Optimal Design of Queuing Systems via Compositional Stochastic Programming

Srujan Teja Thomdapu, *Student Member, IEEE* and Ketan Rajawat, *Member, IEEE,*

Abstract

Well-designed queuing systems form the backbone of modern communications, distributed computing, and content delivery architectures. Designs balancing infrastructure costs and user experience indices require tools from teletraffic theory and operations research. A standard approach to designing such systems involves formulating optimization problems that strive to maximize the pertinent utility functions while adhering to quality-of-service and other physical constraints. In many cases, formulating such problems necessitates making simplistic assumptions on arrival and departure processes to keep the problem simple.

This work puts forth a stochastic optimization framework for designing queuing systems where the exogenous processes may have arbitrary and unknown distributions. We show that many such queuing design problems can generally be formulated as stochastic optimization problems where the objective and constraint are non-linear functions of expectations. The compositional structure obviates the use of classical stochastic approximation approaches where the stochastic gradients are often required to be unbiased. To this end, a constrained stochastic compositional gradient descent algorithm is proposed that utilizes a tracking step for the expected value functions. The non-asymptotic performance of the proposed algorithm is characterized via its iteration complexity. Numerical tests allow us to validate the theoretical results and demonstrate the efficacy of the proposed algorithm.

I. INTRODUCTION

As the Internet traffic continues to grow, meeting service-level agreement (SLA) targets necessitates proactive capacity planning and resource provisioning [1–3]. Sustaining the users' quality of experience (QoE), for instance, requires the backbone and access networks to be designed

---

A part of this work has been submitted in Proceedings of the Asilomar Conference on Signals, Systems, and Computers, 2019.

S. T. Thomdapu, and K. Rajawat are with the Department of Electrical Engineering, Indian Institute of Technology Kanpur, Kanpur-208016, India, e-mail: ({srujant, ketan}@iitk.ac.in).

with sufficiently high available capacity. Virtual machine resources, such as CPU, memory, and storage, must likewise be provisioned to tackle the demand variability in data centers. On the other hand, economic considerations generally dictate that resources be utilized efficiently and over-provisioning be minimized. Well-designed systems must strive to achieve a balance between resource utilization and excess capacity.

Teletraffic theory, utilizing advances in queuing theory and operations research, has been widely used for the design of communication networks [4]. Within this framework, the random exogenous factors such as resource demands and channel gains are studied in isolation and statistically modeled a priori (see Chapter 3 of [5]). Simple models are generally preferred, so as to yield closed-form expressions for various performance metrics such as waiting times, throughput, delay, waiting costs, provisioning costs, and quality-of-service (QoS). Subsequently, the underlying decision variables, such as the arrival rates or link capacities, are tuned to maximize the performance or minimize the costs, while adhering to the various QoS constraints. The latter step may be accomplished by solving a constrained optimization problem; see [4] and references therein.

The classical approach is however not suited to systems where the distribution of the exogenous variables is either unknown or too complicated to allow closed-form expressions. As a simple example, the average ergodic capacity of a single antenna wireless fading channel is not expressible in closed-form [6]. Consequently, existing approaches for wireless system design either ignore the delay or rely on approximations even for simplistic fading models [7–11]. More generally, the performance metrics of almost all such queuing systems can be expressed in terms of expectations involving the random variables. This work considers the design of queuing systems from a stochastic optimization perspective. The network design problem is formulated as a constrained optimization problem where the objective and constraint functions involve expectations with respect to these random variables. We consider handling such problems within the stochastic approximation rubric where a few sequentially available independent realizations of the random process are sufficient to yield near-optimal designs. Of particular interest are low-complexity first-order algorithms that are capable of running in real-time. A supervised learning approach is adopted where the parameters are learned during a training phase that is separate from the significantly longer operational phase. The proposed framework is general, and finds applications in communication networks, traffic systems, data centers, and supply chain.

From a theoretical vantage point, the objective function and constraints in these design problems usually depend on quantities such as queuing delay and throughput, that are themselves non-linear functions of expectations. The compositional structure of these problems obviates the use of standard stochastic gradient descent (SGD) methods where the gradient estimates are required to be unbiased or at least strongly consistent (see e.g., [12–14]). The more general case involving non-linear functions of expectations has only been recently considered in [15], where the stochastic compositional gradient descent (SCGD) algorithm was proposed. However, compositional problems involving expectation constraints have never been considered in the literature, and the proposed SCGD method for constrained problems is the first such algorithm. The non-asymptotic optimality gap and constraint violation incurred by the proposed algorithm is characterized in terms of the number of samples, and constitutes the key theoretical contribution of the present work. The sample error complexity results developed here are of independent interest and can be applied to any setting involving constrained stochastic optimization problems.

The rest of the paper is organized as follows. The next subsection briefly reviews the related work on optimization of queuing systems and on algorithms for solving constrained stochastic problems. Sec. II formulates the general constrained problem and discusses several examples from queuing literature. Sec. III details the proposed algorithm and provides the relevant theoretical guarantees. The simulation results are presented in Sec. IV and finally, Sec. V concludes the paper.

A. Related Work

General compositional stochastic problems, where the goal is to minimize compositions of expected-value functions, were first discussed in [15] and solved via the SCGD algorithm. The corresponding finite-sum variant of the problem have subsequently been considered in [16] and solved via the variance-reduced SCGD. Other variants include the accelerated SCGD [17], those utilizing Markov samples [18], and generalized versions involving multiple nested compositions [19]. A functional variant of SCGD has recently been proposed in [20]. Compositional stochastic problems involving expectation constraints were mentioned in [21, Sec. 6.10] but have not been thoroughly studied.

Within the classical stochastic optimization framework, where the compositional form does not arise, problems with stochastic constraints have been considered. Relevant algorithms include the stochastic dual-descent algorithm [22–25] and the stochastic variant of the Arrow–Hurwicz saddle point method [26]. The other problem formulations with infinite time horizon objective and constraint functions are discussed in [27]. A comprehensive article on constrained stochastic problems of large data can be found in [28]. It is remarked that all these classes of algorithms apply to problems with a specific structure, and sample complexity results for the general versions have not been developed. Further, these algorithms cannot be directly applied to problems involving compositional forms. Interestingly, the compositional stochastic optimization problem considered here subsumes the setting with only expectation constraints.

Optimization approaches for the design of queuing systems have a long history; see [4] and references therein. Formulated as deterministic optimization problems, these network designs consider a wide variety of objectives and systems. As stated earlier, the existing optimization-based approaches do not suffice when the problem involves random variables with unknown distributions. Optimization of queuing systems where the virtual network functions to be executed in an order are in queue are discussed in [29]. The objective function is formulated as a time average processing cost.

The system design problems considered here are static in nature, and the optimal decision variable obtained at the end of the training phase is utilized during the operational phase. Such a setting is in contrast to that of network control, where the goal is to learn optimal policies that depend on the observed environmental state or queue lengths. Utility-optimal control algorithms such as the backpressure algorithm have been extensively studied [30–36]. In network control, the optimal policy is updated at every time instant, and the algorithm continues to run forever. However, the focus here is on determining the optimal provisioning or dimensioning parameters that must be set a priori. Accordingly, the performance of the proposed approach is studied in terms of its sample complexity bounds.

II. PROBLEM FORMULATION AND EXAMPLES

This section formulates the constrained stochastic optimization problem at hand and provides motivating examples within the context of queuing theory. Of particular interest are problems arising in queuing system design, where the distribution of the service times are general and a priori unknown.

A. Problem Formulation

We begin with describing a general stochastic optimization problem where the formulation involves expectations with respect to a collection of k random variables (collected into the vector

$\boldsymbol{\zeta} \in \mathbb{R}^k$) with unknown distributions. Going beyond the classical stochastic optimization setting where the goal is to minimize the expectation of a function, we consider the more general case where the expectation appears as an argument of an outer non-linear function. In general, such a composite structure makes it difficult to obtain unbiased samples of the objective gradient, rendering the SGD and its variants inapplicable.

Consider the constrained optimization problem

$$\mathbf{x}^\star = \arg\min_{\mathbf{x}\in\mathcal{X}} \; f(\mathbb{E}[\mathbf{g}(\mathbf{x},\boldsymbol{\zeta})]) \quad \text{s.t.} \quad \mathbf{q}(\mathbb{E}[\mathbf{h}(\mathbf{x},\boldsymbol{\zeta})]) \leq 0 \tag{$\mathcal{P}$}$$

where the expectation is taken with respect to $\boldsymbol{\zeta}$. Here, $f : \mathbb{R}^m \to \mathbb{R}$, $\mathbf{g} : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}^m$, $\mathbf{h} : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}^d$, and $\mathbf{q} : \mathbb{R}^d \to \mathbb{R}^J$ are continuous functions. The optimization variable $\mathbf{x}$ belongs to the convex set $\mathcal{X} \subset \mathbb{R}^n$, described by simple constraints that are easy to project onto; examples include box or norm-ball constraints. Specific assumptions regarding convexity or smoothness of various functions will be explicitly stated in Sec III. In order to avoid trivial scenarios however, we assume that the problem in ($\mathcal{P}$) is feasible and has a finite solution.

When the distribution of $\boldsymbol{\zeta}$ is not known a priori, the expectations appearing in ($\mathcal{P}$) cannot be evaluated in closed-form. Motivated by classical stochastic approximation methods, the goal is to solve ($\mathcal{P}$) in an online fashion using only independent realizations $\boldsymbol{\zeta}_1$, $\boldsymbol{\zeta}_2$, $\ldots$ that are revealed sequentially. For many problems, such samples yield unbiased (or at least strongly consistent) estimates of the objective gradient, paving the way for the application of the popular SGD algorithm and its variants [37, 38]. In the present case however, unbiased estimates of the gradient $\mathbb{E}[\boldsymbol{\nabla}\mathbf{g}(\mathbf{x},\boldsymbol{\zeta})]\nabla f(\mathbb{E}\mathbf{g}(\mathbf{x},\boldsymbol{\zeta}))$ cannot be obtained from a single sample $\boldsymbol{\zeta}$ due to the presence of the non-linear outer function f. Moreover, projected SGD and its variants cannot be used since the constraint also involves an expectation that cannot be evaluated in closed form.

The set-constrained version of ($\mathcal{P}$) was first considered in [15], where a quasi-gradient approach was proposed. The main idea was to run two iterations: one for performing quasi-gradient steps for estimating $\mathbf{x}^\star$ and another for keeping track of the vector $\mathbb{E}[\mathbf{g}(\mathbf{x}^\star,\boldsymbol{\zeta})]$. The approach in [15] is however still not applicable to ($\mathcal{P}$) due to the presence of stochastic constraints.

Stochastic constraints can be dealt with algorithms operating in the dual domain. For instance, dual descent algorithms proposed in [22, 23] can handle expectation constraints of a specific form. Likewise, primal-dual or saddle point algorithms have also been proposed for problems with expectation constraints [39–41]. However, none of these approaches can handle non-linear functions of expectations either in the objective or constraint functions. In summary, existing

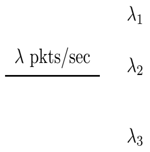
λ_1
λ pkts/sec
λ_2
λ_3

In this section, several queuing system design examples are discussed and the resulting stochastic problems are shown to be of the form in ($\mathcal{P}$). For ease of exposition, we consider the standard utility-minus-penalty objective similar to (1) subject to various functional constraints. Other metrics can also be used, provided that they satisfy the assumptions in Sec. III. Alternatively, all requirements may be expressed as constraints only, so that ($\mathcal{P}$) becomes a feasibility problem.

1) Example 1: Parallel M/G/1 queues: To begin with, consider a set of N parallel $M/G/1$ queues as shown in Fig. 1a. The system consists of a generic source, mimicking a router or an encoder, and capable of supplying at most $\lambda^{\max}$ packets per second on an average. Each packet generated by the source carries a destination address and is accordingly routed to one of the queues. The length of a packet sent to the i-th queue is measured in bits and denoted by the random variable L_i and collected into the random vector $\mathbf{L}$. The lengths are exponentially distributed, but the means $\mathbb{E}L_i$ are unknown and different for each queue. The difference in mean packet lengths may arise due to the heterogeneity in the traffic stream arriving at each queue. The departures from the i-th queue are limited by the link capacity C_i bits per second and the packet service times $X_i := L_i/C_i$.

The optimization variables $\{\lambda_i\}$ are collected into the vector $\boldsymbol{\lambda} \in \mathbb{R}^N_{++}$ and are required to satisfy box constraints $\lambda_i \in [\lambda^{\min}, \lambda^{\max}]$ that may arise from various physical considerations. As in (1), the goal is to maximize the utility function that depends on the throughput, and the delay. While throughput is still given by $\lambda_i \mathbb{E}L_i$, the queuing delay must be calculated using the Pollaczek-Khinchin (P-K) formula, and the objective function becomes $U(\boldsymbol{\lambda}) = \sum_{i=1}^N \left[\psi_i(\lambda_i \mathbb{E}L_i) - \varphi\left(\lambda_i \mathbb{E}L_i^2 / \left(2C_i(C_i - \lambda_i \mathbb{E}L_i)\right)\right)\right]$, where $\psi_i : \mathbb{R} \to \mathbb{R}$ and $\varphi_i : \mathbb{R} \to \mathbb{R}$ are pre-specified functions. Examples include linear utility and cost functions, fairness-inducing log utility $\psi(x) = \bar{\psi}_i \log(x)$, and squared delay penalty $\varphi_i(x) = \bar{\varphi}_i x^2$, where $\bar{\psi}_i$ and $\bar{\varphi}_i$ are positive constants. It can be seen that the objective function depends on a non-linear function of the expectation and can be written in the form required in ($\mathcal{P}$) as

$$[\mathbf{g}(\boldsymbol{\lambda}, \mathbf{L})]_i = \begin{cases} \lambda_i L_i & 1 \le i \le N \\ \lambda_{i-N} L_{i-N}^2 & N+1 \le i \le 2N \end{cases}, \quad f(\mathbf{y}) = \sum_{i=1}^N \left[\varphi\left(\frac{y_{N+i}}{2C_i(C_i - y_i)}\right) - \psi(y_i)\right] \tag{2}$$

where $m = 2N$ and $n = k = N$. The waiting delay of a packet should not exceed the maximum tolerable value, and hence the following constraint is imposed:

$$\max_{1 \le i \le N} \frac{\lambda_i \mathbb{E}L_i^2}{2C_i(C_i - \lambda_i \mathbb{E}L_i)} \le D^{\max} \tag{3}$$

where $D^{\max}$ is the maximum tolerable delay. Comparing with the formulation in ($\mathcal{P}$), it can be seen that $\mathbf{h}(\boldsymbol{\lambda}, \mathbf{L}) = \mathbf{g}(\boldsymbol{\lambda}, \mathbf{L})$ and $q(\mathbf{z}) = \max_{1\leq i\leq N} z_{i+N} / \left(2C_i(C_i - z_i)\right) - D^{\max}$ so that $d = 2N$. The overall optimization problem can be written as

$$\min_{\boldsymbol{\lambda}\in\boldsymbol{\Lambda}} U(\boldsymbol{\lambda}) \quad \text{s.t.} \quad \max_{1\leq i\leq N} \frac{\lambda_i \mathbb{E}L_i^2}{2C_i(C_i - \lambda_i \mathbb{E}L_i)} \leq D^{\max},$$

where $\boldsymbol{\Lambda} := \{\boldsymbol{\lambda} | \lambda^{\min} \leq \lambda_i \leq \lambda^{\max} \forall i = 1, \ldots, N, \sum_{i=1}^N \lambda_i \leq \lambda^{lim}\}$.

