

# Throughput-Optimal Scheduling via Rate Learning

Panagiotis Promponas, Víctor Valls\*, Konstantinos Nikolakakis,  
 Dionysis Kalogerias, and Leandros Tassiulas  
 Yale University, \*IBM Research Europe – Dublin

**Abstract**—We study the problem of designing scheduling policies for communication networks. This problem is often addressed with max-weight-type approaches since they are throughput-optimal. However, max-weight policies make scheduling decisions based on the network congestion, which can be sometimes unnecessarily restrictive. In this paper, we present a “schedule as you learn” (SYL) approach, where we learn an average rate, and then select schedules that generate such a rate in expectation. This approach is interesting because scheduling decisions do not depend on the size of the queue backlogs, and so it provides increased flexibility to select schedules based on other criteria or rules, such as serving high-priority queues. We illustrate the results with numerical experiments for a cross-bar switch and show that, compared to max-weight, SYL can achieve lower latency to certain flows without compromising throughput optimality.

## I. INTRODUCTION

Scheduling policies in networking have historically focused on maximizing throughput, as this was the main bottleneck in communication networks. This is, however, not the case anymore with modern communication networks (e.g., 5G), and cloud applications (e.g., collaborative LaTeX editors and online video games) often require performance metrics besides throughput to achieve a good quality of experience. For example, a latency above 250 ms is considered not suitable for online gaming [1].

Let us illustrate, with a toy example, the problem of a scheduling policy that focuses only on maximizing throughput. Consider the network illustrated in Fig. 1; a server that receives packets from two different flows, and stores them in separate queues. Flow 1 represents a flow that is insensitive to latency (e.g., FTP traffic), and flow 2 is a flow with low-latency requirements (e.g., an online video game). In this example, the server can transmit at most one packet at a time, so it has to decide which of the queues to serve. In this case, the max-weight policy [2] will transmit a packet from the queue with the largest backlog. Such a strategy is throughput-optimal in the sense that all the packets that get into the queues will eventually get out in a finite amount of time.<sup>1</sup> However, the max-weight policy does not provide low latency to flow 2 (see Fig. 2). Now, consider a policy that transmits a packet from queue 2 if this is not empty, and otherwise it serves a packet from queue 1. This *alternative policy* is also throughput optimal, but this provides, in addition, low latency to the second flow

<sup>1</sup>Or more formally, the waiting times of the packets in the queues does not grow unbounded.




Fig. 1. Server with two queues. The server can only serve one packet from the queues at a time. Flow 1 and 2 are directed to nodes 1 and 2, respectively.




Fig. 2. Probability distribution of the packets waiting times for the network in Fig. 1. The waiting times are normalized to the frequency in which the server selects schedules. The packet arrivals of flow 1 and flow 2 are Bernoulli with mean  $(1 - \epsilon)0.8$  and  $(1 - \epsilon)0.2$  respectively, where  $\epsilon = 5 \cdot 10^{-5}$ .

(see Fig. 2). Hence, it is preferable over max-weight for this scenario.

Designing throughput optimal policies is highly non-trivial, especially when the network has arbitrary topology, the traffic arrival statistics are not known, and the network connectivity varies over time. The max-weight policy [2] and variants [3], [4] have become the *de facto* approaches for maximizing throughput in networks. Their fundamentals are well studied by the networking and control communities, and besides throughput-optimality, they have the appealing property that scheduling decisions can be computed in polynomial time by finding a maximum-weighted matching in a graph [5]. In addition, well-established extensions of max-weight, such as the Lyapunov Drift-Plus-Penalty (LDPP) algorithm [4], can minimize a convex utility function of the average throughput. For example, LDPP can serve all the traffic in a wireless network while minimizing the energy used [6]. However, performance metrics, such as latency, are not a function of the average flow, and selecting schedules to be maximum-weighted matchings in a graph may not always be the best choice, as shown in the example in Fig. 1.

