bounded, dependent on $\alpha(\epsilon)$, while the second moment of $\|\overline{\mathbf{T}}^{(\epsilon)}\|$ is bounded by some constant independent of ϵ.

Proposition 3. *If $\Pr\{C_l[t]=0\} > 0, \forall l \in \mathcal{L}$, then, under RSG Algorithm, there exists a constant $N_{T,2}$, independent of ϵ, such that*

$$\mathbb{E}[\|\overline{\mathbf{Q}}_{\perp}^{(\epsilon)}\|^2] = O\left((\alpha(\epsilon))^4 (\log \alpha(\epsilon))^2\right), \tag{11}$$

$$\mathbb{E}[\|\overline{\mathbf{T}}^{(\epsilon)}\|^2] \le N_{T,2}. \tag{12}$$

We prove Proposition 3 by first studying the drift of the Lyapunov function

$$V_{\perp}(\mathbf{Q}^{(\epsilon)}, \mathbf{T}^{(\epsilon)}) \triangleq \|(\mathbf{Q}_{\perp}^{(\epsilon)}, \sqrt{2\alpha(\epsilon) C_{\max} \mathbf{T}^{(\epsilon)}})\|,$$

and show that when $V_{\perp}(\mathbf{Q}^{(\epsilon)}, \mathbf{T}^{(\epsilon)})$ is sufficiently large, it has a strictly negative drift independent of ϵ, which is characterized in the following key lemma.

Lemma 1. *Under the RSG Algorithm, there exist positive constants d and ς, independent of ϵ, such that whenever $V_{\perp}(\mathbf{Q}^{(\epsilon)}[t], \mathbf{T}^{(\epsilon)}[t]) > d$, we have*

$$\mathbb{E}[\Delta V_{\perp}(\mathbf{Q}^{(\epsilon)}[t], \mathbf{T}^{(\epsilon)}[t]) | \mathbf{Q}^{(\epsilon)}[t], \boldsymbol{T}^{(\epsilon)}[t]] < -\varsigma, \tag{13}$$

where $\Delta V_{\perp}(\mathbf{Q}^{(\epsilon)}[t], \mathbf{T}^{(\epsilon)}[t]) \triangleq V_{\perp}(\mathbf{Q}^{(\epsilon)}[t+1], \mathbf{T}^{(\epsilon)}[t+1]) - V_{\perp}(\mathbf{Q}^{(\epsilon)}[t], \mathbf{T}^{(\epsilon)}[t])$.

The proof of Lemma 1 is available in Appendix A.

Note that the TSLS counters have bounded increment but unbounded decrement, since they can at almost increase by 1 and drop to 0 once their corresponding links are scheduled. Due to this characteristic of TLSL, the absolute value of the drift $\Delta V_{\perp}(\mathbf{Q}^{(\epsilon)}, \mathbf{T}^{(\epsilon)})$ has neither an upper bound nor an exponential tail given the current system state $(\mathbf{Q}^{(\epsilon)}, \mathbf{T}^{(\epsilon)})$. Thus, we cannot directly apply Theorem 2.3 in [4], which requires either boundedness or the exponential tail of the Lyapunov drift to establish the existence of the second moment of the stochastic process. Indeed, for a Markov Chain with a strictly negative drift of Lyapunov function, if its Lyapunov drift has bounded increment but unbounded decrement, its second moment may not exist.

Counterexample: Consider a Markov Chain $\{X[t]\}_{t\ge 0}$ with the following transition probability:

$$P_{j,j+1} = \begin{cases} 1 & \text{if } j=0; \\ \frac{1}{2} & \text{if } j=1; \\ \frac{j-1}{j+1} & \text{if } j \ge 2. \end{cases} \quad P_{j,0} = \begin{cases} \frac{1}{2} & \text{if } j=1; \\ \frac{2}{j+1} & \text{if } j \ge 1. \end{cases}$$

The state transition diagram of Markov Chain $\{X[t]\}_{t\ge 0}$ is shown in Figure 5. Consider a linear Lyapunov function

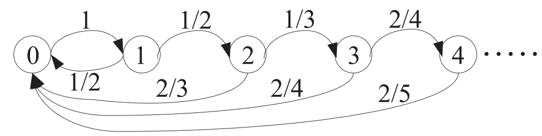

Figure 5: Markov Chain $\{X[t]\}_{t\ge 0}$

X. For any $X \ge 2$, we have

$$\mathbb{E}[X[t+1] - X[t] | X[t] = X] = \frac{X-1}{X+1} - \frac{2X}{X+1} = -1.$$

Thus, the Lyapunov function X has a strictly negative drift when $X \ge 2$ and hence the steady-state distribution of the Markov Chain exists. Recall that its drift increases at almost by 1, but has unbounded decrement, which has similar dynamics with the system under the RSG Algorithm.

Next, we will show that even the first moment of this Markov Chain does not exist, let alone its second moment. Let $\overline{X}$ be the limiting random variable of the Markov Chain and $\pi_j \triangleq \Pr\{\overline{X} = j\}$. According to the global balance equations, we can easily calculate

$$\pi_1 = \pi_0 = \frac{1}{3}, \quad \pi_j = \frac{1}{3j(j-1)}. \tag{14}$$

Thus, we have

$$\mathbb{E}[\overline{X}] = \sum_{j=1}^{\infty} j\pi_j = \frac{1}{3} + \sum_{j=2}^{\infty} \frac{1}{3(j-1)} = \infty.$$

Fortunately, we can establish the boundedness of the second moment of $\|\overline{\mathbf{Q}}_{\perp}^{(\epsilon)}\|$ under the RSG Algorithm by exploiting its unique dynamics under a mild assumption that $\Pr\{C_l[t]=0\} > 0, \forall l \in \mathcal{L}$, which leads to the following lemma that all TSLS counters have an exponential tail independent of ϵ.