2) Example 2: Parallel M/G/1 queues in wireless networks (ergodic capacity): Consider a system of N parallel M/G/1 queues operating within a common base station (BS) and communicating with N users over N orthogonal channels; see Fig. 1b. The arrivals are still Poisson and the arrival rate vector $\boldsymbol{\lambda}$ is an optimization variable. Different from the earlier examples however, the packet lengths are constant and the service time instead depends on the channel capacity and the power allocation. Each queue is allocated with an average power of p_i and the queues adhere to the total power budget $P^{\max}$, that is, $\sum_{i=1}^N p_i \leq P^{\max}$. The allocated powers are collected into the vector $\mathbf{p} \in \mathbb{R}_+^N$. Together, the arrival rates and power allocations belong to the sets $\boldsymbol{\Lambda} := \{\boldsymbol{\lambda} | \lambda^{\min} \leq \lambda_i \leq \lambda^{\max}, i = 1, \ldots, N, \sum_i \lambda_i \leq \lambda^{\max}\}$ and $\boldsymbol{\Pi} := \{\mathbf{p} \mid p_i \geq P^{\min}, \sum_i p_i \leq P^{\max}\}$.

The channel gain from the BS to the i-th user is denoted by ζ_i, collected in the vector $\boldsymbol{\zeta} \in \mathbb{R}^N$, and remains constant over the transmission of a single packet. Assuming sufficiently long coherence times, the transmission rate is close to the ergodic capacity of the channel, given by $b_i(p_i, \zeta_i) = B_i \log(1 + \zeta_i p_i)\ \forall i$, where B_i is the capacity of the i-th channel and is known a priori. If the capacity is measured in packets per second, the service time is the random variable $X_i = 1/b_i(p_i, \zeta_i)$ measured in seconds. The design goal is to determine the optimal values of $\boldsymbol{\lambda}$ and $\mathbf{p}$ so as to maximize the utility function. Applying the PK formula for queuing delay, the utility function can be written as

$$U(\boldsymbol{\lambda}, \mathbf{p}) = \sum_{i=1}^N \left[\psi(\lambda_i) - \varphi_i\left(\frac{\lambda_i \mathbb{E}\left[\frac{1}{b_i^2(p_i,\zeta_i)}\right]}{2\left(1 - \lambda_i \mathbb{E}\left[\frac{1}{b_i(p_i,\zeta_i)}\right]\right)}\right)\right]. \tag{4}$$

It can be seen that the objective function in (4) is of the compositional form in ($\mathcal{P}$) with

$$[\mathbf{g}(\boldsymbol{\lambda}, \mathbf{p}, \boldsymbol{\zeta})]_i = \begin{cases} \lambda_i & 1 \leq i \leq N \\ \frac{\lambda_{i-N}}{b_{i-N}(p_{i-N},\zeta_{i-N})} & N+1 \leq i \leq 2N \\ \frac{\lambda_{i-2N}}{b_{i-2N}^2(p_{i-2N},\zeta_{i-2N})} & 2N+1 \leq i \leq 3N \end{cases}, f(\mathbf{y}) = \sum_{i=1}^N \left[\varphi_i\left(\frac{y_{i+2N}}{2(1 - y_{i+N})}\right) - \psi(y_i)\right], \tag{5}$$

and $m = 3N$, $n = 2N$, and $k = N$. In order to ensure quality of service to every user, it is also required that the physical layer rate for any user does not drop below a specified level, i.e., $\mathbb{E}[\min_i b_i(p_i, \zeta_i)] > R^{\min}$. In other words, on average, the minimum physical layer rate is not below $R^{\min}$. The constraint is again in the form required in ($\mathcal{P}$), with $q(z) = z$, and $h(\boldsymbol{\lambda}, \mathbf{p}, \boldsymbol{\zeta}) = -\min_i b_i(p_i, \zeta_i)$ so that $d = 1$. Unlike in Example 1 however, it can be seen that empirical estimates of the moments of $\boldsymbol{\zeta}$ no longer suffice. Instead, the proposed algorithm will entail tracking $\mathbb{E}[\mathbf{g}(\boldsymbol{\lambda}^\star, \mathbf{p}^\star, \boldsymbol{\zeta})]$ and $\mathbb{E}[\mathbf{h}(\boldsymbol{\lambda}^\star, \mathbf{p}^\star, \boldsymbol{\zeta})]$ explicitly.

3) Example 3: Parallel M/G/1 queues in wireless networks (outage capacity): Ergodic capacity is generally difficult to achieve in practical systems with smaller coherence intervals. We consider the joint design of rate and power for systems where both outage and retransmission events are common. This example builds upon the queuing system described in [43].

Within the outage-limited setting, a packet in the i-th queue is transmitted at a fixed rate R_i regardless of the instantaneous channel conditions. The transmission is deemed successful if $R_i < b_i(p_i, \zeta_i)$ and an outage is declared otherwise. In the case of outage, the server continues to attempt to retransmit a packet until it is successful. Define the indicator function for a successful transmission as

$$r_i(p_i, \zeta_i) = \begin{cases} 1 & R_i < b_i(p_i, \zeta_i) \\ 0 & R_i \geq b_i(p_i, \zeta_i). \end{cases} \tag{6}$$

Then the service times depend on the number of retransmission attempts, which in turn depends on the probability of successful transmission $\rho_i = \mathbb{E}[r_i(p_i, \zeta_i)]$. In particular, it can be shown that for this case, the mean waiting time for a packet is given by $w_i(\lambda_i, p_i) = \lambda_i(1 + \rho_i)/\left(2R_i(1 - \rho_i)(R_i(1 - \rho_i) - \lambda_i)\right)$ while the throughput is given by $R_i(1 - \rho_i)$. The overall utility function can be written as $U(\boldsymbol{\lambda}, \mathbf{p}) = \sum_{i=1}^{N} \left[\psi(R_i(1 - \rho_i)) - \varphi_i\left(w_i(\lambda_i, p_i)\right)\right]$ which is also of the form required in ($\mathcal{P}$) with

$$[\mathbf{g}(\mathbf{p}, \mathbf{h})]_i = \begin{cases} r_i(p_i, \zeta_i) & 1 \leq i \leq N \\ \lambda_{i-N} & N + 1 \leq i \leq 2N \end{cases} \tag{7}$$

$$f(\mathbf{y}) = \sum_{i=1}^{N} \left[\varphi_i\left(\frac{y_{i+N}(1 + y_i)}{2R_i(1 - y_i)(R_i(1 - y_i) - y_{i+N})}\right) - \psi(R_i(1 - y_i))\right] \tag{8}$$

Within the context of M/G/1 queues, we briefly discuss some more examples that serve to illustrate the flexibility and generality of the problem considered here.

1) **Delay variance** captures the jitter and may be an important QoS parameter in streaming and gaming applications. The delay variance can be minimized by adding terms of the form $(\mathbb{E}[w_i(\lambda_i, p_i)])^2 - \mathbb{E}[w_i(\lambda_i, p_i)^2]$ in Example 3. Likewise, various moments of delay or throughput may also be incorporated.
2) More generally, **tail probabilities** of the delay or throughput may be considered. For instance in Example 3, where the service time X_i is a random variable, the tail probability can be written as $-\lim_{T\to\infty}(1/T)\,\mathbb{P}\left((1/T)\sum_{t=1}^{T} X_{i,t} > x\right) = \max_\theta\left[\theta x - \log \mathbb{E}[\exp(\theta X_i)]\right]$, and may therefore be incoporated within the objective function.
3) **Worst-case delay** can be considered by using the term $\max_{1\le i\le N} w_i(\lambda_i, p_i)$ in the utility function. Likewise, other non-linear combinations of the delays and throughput may be used.
4) **Expectation constraints** can also be incorporated via penalty methods. For instance, instead of imposing the constraint $\mathbb{E}[b_i(p_i, \zeta_i)] \ge \epsilon$ in Example II-B2, it is possible to use a log-barrier penalty of the form $\log\left(\mathbb{E}[b_i(p_i, \zeta_i)] - \epsilon\right)$ can be used in the objective function. Likewise, instead of imposing the constraint of maximum delay $w_i(\lambda_i, p_i) \le w^{\max}$, a log-barrier penalty of the form $\log\left(w^{\max} - w_i(\lambda_i, p_i)\right)$ can be used. Log penalties have been widely used in the context of interior point methods.
5) **Probabilistic constraints** can be incorporated much in the same way as in Example II-B3. For instance, a constraint of the form $\mathbb{P}(b_i(p_i, \zeta_i) > R_i) > 0.9$ can be converted into an expectation constraint $\mathbb{E}[r_i(p_i, \zeta_i)] > 0.9$ and dealt with using an appropriate penalty function.

4) Example 4: Parallel G/G/1 queues in wireless networks (effective capacity): This example considers the more complicated G/G/1 queue where the expressions for delay cannot be characterized in closed form. The effective capacity, defined as the minimum rate required to guarantee a certain delay (say W), is an approximation that allows handling queues with arbitrary arrival and service distributions. Effective capacity has been widely studied for fading channels in systems where a maximum delay of W units may be tolerated [7, 9, 10].

Let us consider a similar model as in Fig. 1b but with N G/G/1 queues. The packet lengths are constant and the departures depend on the random channel gains and allocated powers $\{p_i \in \mathbb{R}_+\}_{i=1}^{N}$. As in Examples 2, 3, the power vector $\mathbf{p}$ is required to lie in the compact set $\mathbf{\Pi}$. Different from the earlier examples and for the sake of simplicity, it is assumed that the traffic arrival process at the i-th queue has mean m_i^a and variance $(\sigma_i^a)^2$, both of which are known a priori.

Before formulating the problem, let us briefly review the theory of effective capacity. At any time slot t, let $a_i(t)$ denote the instantaneous arrival rate at the i-th queue and $c_i(t)$ denote the instantaneous channel capacity of the i-th channel. The corresponding accumulated rates are denoted by $A_i(t) := \sum_{\tau=1}^{t} a_i(t)$ and $C_i(t) = \sum_{\tau=1}^{t} c_i(t)$, respectively. The effective capacity of the i-th user is given by [9, Eq. (7)]

$$\alpha_i(\theta_i) = \lim_{t\to\infty} \frac{1}{tu_i} \log \mathbb{E}\left[e^{u_i A_i(t)}\right] |_{u_i=\theta_i} = \lim_{t\to\infty} -\frac{1}{tu_i} \log \mathbb{E}\left[e^{-u_i C_i(t)}\right] |_{u_i=\theta_i}, \tag{9}$$

where θ_i is the QoS exponent and is the solution to the above equation [10, Eq. (3)]. The steady-state approximation is often used to simplify these expressions. Specifically, as $t \to \infty$, the application of the central limit theorem allows us to assume the averages $A_i(t)/t$ and $C_i(t)/t$ as Gaussian random variables with finite means and variances.

In the present case, recall that the instantaneous channel capacity is given by $b_i(p_i, \zeta_i)$ where ζ_i is the instantaneous channel. Therefore the mean and variance of $C_i(t)/t$ can be approximated as $m_i^c(p_i) = \mathbb{E}[b_i(p_i, \zeta_i)]$, and $\sigma_i^c(p_i) = \sqrt{\mathbb{E}[b_i^2(p_i, \zeta_i)] - (\mathbb{E}[b_i(p_i, \zeta_i)])^2}$ respectively. As stated earlier, we assume that the means $\{m_i^a\}$ and variances $\{(\sigma_i^a)^2\}$ of the arrival processes are known. Hence (9) can be rewritten as [10, Eq. (6)]

$$m_i^a + \frac{u_i}{2} (\sigma_i^a)^2 |_{u_i=\theta_i} = m_i^c(p_i) - \frac{u_i}{2} (\sigma_i^c(p_i))^2 |_{u_i=\theta_i}, \tag{10}$$

implying that the QoS quotient θ_i is given by $\theta_i(p_i) = (m_i^c(p_i) - m_i^a) / \left((\sigma_i^a)^2 + (\sigma_i^c(p_i))^2\right)$. where we have explicated the dependence on the optimization variables $\{p_i\}$. Substituting into (9), we obtain the effective capacity as

$$\alpha_i(p_i) = \mathbb{E}[b_i(p_i, \zeta_i)] - \frac{\mathbb{E}[b_i(p_i, \zeta_i)] - m_i^a}{(\sigma_i^a)^2 + \mathbb{E}[b_i^2(p_i, \zeta_i)] - (\mathbb{E}[b_i(p_i, \zeta_i)])^2} \left(\mathbb{E}\left[b_i^2(p_i, \zeta_i)\right] - (\mathbb{E}[b_i(p_i, \zeta_i)])^2\right).$$

Proceeding along similar lines, the approximation for the probability of delay exceeding the pre-specified value of W time slots is given by $\mathbb{P}\{W_i(t)|_{t\to\infty} > W\} \approx \eta e^{-\theta_i(p_i)\alpha_i(p_i)W}$ where η is a normalization constant. So the utility function for this case can be defined as $U(\mathbf{p}) = \sum_{i=1}^{N} \left[\psi_i(\alpha_i(p_i)) - \varphi_i\left(e^{-\theta_i(p_i)\alpha_i(p_i)W}\right)\right]$. The utility is therefore in the compositional form with the inner function

$$[\mathbf{g}(\mathbf{p}, \boldsymbol{\zeta})]_i = \begin{cases} b_i(p_i, \zeta_i) & 1 \le i \le N \\ b_i^2(p_i, \zeta_i) & N+1 \le i \le 2N \end{cases},$$

and outer function is given by $f(\mathbf{y}) = \sum_{i=1}^{N} \left[\psi_i\left(\alpha_i(y_i, y_{i+N})\right) + \varphi_i\left(e^{-\theta_i(y_i,y_{i+N})\alpha_i(y_i,y_{i+N})W}\right)\right]$, where $\theta_i(u,v) = \frac{u - m_i^a}{(\sigma_i^a)^2 + v - u^2}$ and $\alpha_i(u,v) = m_i^a + \frac{\theta_i(u,v)}{2}(\sigma_i^a)^2$. In conclusion, it can be seen that the proposed formulation in $(\mathcal{P})$ is applicable to generalized capacity definitions that also involve non-linear functions of expectations.