In this paper, we study the problem of designing throughput optimal scheduling policies for networks with arbitrary topology when the statistics of the arrival process are unknown. Unlike max-weight approaches, which seek to stabilize a system of queues directly, we see the problem as learning an “average” rate vector that guarantees queue

stability, and then “stack” a (randomized) scheduling policy on top that guarantees such an average rate in expectation. That is, *we select schedules as we learn an average rate*. Such an approach is appealing because it decouples the choice of schedule from the size of the queue backlogs at a given time, thus providing more flexibility as to *how* to select schedules. Max-weight is *a* throughput-optimal policy, but not the only one. Among all throughput-optimal policies, it is interesting to select/design one that boosts other performance metrics important for the problem of interest, e.g., low latency or reduced memory usage. To this end, the main contributions of the paper are:

- A new approach to designing throughput optimal policies (Sec. IV). The strategy consists of learning an average rate that guarantees queue stability/maximum throughput, and then selecting schedules that ensure such an average rate in expectation.
- A new algorithm (Algorithm 3) that is throughput optimal. The algorithm uses Nesterov’s dual-averaging [7] for learning the “average” rate, and a randomized policy for selecting the schedules. In our approach, scheduling decisions do not depend on the size of the queue backlogs.
- Two numerical experiments for a cross-bar switch scheduling problem that illustrate the performance of our approach. In addition, we present a variant of SYL to illustrate how we can “bias” the selection of schedules to reduce the latency of a flow while retaining throughput optimally.

Our approach is novel and provides an alternative perspective on how to tackle queue scheduling problems. Furthermore, we believe it makes the design of *new* throughput-optimal policies conceptually simpler. The learning process can be offloaded to a subroutine, which allows us to focus on designing a scheduling policy for a particular system. Such an approach can be helpful for systems with complex operational constraints, such as quantum networks with qubit memory constrains and decoherence [8].

The rest of the paper is organized as follows. Sec. II gives an overview of previous work and position our contributions. Sec. III introduces the preliminaries: the notation, a sufficient condition for queue stability, and the dual averaging algorithm. In Sec. IV, we formulate the problem and discuss how we can tackle it depending on whether the statistics of the packet arrivals are known, or not. Sec. V presents the main technical results (Algorithm 3 and Theorem 1) and discusses the limitations of our approach. Finally, Sec. VI presents the numerical experiments for a cross-bar switch, and Sec. VII concludes. All the proofs are in the Appendix.

## II. BACKGROUND AND RELATED WORK

A significant amount of research on throughput-optimal policies started as result of the paper in [2]. In brief, [2] characterized the set of rates for which throughput-optimal policies exist and showed how to design a throughput-optimal policy using a discrete-time Lyapunov control approach. The throughput-optimal policy presented in [2] is known

as max-weight, as this selects schedules that correspond to maximum-weighted matchings in the connectivity graph—the weights of edges are the sizes of the queue backlogs. This line of work has inspired several extensions of the original algorithm with applications to many networking and resource allocation problems. Some notable extension include utility functions [4], constrained schedules [9], heavy-tailed traffic [10], and online learning [11].

The appealing features of max-weight algorithms<sup>2</sup> motivated researchers to develop convex optimization algorithms for network resource allocation problems with similar features. In short, the approaches often consists of connecting Lagrange multipliers with scaled queues’ occupancy in a specific convex optimization method, and then show, via the Lagrange multipliers, that the queues are stable. See, for example, [12, Sec. 4.7], [13, Sec. V], [14, Sec. IV-B].

The approach in this paper is conceptually different from max-weight and convex optimization approaches. We show that satisfying an *average rate* vector condition is sufficient for obtaining queue stability and then aim to find/learn such a rate. We use dual-averaging (i.e., a convex optimization algorithm) to learn the rate vector, but other techniques are possible. Notably, our learning algorithm does not connect dual variables with queues, which is the distinctive characteristic of previous convex optimization works [12], [13], [14]. Furthermore, we use *diminishing* step sizes to learn the rate vector (i.e.,  $\alpha \rightarrow 0$ ), which is in contrast to the *constant* step size employed in previous approaches to ensure queue stability (see, for example, [3, Sec. 4.10] and [13, Sec. V]).

## III. PRELIMINARIES

This section introduces the notation, a sufficient condition for strong stability in queuing systems, and the dual-averaging algorithm that we will use in Sec. V to learn a rate vector that ensures queue stability.