Lemma 2. *If $p_l \triangleq \Pr\{C_l[t]=0\} > 0, \forall l \in \mathcal{L}$, then, under the RSG Algorithm, there exists a $\vartheta \in (0,1)$, independent of ϵ, such that*

$$\Pr\{\overline{T}_l^{(\epsilon)} = m\} \le 2\vartheta^m, \quad \forall l \in \mathcal{L}. \tag{15}$$

The proof of Lemma 2 is available in Appendix B.

Remark: We can also show that all TLSL counters still have an exponential tail independent of ϵ in non-fading single-hop network topologies. The extension to the more general setup is left for future search.

Lemma 2 directly implies (12). The rest of proof mainly builds on the analytical technique in [4], while it requires carefully partitioning the space $(\mathbf{Q}_{\perp}^{(\epsilon)}, \mathbf{T}^{(\epsilon)})$. The detailed outline can be found in Appendix C.

6.2 Proof of Main Result

We first give an upper bound on $\mathbb{E}[\langle \mathbf{c}, \overline{\mathbf{Q}}^{(\epsilon)}\rangle]$ by using the methodology of "setting the drift of a Lyapunov function equal to zero". We will omit the superscript ϵ associated with the queue lengths and TSLS counters for brevity in the rest of proof. To derive an upper bound, we need the following fundamental identity (see Lemma 8 in [2]):

$$\frac{\mathbb{E}\left[\langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle^2\right]}{2} + \frac{\mathbb{E}\left[\langle \mathbf{c}, \mathbf{A} - \mathbf{S}^*(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle^2\right]}{2}$$
$$+\mathbb{E}\left[\langle \mathbf{c}, \overline{\mathbf{Q}} + \mathbf{A} - \mathbf{S}^*(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle \langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle\right]$$
$$=\mathbb{E}\left[\langle \mathbf{c}, \overline{\mathbf{Q}}\rangle \langle \mathbf{c}, \mathbf{S}^*(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J) - \mathbf{A}\rangle\right], \tag{16}$$

which is derived through setting $\mathbb{E}[\Delta W_{\parallel}(\mathbf{Q}, \mathbf{T})] = 0$.

Next, we give upper bounds for each individual term in the left hand side of (16) and a lower bound for the right hand side of (16). Due to the space limitations, we omit the details and directly give results with some simple explanations.

By setting the mean drift of $\langle \mathbf{c}, \mathbf{Q}\rangle$ equal to zero and using the fact that $U_l \leq C_{\max}$ for all l, we have

$$\frac{1}{2}\mathbb{E}\left[\langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle^2\right] \leq \frac{\epsilon}{2}\langle \mathbf{c}, C_{\max}\mathbf{1}\rangle. \tag{17}$$

This means that there is almost no unused services under heavy-traffic conditions.

By observing that the RSG Algorithm selects the schedule $\mathbf{S}$ which maximizes $\langle \mathbf{c}, \mathbf{S}\rangle$ with high probability, we can show

$$\mathbb{E}\left[\langle \mathbf{c}, \mathbf{A} - \mathbf{S}^*(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle^2\right]$$
$$\leq \zeta^{(\epsilon)} + 2b\sum_{j=0}^{M} \frac{\epsilon}{\gamma_j} b_j + \sum_{j=0}^{M} \frac{\epsilon}{\gamma_j}\left((b_j)^2 + \langle \mathbf{c}, C_{\max}\mathbf{1}\rangle^2\right), \tag{18}$$

where we recall that $\zeta^{(\epsilon)}$ is defined in Proposition 2, and

$$\pi_j \triangleq \Pr\left\{\langle \mathbf{c}, \mathbf{S}^*(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle = b_j \middle| J = j\right\},$$
$$\gamma_j \triangleq \min\left\{b_j - \langle \mathbf{c}, \mathbf{r}\rangle : \text{for all } \mathbf{r} \in \mathcal{S}^j \setminus \mathcal{H}^{(\mathbf{c})}\right\}.$$

Inequality (18) indicates that the second moment of $\mathbf{c}$-weighted difference between arrivals and services is dominated by the $\mathbf{c}^2$-weighted variance of the arrival process and the variance of the channel fading process in the heavy-traffic limit.

In addition, by using similar arguments as in the proof for Proposition 4 in [2], we have

$$\mathbb{E}\left[\langle \mathbf{c}, \overline{\mathbf{Q}} + \mathbf{A} - \mathbf{S}^*(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle \langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J)\rangle\right]$$
$$\leq \sqrt{\epsilon \mathbb{E}[\|\overline{\mathbf{Q}}_{\perp}\|^2] \frac{C_{\max}}{c_{\min}}}, \tag{19}$$

where $c_{\min} \triangleq \min_{m \in \{l : c_l > 0\}} c_m$.

Finally, by using the definition of the RSG Algorithm and Proposition 3, we have

$$\mathbb{E}\left[\langle \mathbf{c}, \overline{\mathbf{Q}}\rangle \langle \mathbf{c}, \mathbf{S}^*(\overline{\mathbf{Q}}, \overline{\mathbf{T}}, J) - \mathbf{A}\rangle\right]$$
$$\geq \epsilon \mathbb{E}\left[\|\overline{\mathbf{Q}}_{\parallel}\|\right] - \cot(\theta)\sqrt{2\left(\mathbb{E}[\|\overline{\mathbf{Q}}_{\perp}\|^2] + (\alpha(\epsilon))^2 N_{\mathbf{T},2}\right)\epsilon}$$
$$\times \sqrt{\sum_{j=0}^{M} \frac{1}{\gamma_j}\left((b_j)^2 + \langle \mathbf{c}, C_{\max}\mathbf{1}\rangle^2\right)}, \tag{20}$$

where $\theta \in (0, \frac{\pi}{2}]$ is an angle such that $\langle \mathbf{c}, \mathbf{R}^*(\mathbf{Q}, \mathbf{T})\rangle = b$, for all $\mathbf{Q}$ and $\mathbf{T}$ satisfying $\frac{\|(\mathbf{Q}+\alpha\mathbf{T})_{\parallel}\|}{\|\mathbf{Q}+\alpha\mathbf{T}\|} \geq \cos(\theta)$, and $\mathbf{R}^*(\mathbf{Q}, \mathbf{T}) \triangleq \mathbb{E}[\mathbf{S}^*(\mathbf{Q}, \mathbf{T}, J)|\mathbf{Q}, \mathbf{T}]$.