5) Example 5: Resource provisioning in cloud services: Cloud computing service companies seek to ensure that resources such as power, CPU, and storage are made available to the users as per their demands. Given the uncertainties in the user demands however, the task of distributing the resources to the users is challenging. Ideally, the resources must be provisioned so as to minimize the probability of SLA violation. This example formulates the resource provisioning problem from [2, 3] for the case when demand distributions are unknown.

Consider a service provider that offers N different classes of service and offers a total of C resource units. A customer subscribes to the i-th class by paying a service charge of p_i units in return for r_i units of resources. The service provider also incurs a maintenance charge of χ units per unit time, and per unit resource. An incoming customer may attempt to subscribe to class i of its choice, but may be blocked if the available resources are less than r_i units. The goal is to determine the optimal number of resource units r_i corresponding to each resource class and the total number of resource units to be provisioned.

Let ζ_i denote the load for class i, which represents the number of customers willing to subscribe to class i if the service provider was not resource-constrained. Then the blocking probability for the i-th class is given by [3, Eq. 1] $P_i^B(\mathbf{r}, C) = \mathbb{P}(C - r_i < \sum_{j=1}^{N} \zeta_j r_j \leq C)/\mathbb{P}(\sum_{j=1}^{N} \zeta_j r_j \leq C)$. Roughly, the net revenue obtained by the service provider is given by $R = \sum_i p_i(1 - P_i^B(\mathbf{r}, C))$. On the other hand, the maintenance costs incurred by the service provider amount to χC units per unit time. Consequently, the provider's average profit is given by $U(\mathbf{r}, C) = \sum_{i=1}^{N} p_i N_i (1 - P_i^B(\mathbf{r}, C)) - \chi C$, where N_i is the average number of customers actually subscribing to class i per unit time. Finally, the revenue offered for the i-th class should be a function of the price p_i for that class. Without loss of generality, let $p_1 < p_2 < \ldots < p_N$. Then, the resource units offered for class $i+1$ should be sufficiently different from that offered for class i, and may be restricted to

$$r_i + l(p_{i+1} - p_i) \leq r_{i+1} \leq r_i + u(p_{i+1} - p_i). \tag{11}$$

where l and u are constants. The goal is to maximize the utility $U(\mathbf{r}, C)$ subject to constraints in (11). The problem can be written in the compositional form by defining

$$a_i(\mathbf{r}, C, \boldsymbol{\zeta}) = \begin{cases} 1 & \text{if } C - r_i < \boldsymbol{\zeta}^T \mathbf{r} \le C \\ 0 & \text{Otherwise} \end{cases} \qquad b_i(\mathbf{r}, C, \boldsymbol{\zeta}) = \begin{cases} 1 & \text{if } \boldsymbol{\zeta}^T \mathbf{r} \le C \\ 0 & \text{Otherwise} \end{cases} \tag{12}$$

so that the inner function is given by

$$[\mathbf{g}(\mathbf{r}, C, \boldsymbol{\zeta})]_i = \begin{cases} a_i(\mathbf{r}, C, \boldsymbol{\zeta}) & 1 \le i \le N \\ b_i(\mathbf{r}, C, \boldsymbol{\zeta}) & N+1 \le i \le 2N \\ C & i = 2N+1 \end{cases} \tag{13}$$

and the outer function is $f(\mathbf{y}) = \sum_{i=1}^{N} p_i N_i \left(1 - \frac{y_i}{y_{N+i}}\right) - \chi y_{2N+1}$. This example demonstrates the applicability of the compositional form to problems involving probabilistic functions such as $P_i^B(\mathbf{r}, C)$.

Before concluding, it is remarked that the objective function and constraints in the examples are specifically selected for simplifying the exposition. More general utility functions and constraints are also possible, as long as they adhere to the compositional structure of ($\mathcal{P}$).

III. STOCHASTIC COMPOSITIONAL GRADIENT DESCENT FOR CONSTRAINED PROBLEMS

This section details the proposed algorithm for solving ($\mathcal{P}$) and provides the corresponding convergence rates. For the sake of brevity, we define $F(\mathbf{x}) := f(\mathbb{E}[\mathbf{g}(\mathbf{x}, \boldsymbol{\zeta})])$ and $\mathbf{Q}(\mathbf{x}) := \mathbf{q}(\mathbb{E}[\mathbf{h}(\mathbf{x}, \boldsymbol{\zeta})])$. In the unconstrained case, ($\mathcal{P}$) can be solved using the SCGD algorithm proposed in [15] for convex non-smooth or smooth non-convex F. This section focuses on the constrained case when both F and Q are convex but possibly non-smooth.

A. Assumptions

We begin with discussing the necessary assumptions on the functions f, $\mathbf{g}$, $\mathbf{q}$ and $\mathbf{h}$.

A1. The outer functions f, $\mathbf{q}$ are continuously differentiable and the inner functions $\mathbf{g}$, $\mathbf{h}$ are continuous. Consequently, the sub-gradients of the objective and constraint functions are well-defined, with

$$\nabla F(\mathbf{x}) := \mathbb{E}\widetilde{\nabla}\mathbf{g}(\mathbf{x}, \boldsymbol{\zeta}) \nabla f\left(\mathbb{E}\mathbf{g}(\mathbf{x}, \boldsymbol{\zeta})\right), \quad \nabla Q_j(\mathbf{x}) := \mathbb{E}\widetilde{\nabla}\mathbf{h}(\mathbf{x}, \boldsymbol{\zeta}) \nabla q_j\left(\mathbb{E}\mathbf{h}(\mathbf{x}, \boldsymbol{\zeta})\right) \tag{14}$$

where $[\widetilde{\nabla}\mathbf{g}(\mathbf{x}, \boldsymbol{\zeta})]_{ij} := \partial g_j(\mathbf{x}, \boldsymbol{\zeta})/\partial x_i$ and $[\widetilde{\nabla}\mathbf{h}(\mathbf{x}, \boldsymbol{\zeta})]_{ij} := \partial h_j(\mathbf{x}, \boldsymbol{\zeta})/\partial x_i$ with ∂ denoting the generalized gradient operation.

A2. The problem $(\mathcal{P})$ is a convex optimization problem and the constraints can be satisfied strictly, i.e., there exists $\widetilde{\mathbf{x}}$ such that $\max_j Q_j(\widetilde{\mathbf{x}}) + \sigma \le 0$ for some constant $\sigma > 0$. The set $\mathcal{X}$ is closed and compact, i.e., $\sup_{\mathbf{x},\mathbf{x}' \in \mathcal{X}} \|\mathbf{x} - \mathbf{x}'\|^2 \le D_x < \infty$.

A3. The random variables $\boldsymbol{\zeta}_1, \boldsymbol{\zeta}_2, \ldots$ are independent and identically distributed.

A4. The functions $\mathbf{g}$, $\mathbf{h}$ are Lipschitz continuous in expectation and have bounded second order moments, i.e.,

$$\mathbb{E}\left[\left\|\widetilde{\nabla}\mathbf{g}(\mathbf{x},\boldsymbol{\zeta})\right\|^2\right] \le C_g, \qquad \mathbb{E}\left[\|\mathbf{g}(\mathbf{x},\boldsymbol{\zeta}) - \mathbb{E}\mathbf{g}(\mathbf{x},\boldsymbol{\zeta})\|^2\right] \le V_g \qquad \forall \mathbf{x} \in \mathcal{X} \tag{15}$$

$$\mathbb{E}\left[\left\|\widetilde{\nabla}\mathbf{h}(\mathbf{x},\boldsymbol{\zeta})\right\|^2\right] \le C_h, \qquad \mathbb{E}\left[\|\mathbf{h}(\mathbf{x},\boldsymbol{\zeta}) - \mathbb{E}\mathbf{h}(\mathbf{x},\boldsymbol{\zeta})\|^2\right] \le V_h \qquad \forall \mathbf{x} \in \mathcal{X} \tag{16}$$

A5. The functions f, q are smooth and have bounded gradients, i.e.,

$$\|\nabla f(\mathbf{y})\|^2 \le C_f, \qquad \|\nabla f(\mathbf{y}) - \nabla f(\bar{\mathbf{y}})\| \le L_f \|\mathbf{y} - \bar{\mathbf{y}}\| \qquad \forall \mathbf{y}, \bar{\mathbf{y}} \in \mathbb{R}^m \tag{17}$$

$$\|\nabla \mathbf{q}(\mathbf{z})\|^2 \le C_q, \qquad \|\nabla \mathbf{q}(\mathbf{z}) - \nabla \mathbf{q}(\bar{\mathbf{z}})\| \le L_q \|\mathbf{z} - \bar{\mathbf{z}}\| \qquad \forall \mathbf{z}, \bar{\mathbf{z}} \in \mathbb{R}^d. \tag{18}$$

These assumptions are standard and hold for most problems at hand. In particular, we do not require $\mathbf{g}$ and $\mathbf{h}$ to be smooth or convex. Only the outer functions f, q are required to be smooth and only the composite functions F and $\mathbf{Q}$ are required to be convex. As in [15], the matrices $\widetilde{\nabla}\mathbf{g}(\mathbf{x},\boldsymbol{\zeta})$ and $\widetilde{\nabla}\mathbf{h}(\mathbf{x},\boldsymbol{\zeta})$ may consist of gradients, sub-gradients, or other directional gradients satisfying (14). Further, none of the functions are required to be strongly convex. Note that the Lipschitz continous function $\mathbf{Q}$ is defined over a compact domain, and hence is bounded.

B. Proposed Algorithm

In order to describe the proposed algorithm, let us also define the smooth and convex function $\ell(\mathbf{w}) = \sum_{j=1}^{J} \ell_j(w_j + \gamma)$ where ℓ_j is defined as

$$\ell_j(x) := \begin{cases} \frac{1}{2}x^2 & 0 \le x \le C_\ell \\ C_\ell x - \frac{C_\ell^2}{2} & x > C_\ell \\ 0 & x < 0 \end{cases} \tag{19}$$

and γ and C_ℓ are algorithm parameters. It can be seen that $\nabla \ell_j(x) = \max\{x, C_\ell\}$ for $x > 0$ and zero otherwise. Such a function allows us to define a convex penalty function $L(\mathbf{x}) := \ell(\mathbf{q}(\mathbb{E}[\mathbf{h}(\mathbf{x},\boldsymbol{\zeta})]))$ that quantifies the constraint violation and allows us to control it through the choice of the parameter γ. The penalty function defined in (19) is quadratic and imposes square penalty on the constraint violation. Intuitively, the update takes the iterate towards the descent direction of optimal function F as well as towards the feasible region $\{\mathbf{x}|\mathbf{Q}(\mathbf{x}) \leq 0\}$.

Recall that the SCGD algorithm for the unconstrained version of $\mathcal{P}$ entails carrying out updates of the form $\mathbf{x}_{t+1} = \mathbf{x}_t - \alpha_t \hat{\nabla} F(\mathbf{x}_t)$ where $\hat{\nabla} F(\mathbf{x}_t)$ is an approximation of the actual gradient $\nabla F(\mathbf{x}_t)$ and α_t is the step size. To handle the constrained case, we consider the iterations $\mathbf{x}_{t+1} = \mathbf{x}_t - \alpha_t \hat{\nabla} F(\mathbf{x}_t) - \delta_t \hat{\nabla} L(\mathbf{x}_t)$ where δ_t is also a step size parameter and $\hat{\nabla} L(\mathbf{x}_t)$ approximates the true gradient $\nabla L(\mathbf{x}_t)$. Finally, the approximate gradient $\hat{\nabla} L(\mathbf{x}_t)$, and $\hat{\nabla} F(\mathbf{x}_t)$ are calculated by introducing new variables $\mathbf{z}_t$, and $\mathbf{y}_t$ that track the quantities $\mathbb{E}[\mathbf{h}(\mathbf{x}^\star, \boldsymbol{\zeta})]$, and $\mathbb{E}[\mathbf{g}(\mathbf{x}^\star, \boldsymbol{\zeta})]$ are respectively. The step sizes α_t, and δ_t are chosen such that the updates converge to a point that is not only optimal but also feasible. The full updates are summarized in Algorithm 1. The operation $\Pi_{\mathcal{X}}$ denotes projection onto the set $\mathcal{X}$. The computational complexity of the proposed algorithm is only slightly higher than that of the classical stochastic gradient descent method. At each time instant, only gradients of $\mathbf{g}$ and $\mathbf{h}$ are needed in order to carry out the updates. As compared to the SCGD algorithm for unconstrained setting [15], the proposed algorithm incurs only a single additional update. Finally, the iteration complexity results specify the precise number of data points needed in order to achieve a specified performance, making the algorithm readily applicable to any real-world application.