### A. Notation

We use  $\mathbf{R}^n$  to denote the set of  $n$ -dimensional vectors, and  $\mathbf{R}_+^n$  the set of vectors with non-negative entries. All vectors are in column form, and we use  $\mathbf{1}$  to denote the all ones column vector—its dimension will be clear from the context. For two vectors  $x$  and  $y$ , we use  $\langle x, y \rangle$  to denote their inner product. We use  $x \preceq (\succeq) y$  to indicate that  $x$  is component-wise smaller (greater) than  $y$ . The function  $[x]^+ : \mathbf{R}^n \rightarrow \mathbf{R}^n$  returns a vector where its  $j$ -th component is equal to the  $j$ -th component of  $x$  if this is greater than or equal to zero, or zero otherwise. Finally, we use  $\|\cdot\|$  to denote the  $\ell_2$ -norm, and  $\nabla f(x)$  to denote a (sub)gradient of a function  $f$  at  $x$ .

### B. Strong stability of a queuing system

Consider a system of  $n$  queues that operates in slotted time. In each time slot  $k = 1, 2, \dots$ , the queues evolve according to the following recursion:

$$Q_{k+1} = [Q_k + Z_k]^+ \quad k = 1, 2, \dots \quad (1)$$

<sup>2</sup>Throughput optimality; it does not require to know the statistics of the arrival process; and schedules can be computed in polynomial time.

---

**Algorithm 1** Dual averaging algorithm
 

---

**Input:**  $f$  and  $C$  convex,  $\phi$  is  $\sigma$ -strongly convex  
**Set:**  $s_0 = 0, k = 1$   
**while** termination condition is not met **do**  
    $\alpha_k > 0$   
    $x_k = \arg \max_{u \in C} \{s_{k-1}, u\} - \phi(u)$   
    $s_k \leftarrow s_{k-1} - \alpha_k \widetilde{\nabla} f(x_k)$   
    $k \leftarrow k + 1$   
**end while**

---

where  $Q_1 \in \mathbf{R}_+^n$ , and  $Z_k \in \mathbf{R}^n$  indicates the net increment of items in each of the queues.

A queuing system is strongly stable when the expected backlog of all queues is finite [4], i.e.,

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{i=1}^k \mathbf{E}[\|Q_i\|] < \infty. \quad (2)$$

Strong stability is an important property to ensure throughput-optimality. Informally, by ensuring that the queues do not overflow, we guarantee that all the packets that get into the queues will eventually get out. A queueing system will be strongly stable depending on the properties of the sequence  $\{Z_i\}_{i=1}^k$ . The following proposition establishes a simple condition to have strongly stable queues.

**Proposition 1.** Consider the update in Eq. (1) with  $Q_1 \in \mathbf{R}_+^n$ . Suppose that

$$\mathbf{E}[Z_k] \preceq -\eta \mathbf{1} \quad \text{for all } k = 1, 2, \dots$$

with  $\eta > 0$  and  $\|Z_k\| \leq \sigma$  for some constant  $\sigma < \infty$ . Also, suppose that  $Z_k$  is independent of  $Q_k$ . Then, the queuing system is strongly stable.

The result above holds also when the condition  $\mathbf{E}[Z_k] \preceq -\eta \mathbf{1}$  is only satisfied for  $k \geq \tau$ , where  $\tau \in \{1, 2, \dots\}$  and  $Q_\tau \in \mathbf{R}_+^n$ . We will use this fact in Lemma 3 in Sec. V.

### C. Dual averaging

We present Nesterov's dual averaging algorithm [7] in generic convex optimization form as this will allow us to streamline the proofs and exposition in Sec. V.

Consider the convex optimization problem:

$$\underset{x \in C}{\text{minimize}} \quad f(x) \quad (3)$$

where  $f : \mathbf{R}^n \rightarrow \mathbf{R}$  and  $C \subseteq \mathbf{R}^n$  are convex. The objective function does not need to be differentiable. To minimize  $f$ , we generate a sequence  $\{x_i \in C\}_{i=1}^k$  to query a stochastic first-order oracle. On querying the oracle at point  $x_k$ , this returns vectors of the form

$$\widetilde{\nabla} f(x_k) = \nabla f(x_k) + \xi_k,$$

where  $\xi_k$  is a bounded random variable with  $\mathbf{E}[\xi_k] = 0$  for all  $k$ . With such an oracle, we can design an algorithm that solves (3). The following lemma is a special case of [7, Theorem 1], and we will use it to prove Lemma 2 in Sec. V.