By substituting bounds (17), (18), (19) and (20) into identity (16), we have

$$\epsilon \mathbb{E}\left[\|\overline{\mathbf{Q}}_{\parallel}\|\right] \leq \frac{\zeta^{(\epsilon)}}{2} + \overline{B}^{(\epsilon)}, \tag{21}$$

where

$$\overline{B}^{(\epsilon)} \triangleq \frac{\epsilon}{2}\langle \mathbf{c}, C_{\max}\mathbf{1}\rangle + \sqrt{\epsilon \mathbb{E}[\|\overline{\mathbf{Q}}_{\perp}\|^2] \frac{C_{\max}}{c_{\min}}}$$
$$+ \frac{1}{2}\sum_{j=0}^{M} \frac{\epsilon}{\gamma_j}\left((b_j)^2 + \langle \mathbf{c}, C_{\max}\mathbf{1}\rangle^2\right)$$
$$+ b\sum_{j=0}^{M} \frac{\epsilon}{\gamma_j} b_j + \cot(\theta)\sqrt{2\left(\mathbb{E}[\|\overline{\mathbf{Q}}_{\perp}\|^2] + (\alpha(\epsilon))^2 N_{\mathbf{T},2}\right)\epsilon}$$
$$\times \sqrt{\sum_{j=0}^{M} \frac{1}{\gamma_j}\left((b_j)^2 + \langle \mathbf{c}, C_{\max}\mathbf{1}\rangle^2\right)}. \tag{22}$$

Thus, if $\lim_{\epsilon \downarrow 0} \overline{B}^{(\epsilon)} = 0$, then RSG Algorithm is heavy-traffic optimal. Noting that $N_{\mathbf{T},2}$ is independent of ϵ, to satisfy $\lim_{\epsilon \downarrow 0} \overline{B}^{(\epsilon)} = 0$, it is enough to have

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}[\|\overline{\mathbf{Q}}_{\perp}\|^2] = 0 \text{ and } \lim_{\epsilon \downarrow 0} \epsilon \left(\alpha(\epsilon)\right)^2 = 0. \tag{23}$$

By using Proposition 3, it is easy to see that $\alpha(\epsilon) = O(\frac{1}{\sqrt[5]{\epsilon}})$ meets the above requirements.

7. CONCLUSION

In this paper, we studied the heavy-traffic behavior of the recently proposed maximum-weight type scheduling algorithm, called Regular Service Guarantee (RSG) Algorithm, that not only achieves throughput optimality but also provides regular services through the control parameter $\alpha \geq 0$. We showed that the RSG Algorithm is heavy-traffic optimal as long as $\alpha = O(\frac{1}{\sqrt[5]{\epsilon}})$, where ϵ is the heavy-traffic parameter characterizing the closeness of the arrival rate vector to the boundary of the capacity region. Noting that the service regularity improves with increasing α, our result reveals that the RSG Algorithm with a carefully selected parameter α can achieve the best service regularity performance among the class of the RSG Algorithms without sacrificing the mean delay optimality under heavy-traffic conditions.

8. ACKNOWLEDGMENTS

This work is supported by the NSF grants: CAREER-CNS-0953515 and CCF-0916664. Also, the work of A. Eryilmaz was in part supported by the QNRF grant number NPRP 09-1168-2-455.

9. REFERENCES

[1] M. Bramson. State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems*, 30(1):89–140, 1998.

[2] A. Eryilmaz and R. Srikant. Asymptotically tight steady-state queue length bounds implied by drift conditions. *Queueing Systems*, 72:311–359, 2012.

[3] G. Foschini and J. Salz. A basic dynamic routing problem and diffusion. *IEEE Transactions on Communications*, 26(3):320–327, 1978.

[4] B. Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advances in Applied Probability*, 14(3):502–525, 1982.

[5] I. Hou, V. Borkar, and P. R. Kumar. A theory of QoS for wireless. In *Proceedings of IEEE INFOCOM*, Rio de Janeiro, Brazil, April 2009.

[6] I. Hou and P. R. Kumar. Scheduling heterogeneous real-time traffic over fading wireless channels. In *Proceedings of IEEE INFOCOM*, San Diego, CA, March 2010.

[7] J. Jaramillo, R. Srikant, and L. Ying. Scheduling for optimal rate allocation in ad hoc networks with heterogeneous delay constraints. *IEEE Journal on Selected Areas in Communications*, 29(5):979–987, 2011.

[8] R. Li, B. Li, and A. Eryilmaz. Throughput-optimal scheduling with regulated inter-service times. In *Proceedings of IEEE INFOCOM*, Turin, Italy, April 2013.

[9] D. Shah and D. Wischik. Switched networks with maximum weight policies: Fluid approximation and multiplicative state space collapse. *Annals of Applied Probability*, 22(1):70–127, 2012.

[10] A. Stolyar. Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *The Annals of Applied Probability*, 14(1):1–53, 2004.

[11] L. Tassiulas. Scheduling and performance limits of networks with constantly changing topology. *IEEE Transactions on Information Theory*, 43(3):1067–1073, 1997.

[12] 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*, 36(12):1936–1948, 1992.

[13] W. Whitt. Weak convergence theorems for priority queues: Preemptive-resume discipline. *Journal of Applied Probability*, 8(1):74–94, 1971.

[14] R. Williams. Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing Systems*, 30(1):27–88, 1998.