Algorithm 1 Constrained SCGD

```
**Input:** $\mathbf{x}_1 \in \mathbb{R}^n$, $\mathbf{y}_1 \in \mathbb{R}^m$, $\mathbf{z}_1 \in \mathbb{R}^d$ step sizes $\alpha_t, \beta_t, \delta_t \subset (0, 1]$.
**for** $t = 1, 2, ...$
    Observe the random variable $\boldsymbol{\zeta}_t$, and update
        $\mathbf{y}_{t+1} = (1 - \beta_t)\mathbf{y}_t + \beta_t \mathbf{g}(\mathbf{x}_t, \boldsymbol{\zeta}_t), \qquad \mathbf{z}_{t+1} = (1 - \beta_t)\mathbf{z}_t + \beta_t \mathbf{h}(\mathbf{x}_t, \boldsymbol{\zeta}_t)$    (20)
        $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}} \left\{ \mathbf{x}_t - \alpha_t \widetilde{\nabla} \mathbf{g}(\mathbf{x}_t, \boldsymbol{\zeta}_t) \nabla f(\mathbf{y}_{t+1}) - \delta_t \widetilde{\nabla} \mathbf{h}(\mathbf{x}_t, \boldsymbol{\zeta}_t) \nabla \mathbf{q}(\mathbf{z}_{t+1}) \nabla \ell(\mathbf{q}(\mathbf{z}_{t+1})) \right\}$    (21)
**end**
**Output:** $\hat{\mathbf{x}} = \frac{2}{T} \sum_{t=T/2}^{T} \mathbf{x}_t$.
```

C. Performance Analysis

This section provides the performance analysis of the proposed algorithm. We begin with describing the surrogate problem

$$\mathbf{x}^\gamma = \arg\min_{\mathbf{x}\in\mathcal{X}} F(\mathbf{x}), \quad \text{s.t. } \mathbf{Q}(\mathbf{x}) + \gamma\mathbf{1} \leq 0 \tag{22}$$

for a given value of γ. The subsequent analysis utilizes the following key steps (a) bounding the optimality gap $F(\hat{\mathbf{x}}) - F(\mathbf{x}^\gamma)$ and the penalty $L(\hat{\mathbf{x}})$; (b) relating the solutions $\mathbf{x}^\star$ and $\mathbf{x}^\gamma$; and (c) obtaining the corresponding bounds on the optimality gap $F(\hat{\mathbf{x}}) - F(\mathbf{x}^\star)$ and the constraint violation $\sum_j \max\{Q_j(\hat{\mathbf{x}}), 0\}$.

For the first step, we establish the following Lemma, whose proof is provided in Appendix A.

Lemma 1. Under Assumptions **A1**-**A5**, the iterates generated by Algorithm 1 satsify

$$\sum_{t=T/2}^{T} \mathbb{E}\left[F(\mathbf{x}_t) + \frac{\delta_t}{\alpha_t}L(\mathbf{x}_t) - F(\mathbf{x}^\gamma)\right] \leq \mathcal{O}\left(\frac{1}{\alpha_T} + \sum_{t=T/2}^{T}\left[\frac{\delta_t^2}{\alpha_t} + \frac{\beta_t^2}{\alpha_t} + \frac{\delta_t^2}{\alpha_t\beta_t}\right]\right). \tag{23}$$

The constants within $\mathcal{O}(\cdot)$ in (23) depend only on the $D_x, C_q, C_h, C_g, J, C_\ell, L_f, L_q, V_g$, and V_h. The proof of Lemma 1 follows along the lines of the SCGD proof in [15], with appropriate modifications incorporated to handle the two variables $\mathbf{y}_t$ and $\mathbf{z}_t$, and the required changes in obtaining the bound. Specifically, the proof requires us to re-define appropriate sequences that allow the telescopic sums to be taken.

For the next step, we relate the solutions of (22) and ($\mathcal{P}$). The proof is provided in Appendix B.

Lemma 2. Under Assumption (**A2**) and for $0 < \gamma < \sigma/2$, it holds that $F(\mathbf{x}^\gamma) - F(\mathbf{x}^\star) \leq 2\sqrt{C_g C_f} D_x \gamma/\sigma$.

Lemma 2 follows from the strong duality of (22) and from standard results in duality theory. An implication of Lemma 2 is that the gap between the solutions of ($\mathcal{P}$) and (22) cannot be large if γ is small. Having stated the preliminary results, we now state the main result of the paper as the following theorem. For the sake of brevity, let us define

$$\omega = 2\frac{D_1}{T\alpha_T} + 2\frac{D_2}{T}\sum_{t=T/2}^{T}\left[\frac{\delta_t^2}{\alpha_t} + \frac{\beta_t^2}{\alpha_t} + \frac{\delta_t^2}{\alpha_t\beta_t}\right] \tag{24}$$

where D_1 and D_2 are constants as defined in Appendix A.

Theorem 1. Let the step-sizes be chosen as

1) either diminishing with $\alpha_t = t^{-a}$, $\beta_t = t^{-b}$, $\delta_t = t^{-c}$ for $1 \leq t \leq T$;
2) or constant with $\alpha_t = T^{-a}$, $\beta_t = T^{-b}$, $\delta_t = T^{-c}$ for all $1 \leq t \leq T$,

with $1 > a \geq c \geq b > 0$. Under **(A1)**-**(A5)** and for C_ℓ chosen such that $\max_{\mathbf{x}\in\mathcal{X}} Q_j(\mathbf{x}) \leq C_\ell - \gamma$, we have the following bounds

$$\mathbb{E}\left[F(\hat{\mathbf{x}})\right] - F^\star \leq \omega + \frac{2\sqrt{C_g C_f} D_x}{\sigma}\gamma, \quad \max_j Q_j(\hat{\mathbf{x}}) \leq \sqrt{\frac{J T^{c-a}\left(\omega + \sqrt{C_f C_g} D_x\right)}{2^{c-a}}} - \gamma$$

where $\omega = \mathcal{O}\left(T^{a-1} + T^{a-2c} + T^{a-2b} + T^{a+b-2c}\right)$.

Proof: The proof is written in two parts. We first begin with analyzing the optimality gap and subsequently bound the constraint violation.

Since the right-hand side of (23) is equal to $\omega T/2$ (cf. (68)), we have that

$$\sum_{t=T/2}^{T} \mathbb{E}\left[F(\mathbf{x}_t) + \frac{\delta_t}{\alpha_t} L(\mathbf{x}_t) - F(\mathbf{x}^\gamma)\right] \leq \omega \frac{T}{2}. \tag{25}$$

Since $\mathbb{E}\left[L(\mathbf{x}_t)\right] \geq 0$, it follows that $\sum_{t=T/2}^{T} \mathbb{E}\left[F(\mathbf{x}_t) - F(\mathbf{x}^\gamma)\right] \leq \omega \frac{T}{2}$. Therefore the optimality gap can be bounded as

$$\frac{2}{T}\sum_{t=T/2}^{T} \mathbb{E}\left[F(\mathbf{x}_t) - F(\mathbf{x}^\star)\right] \leq \omega + \frac{2}{T}\sum_{t=T/2}^{T} \mathbb{E}\left[F(\mathbf{x}^\gamma) - F(\mathbf{x}^\star)\right] \leq \omega + \frac{2\sqrt{C_g C_f} D_x}{\sigma}\gamma. \tag{26}$$

Since F is convex, the bound in (26) implies that $\mathbb{E}\left[F(\hat{\mathbf{x}})\right] - F^\star \leq \mathcal{O}(\omega + \gamma)$ where both ω and γ may depend on T.

Next observe that since F is Lipschitz with parameter $\sqrt{C_f C_g}$ and the set $\mathcal{X}$ is compact, it follows that $F(\mathbf{x}_t) - F(\mathbf{x}^\gamma) \geq -\sqrt{C_f C_g}\left\|\mathbf{x}_t - \mathbf{x}^\gamma\right\| \geq -\sqrt{C_f C_g} D_x$. Substituting into (25) and rearranging yields

$$\frac{2}{T}\sum_{t=T/2}^{T} \frac{\delta_t}{\alpha_t}\mathbb{E}\left[L(\mathbf{x}_t)\right] \leq \omega + \sqrt{C_f C_g} D_x. \tag{27}$$

Since the sequence $\{\frac{\alpha_t}{\delta_t}\}$ is either diminishing or constant, it holds that $\frac{\alpha_t}{\delta_t} \leq \frac{\alpha_{T/2}}{\delta_{T/2}}$ for all $T/2 \leq t \leq T$. Recalling that $\hat{\mathbf{x}} = \frac{2}{T}\sum_{t=T/2}^{T} \mathbf{x}_t$ and that the function L is convex, the left-hand side of (27) can be lower bounded as

$$\frac{2}{T}\sum_{t=T/2}^{T} \frac{\delta_t}{\alpha_t}\mathbb{E}\left[L(\mathbf{x}_t)\right] \geq \frac{\delta_{T/2}}{\alpha_{T/2}}\frac{2}{T}\sum_{t=T/2}^{T} \mathbb{E}\left[L(\mathbf{x}_t)\right] \geq \frac{\delta_{T/2}}{\alpha_{T/2}}\mathbb{E}\left[L(\hat{\mathbf{x}})\right]. \tag{28}$$

TABLE I: Summery of bounds on optimality gap and constraint violation

Choice of constants a, b, c	Optimality gap	Constraint violation	γ
$a = 0.9167$, $b = 0.5$, $c = 0.75$	$\mathcal{O}\left(T^{-1/12}\right)$	0	$\mathcal{O}\left(T^{-1/12}\right)$
$a = 0.9167$, $b = 0.5$, $c = 0.75$	$\mathcal{O}\left(T^{-1/12}\right)$	$\mathcal{O}\left(T^{-1/12}\right)$	0
$a = 0.875$, $b = 0.5$, $c = 0.75$	$\mathcal{O}\left(T^{-1/8}\right)$	$\mathcal{O}\left(T^{-1/16}\right)$	0
$a = 0.75$, $b = 0.5$, $c = 0.75$	$\mathcal{O}\left(T^{-1/4}\right)$	$\mathcal{O}(1)$	0

Recall that Q is bounded and therefore there exists C_ℓ such that $\max_j Q_j(\mathbf{x}) \leq C_\ell - \gamma$ for $\gamma \in (0,1)$ or equivalently, $L(\hat{\mathbf{x}}) = \sum_j [Q_j(\hat{\mathbf{x}}) + \gamma]_+^2$. Therefore, the lower bound in (28) becomes

$$\frac{2}{T}\sum_{t=T/2}^{T}\frac{\delta_t}{\alpha_t}\mathbb{E}\left[L(\mathbf{x}_t)\right] \geq \frac{\delta_{T/2}}{\alpha_{T/2}}\sum_{j=1}^{J}[Q_j(\hat{\mathbf{x}}) + \gamma]_+^2 \geq \frac{\delta_{T/2}}{J\alpha_{T/2}}\left(\sum_{j=1}^{J}[Q_j(\hat{\mathbf{x}}) + \gamma]_+\right)^2. \tag{29}$$

Substituting (29) into (27) and taking square root, we obtain

$$\sum_{j=1}^{J}[Q_j(\hat{\mathbf{x}}) + \gamma]_+ \leq \sqrt{\frac{J\alpha_{T/2}\left(\omega + \sqrt{C_f C_g} D_x\right)}{\delta_{T/2}}}$$

$$\Rightarrow \max_j Q_j(\hat{\mathbf{x}}) \leq \sqrt{\frac{JT^{c-a}\left(\omega + \sqrt{C_f C_g} D_x\right)}{2^{c-a}}} - \gamma.$$

Finally the required order bound on ω can be obtained simply by substituting the constant/diminishing step-sizes and ignoring all constant terms. ■

Theorem 1 provides a general bound on the optimality gap and the constraint violation that depends on the choice of constants a, b, c, and γ. For instance, it can be seen that if $\gamma = \sqrt{\frac{JT^{c-a}\left(\omega+\sqrt{C_f C_g} D_x\right)}{2^{c-a}}}$ the proposed algorithm incurs zero constraint violation. By carefully choosing the parameters a, b, and c, it is then possible to minimize the bounds on the optimality gap and the constraint violation, as shown in Table I. As expected, the optimality gap decays slowly when stricter requirements are imposed on the constraint violation. It is remarked that the SCGD algorithm proposed in [15] for unconstrained problems also incurs an optimality gap of $\mathcal{O}(T^{-1/4})$, implying that the bounds obtained are tight in the limiting case.

IV. SIMULATION RESULTS

This section provides numerical tests for all examples discussed in Sec. II-B. For each example, Algorithm 1 is implemented over a network of $N = 3$ parallel queues for $\psi_i(x) = \bar{\psi}_i \log(x)$, and

$\varphi_i(x) = \bar{\varphi}_i x$ for $1 \leq i \leq N$. We use the values $\bar{\psi}_1 = 1$, $\bar{\psi}_2 = 1.5$, $\bar{\psi}_3 = 2$, $\bar{\varphi}_1 = 10$, $\bar{\varphi}_2 = 15$, and $\bar{\varphi}_3 = 20$ in all the examples. The step sizes are chosen as constant with the selection $a = 0.9167$, $b = 0.5$, and $c = 0.75$. For all cases, the evolution of the optimality gap $F(\hat{\mathbf{x}}) - F^\star$ as well as the constraint violation $\max_j Q_j(\hat{\mathbf{x}})$ with iterations is shown.

A. Example 1: Parallel M/G/1 Queues

For the problem considered in Example 1, it is clear that outer functions f and q are differentiable (cf. (2) and (3)). Likewise, the inner functions $\mathbf{g}$ and $\mathbf{h}$ are continuous, as required in (**A1**). Further, observe that both the objective function as well as the constraint functions are convex in $\boldsymbol{\lambda}$. In its current form however, the gradient of the outer functions f and q is unbounded. To this end, let $L^{\max}$ denote the maximum length of the packet and let $\lambda^{\max}$ be such that $\lambda^{\max} \leq \epsilon C / L^{\max}$ for some given constant $\epsilon \in (0,1)$. Such a requirement ensures that arrival rate is not arbitrarily close to the system capacity.

Given $0 < \sigma < D^{\max}$, it can be seen that any $\lambda_i < 2C_i^2(D^{\max} - \sigma)/(\mathbb{E}L_i^2 + 2C_i(D^{\max} - \sigma)\mathbb{E}L_i)$ is a strictly feasible solution to (3), as required in (**A2**). Next, (**A3**) holds if the packet lengths are i.i.d. and (**A4**) holds with $C_g = V_g = C_h = V_h = \sum_{i=1}^{N} \lambda^{\max} \left(\mathbb{E}L_i^2 + \mathbb{E}L_i^4\right)$ where the right-hand side is bounded since $L^{\max} < \infty$. Finally, (**A5**) holds with constants

$$C_f = \sum_{i=1}^{N} \left[\left(\bar{\varphi}_i \frac{1 - \epsilon + \epsilon L_i^{\max}}{4C_i \left(1 - \epsilon\right)^2} \right)^2 + \left(\frac{\bar{\psi}_i}{\lambda^{\min}} \right)^2 \right],$$

$$L_f = \max_{1 \leq i \leq N} \left\{ \bar{\varphi}_i \frac{1 - \epsilon + \epsilon L_i^{\max}}{C_i^3 (1-\epsilon)^3} + \bar{\varphi}_i \frac{\epsilon C_i L^{\max} - 1}{2C_i^3 (1-\epsilon)^2} + 2\bar{\psi}_i \left(\frac{1}{\lambda^{\min}} \right)^2, \bar{\varphi}_i \frac{1}{4C_i^2 (1-\epsilon)} \right\}.$$

It can be seen that the constants depend on the specified value of ϵ. In other words, attempt to operate the system very close to the boundary may lead to larger constants and slower rates of convergence.