**Lemma 1.** Consider the optimization in (3). Let  $\phi$  be a  $\sigma$ -strongly convex function such that  $\phi(u) \geq 0$  for all  $u \in C$ . Suppose that  $\mathbf{E}[\widetilde{\nabla} f(x_k)] = \nabla f(x_k)$ . For any  $w \in C$ , Algorithm 1 ensures that

$$\begin{aligned} & \mathbf{E} \left[ \sum_{i=1}^k \alpha_i f(x_i) - kf(w) \right] \\ & \leq \mathbf{E} \left[ \sum_{i=1}^k \alpha_i \langle \nabla f(x_i), x_i - w \rangle \right] \\ & \leq \phi(w) + \frac{1}{2\sigma} \sum_{i=1}^k \alpha_i^2 \|\widetilde{\nabla} f(x_i)\|^2. \end{aligned} \quad (4)$$

An example of function  $\phi(u)$  is  $\frac{1}{2}\|u - x_0\|^2$ , i.e., a function that measures the distance between  $u$  and a vector  $x_0 \in C$ . The bound in Eq. (4) is closely connected to the convergence of subgradient methods [15]. If we set  $w = x^* := \arg \min_{u \in C} f(u)$  and assume that  $\|\widetilde{\nabla} f(x_k)\|$  are uniformly upper bounded for all  $k$ , one can show that  $f(x_k)$  converges, in probability, to  $f(x^*)$  (see [15] for the details).

## IV. PROBLEM FORMULATION

Consider a system of queues that evolve as follows:

$$Q_{k+1} = [Q_k + A_k - S_k]^+ \quad k = 1, 2, \dots$$

where  $Q_k \in \mathbf{R}_+^n$ ,  $\{A_i \in \mathcal{A}\}_{i=1}^k$  and  $\{S_i \in \mathcal{S}\}_{i=1}^k$  are the packet arrivals and service in a time slot  $t$ . Without loss of generality, we assume that  $\mathcal{A} := \{a_1, \dots, a_{|\mathcal{A}|}\}$  and  $\mathcal{S} := \{s_1, \dots, s_{|\mathcal{S}|}\}$  are discrete sets in  $\mathbf{R}^n$ , e.g.,  $\mathcal{A}, \mathcal{S} \subseteq \{0, 1\}^n$ . Also, we will refer to the elements in  $\mathcal{S}$  as schedules.

For a given arrival process  $\{A_i \in \mathcal{A}\}_{i=1}^k$ , our goal is to design a policy that generates a sequence  $\{S_i \in \mathcal{S}\}_{i=1}^k$  that stabilizes the system of queues (see Sec. III-B). Such a problem is feasible when

$$\lambda := \lim_{k \rightarrow \infty} \frac{1}{k} \sum_{i=1}^k A_i, \quad A_i \in \mathcal{A}$$

is in the set

$$\Lambda := \{\lambda \in \mathbf{R}_+^n : \lambda \preceq \mu \text{ for some } \mu \in \text{conv}(\mathcal{S})\},$$

which is known in the networking literature as the *network capacity region* [2]. In short, the network capacity region contains all the vectors  $\lambda$  such that  $\lambda \preceq \mu$  for a vector  $\mu$  in the convex hull of  $\mathcal{S}$ . Also, note that, by construction, we can write any vector  $\mu \in \text{conv}(\mathcal{S})$  as the running average of service vectors in  $\mathcal{S}$ , i.e.,  $\mu = \lim_{k \rightarrow \infty} \frac{1}{k} \sum_{i=1}^k S_i$ , with  $S_i \in \mathcal{S}$ .

In the rest of the paper, we will make the following assumption, which ensures that the problem is *strictly feasible* (see, for example, [3, pp. 104]).

**Assumption 1.**  $\lambda$  is in the interior of the network capacity region. That is, for a given vector  $\lambda$ , there exists a vector  $\mu \in \text{conv}(\mathcal{S})$  and a  $\eta > 0$  such that  $\lambda + \eta \mathbf{1} \preceq \mu$ .