APPENDIX

A. PROOF OF LEMMA 1

We assume $\boldsymbol{\lambda}^{(\epsilon)} \succ \mathbf{0}$. Indeed, if $\lambda_l^{(\epsilon)} = 0$ for some link l, then no arrivals occur in the link l. Thus, we do not need to consider such links. Since normal vector $\mathbf{c} \succeq \mathbf{0}$, we have $\boldsymbol{\lambda}^{(0)} \succ 0$. In addition, since $\boldsymbol{\lambda}^{(0)}$ is a relative interior point of dominant hyperplane $\mathcal{H}^{(\mathbf{c})}$, there exists a small enough $\delta > 0$ such that

$$\mathcal{B}_\delta \triangleq \mathcal{H}^{(\mathbf{c})} \bigcap \left\{ \mathbf{r} \succ \mathbf{0} : \|\mathbf{r} - \boldsymbol{\lambda}^{(0)}\| \leq \delta \right\}, \qquad (24)$$

representing the set of vectors on the hyperplane $\mathcal{H}^{(\mathbf{c})}$ that are within δ distance from $\boldsymbol{\lambda}^{(0)}$, lies strictly within the face $\mathcal{F}^{(\mathbf{c})} \triangleq \mathcal{H}^{(\mathbf{c})} \bigcap \mathcal{R}$.

In the rest of proof, we will omit ϵ associated with the queue length processes, the TSLS counters and parameter $\alpha(\epsilon)$ for brevity. Noting the difficulty to directly study the drift of Lyapunov function $V_\perp(\mathbf{Q}, \mathbf{T})$, we relate it with the drift of other proper Lyapunov functions, which is characterized in the following lemma.

LEMMA 3. *Define the following Lyapunov functions:*

$$W(\mathbf{Q}, \mathbf{T}) \triangleq \|(\mathbf{Q}, \sqrt{2\alpha C_{\max}}\mathbf{T})\|^2, \qquad (25)$$

$$W_\parallel(\mathbf{Q}, \mathbf{T}) \triangleq \|\mathbf{Q}_\parallel\|^2. \qquad (26)$$

Then, given $\mathbf{Q}[t] = \mathbf{Q}$ *and* $\mathbf{T}[t] = \mathbf{T}$*, their one-step drifts denoted by:*

$$\Delta W(\mathbf{Q}, \mathbf{T}) \triangleq \left[W(\mathbf{Q}[t+1], \mathbf{T}[t+1]) - W(\mathbf{Q}[t], \mathbf{T}[t])\right],$$

$$\Delta W_\parallel(\mathbf{Q}, \mathbf{T}) \triangleq \left[W_\parallel(\mathbf{Q}[t+1], \mathbf{T}[t+1]) - W_\parallel(\mathbf{Q}[t], \mathbf{T}[t])\right],$$

satisfy the following inequality:

$$\Delta V_\perp(\mathbf{Q}, \mathbf{T}) \leq \frac{\Delta W(\mathbf{Q}, \mathbf{T}) - \Delta W_\parallel(\mathbf{Q}, \mathbf{T})}{2\|(\mathbf{Q}_\perp, \sqrt{2\alpha C_{\max}}\mathbf{T})\|}. \qquad (27)$$

The proof of Lemma 3 is similar to that in [2] and is omitted here for brevity.

The rest of proof follows from Lemma 3 by studying the conditional expectation of $\Delta W(\mathbf{Q}, \mathbf{T})$ and $\Delta W_\parallel(\mathbf{Q}, \mathbf{T})$. We will omit the time reference $[t]$ without confusion.

We first consider $\mathbb{E}\left[\Delta W(\mathbf{Q}, \mathbf{T}) | \mathbf{Q}[t] = \mathbf{Q}, \mathbf{T}[t] = \mathbf{T}\right]$. It is not hard to show that

$$\begin{aligned}&\mathbb{E}\left[\Delta W(\mathbf{Q}, \mathbf{T}) | \mathbf{Q}[t] = \mathbf{Q}, \mathbf{T}[t] = \mathbf{T}\right] \\ \leq& 2\mathbb{E}\left[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}\right] + K_1 - 2\alpha \mathbb{E}\left[\langle \mathbf{T}, C_{\max}\mathbf{S}^*\rangle | \mathbf{Q}, \mathbf{T}\right],\end{aligned} \qquad (28)$$

where $K_1 \triangleq L \max\{A_{\max}^2, C_{\max}^2\} + 2\alpha L C_{\max}$, and we use the fact that $\sum_{l=1}^L T_l[t+1] - \sum_{l=1}^L T_l[t] = L - |\mathbf{H}^*| - \sum_{l \in \mathbf{H}^*} T_l[t]$, where $\mathbf{H}^* \triangleq \{l \in \mathcal{L} : S_l^*[t] C_l[t] > 0\}$.

Next, we consider $\mathbb{E}\left[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}\right]$. By using the definition of projection $\boldsymbol{\lambda}^{(0)}$, we have

$$\begin{aligned}&\mathbb{E}\left[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}\right] \\ =& \langle \mathbf{Q}, \boldsymbol{\lambda}^{(0)} - \epsilon \mathbf{c}\rangle - \mathbb{E}[\langle \mathbf{Q}, \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}] \\ =& -\epsilon \|\mathbf{Q}_\parallel\| + \langle \mathbf{Q}, \boldsymbol{\lambda}^{(0)}\rangle - \mathbb{E}[\langle \mathbf{Q} + \alpha \mathbf{T}, \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}] \\ &+ \alpha \mathbb{E}[\langle \mathbf{T}, \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}].\end{aligned} \qquad (29)$$

Given the queue-length vector $\mathbf{Q}[t]$, TSLS vector $\mathbf{T}[t]$ and the global channel state $J[t]$ at the beginning of slot t, according to the definition of the RSG Algorithm, we have

$$\langle \mathbf{Q}[t] + \alpha \mathbf{T}[t], \mathbf{S}^*[t] \cdot \mathbf{C}[t]\rangle = \max_{\mathbf{S} \in \mathcal{S}^{J[t]}} \langle \mathbf{Q}[t] + \alpha \mathbf{T}[t], \mathbf{S} \cdot \mathbf{C}[t]\rangle,$$

which implies

$$\langle \mathbf{Q} + \alpha \mathbf{T}, \mathbb{E}[\mathbf{S}^* \cdot \mathbf{C} | \mathbf{Q}, \mathbf{T}]\rangle = \max_{\mathbf{r} \in \mathcal{R}} \langle \mathbf{Q} + \alpha \mathbf{T}, \mathbf{r}\rangle. \qquad (30)$$