Fig. 2 considers an example system with $D^{\max} = 50ms$, $C = [100\ 200\ 500]$ kbps, $L^{\max} = [20\ 30\ 60]$ kb, $\lambda^{\min} = 0.1$, $\lambda = 15$, $\lambda^{\max} = [5\ 7\ 9]$, and mean packet lengths $\bar{L} = [15\ 20\ 35]$ bits. The simulation is run for different packet length distributions, all truncated to respect the maximum packet length requirements. It is remarked that Algorithm 1 needs one sample of L_i at each time instant, in addition to the memory required to store the different iterates. As can be seen from Fig. 2, 100 samples are enough for the algorithm to converge. Note that the problem is simpler if the distribution of packet lengths is known and $F^\star$ can be directly calculated as a

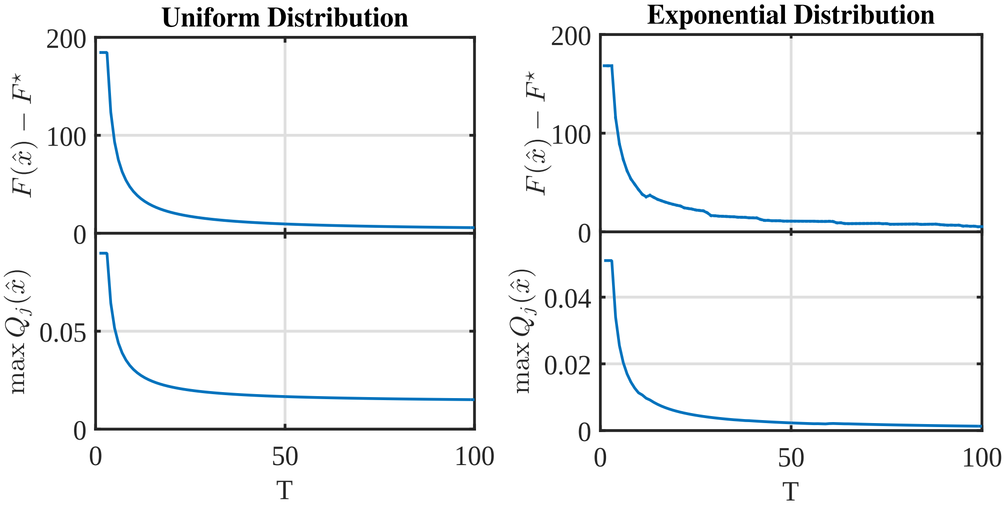

Uniform Distribution

Exponential Distribution

Fig. 2: Evolution of the objective function and the constraint violation for Example 1

function of the first and second moments of L_i. As expected, the optimality gap and the constraint violations decay with iterations.

B. Example 2: Parallel M/G/1 queues in wireless networks (ergodic capacity)

It can be seen that the inner functions $\mathbf{g}$ and h are continuous as required in (**A1**). For ease of analysis, let us assume that the support of the random variable ζ_i is bounded away from zero, i.e., $\zeta_i \geq G$. Coupled with the fact that $p_i \geq P^{\min}$, it follows that $b_i(p_i, \zeta_i) > B_i \log(1 + P^{\min} G)$. Therefore, it is always possible to find $p_i \in (P^{\min}, P^{\max})$ such that $\mathbb{E}\left[\min_i b_i(p_i, \zeta_i)\right] > R^{\min} - \sigma$ as required in (**A2**). Further, since $\mathbf{g}$ is locally Lipschitz on compact sets $\mathbf{\Lambda}$ and $\mathbf{\Pi}$, it follows that both $\mathbf{g}(\boldsymbol{\lambda}, \mathbf{p}, \boldsymbol{\zeta})$ and $\nabla \mathbf{g}(\boldsymbol{\lambda}, \mathbf{p}, \boldsymbol{\zeta})$ are bounded for all $\mathbf{p} \in \mathbf{\Pi}$ and $\boldsymbol{\lambda} \in \mathbf{\Lambda}$. Indeed, as required in (**A4**), we have that

$$C_g = V_g = N + \sum_{i=1}^{N} \left[\frac{1}{b_i^2(P^{\min}, G)} \left(1 + \frac{\lambda^{\max}}{(1 + P^{\min} G)^2 \left(b_i^2(P^{\min}, G)\right)}\right) + \frac{1}{b_i^4(P^{\min}, G)} \left(1 + \frac{2\lambda^{\max}}{(1 + P^{\min} G)^2 \left(b_i^2(P^{\min}, G)\right)}\right) \right],$$

$C_h = V_h = \frac{\max_i B_i}{(1+P^{\min}G)^2}$ where B_i is the capacity of the i-th channel.

Observe that the outer function f as defined in Example 2 is unbounded and has a restricted domain. Instead, we redefine the outer functions as $\tilde{f}$ that approximates f while also satisfying the regularity conditions. Let $\tilde{f}(\mathbf{y}) = \sum_i \varphi_i(\varsigma_i(y_{i+2N}, y_{i+N})) - \psi(y_i)$ where

$$\varsigma_i(x, y) = \begin{cases} \frac{x}{1-y} & y < 1 - \epsilon \\ x^{\frac{y-1+2\epsilon}{\epsilon^2}} & y \geq 1 - \epsilon. \end{cases}$$

Note that since $b_i(p_i, \zeta_i) > B_i \log(1 + P^{\min} G)$, the iterate $\mathbf{y}_t$ takes values in a compact domain, and consequently, it suffices to ensure that $\tilde{f}(\mathbf{y})$ is locally Lipschitz, implying that $\nabla \tilde{f}(\mathbf{y})$ is bounded. As required in (**A5**), we have that

$$C_f = \sum_{i=1}^{N} \left[\frac{\bar{\varphi}_i^2}{4\left(1-\epsilon\right)^2} \left(1 + \frac{1}{b_i^2(P^{\min}, G)\left(1-\epsilon\right)^2} \right) + \left(\frac{\bar{\psi}_i}{\lambda^{\min}} \right)^2 \right]$$

$$L_f = \max_{1 \leq i \leq N} \left\{ \frac{\bar{\varphi}_i}{\left(1-\epsilon\right)^2} \left(\frac{1}{2} + \frac{1}{b_i(P^{\min}, G)\left(1-\epsilon\right)} \right) + 2\bar{\psi}_i \left(\frac{1}{\lambda^{\min}} \right)^2 \right\}.$$

As $q(z) = z$, it is easy to see that $L_q = 0$, and $C_q = 1$. Finally, (**A3**) holds if the channel gains are i.i.d.

Verifying the convexity of the redefined objective function is challenging, as the expression for the Hessian is quite complicated. Instead, we resort to a numerical verification by observing that the objective function is smooth in $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}$ and $\mathbf{p} \in \boldsymbol{\Pi}$. The Hessian for a summand of $\tilde{f}$ is evaluated over a two dimensional grid in $\boldsymbol{\Lambda} \times \boldsymbol{\Pi}$ and is found to be always strictly positive definite for the parameters considered in the simulations. Numerical verification is deferred to Appendix C.

For experiments we consider a setting with $B_i = 10$ kbps for $i = 1, \ldots, N$, $\lambda = 37$, $\epsilon = 0.95$, $G = 0.25$, $R^{\min} = 35$ kbps, $P^{\min} = 14$ units, $P^{\max} = 100$ units, $\lambda^{\min} = 0.1$, $\lambda^{\max} = 15$. Rayleigh fading is assumed so that ζ_i are chi-squared distributed with support truncated to G. Two sets of simulations are carried out with 5, and 10 antennas at the BS, and the performance is shown in Fig. 3. It can be seen that the optimality gap decays to zero as expected. Interestingly, observe that the constraints are never violated when the number of antennas are large, likely due to the diversity gains available to the system. Different from Example 1 however, $F^\star$ cannot be obtained in closed form but requires brute force search, since the moments of the inner function are themselves functions of the powers. From the Fig. 3, it is evident that for this example, 300 samples are required to obtain a feasible solution that is close to optimum.

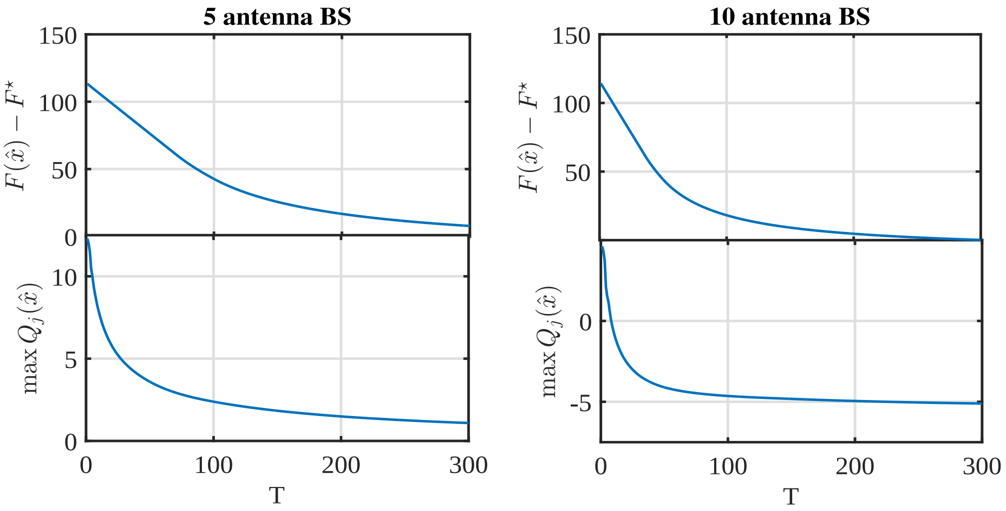

Fig. 3: Evolution of the objective function and the constraint violation for Example 2 (ergodic capacity) for 5 antennas and 10 antennas at the BS.

C. Example 3: Parallel M/G/1 queues in wireless networks (outage capacity)

The objective function in Example 3 is not convex, and therefore the results in Theorem III do not apply. Nevertheless, the results from [15] are still applicable under (**A1**), (**A3**), (**A4**), and (**A5**). First, note that $\mathbf{\Lambda}$ and $\mathbf{\Pi}$ are defined as in Example 2. The channel is also assumed to have a support bounded away from zero, i.e., $\zeta_i \geq G$. Since the domain of f is restricted, it is possible to utilize an approximate version as in Example 2. Alternatively, if the problem parameters satisfy $\epsilon R_i \left(1 - r_i(P^{\min}, G)\right) \geq \lambda^{\max}$ for some ϵ, then the outer function's gradient is bounded. In the current form, inner function g (see. (7)) is also not continuous as required in (**A1**). To this end we approximate the indicator function in (6) with the sigmoid function $r_i(p_i, \zeta_i) = 1/\left(1 + e^{-\eta(R_i - b_i(p_i,\zeta_i))}\right)$ where $\eta \geq 1$ is a parameter that controls the tightness of the approximation. Outer function f in (8) is clearly continuously differentiable (as required in (**A1**)) and (**A3**) holds if the channel gains are i.i.d. Further (**A4**) holds with

$$C_g = V_g = N + \sum_{i=1}^{N} \frac{1}{2} \left(1 + \left(\frac{B_i G}{1 + P_i^{\min} G}\right)^2\right). \tag{30}$$

Since the domain of $(\mathbf{p}, \boldsymbol{\lambda})$ is compact and f is locally Lipschitz, appropriate constants C_f and L_f can likewise be found for (**A5**) as well.

For experiments we consider a setting with $B_i = 100$ kbps, $\lambda = 45$, $\epsilon = 0.95$, $G = 0.25$, $R = [30, 35, 40]$ kbps, $P^{\min} = 10$ units, $P^{\max} = 100$ units, $\lambda^{\min} = 0.1$, and $\lambda_i^{\max} = 25$. Channel

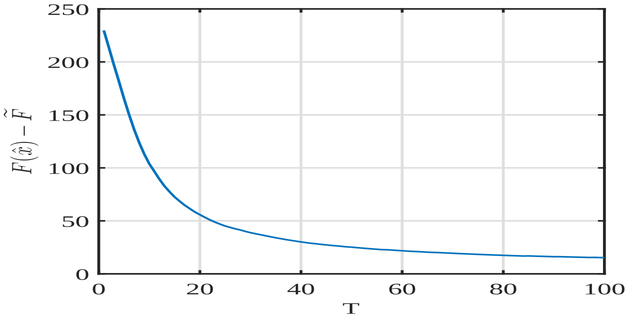

Fig. 4: Evolution of the objective function of Example 3 (outage capacity).

gains are exponentially distributed with mean 1 and the base station has a single antenna. The SCGD algorithm [15] is implemented for this example and the simulations are shown in Fig. 4. For this case, the algorithm is only guaranteed to converge to a stationary point. However, in order to also evaluate the quality of the solution, we minimize the objective using the trust-region reflective method as implemented in the fmincon function of MATLAB. It is remarked that solving the problem using such an algorithm is generally impractical since it requires carrying out simulations in order to calculate the expected gradient and Hessian at every iteration. As evident from Fig. 4, it can be seen that the objective function values converge to a local minimum value $\tilde{F}$ obtained from MATLAB, while incurring significantly lower computational cost.