The problem of designing a service policy that stabilizes the queuing system can be divided into two cases, depending on whether  $\lambda$  is known, or not. We are interested in solving the problem when  $\lambda$  is not known (Sec. IV-B), but present the case where  $\lambda$  is known since both cases are closely connected in our work.

#### A. $\lambda$ is known

In this case, we can design a service policy in two steps:

- (i) Find a vector  $\mu \in \text{conv}(\mathcal{S})$  such that  $\lambda + \eta \mathbf{1} \preceq \mu$  for some  $\eta > 0$ .
- (ii) Decompose  $\mu$  as the convex combination of elements in  $\mathcal{S}$ , i.e., write  $\mu = \sum_{j=1}^{|\mathcal{S}|} \theta_j s_j$  for some  $\theta_j \geq 0$  with  $\sum_{j=1}^{|\mathcal{S}|} \theta_j = 1$ . Such a decomposition is always possible because  $\mu$  is, by construction, in the convex hull of  $\mathcal{S}$ .

With the decomposition of  $\mu$ , we can select an element in  $\mathcal{S}$  with probability equal to its weight, i.e.,

$$\text{Prob}(S_k = s_j) = \theta_j. \quad (5)$$

Such a policy is often known as the *randomized policy*; it will generate a sequence  $\{S_i \in \mathcal{S}\}_{i=1}^k$  where  $\mathbf{E}[S_k] = \mu$  for all  $k \in \{1, 2, \dots\}$ . Thus, we have that  $\mathbf{E}[A_k] - \mathbf{E}[S_k] \preceq -\eta \mathbf{1}$  and so the queues will be strongly stable by Proposition 1.

Finding a vector  $\mu$  that ensures  $\lambda + \eta \mathbf{1} \preceq \mu$  is easy if we can check whether a vector is in the convex hull of  $\mathcal{S}$ . For instance, we could solve the optimization  $\max_{\eta \geq 0} \eta$  s.t.  $\lambda + \eta \mathbf{1} \in \text{conv}(\mathcal{S})$  with the bisection method. The complexity of decomposing a vector  $\mu$  as a convex combination of vectors in  $\mathcal{S}$  depends on the problem. For example, if the number of elements in  $\mathcal{S}$  is small, we can obtain such a decomposition by solving the convex optimization problem:  $\min_{\theta_j \geq 0} \|\sum_{j=1}^{|\mathcal{S}|} \theta_j s_j - \mu\|^2$  s.t.  $\sum_{j=1}^{|\mathcal{S}|} \theta_j = 1$ . However, that is not always possible. For example, in the cross-bar switches, the number of actions (i.e.,  $|\mathcal{S}|$ ) increases factorially with the number of input/output ports [16].<sup>3</sup>

To conclude this section, we present an example of how to decompose a  $3 \times 3$  traffic matrix of a cross-bar switch with the approach described above.

**Example 1.** Consider a  $3 \times 3$  cross-bar switch with arrival matrix (i.e., a 9-dimensional vector):

$$\lambda = \begin{bmatrix} 0.6 & 0.3 & 0.0 \\ 0.1 & 0.0 & 0.8 \\ 0.2 & 0.6 & 0.1 \end{bmatrix}. \quad (6)$$

Set  $\mathcal{S}$  corresponds to the permutation matrices plus the null action (i.e., the all-zeros matrix). Matrix  $\lambda$  is in the interior of the capacity region since the sum of the rows and columns is equal to 0.9. A vector  $\mu$  that satisfies  $\lambda + \eta \mathbf{1} \preceq \mu$  is

$$\mu = \begin{bmatrix} 0.633 & 0.333 & 0.033 \\ 0.133 & 0.033 & 0.833 \\ 0.233 & 0.633 & 0.133 \end{bmatrix}.$$

<sup>3</sup>e.g., a switch with 16 ports has over 20 trillion schedules/actions.