Thus, we have

$$\begin{aligned}\langle \mathbf{Q} + \alpha \mathbf{T}, \mathbb{E}[\mathbf{S}^* \cdot \mathbf{C} | \mathbf{Q}, \mathbf{T}]\rangle &= \max_{\mathbf{r} \in \mathcal{R}} \langle \mathbf{Q} + \alpha \mathbf{T}, \mathbf{r}\rangle \\ &\geq \max_{\mathbf{r} \in \mathcal{B}_\delta} \langle \mathbf{Q} + \alpha \mathbf{T}, \mathbf{r}\rangle \geq \max_{\mathbf{r} \in \mathcal{B}_\delta} \langle \mathbf{Q}, \mathbf{r}\rangle + \langle \alpha \mathbf{T}, \mathbf{r}^*\rangle,\end{aligned}$$

where $\mathbf{r}^* \in \arg\max_{\mathbf{r} \in \mathcal{B}_\delta} \langle \mathbf{Q}, \mathbf{r}\rangle$. Since $\boldsymbol{\lambda}^{(0)} \succ \mathbf{0}$, we can find a $\delta > 0$ sufficiently small such that $r_l \geq r_{\min}$ for all $\mathbf{r} = (r_l)_{l \in \mathcal{L}} \in \mathcal{B}_\delta$ and some $r_{\min} > 0$. Hence, we have

$$\langle \mathbf{Q} + \alpha \mathbf{T}, \mathbb{E}[\mathbf{S}^* \cdot \mathbf{C} | \mathbf{Q}, \mathbf{T}]\rangle \geq \max_{\mathbf{r} \in \mathcal{B}_\delta} \langle \mathbf{Q}, \mathbf{r}\rangle + \alpha r_{\min} \|\mathbf{T}\|_1.$$

By substituting above inequality into (29), we have

$$\begin{aligned}\mathbb{E}\left[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}\right] \leq& -\epsilon \|\mathbf{Q}_\parallel\| + \min_{\mathbf{r} \in \mathcal{B}_\delta} \langle \mathbf{Q}, \boldsymbol{\lambda}^{(0)} - \mathbf{r}\rangle \\ &- \alpha r_{\min} \|\mathbf{T}\|_1 + \alpha \mathbb{E}[\langle \mathbf{T}, \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}].\end{aligned}$$

Since $\boldsymbol{\lambda}^{(0)} - \mathbf{r}$ is perpendicular to the normal vector $\mathbf{c}$ for $\mathbf{r} \in \mathcal{B}_\delta$, we have

$$\min_{\mathbf{r}\in\mathcal{B}_\delta} \langle \mathbf{Q}, \boldsymbol{\lambda}^{(0)} - \mathbf{r}\rangle = \min_{\mathbf{r}\in\mathcal{B}_\delta} \langle \mathbf{Q}_\perp, \boldsymbol{\lambda}^{(0)} - \mathbf{r}\rangle = -\delta\|\mathbf{Q}_\perp\|.$$

Hence, we have

$$\begin{aligned}\mathbb{E}\left[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}\right] \le& -\epsilon\|\mathbf{Q}_\parallel\| - \delta\|\mathbf{Q}_\perp\| \\ &- \alpha r_{\min}\|\mathbf{T}\|_1 + \alpha\mathbb{E}[\langle \mathbf{T}, \mathbf{S}^* \cdot \mathbf{C}\rangle | \mathbf{Q}, \mathbf{T}]. \end{aligned} \tag{31}$$

Thus, by substituting (31) into (28), we have

$$\begin{aligned}\mathbb{E}[\Delta W(\mathbf{Q}, \mathbf{T}) | \mathbf{Q}, \mathbf{T}] \le& -2\epsilon\|\mathbf{Q}_\parallel\| - 2\delta\|\mathbf{Q}_\perp\| \\ &- 2\alpha r_{\min}\|\mathbf{T}\|_1 + K_1. \end{aligned} \tag{32}$$

Using techniques in showing (32) in [2], we have

$$\mathbb{E}[\Delta W_\parallel(\mathbf{Q}, \mathbf{T}) | \mathbf{Q}, \mathbf{T}] \ge -2\epsilon\|\mathbf{Q}_\parallel\| - K_2, \tag{33}$$

where $K_2 \triangleq 2LC_{\max}^2$.

By using the bounds (32), (33) and Lemma 3, we have

$$\begin{aligned}&\mathbb{E}[\Delta V_\perp(\mathbf{Q}, \mathbf{T}) | \mathbf{Q}, \mathbf{T}] \\ \le& \frac{\mathbb{E}\left[\Delta W(\mathbf{Q}, \mathbf{T}) - \Delta W_\parallel(\mathbf{Q}, \mathbf{T}) | \mathbf{Q}, \mathbf{T}\right]}{2\|(\mathbf{Q}_\perp, \sqrt{2\alpha C_{\max}\mathbf{T}})\|} \\ \le& \frac{-2\delta\|\mathbf{Q}_\perp\| - 2\alpha r_{\min}\sum_{l=1}^L T_l + K_1(\alpha) + K_2}{2\|(\mathbf{Q}_\perp, \sqrt{2\alpha C_{\max}\mathbf{T}})\|}. \end{aligned}$$