D. Example 4: Parallel G/G/1 queues in wireless networks (effective capacity)

Similar to the previous example, the objective function here is not convex, and the results from [15] apply under assumptions (**A1**), (**A3**), (**A4**), and (**A5**). Channel gains at consecutive update intervals are considered to be i.i.d., and therefore (**A3**) holds. As defined in Example 2, 3, $\Pi(\cdot)$ acts as the ball constraint, and because of which the domain of f is restricted. As required in (**A5**), the outer function's gradient is bounded as long as $\mathbb{E}\left[b_i(P^{\max}, \zeta_i)\right]$ and $\mathbb{E}\left[\left(b_i(P^{\max}, \zeta_i)\right)^2\right]$ are finite. Additionally, outer function f is differentiable as required in (**A1**). Further (**A4**) holds with

$$C_g = \left(1 + 4B_i^2\right) B_i^2 \left(P^{\max}\right)^2, \qquad V_g = \left(1 + 4B_i^2\right) B_i^2 \left(P^{\max}\right)^4. \tag{31}$$

For experiments, we consider a setting with $B_i = 100$ kbps, $P^{\min} = 0.1$ units, $P^{\max} = 0.9$ units, and $W = 500$ ms. Channel gains are exponentially distributed with means $[0.8, 0.9, 1]$, and the

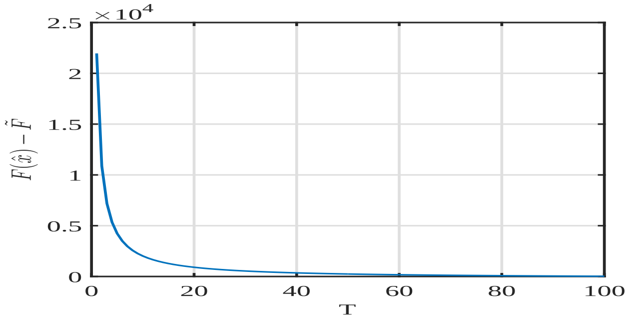

Fig. 5: Evolution of the objective function of Example 4 (effective capacity).

base station has a single antenna. Simulations for this example are shown in Fig. 5. Since the problem is non convex, the algorithm [15] is only guaranteed to converge to a stationary point. The local optimum objective $\tilde{F}$ is calculated by using the approximations in [10]. Since the objective function is of composition form, the inner functions $\mathbb{E}\left[b_i(p_i, \zeta_i)\right]$, $\mathbb{E}\left[(b_i(p_i, \zeta_i))^2\right]$ are approximated to a deterministic functions of p_i, and $\mathbb{E}\left[\zeta\right]$, for Rayleigh channels. In other words, composition form can be avoided if the channel gain ζ_i is of exponentially distributed. However the approximations are of some complicated functions, and valid only for Rayleigh channels. In contrast, the SCGD algorithm is applicable to any channel with arbitrary distributions. As evident from Fig. 5, the objective function value converges to a local minimum $\tilde{F}$.

V. CONCLUSION

The problem of queuing system design is formulated as a constrained stochastic optimization problem. As the objective and the constraints of the problem are non-linear functions of expectations, classical stochastic approximation methods cannot be used. Instead a constrained variant of the stochastic compositional gradient descent method is proposed where the expected value functions are tracked explicitly. The sample complexity of the proposed algorithm is characterized and it is shown that after T iterations, the average optimality gap and constraint violation are at most $\mathcal{O}\left(T^{-1/12}\right)$. Alternative choices of the parameters also yield zero constraint violation but a larger optimality gap of $\mathcal{O}\left(T^{-1/12}\right)$ or an $\mathcal{O}(1)$ constraint violation and the best possible optimality gap of $\mathcal{O}\left(T^{-1/4}\right)$. The proposed algorithm is tested over the different queuing design examples and it is shown that the optimality gap and the constraint violation decay as expected.

APPENDIX A

PROOF OF LEMMA 1

The overall proof is split into four lemmas. For the sake of brevity, let us define the function $p(\mathbf{z}) := \ell\left(\mathbf{q}(\mathbf{z})\right)$. The following preliminary lemma establishes the necessary properties of p.

Lemma 3. The function $p(\mathbf{z})$ has Lipschitz continuous and bounded gradients,

$$\left\|\nabla p(\mathbf{z})\right\|^2 \leq JC_q C_\ell^2 \qquad \left\|\nabla p(\mathbf{z}) - \nabla p(\bar{\mathbf{z}})\right\| \leq \left(\sqrt{J}L_q C_\ell + C_q\right)\left\|\mathbf{z} - \bar{\mathbf{z}}\right\| \quad \forall \mathbf{z}, \bar{\mathbf{z}} \in \mathbb{R}^d.$$

Proof: From the chain rule, we have that $\nabla p(\mathbf{z}) = \nabla \mathbf{q}(\mathbf{z}) \nabla \ell(\mathbf{q}(\mathbf{z}))$. Using triangle inequality, Assumption (**A5**), and the definition of ℓ, we obtain $\left\|\nabla p(\mathbf{z})\right\|^2 \leq JC_q C_\ell^2$ for all $\mathbf{z}$. In a similar vein, it can be seen that p is smooth since both ℓ and $\mathbf{q}$ are smooth functions. In particular, we have that

$$\begin{aligned}\nabla p(\mathbf{z}) - \nabla p(\bar{\mathbf{z}}) &= \nabla \mathbf{q}(\mathbf{z}) \nabla \ell\left(\mathbf{q}(\mathbf{z})\right) - \nabla \mathbf{q}(\bar{\mathbf{z}}) \nabla \ell\left(\mathbf{q}(\bar{\mathbf{z}})\right) \\ &= \nabla \mathbf{q}(\mathbf{z})\left(\nabla \ell\left(\mathbf{q}(\mathbf{z})\right) - \nabla \ell\left(\mathbf{q}(\bar{\mathbf{z}})\right)\right) + \left(\nabla \mathbf{q}(\mathbf{z}) - \nabla \mathbf{q}(\bar{\mathbf{z}})\right) \nabla \ell\left(\mathbf{q}(\bar{\mathbf{z}})\right).\end{aligned}$$

Taking norm, and applying the Cauchy-Schwartz inequality, we obtain

$$\begin{aligned}\left\|\nabla p(\mathbf{z}) - \nabla p(\bar{\mathbf{z}})\right\| &\leq \left\|\nabla \mathbf{q}(\mathbf{z})\right\| \left\|\nabla \ell\left(\mathbf{q}(\mathbf{z})\right) - \nabla \ell\left(\mathbf{q}(\bar{\mathbf{z}})\right)\right\| + \left\|\nabla \mathbf{q}(\mathbf{z}) - \nabla \mathbf{q}(\bar{\mathbf{z}})\right\| \left\|\nabla \ell\left(\mathbf{q}(\bar{\mathbf{z}})\right)\right\| \\ &\leq \sqrt{C_q}\left\|\mathbf{q}(\mathbf{z}) - \mathbf{q}(\bar{\mathbf{z}})\right\| + L_q \sqrt{J} C_\ell \left\|\mathbf{z} - \bar{\mathbf{z}}\right\| \leq \left(C_q + \sqrt{J} L_q C_\ell\right)\left\|\mathbf{z} - \bar{\mathbf{z}}\right\|\end{aligned}$$

where we have used (**A5**) and the implication that $\mathbf{q}$ is Lipschitz continuous with parameter $\sqrt{C_q}$. ■

Building upon Lemma 3, we now characterize the consecutive iterate difference and show that it is $\mathcal{O}(\alpha_t^2 + \delta_t^2)$.

Lemma 4. The updates in (21) yields

$$\mathbb{E}\left[\left\|\mathbf{x}_{t+1} - \mathbf{x}_t\right\|^2\right] \leq 2\alpha_t^2 C_f C_g + 2\delta_t^2 J C_\ell^2 C_q C_h. \tag{32}$$

for all $t \geq 1$.

Proof: From Algorithm 1, since $\mathbf{x}_t \in \mathcal{X}$, the non-expansiveness property of the projection operation implies that

$$\|\mathbf{x}_{t+1}-\mathbf{x}_t\|^2 = \left\|\Pi_{\mathcal{X}}\left\{\mathbf{x}_t - \alpha_t\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\mathbf{y}_{t+1}) - \delta_t\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\mathbf{z}_{t+1})\right\} - \mathbf{x}_t\right\|^2 \tag{33}$$

$$\leq \left\|\alpha_t\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\mathbf{y}_{t+1}) + \delta_t\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\mathbf{z}_{t+1})\right\|^2 \tag{34}$$

$$\leq 2\alpha_t^2\left\|\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\mathbf{y}_{t+1})\right\|^2 + 2\delta_t^2\left\|\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\mathbf{z}_{t+1})\right\|^2 \tag{35}$$

$$\leq 2\alpha_t^2\left\|\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\right\|^2\|\nabla f(\mathbf{y}_{t+1})\|^2 + 2\delta_t^2\left\|\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\right\|^2\|\nabla p(\mathbf{z}_{t+1})\|^2 \tag{36}$$

$$\leq 2\alpha_t^2 C_f\left\|\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\right\|^2 + 2J\delta_t^2 C_\ell^2 C_q\left\|\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\right\|^2, \tag{37}$$

where (35) and (36) follow from the Peter-Paul and Cauchy-Schwartz inequalities, respectively. The last inequality follows from Lemma 3 and (**A5**). Taking expectations and using (**A4**), we obtain

$$\mathbb{E}\left[\|\mathbf{x}_{t+1}-\mathbf{x}_t\|^2\right] \leq 2\alpha_t^2 C_f\mathbb{E}\left[\left\|\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\right\|^2\right] + 2\delta_t^2 JC_\ell^2 C_q\mathbb{E}\left[\left\|\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\right\|^2\right] \tag{38}$$

$$\leq 2\alpha_t^2 C_f C_g + 2\delta_t^2 JC_\ell^2 C_q C_h \tag{39}$$

which is the required result. ■

Next, we establish a key result that legitimizes the tracking steps in (20). To this end, define functions $\bar{\mathbf{g}}(\mathbf{x}) := \mathbb{E}_{\boldsymbol{\zeta}}[\mathbf{g}(\mathbf{x},\boldsymbol{\zeta})]$ and $\bar{\mathbf{h}}(\mathbf{x}) := \mathbb{E}_{\boldsymbol{\zeta}}[\mathbf{h}(\mathbf{x},\boldsymbol{\zeta})]$ for all $\mathbf{x}\in\mathcal{X}$. Also, let $\mathcal{F}_t$ be the sigma algebra generated by the random variables $\{\boldsymbol{\zeta}_1,\ldots,\boldsymbol{\zeta}_{t-1}\}$ so that $\mathbb{E}[\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\mid\mathcal{F}_t] = \bar{\mathbf{g}}(\mathbf{x}_t)$ from (**A3**).

Lemma 5. For the choice $\alpha_t\leq\delta_t$ and $\delta_t\leq\beta_t$, the updates (20) in Algorithm 1 yield

$$\mathbb{E}\left[\|\mathbf{y}_{t+1}-\bar{\mathbf{g}}(\mathbf{x}_t)\|^2\right] \leq (1-\beta_t)\mathbb{E}\left[\|\mathbf{y}_t-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\|^2\right] + 2V_g\beta_t^2 + 2\frac{\alpha_t^2}{\beta_t}C_fC_g^2 + 2\frac{\delta_t^2}{\beta_t}JC_\ell^2C_gC_qC_h \leq D_y \tag{40}$$

$$\mathbb{E}\left[\left\|\mathbf{z}_{t+1}-\bar{\mathbf{h}}(\mathbf{x}_t)\right\|^2\right] \leq (1-\beta_t)\mathbb{E}\left[\left\|\mathbf{z}_t-\bar{\mathbf{h}}(\mathbf{x}_{t-1})\right\|^2\right] + 2V_h\beta_t^2 + 2\frac{\alpha_t^2}{\beta_t}C_fC_gC_h + 2\frac{\delta_t^2}{\beta_t}JC_\ell^2C_qC_h^2 \leq D_z \tag{41}$$

where D_y and D_z are initialization dependent constants.

Proof: We provide the proof of (40) while observing that (41) follows along similar lines. Let $\mathbf{e}_t := (1-\beta_t)\left(\bar{\mathbf{g}}(\mathbf{x}_t)-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\right)$ be the difference between evaluations of $\bar{\mathbf{g}}$ at consecutive iterates, scaled by $(1-\beta_t)$. From the Peter-Paul inequality, we have that

$$\|\mathbf{y}_{t+1}-\bar{\mathbf{g}}(\mathbf{x}_t)\|^2 \leq (1+\beta_t)\,\|\mathbf{y}_{t+1}-\bar{\mathbf{g}}(\mathbf{x}_t)+\mathbf{e}_t\|^2+(1+\frac{1}{\beta_t})\|\mathbf{e}_t\|^2. \tag{42}$$

Since $\bar{\mathbf{g}}$ is Lipschitz (cf. (**A4**)), the second term of (42) can be bounded as

$$\|\mathbf{e}_t\| \leq (1-\beta_t)\sqrt{C_g}\,\|\mathbf{x}_t-\mathbf{x}_{t-1}\|\,. \tag{43}$$

For the first term of (42), we have that

$$\mathbf{y}_{t+1}-\bar{\mathbf{g}}(\mathbf{x}_t)+\mathbf{e}_t=(1-\beta_t)\left(\mathbf{y}_t-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\right)+\beta_t\left(\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)-\bar{\mathbf{g}}(\mathbf{x}_t)\right). \tag{44}$$

Taking squared norm, applying conditional expectation given $\mathcal{F}_t$, and using the definition of $\bar{\mathbf{g}}$, we obtain

$$\begin{aligned}\mathbb{E}\left[\|\mathbf{y}_{t+1}-\bar{\mathbf{g}}(\mathbf{x}_t)+\mathbf{e}_t\|^2\mid\mathcal{F}_t\right]&=\mathbb{E}\left[\|(1-\beta_t)\left(\mathbf{y}_t-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\right)+\beta_t\left(\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)-\bar{\mathbf{g}}(\mathbf{x}_t)\right)\|^2\mid\mathcal{F}_t\right]\\&=(1-\beta_t)^2\|\mathbf{y}_t-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\|^2+\beta_t^2\mathbb{E}\left[\|\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)-\bar{\mathbf{g}}(\mathbf{x}_t)\|^2|\mathcal{F}_t\right]\\&\quad+2(1-\beta_t)\beta_t\left(\mathbf{y}_t-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\right)^T\mathbb{E}\left[\left(\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)-\bar{\mathbf{g}}(\mathbf{x}_t)\right)|\mathcal{F}_t\right]\\&\leq(1-\beta_t)^2\|\mathbf{y}_t-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\|^2+\beta_t^2V_g.\end{aligned} \tag{45}$$

where the last inequality follows from (**A4**). Taking full expectation in (45) and plugging into (42), we obtain the required result as

$$\mathbb{E}\left[\|\mathbf{y}_{t+1}-\bar{\mathbf{g}}(\mathbf{x}_t)\|^2\right] \tag{46}$$

$$\leq(1+\beta_t)(1-\beta_t)^2\mathbb{E}\left[\|\mathbf{y}_t-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\|^2\right]+(1+\beta_t)\beta_t^2V_g+\frac{1-\beta_t^2}{\beta_t}C_g\mathbb{E}\left[\|\mathbf{x}_t-\mathbf{x}_{t-1}\|^2\right] \tag{47}$$

$$\leq(1-\beta_t)\mathbb{E}\left[\|\mathbf{y}_t-\bar{\mathbf{g}}(\mathbf{x}_{t-1})\|^2\right]+2V_g\beta_t^2+2\frac{\alpha_t^2}{\beta_t}C_fC_g^2+2\frac{\delta_t^2}{\beta_t}JC_\ell^2C_gC_qC_h. \tag{48}$$

where we have also used the fact that $\beta_t<1$. To establish the upper bound, take $D_y:=\mathbb{E}\left[\|\mathbf{y}_2-\bar{\mathbf{g}}(\mathbf{x}_1)\|^2\right]+2V_g+2C_g\left(C_fC_g+JC_\ell^2C_qC_h\right)$ and observe that $\mathbb{E}\left[\|\mathbf{y}_3-\bar{\mathbf{g}}(\mathbf{x}_2)\|^2\right]\leq(1-\beta_t)D_y+2V_g\beta_t^2+2\frac{\alpha_t^2}{\beta_t}C_fC_g^2+2\frac{\delta_t^2}{\beta_t}JC_\ell^2C_gC_qC_h\leq D_y$ since $\alpha_t/\beta_t\leq1$ and $\delta_t/\beta_t\leq1$. The argument can inductively be applied to all values of t.