---

**Algorithm 2** Dual averaging algorithm with stochastic (sub)gradients for the Lagrange dual problem (8)

---

**Input:**  $f$  and  $C$  convex,  $\phi(u) = \frac{1}{2}\|u\|^2$   
**Set:**  $s_1 = 0$ ,  $x_1 \in C$ ,  $k = 1$   
**while** termination condition is not met **do**  
 $\alpha_k \leftarrow 1/\sqrt{k}$   
 $y_k = \arg \max_{v \geq 0} \{\langle s_{k-1}, v \rangle - \phi(v)\}$   
 $\mu_k = \arg \min_{u \in C} \{f(u) + \langle y_k, \lambda - u \rangle\}$   
 $s_{k+1} \leftarrow s_k + \alpha_k (A_k - \mu_k)$   
 $k \leftarrow k + 1$   
**end while**

---

That is, we have added 0.033 to each entry of  $\lambda$ . We can decompose  $\mu$  as follows:

$$\frac{19}{30} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} + \frac{6}{30} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} + \frac{4}{30} \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} + \frac{1}{30} \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}.$$

A randomized policy will select each permutation matrix, respectively, with probability  $\frac{19}{30}$ ,  $\frac{6}{30}$ ,  $\frac{4}{30}$ , and  $\frac{1}{30}$ .

#### B. $\lambda$ is not known

A straightforward approach to tackle this case is to observe the arrivals process  $\{A_i \in \mathcal{A}\}_{i=1}^k$  for long enough to ensure that  $\frac{1}{k} \sum_{i=1}^k A_i$  is sufficiently close to  $\lambda$  to estimate a vector  $\mu$  such that  $\lambda + \eta \mathbf{1} \preceq \mu$  for some  $\eta > 0$ . With that, we can decompose  $\mu$  as indicated in the previous section. However, such an approach is typically not desirable because it requires to remain idle (i.e., not serve packets) for a long period of time.

An alternative approach is the following. Suppose that we have access to a sequence  $\{\bar{\mu}_i \in \text{conv}(\mathcal{S})\}_{i=1}^k$  such that  $\lambda + \eta \mathbf{1} \preceq \mathbf{E}[\bar{\mu}_k]$  for all  $k \geq \tau$  for some constant  $\tau \in \{1, 2, \dots\}$ . Then, we can use  $\bar{\mu}_k$  to construct a “randomized” policy as explained in Sec. IV-A. In particular, at time  $k$ , we will write  $\bar{\mu}_k$  as the convex combination of elements in  $\mathcal{S}$ , and select one of the elements with probability proportional to the weights. The difference with the approach mentioned in the paragraph above is that we do not wait to learn a  $\mu$  that satisfies  $\lambda + \eta \mathbf{1} \preceq \mu$  to start the scheduling process. Instead, we schedule as we learn. In the next section, we present an algorithm that learns  $\bar{\mu}_k$  with Nesterov’s dual averaging algorithm but other optimization techniques could be used as long as the necessary conditions are met.

## V. MAIN RESULTS

In this section, we present an approach for constructing a sequence  $\{\bar{\mu}_i \in C\}_{i=1}^k$  that satisfies  $\lambda + \eta \mathbf{1} \preceq \mathbf{E}[\bar{\mu}_k]$  for all  $k \geq \tau$  for some constant  $\tau \in \{1, 2, \dots\}$ .

To start, consider the constrained optimization problem:

$$\begin{aligned} & \underset{\mu \in C}{\text{minimize}} && f(\mu) \\ & \text{subject to} && \lambda \preceq \mu \end{aligned} \quad (7)$$

where  $\lambda$  is the average of the arrivals,  $\mu$  the average service, and  $C := \text{conv}(\mathcal{S})$ . Without loss of generality, we assume that  $0 \in C$ , i.e., there is the option to not serve any packets from the queues.<sup>4</sup> The objective function just expresses a

<sup>4</sup>We will use the fact that  $0 \in C$  in the proof of Lemma 2.

preference on the average service rate  $\mu$  that satisfy the constraints. In the schematic example in Fig. 3, the objective  $f$  controls where  $\mu^* := \arg \min_{\mu \in C, \lambda \preceq \mu} f(\mu)$  will be in the yellow region (i.e., the set of vectors where  $\mu \succeq \lambda$ ).

To tackle (7), we formulate the Lagrange dual problem, which allow us to relax the knowledge of  $\lambda$  by using stochastic (sub)gradients. This will be clear shortly. For now, consider the Lagrange dual problem of (7):

$$\underset{y \succeq 0}{\text{maximize}} \quad h(y) \quad (8)$$

where

$$h(y) := \min_{\mu \in C} \{f(\mu) + \langle y, \lambda - \mu \rangle\}.$$

The Lagrange dual function  $h$  is concave [17, Chapter 5]. Thus, we can use Algorithm 1 since  $-h$  is convex.