Note that $\alpha T_l \ge \alpha T_l \mathbb{1}_{\{\alpha T_l \ge 1\}} \ge \sqrt{\alpha T_l}\mathbb{1}_{\{\alpha T_l \ge 1\}} = \sqrt{\alpha T_l} - \sqrt{\alpha T_l}\mathbb{1}_{\{\alpha T_l < 1\}} \ge \sqrt{\alpha T_l} - 1$, and $\|\mathbf{Q}_\perp\| \ge \frac{1}{\sqrt{L}}\|\mathbf{Q}_\perp\|_1$, where $\mathbb{1}_{\{\cdot\}}$ is an indicator function. Thus, we have

$$\begin{aligned}&\mathbb{E}[\Delta V_\perp(\mathbf{Q}, \mathbf{T}) | \mathbf{Q}, \mathbf{T}] \\ \le& \frac{-\frac{2\delta}{\sqrt{L}}\|\mathbf{Q}_\perp\|_1 - 2r_{\min}\sum_{l=1}^L \sqrt{\alpha T_l} + K_1 + K_2 + 2Lr_{\min}}{2\|(\mathbf{Q}_\perp^{(k)}, \sqrt{2\alpha C_{\max}\mathbf{T}})\|} \\ \le& -\min\left\{\frac{\delta}{\sqrt{L}}, \frac{r_{\min}}{\sqrt{2C_{\max}}}\right\} + \frac{K_1 + K_2 + 2Lr_{\min}}{2\|(\mathbf{Q}_\perp, \sqrt{2\alpha C_{\max}\mathbf{T}})\|}. \end{aligned}$$

Hence, for any $0 < \varsigma < \min\left\{\frac{\delta}{\sqrt{L}}, \frac{r_{\min}}{\sqrt{2C_{\max}}}\right\}$, by taking

$$d \triangleq \frac{K_1 + K_2 + 2Lr_{\min}}{2\left(\min\left\{\frac{\delta}{\sqrt{L}}, \frac{r_{\min}}{\sqrt{2C_{\max}}}\right\} - \varsigma\right)}, \tag{34}$$

we have the desired result.

B. PROOF OF LEMMA 2

If the event $\mathcal{E}_j \triangleq \{C_l[j] > 0, C_i[j] = 0, \forall i \neq l\}$ always happens from $j = t - m + 1$ to t and there is at least one packet arriving at link l in this time duration, then under the RSG Algorithm, link l should be scheduled at least once during the past m slots and thus $T_l^{(\epsilon)}[t] < m$. This implies

$$\begin{aligned}&\Pr\{T_l^{(\epsilon)}[t] = m\} \\ \le& \Pr\{\mathcal{E}_j \text{ didn't happen for some } j \in [t-m+1, t]\} \\ &+ \Pr\{\text{No packet arrived at link } l \text{ from } t-m+1 \text{ to } t\}. \end{aligned}$$

Hence, we have $\Pr\{T_l^{(\epsilon)}[t] = m\} \le (1 - (1-p_l)\Pi_{i\neq l}p_i)^m + q_l^m \le 2\vartheta_l^m$, where $q_l \triangleq \Pr\{A_l[t] = 0\} < 1, \forall l \in \mathcal{L}$, and $\vartheta_l \triangleq \max\{1 - (1-p_l)\Pi_{i\neq l}p_i, q_l\} \in (0,1)$. Hence, we have $\Pr\{\overline{T}_l^{(\epsilon)} = m\} < 2\vartheta_l^m$. Thus, by taking $\vartheta \triangleq \max_{l\in\mathcal{L}}\vartheta_l$, we have the desired result.

C. PROOF OUTLINE OF PROPOSITION 3

In the rest of proof, we will omit ϵ associated with the queue length processes, the TSLS counters and parameter $\alpha(\epsilon)$ for brevity. It is quite challenging to directly give an upper bound on $\mathbb{E}[\|\overline{\mathbf{Q}}_\perp\|^2]$. Instead, we upper-bound the moment generation function of $\|\overline{\mathbf{Q}}_\perp\|$, and use the relationship between the moments of a random variable and its moment generation function to upper-bound $\mathbb{E}[\|\overline{\mathbf{Q}}_\perp\|^2]$ as shown in the following lemma.

LEMMA 4. *For a random variable X with $\mathbb{E}[e^{\eta X}] < \infty$ for some $\eta > 0$, we have*

$$\mathbb{E}[X^n] \le \frac{1}{\eta^n}\left(\log\left(e^{n-1}\mathbb{E}[e^{\eta X}]\right)\right)^n, \tag{35}$$

for $n = 1, 2, 3\cdots$.

Please see the Appendix D for the proof of Lemma 4.

Let $\mathbf{Z}[t] \triangleq \left(\mathbf{Q}_\perp[t], \sqrt{2\alpha C_{\max}\mathbf{T}[t]}\right)$. We first give an upper bound on $\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|}\middle|\mathbf{Q}[t], \mathbf{T}[t]\right]$. To that end, let $l^*[t] \in \arg\max_l T_l[t]$. We partition $(\mathbf{Q}_\perp[t], \mathbf{T}[t])$ into sets $\mathcal{F}_1$, $\mathcal{F}_2$ and $\mathcal{F}_3$, where

$$\begin{aligned}&\mathcal{F}_1 \triangleq \{\|\mathbf{Z}[t]\| \le d\}; \mathcal{F}_2 \triangleq \left\{\|\mathbf{Z}[t]\| > d, \|\mathbf{Q}_\perp[t]\| > T_{l^*[t]}[t]\right\}; \\ &\mathcal{F}_3 \triangleq \left\{\|\mathbf{Z}[t]\| > d, \|\mathbf{Q}_\perp[t]\| \le T_{l^*[t]}[t]\right\}.\end{aligned}$$

Then, we have

$$\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|}\middle|\mathbf{Q}[t], \mathbf{T}[t]\right] = \sum_{i=1}^3 \mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|}; \mathcal{F}_i\middle|\mathbf{Q}[t], \mathbf{T}[t]\right]. \tag{36}$$

Next, we consider each term in (36) individually.

(i) On event $\mathcal{F}_1$, we can show that if $\|\mathbf{Z}[t]\| \le d$, then

$$\begin{aligned}\|\mathbf{Z}[t+1]\|^2 \le& d^2 + 2\left(A_{\max} + 2\sqrt{L}C_{\max}\right)d\sqrt{L} \\ &+ L\left(A_{\max} + 2\sqrt{L}C_{\max}\right)^2 + 2L\alpha C_{\max} \triangleq G_1^2, \end{aligned} \tag{37}$$

Hence, we have

$$\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|}; \mathcal{F}_1\middle|\mathbf{Q}[t], \mathbf{T}[t]\right] \le e^{\eta G_1} \tag{38}$$

To consider other two terms in (36), we need the following lemma.