The derivation of (41) follows along similar lines. For the upper bound, we set $D_z=\mathbb{E}\left[\left\|\mathbf{z}_2-\bar{\mathbf{h}}(\mathbf{x}_1)\right\|^2\right]+2V_h+2C_h\left(C_fC_g+JC_\ell^2C_qC_h\right)$ and proceed inductively as before.

■

Next, we analyze the distance from the optimal set $\|\mathbf{x}_t-\mathbf{x}^\gamma\|^2$, and establish its relation with the optimality gap and the constraint violation.

Lemma 6. For convex functions F and $\mathbf{Q}$, we have the following bound.

$$\begin{align}
\mathbb{E}\left[\|\mathbf{x}_{t+1}-\mathbf{x}^{\gamma}\|^2\right] \le{}& \mathbb{E}\left[\|\mathbf{x}_{t}-\mathbf{x}^{\gamma}\|^2\right] - 2\alpha_t\mathbb{E}\left[F(\mathbf{x}_t)+\frac{\delta_t}{\alpha_t}L(\mathbf{x}_t)-F(\mathbf{x}^{\gamma})\right] + 2\alpha_t^2 C_f C_g \nonumber\\
&+ 2\delta_t^2 J C_\ell^2 C_q C_h + \beta_t\mathbb{E}\left[\|\mathbf{y}_{t+1}-\bar{\mathbf{g}}(\mathbf{x}_t)\|^2\right] + \beta_t\mathbb{E}\left[\left\|\mathbf{z}_{t+1}-\bar{\mathbf{h}}(\mathbf{x}_t)\right\|^2\right] \nonumber\\
&+ \frac{\alpha_t^2}{\beta_t}L_f D_x C_g + \frac{\delta_t^2}{\beta_t}\left(\sqrt{J}L_q C_\ell + C_q\right) D_x C_h. \tag{49}
\end{align}$$

Proof: Since $\mathbf{x}^{\gamma}\in\mathcal{X}$, the non-expansiveness of the $\Pi_{\mathcal{X}}$ operator implies that

$$\begin{align}
\|\mathbf{x}_{t+1}-\mathbf{x}^{\gamma}\|^2 &= \left\|\Pi_{\mathcal{X}}\left\{\mathbf{x}_t-\alpha_t\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\mathbf{y}_{t+1})-\delta_t\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\mathbf{z}_{t+1})\right\}-\mathbf{x}^{\gamma}\right\|^2 \tag{50}\\
&\le \left\|\mathbf{x}_t-\mathbf{x}^{\gamma}-\left(\alpha_t\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\mathbf{y}_{t+1})+\delta_t\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\mathbf{z}_{t+1})\right)\right\|^2 \tag{51}\\
&= \|\mathbf{x}_t-\mathbf{x}^{\gamma}\|^2 + \left\|\alpha_t\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\mathbf{y}_{t+1})+\delta_t\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\mathbf{z}_{t+1})\right\|^2 \tag{52}\\
&\quad - 2\left(\mathbf{x}_t-\mathbf{x}^{\gamma}\right)^T\left(\alpha_t\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\mathbf{y}_{t+1})+\delta_t\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\mathbf{z}_{t+1})\right) \tag{53}\\
&= \|\mathbf{x}_t-\mathbf{x}^{\gamma}\|^2 + \|\mathbf{x}_{t+1}-\mathbf{x}_t\|^2 + u_t + v_t \tag{54}\\
&\quad - 2\left(\mathbf{x}_t-\mathbf{x}^{\gamma}\right)^T\left(\alpha_t\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\bar{\mathbf{g}}(\mathbf{x}_t))+\delta_t\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\bar{\mathbf{h}}(\mathbf{x}_t))\right) \tag{55}
\end{align}$$

where $u_t = 2\alpha_t\left(\mathbf{x}_t-\mathbf{x}^{\gamma}\right)^T\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\left(\nabla f(\bar{\mathbf{g}}(\mathbf{x}_t))-\nabla f(\mathbf{y}_{t+1})\right)$, and $v_t = 2\delta_t\left(\mathbf{x}_t-\mathbf{x}^{\gamma}\right)^T\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\left(\nabla p(\bar{\mathbf{h}}(\mathbf{x}_t))-\nabla p(\mathbf{z}_{t+1})\right)$. Taking conditional expectation given $\mathcal{F}_t$ in (55), we obtain

$$\begin{align}
\mathbb{E}\left[\|\mathbf{x}_{t+1}-\mathbf{x}^{\gamma}\|^2|\mathcal{F}_t\right] \le{}& \|\mathbf{x}_t-\mathbf{x}^{\gamma}\|^2 + \mathbb{E}\left[\|\mathbf{x}_{t+1}-\mathbf{x}_t\|^2|\mathcal{F}_t\right] + \mathbb{E}[u_t|\mathcal{F}_t] + \mathbb{E}[v_t|\mathcal{F}_t] \nonumber\\
&- 2\left(\mathbf{x}_t-\mathbf{x}^{\gamma}\right)^T\mathbb{E}\left[\alpha_t\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla f(\bar{\mathbf{g}}(\mathbf{x}_t))+\delta_t\widetilde{\nabla}\mathbf{h}(\mathbf{x}_t,\boldsymbol{\zeta}_t)\nabla p(\bar{\mathbf{h}}(\mathbf{x}_t))|\mathcal{F}_t\right] \nonumber\\
={}& \|\mathbf{x}_t-\mathbf{x}^{\gamma}\|^2 + 2\alpha_t^2 C_f C_g + 2\delta_t^2 J C_\ell^2 C_q C_h + \mathbb{E}[u_t|\mathcal{F}_t] + \mathbb{E}[v_t|\mathcal{F}_t] \nonumber\\
&- 2\left(\mathbf{x}_t-\mathbf{x}^{\gamma}\right)^T\left(\alpha_t\nabla F(\mathbf{x}_t)+\delta_t\nabla L(\mathbf{x}_t)\right) \tag{56}\\
\le{}& \|\mathbf{x}_t-\mathbf{x}^{\gamma}\|^2 + 2\alpha_t^2 C_f C_g + 2\delta_t^2 J C_\ell^2 C_q C_h + \mathbb{E}[u_t|\mathcal{F}_t] + \mathbb{E}[v_t|\mathcal{F}_t] \nonumber\\
&- 2\alpha_t\left(F(\mathbf{x}_t)-F(\mathbf{x}^{\gamma})\right) - 2\delta_t(L(\mathbf{x}_t)-L(\mathbf{x}^{\gamma})). \tag{57}
\end{align}$$

where (57) follows from Lemma 4 and the convexity of functions F and L. Next, noting that $L(\mathbf{x}^{\gamma})=0$ and taking full expectation in (57), we obtain

$$\begin{align}
\mathbb{E}\left[\|\mathbf{x}_{t+1}-\mathbf{x}^{\gamma}\|^2\right] \le{}& \mathbb{E}\left[\|\mathbf{x}_t-\mathbf{x}^{\gamma}\|^2\right] - 2\alpha_t\left(F(\mathbf{x}_t)-F(\mathbf{x}^{\gamma})\right) - 2\delta_t L(\mathbf{x}_t) + 2\alpha_t^2 C_f C_g \nonumber\\
&+ 2\delta_t^2 J C_\ell^2 C_q C_h + \mathbb{E}\left[u_t\right] + \mathbb{E}\left[v_t\right]. \tag{58}
\end{align}$$

In order to bound u_t, we utilize the Cauchy-Schwartz and Peter-Paul inequalities as

$$
\begin{aligned}
u_t &= 2\alpha_t \left(\mathbf{x}_t - \mathbf{x}^\gamma\right)^T \widetilde{\nabla}\mathbf{g}(\mathbf{x}_t, \boldsymbol{\zeta}_t) \left(\nabla f(\bar{\mathbf{g}}(\mathbf{x}_t)) - \nabla f(\mathbf{y}_{t+1})\right) && (59)\\
&\le 2\alpha_t L_f \left\|\mathbf{x}_t - \mathbf{x}\right\| \left\|\mathbf{y}_{t+1} - \bar{\mathbf{g}}(\mathbf{x}_t)\right\| \left\|\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t, \boldsymbol{\zeta}_t)\right\| && (60)\\
&\le \beta_t \left\|\mathbf{y}_{t+1} - \bar{\mathbf{g}}(\mathbf{x}_t)\right\|^2 + \frac{\alpha_t^2}{\beta_t} L_f \left\|\mathbf{x}_t - \mathbf{x}\right\|^2 \left\|\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t, \boldsymbol{\zeta}_t)\right\|^2 && (61)\\
&\le \beta_t \left\|\mathbf{y}_{t+1} - \bar{\mathbf{g}}(\mathbf{x}_t)\right\|^2 + \frac{\alpha_t^2}{\beta_t} L_f D_x \left\|\widetilde{\nabla}\mathbf{g}(\mathbf{x}_t, \boldsymbol{\zeta}_t)\right\|^2. && (62)
\end{aligned}
$$

where (62) follows from (**A2**). Taking expectation on both sides of (62), we obtain $Eu_t \le \beta_t \mathbb{E}\left[\left\|\mathbf{y}_{t+1} - \bar{\mathbf{g}}(\mathbf{x}_t)\right\|^2\right] + \frac{\alpha_t^2}{\beta_t} L_f D_x C_g$, and we can also obtain $\mathbb{E}[v_t] \le \beta_t \mathbb{E}\left[\left\|\mathbf{z}_{t+1} - \bar{\mathbf{h}}(\mathbf{x}_t)\right\|^2\right] + \frac{\delta_t^2}{\beta_t}\left(\sqrt{J} L_q C_\ell + C_q\right) D_x C_h$. The required result follows from substituting into (58). ■

Having established the required recursive relationships, we proceed with applying telescopic sums in order to obtain the required bounds. Let

$$
I_t := \mathbb{E}\left[\left\|\mathbf{x}_t - \mathbf{x}^\gamma\right\|^2 + \left\|\mathbf{y}_t - \bar{\mathbf{g}}(\mathbf{x}_{t-1})\right\|^2 + \left\|\mathbf{z}_t - \bar{\mathbf{h}}(\mathbf{x}_{t-1})\right\|^2\right] \tag{63}
$$

and observe that $I_t \le D_x + D_y + D_z =: D_1$. Multiplying (40) and (41) with $(1 + \beta_t)$ and substituting in (49), we obtain

$$
\begin{aligned}
I_{t+1} \le{}& I_t - 2\alpha_t \mathbb{E}\left[F(\mathbf{x}_t) + \frac{\delta_t}{\alpha_t} L(\mathbf{x}_t) - F(\mathbf{x}^\gamma)\right] + 2\delta_t^2\left(C_f C_g + J C_\ell^2 C_q C_h\right) + 4\left(V_g + V_h\right)\beta_t^2\\
&+ \frac{\delta_t^2}{\beta_t}\left(4\left(C_g + C_h\right) C_f C_g + 4\left(C_g + C_h\right) J C_\ell^2 C_q C_h + \left(L_f C_g + \sqrt{J} L_q C_h C_\ell + C_h C_q\right) D_x\right)\\
\le{}& I_t - 2\alpha_t \mathbb{E}\left[F(\mathbf{x}_t) + \frac{\delta_t}{\alpha_t} L(\mathbf{x}_t) - F(\mathbf{x}^\gamma)\right] + D_2\left(\delta_t^2 + \beta_t^2 + \frac{\delta_t^2}{\beta_t}\right). && (64)
\end{aligned}
$$

where D_2 is the maximum of $4\left(C_g + C_h\right) C_f C_g + 4\left(C_g + C_h\right) J C_\ell^2 C_q C_h$
$+ \left(L_f C_g + \sqrt{J} L_q C_h C_\ell + C_h C_q\right) D_x$, $2\left(C_f C_g + C_\ell^2 C_q C_h\right)$, and $4\left(V_g + V_h\right)$. Summing over

$t = T/2$ to T, noting that $\alpha_t \leq \alpha_{T/2}$, and rearranging, we obtain the required result as

$$\sum_{t=T/2}^{T} \mathbb{E}\left[F(\mathbf{x}_t) + \frac{\delta_t}{\alpha_t} L(\mathbf{x}_t) - F(\mathbf{x}^\gamma)\right] \tag{65}$$

$$\leq \frac{I_{T/2}}{\alpha_{T/2}} + \sum_{t=T/2}^{T} I_t \left(\frac{1}{\alpha_t} - \frac{1}{\alpha_{t-1}}\right) - \frac{I_{T+1}}{\alpha_{T+1}} + D_2 \left(\sum_{t=T/2}^{T} \left[\frac{\delta_t^2}{\alpha_t} + \frac{\beta_t^2}{\alpha_t} + \frac{\delta_t^2}{\alpha_t \beta_t}\right]\right) \tag{66}$$

$$\leq \frac{D_1}{\alpha_{T/2}} + D_1 \sum_{t=T/2+1}^{T} \left(\frac{1}{\alpha_t} - \frac{1}{\alpha_{t-1}}\right) + D_2 \left(\sum_{t=T/2}^{T} \left[\frac{\delta_t^2}{\alpha_t} + \frac{\beta_t^2}{\alpha_t} + \frac{\delta_t^2}{\alpha_t \beta_t}\right]\right) \tag{67}$$

$$\leq \frac{D_1}{\alpha_T} + D_2 \left(\sum_{t=T/2}^{T} \left[\frac{\delta_t^2}{\alpha_t} + \frac{\beta_t^2}{\alpha_t} + \frac{\delta_t^2}{\alpha_t \beta_t}\right]\right) \tag{68}$$

where we have dropped the negative terms from the right-hand side. which is the required result.