LEMMA 5. *Under the RSG Algorithm, if $\|\mathbf{Z}[t]\| > d$, then*

$$\begin{aligned}&|\|\mathbf{Z}[t+1]\| - \|\mathbf{Z}[t]\|| \\ \le& 2L\max\{A_{\max}, C_{\max}\} + \frac{2\alpha C_{\max}L}{d} \\ &+ 2\alpha C_{\max}\frac{\sum_{l\in\mathbf{H}^*} T_l[t]}{\sqrt{\|\mathbf{Q}_\perp[t]\|^2 + 2\alpha C_{\max}\sum_{l=1}^L T_l[t]}}, \end{aligned} \tag{39}$$

where $\mathbf{H}^* \triangleq \{l : S_l^*[t]C_l[t] > 0\}$.

The proof is omitted due to space limitations.

(ii) On event $\mathcal{F}_2$, we have

$$\frac{\sum_{l\in\mathbf{H}^*} T_l[t]}{\sqrt{\|\mathbf{Q}_\perp[t]\|^2 + 2\alpha C_{\max}\sum_{l=1}^L T_l[t]}} \le \frac{LT_{l^*[t]}[t]}{\|\mathbf{Q}_\perp[t]\|} \le L.$$

By substituting above inequality into (39), we get

$$|\|\mathbf{Z}[t+1]\| - \|\mathbf{Z}[t]\|| \le G_2, \tag{40}$$

where $G_2 \triangleq 2L\max\{A_{\max}, C_{\max}\} + \frac{2\alpha C_{\max} L}{d} + 2\alpha C_{\max} L$. Noting that (13) and (40) satisfy conditions of Lemma 2.2 in [4], there exists $\eta_1 > 0$, and $\rho = e^{\eta G_2} - \eta(G_2 + \varsigma) \in (0, 1)$, independent of ϵ, such that

$$\mathbb{E}\left[e^{\eta(\|\mathbf{Z}[t+1]\| - \|\mathbf{Z}[t]\|)}; \mathcal{F}_2 \Big| \mathbf{Q}[t], \mathbf{T}[t]\right] \le \rho, \forall 0 < \eta < \eta_1.$$

Thus, we have

$$\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|}; \mathcal{F}_2 \Big| \mathbf{Q}[t], \mathbf{T}[t]\right] \le \rho e^{\eta\|\mathbf{Z}[t]\|}. \tag{41}$$

(iii) On event $\mathcal{F}_3$, we have

$$\begin{aligned}&\frac{\sum_{l\in\mathbf{H}^*} T_l[t]}{\sqrt{\|\mathbf{Q}_\perp[t]\|^2 + 2\alpha C_{\max}\sum_{l=1}^L T_l[t]}}\\ \le&\frac{L T_{l^*[t]}[t]}{\sqrt{2\alpha C_{\max} T_{l^*[t]}[t]}} = \frac{L}{\sqrt{2\alpha C_{\max}}}\sqrt{T_{l^*[t]}[t]}.\end{aligned} \tag{42}$$

By substituting (42) into (39), we get

$$\begin{aligned}&|\|\mathbf{Z}[t+1]\| - \|\mathbf{Z}[t]\||\\ \le& 2L\max\{A_{\max}, C_{\max}\} + \frac{2\alpha C_{\max} L}{d} + L\sqrt{2\alpha C_{\max}}\sqrt{T_{l^*[t]}[t]}.\end{aligned}$$

In addition, on event $\mathcal{F}_3$, we have

$$\|\mathbf{Z}[t]\| \le \sqrt{T^2_{l^*[t]}[t] + 2\alpha C_{\max} L T_{l^*[t]}[t]}. \tag{43}$$

Hence, we have

$$\begin{aligned}\|\mathbf{Z}[t+1]\| \le& \|\mathbf{Z}[t]\| + |\|\mathbf{Z}[t+1]\| - \|\mathbf{Z}[t]\||\\ \le& F_1 T_{l^*[t]}[t] + F_2,\end{aligned} \tag{44}$$

where $F_1 \triangleq L\sqrt{2\alpha C_{\max}} + \sqrt{1 + 2\alpha C_{\max} L}$ and $F_2 \triangleq \frac{2\alpha C_{\max} L}{d} + 2L\max\{A_{\max}, C_{\max}\}$. Thus, we have

$$\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|}; \mathcal{F}_3 \Big| \mathbf{Q}[t], \mathbf{T}[t]\right] \le e^{\eta F_2} e^{\eta F_1 T_{l^*[t]}[t]}. \tag{45}$$

By substituting (38), (41) and (45) into (36), we have

$$\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|} | \mathbf{Q}[t], \mathbf{T}[t]\right] \le e^{\eta G_1} + \rho e^{\eta\|\mathbf{Z}[t]\|} + e^{\eta F_2} e^{\eta F_1 T_{l^*[t]}[t]}.$$

By taking expectation on both sides, we have

$$\begin{aligned}&\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|}\right] \le e^{\eta G_1} + \rho\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t]\|}\right] + e^{\eta F_2}\mathbb{E}\left[e^{\eta F_1 T_{l^*[t]}[t]}\right]\\ &\le e^{\eta G_1} + \rho\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t]\|}\right] + e^{\eta F_2}\sum_{l=1}^L \mathbb{E}\left[e^{\eta F_1 T_l[t]}\right].\end{aligned} \tag{46}$$

By Lemma 2, there exist $\eta_2 > 0$ such that $e^{\eta_2 F_1}\vartheta < 1$ and for $0 < \eta < \eta_2$, we have

$$\mathbb{E}\left[e^{\eta F_1 T_l[t]}\right] \le 2\sum_{m=0}^{\infty} e^{\eta F_1 m}\vartheta^m = \frac{2}{1 - e^{\eta F_1}\vartheta} \tag{47}$$