APPENDIX B
PROOF OF LEMMA 2

Let us associate dual variable $\boldsymbol{\lambda}$ with the constraint in (22) and let $\boldsymbol{\lambda}^\gamma \geq 0$ be the dual optimal solution. Since $\gamma < \sigma/2 < \sigma$, there exists a strictly feasible $\widetilde{\mathbf{x}}$ such that $Q_j(\widetilde{\mathbf{x}}) + \gamma < Q_j(\widetilde{\mathbf{x}}) + \sigma \leq 0$, and consequently strong duality holds for (22). Therefore, we have that

$$F(\mathbf{x}^\gamma) = \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}) + \langle \boldsymbol{\lambda}^\gamma, \mathbf{Q}(\mathbf{x}) + \gamma \mathbf{1} \rangle \leq F(\mathbf{x}^\star) + \langle \boldsymbol{\lambda}^\gamma, \mathbf{Q}(\mathbf{x}^\star) + \gamma \mathbf{1} \rangle \leq F^\star + \gamma \mathbf{1}^T \boldsymbol{\lambda}^\gamma, \tag{69}$$

where we have used the fact that $\mathbf{Q}(\mathbf{x}^\star) \leq 0$. Let $\widetilde{\mathbf{x}}$ be a strictly feasible point specified in (**A2**). Then from (69), it also follows that

$$F(\mathbf{x}^\gamma) \leq F(\widetilde{\mathbf{x}}) + \langle \boldsymbol{\lambda}^\gamma, \mathbf{Q}(\widetilde{\mathbf{x}}) + \gamma \mathbf{1} \rangle \leq F(\widetilde{\mathbf{x}}) + (\gamma - \sigma) \mathbf{1}^T \boldsymbol{\lambda}^\gamma \tag{70}$$

or equivalently

$$\mathbf{1}^T \boldsymbol{\lambda}^\gamma \leq \frac{F(\widetilde{\mathbf{x}}) - F(\mathbf{x}^\gamma)}{\sigma - \gamma} \leq \frac{2\sqrt{C_g C_f} D_x}{\sigma} \tag{71}$$

where the last inequality follows since F is Lipschitz continuous with parameter $\sqrt{C_g C_g}$ and since $\mathcal{X}$ is compact. The required bound follows from substituting (71) into (69).

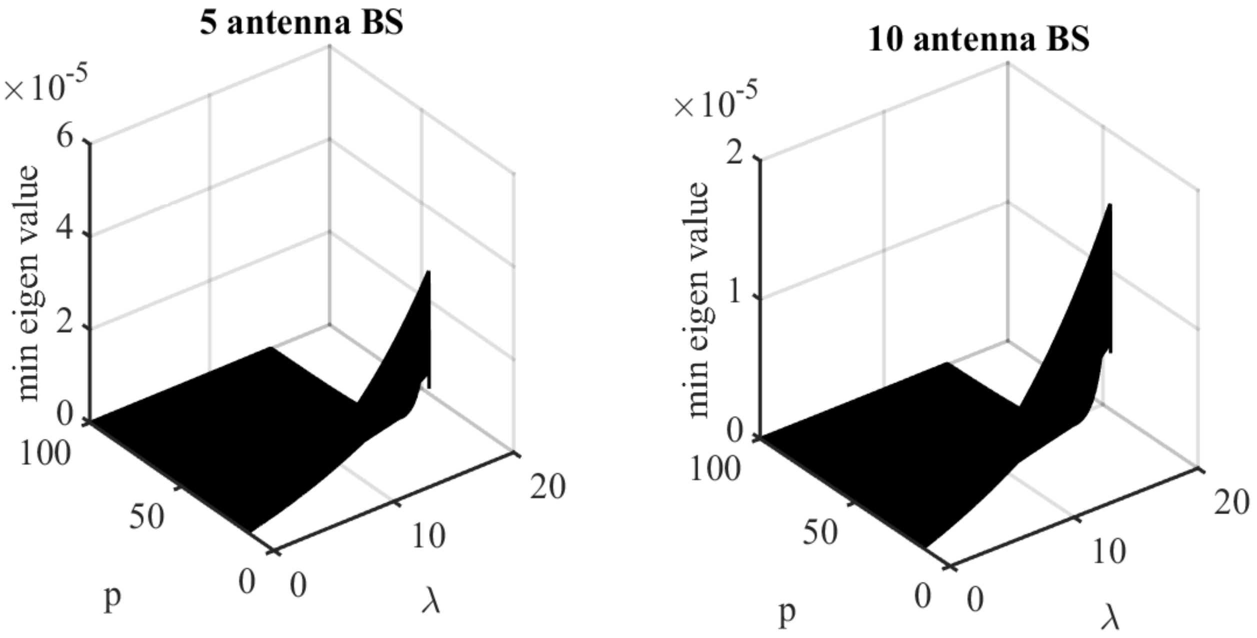

5 antenna BS

10 antenna BS

Fig. 6: minimum eigenvalue of the Hessian of (72) with K = 5, and K = 10 antennas at the BS.

APPENDIX C

NUMERICAL VERIFICATION OF CONVEXITY OF THE FORMULATION IN EXAMPLE 2

Lemma 7. The following function is convex with respect to λ, and p for the box constraints $0.1 \leq \lambda \leq 15$, and $14 \leq p \leq 100$

$$U(\lambda, p) = \frac{\lambda \mathbb{E}\left[\left(\frac{1}{B \log(1+p\zeta)}\right)^2\right]}{2\left(1 - \lambda \mathbb{E}\left[\frac{1}{B \log(1+p\zeta)}\right]\right)} - 0.1 \log(\lambda), \tag{72}$$

where ζ is a Chi-squared distributed random variable with $2K$ degrees of freedom.

Proof: Since the expression in (72) is of non linear expectation function, it is more difficult to get an analytical proof. Instead, we try to derive the Hessian of (72) numerically, and show that the minimum eigenvalue of Hessian is nonnegative. To do that, we take help of two MATLAB functions namely ‘numerical gradient’, and ‘numerical integration’. The results in Fig. 6 confirm that the Hessian of the expression in (72) is positive semidefinite matrix.

■

REFERENCES

[1] R. Buyya, S. K. Garg, and R. N. Calheiros, “SLA-oriented resource provisioning for cloud computing: challenges, architecture, and solutions,” in *Proc. of the IEEE SC2*, Dec. 2011.

[2] Y. Tan, Y. Lu, and C. H. Xia, "Provisioning for large scale cloud computing services," *ACM Sigmetrics Perf. Eval. Review*, vol. 40, no. 1, pp. 407–408, 2012.

[3] Y. Tan and C. H. Xia, "An adaptive learning approach for efficient resource provisioning in cloud services," *ACM Sigmetrics Perf. Eval. Review*, vol. 42, no. 4, pp. 3–11, 2015.

[4] S. Stidham Jr, *Optimal Design of Queueing Systems*. Chapman and Hall/CRC, 2009.

[5] D. P. Bertsekas, R. G. Gallager, and P. Humblet, *Data Networks*. Prentice-Hall International New Jersey, 1992, vol. 2.

[6] S. Lee, M. Han, and D. Hong, "Average SNR and ergodic capacity analysis for opportunistic DF-relaying with outage over rayleigh fading channels," *IEEE Trans. Wireless Commun.*, vol. 8, no. 6, pp. 2807–2812, 2009.

[7] C.-S. Chang and J. A. Thomas, "Effective bandwidth in high-speed digital networks," *IEEE J. Sel. Areas Commun.*, vol. 13, no. 6, pp. 1091–1100, 1995.

[8] G. L. Choudhury, D. M. Lucantoni, and W. Whitt, "Squeezing the most out of atm," *IEEE Trans. Commun.*, vol. 44, no. 2, pp. 203–217, 1996.

[9] D. Wu and R. Negi, "Effective capacity: a wireless link model for support of quality of service," *IEEE Trans. Wireless Commun.*, vol. 2, no. 4, pp. 630–643, 2003.

[10] B. Soret, M. C. Aguayo-Torres, and J. T. Entrambasaguas, "Capacity with explicit delay guarantees for generic sources over correlated rayleigh channel," *IEEE Trans. Wireless Commun.*, vol. 9, no. 6, pp. 1901–1911, 2010.

[11] J. P. Champati, H. Al-Zubaidy, and J. Gross, "Statistical guarantee optimization for age of information for the D/G/1 queue," in *Proc. of the IEEE Infocom Workshops*, Apr. 2018, pp. 130–135.

[12] A. Benveniste, M. Métivier, and P. Priouret, *Adaptive Algorithms and Stochastic Approximations*. Springer Science & Business Media, 2012, vol. 22.

[13] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Prentice hall Englewood Cliffs, NJ, 1989, vol. 23.

[14] V. S. Borkar, *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge Univ. Press, 2008.

[15] M. Wang, E. X. Fang, and H. Liu, "Stochastic compositional gradient descent: algorithms for minimizing compositions of expected-value functions," *Mathematical Programming*, vol. 161, no. 1-2, pp. 419–449, 2017.

[16] X. Lian, M. Wang, and J. Liu, "Finite-sum composition optimization via variance reduced gradient descent," in *Proc. of the Artificial Intel. and Stats.*, Apr. 2017, pp. 1159–1167.

[17] M. Wang, J. Liu, and E. Fang, "Accelerating stochastic composition optimization," in *Proc. of NIPS*, Dec. 2016, pp. 1714–1722.

[18] M. Wang and J. Liu, "A stochastic compositional gradient method using markov samples," in *Proc. of the IEEE WSC*, Dec. 2016, pp. 702–713.

[19] S. Yang, M. Wang, and E. X. Fang, "Multilevel stochastic gradient methods for nested composition optimization," *SIAM J. on Optim.*, vol. 29, no. 1, pp. 616–659, 2019.

[20] A. Koppel, A. S. Bedi, and K. Rajawat, "Controlling the the bias-variance tradeoff via coherent risk for robust learning with kernels," in *Proc. of the ACC*, May 2019.

[21] Y. M. Ermoliev and R.-B. Wets, *Numerical Techniques for Stochastic Optimization*. Springer-Verlag, 1988.

[22] X. Wang and G. B. Giannakis, "Resource allocation for wireless multiuser OFDM networks," *IEEE Trans. Inf. Theory*, vol. 57, no. 7, pp. 4359–4372, 2011.

[23] A. Ribeiro, "Ergodic stochastic optimization algorithms for wireless communication and networking," *IEEE Trans. Signal Process.*, vol. 58, no. 12, pp. 6369–6386, 2010.

[24] A. S. Bedi and K. Rajawat, "Asynchronous incremental stochastic dual descent algorithm for network resource allocation,"

IEEE Trans. Signal Process., vol. 66, no. 9, pp. 2229–2244, 2018.

[25] T. Chen, A. Mokhtari, X. Wang, A. Ribeiro, and G. B. Giannakis, “Stochastic averaging for constrained optimization with application to online resource allocation,” *IEEE Trans. Signal Process.*, vol. 65, no. 12, pp. 3078–3093, 2017.

[26] A. S. Bedi, A. Koppel, and K. Rajawat, “Asynchronous saddle point algorithm for stochastic optimization in heterogeneous networks,” *IEEE Trans. Signal Process.*, vol. 67, no. 7, pp. 1742–1757, 2019.

[27] X. Wang, X. Chen, T. Chen, L. Huang, and G. B. Giannakis, “Two-scale stochastic control for integrated multipoint communication systems with renewables,” *IEEE Trans. Smart Grid*, vol. 9, no. 3, pp. 1822–1834, 2016.

[28] T. Chen, S. Barbarossa, X. Wang, G. B. Giannakis, and Z.-L. Zhang, “Learning and management for internet of things: Accounting for adaptivity and scalability,” in *Proc. of the IEEE*, 2019.

[29] X. Chen, W. Ni, I. B. Collings, X. Wang, and S. Xu, “Automated function placement and online optimization of network functions virtualization,” *IEEE Trans. Commun.*, vol. 67, no. 2, pp. 1225–1237, 2018.

[30] L. Tassiulas, “Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks,” *IEEE Trans. Autom. Control*, vol. 31, no. 12, 1992.

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

[32] P. Kumar and S. P. Meyn, “Stability of queueing networks and scheduling policies,” *IEEE Trans. Autom. Control*, vol. 40, no. 2, pp. 251–260, 1995.

[33] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand, “Achieving 100% throughput in an input-queued switch,” *IEEE Trans. Commun.*, vol. 47, no. 8, pp. 1260–1267, 1999.

[34] M. J. Neely and R. Urgaonkar, “Optimal backpressure routing for wireless networks with multi-receiver diversity,” *Ad Hoc Networks*, vol. 7, no. 5, pp. 862–881, 2009.

[35] A. Gupta, X. Lin, and R. Srikant, “Low-complexity distributed scheduling algorithms for wireless networks,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 6, pp. 1846–1859, 2009.

[36] L. Ying, S. Shakkottai, A. Reddy, and S. Liu, “On combining shortest-path and back-pressure routing over multihop wireless networks,” *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 841–854, 2011.

[37] S. Ghadimi and G. Lan, “Stochastic first-and zeroth-order methods for nonconvex stochastic programming,” *SIAM J. on Optim.*, vol. 23, no. 4, pp. 2341–2368, 2013.

[38] J. Chen and R. Luss, “Stochastic gradient descent with biased but consistent gradient estimators,” *arXiv preprint arXiv:1807.11880*, 2018.

[39] A. Nedić and A. Ozdaglar, “Subgradient methods for saddle-point problems,” *J. of Opt. Theory and Applications.*, vol. 142, no. 1, pp. 205–228, 2009.

[40] A. Koppel, F. Y. Jakubiec, and A. Ribeiro, “A saddle point algorithm for networked online convex optimization,” *IEEE Trans. Signal Process.*, vol. 63, no. 19, pp. 5149–5164, 2015.

[41] A. S. Bedi, A. Koppel, and K. Rajawat, “Beyond consensus and synchrony in decentralized online optimization using saddle point method,” in *Proc. of the IEEE Asilomar*, 2017, pp. 293–297.

[42] M. Chiang, “Geometric programming for communication systems,” *Found. and Trends in Comm. and Inf. Theory*, vol. 2, no. 1–2, pp. 1–154, 2005.

[43] N. Ahmed and R. G. Baraniuk, “Throughput maximization for ARQ-like systems in fading channels with coding and queuing delay constraints,” in *Proc. of the IEEE Asilomar*, vol. 2, Nov. 2004, pp. 1463–1467.