By substituting (47) into (46), we have

$$\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t+1]\|}\right] \le \rho\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t]\|}\right] + G, \tag{48}$$

holding for $0 < \eta < \eta_0 \triangleq \min\{\eta_1, \eta_2\}$, where $G \triangleq e^{\eta G_1} + \frac{2Le^{\eta F_2}}{1 - e^{\eta F_1}\vartheta}$. By using inequality (48) and iterating over t, we have

$$\mathbb{E}\left[e^{\eta\|\mathbf{Z}[t]\|}\right] \le \rho^t e^{\eta\|\mathbf{Z}[0]\|} + \frac{1 - \rho^t}{1 - \rho} G \le e^{\eta\|\mathbf{Z}[0]\|} + \frac{G}{1 - \rho},$$

which implies $\mathbb{E}\left[e^{\eta\|\mathbf{Q}_\perp[t]\|}\right] \le e^{\eta\|\mathbf{Z}[0]\|} + \frac{G}{1-\rho}$.

Thus, we have

$$\mathbb{E}\left[e^{\eta\|\overline{\mathbf{Q}}_\perp[t]\|}\right] \le e^{\eta\|\mathbf{Z}[0]\|} + \frac{G}{1 - \rho} \tag{49}$$

where $G \triangleq e^{\eta G_1} + \frac{2Le^{\eta F_2}}{1 - e^{\eta F_1}\vartheta}$, $\rho \triangleq e^{\eta G_2} - \eta(G_2 + \varsigma) \in (0, 1)$, $d = O(\alpha)$, $G_1 = O(\alpha)$, $G_2 = O(\alpha)$, $F_1 = O(\sqrt{\alpha})$ and $F_2 = O(1)$. Note that we need to choose a $\eta > 0$ such that

$$1 - e^{\eta F_1}\vartheta < 1 \tag{50}$$

$$e^{\eta G_2} - \eta(G_2 + \varsigma) < 1. \tag{51}$$

It is not hard to verify that

$$0 < \eta \le \frac{1}{2}\min\left\{\frac{1}{F_1}\ln\frac{1}{\vartheta}, \frac{1}{G_2}\ln\frac{G_2 + \varsigma}{G_2}\right\} \tag{52}$$

satisfies above requirements. If α is large enough such that $\frac{\varsigma}{G_2} < 1$ and $G_2 \gg F_1$, then we have

$$\frac{1}{G_2}\ln\frac{G_2 + \varsigma}{G_2} \le \frac{1}{F_1}\ln\frac{1}{\vartheta}. \tag{53}$$

Thus, we can take $\eta^* \triangleq \frac{1}{2G_2}\ln\frac{G_2+\varsigma}{G_2}$ to meet the above requirements, and hence $\eta^* = O(\frac{1}{\alpha^2})$.

Taking $\eta = \eta^*$ and noting that $\eta^* < \frac{1}{2F_1}\ln\frac{1}{\vartheta}$, we have

$$\begin{aligned}&\frac{G}{1 - \rho} = \frac{e^{\eta^* G_1} + \frac{2Le^{\eta^* F_2}}{1 - \vartheta e^{\eta^* F_1}}}{1 - \left(e^{\eta^* G_2} - \eta^*(G_2 + \varsigma)\right)}\\ \le&\frac{e^{\eta^* G_1} + \frac{2Le^{\eta^* F_2}}{1 - \sqrt{\vartheta}}}{1 - \left(e^{\eta^* G_2} - \eta^*(G_2 + \varsigma)\right)}\\ =&\frac{e^{\eta^* G_1} + \frac{2Le^{\eta^* F_2}}{1 - \sqrt{\vartheta}}}{1 - \left(1 + \frac{\varsigma}{G_2}\right)^{\frac{1}{2}} + \frac{1}{2}\left(1 + \frac{\varsigma}{G_2}\right)\ln\left(1 + \frac{\varsigma}{G_2}\right)}\\ \overset{(a)}{=}& O\left(\frac{1}{1 - \left(1 + \frac{\varsigma}{2G_2}\right) + \frac{1}{2}\left(1 + \frac{\varsigma}{G_2}\right)\frac{\varsigma}{G_2}}\right) = O\left(\alpha^2\right),\end{aligned}$$

where the step (a) uses $\eta^* G_1 = O\left(\frac{1}{\alpha}\right)$ and $\eta^* F_2 = O\left(\frac{1}{\alpha^2}\right)$. Thus, we have $\mathbb{E}\left[e^{\eta^*\|\overline{\mathbf{Q}}_\perp\|}\right] = O(\alpha^2)$. By using Lemma 4, we have

$$\begin{aligned}\mathbb{E}[\|\overline{\mathbf{Q}}_\perp\|^2] \le& \frac{1}{(\eta^*)^2}\left(\log\left(e\mathbb{E}\left[e^{\eta^*\|\overline{\mathbf{Q}}_\perp[t]\|}\right]\right)\right)^2\\ =& O\left(\alpha^4(\log\alpha)^2\right).\end{aligned}$$

D. PROOF OF LEMMA 4

$$\begin{aligned}\mathbb{E}[X^n] =& \frac{1}{\eta^n}\mathbb{E}\left[\left(\log e^{\eta X}\right)^n\right]\\ \overset{(a)}{\le}& \frac{1}{\eta^n}\mathbb{E}\left[\left(\log\left(e^{n-1}e^{\eta X}\right)\right)^n\right]\\ \overset{(b)}{\le}& \frac{1}{\eta^n}\left(\log\left(e^{n-1}\mathbb{E}[e^{\eta X}]\right)\right)^n,\end{aligned} \tag{54}$$

where the step (a) follows from the fact that $f(y) = (\log y)^n$ is increasing in $y \in [1, \infty)$ for $n = 1, 2, \cdots$; (b) uses the fact that $g(y) = \left(\log\left(e^{n-1}y\right)\right)^n$ is concave in $[1, \infty)$ for $n = 1, 2, \cdots$, and Jensen's Inequality.