Next, consider that we minimize  $-h$  by querying a stochastic first-order oracle at  $y_k$  (see Sec. III-C). The oracle returns vectors of the form

$$\begin{aligned} -\tilde{\nabla} h(y_k) &= -\nabla h(y_k) + \xi_k \\ &= \mu_k - \lambda + \xi_k & (-\nabla h(y_k) &= \mu_k - \lambda) \\ &= \mu_k - A_k & (\xi_k &:= \lambda - A_k) \end{aligned}$$

where

$$\mu_k \in \arg \min_{u \in C} \{f(u) + \langle y_k, \lambda - u \rangle\}. \quad (9)$$

There are two important points to note. First, if  $\mathbf{E}[A_k] = \lambda$ , we have that  $\mathbf{E}[\xi_k] = 0$  for all  $k$ , and so we satisfy the condition for dual averaging with stochastic gradients (see Sec. III-C). The second point is that we do not need to know  $\lambda$  to compute  $\mu_k$ ; see Eq. (9). Thus, if  $\mathbf{E}[A_k] = \lambda$ , we can construct such an oracle by solving a convex optimization and “replacing”  $\lambda$  with  $A_k$ . We have the following result.

**Lemma 2.** *Suppose  $f$  is  $m$ -strongly convex and that Slater’s condition holds, i.e., there exists a  $\eta > 0$  such that  $\lambda + \eta \mathbf{1} \preceq \mu$  for some  $\mu \in C$ . Algorithm 2 (i.e., Algorithm 1 applied to problem (8)) with  $\phi(u) = \frac{1}{2}\|u\|^2$  ensures:*

$$\mathbf{E} \left[ \|\bar{\mu}_k - \mu^*\|^2 \right] \leq \frac{B}{\sigma m} \frac{\log(k) + 1}{\sqrt{k}}$$

where  $\mu^* := \arg \min_{\mu \in C, \lambda \preceq \mu} f(\mu)$ ,

$$\bar{\mu}_k := \frac{1}{\sum_{i=1}^k \alpha_i} \sum_{i=1}^k \alpha_i \mu_i, \quad (10)$$

and  $B := \frac{1}{2} \max_{\mu \in C, a \in A} \|\mu - a\|^2$ .

The result ensures that  $\mathbf{E}[\|\bar{\mu}_k - \mu^*\|^2] \rightarrow 0$  as  $k \rightarrow \infty$  at a rate of  $\tilde{O}(1/\sqrt{k})$ . Thus, for large enough  $k$ ,  $\bar{\mu}_k$  will be within a ball around  $\mu^*$  with high probability. Fig. 3 shows, schematically, the intuition behind this idea. The ball where  $\bar{\mu}_k$  lives shrinks with  $k$ , and therefore, for  $k$  sufficiently large, it will eventually happen that  $\lambda \prec \mathbf{E}[\bar{\mu}_k]$  if  $\lambda \prec \mu^*$ . The following lemma establishes a time  $\tau \in \{1, 2, \dots\}$  for which  $\mathbf{E}[\bar{\mu}_k]$  will be strictly larger than  $\lambda$  for all  $k \geq \tau$ .

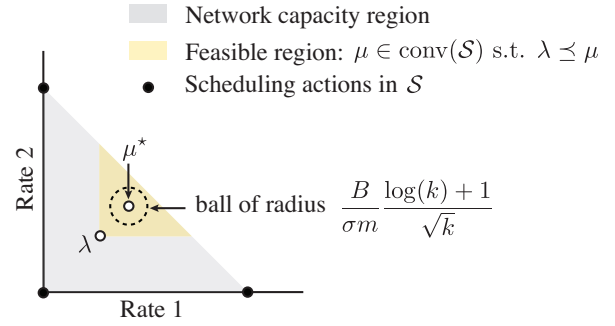


Fig. 3. Schematic illustration of the convergence result in Lemma 2. Vector  $\bar{\mu}_k$  is within a ball of radius  $B(\log(k) + 1)/(\sigma m \sqrt{k})$  centered at  $\mu^*$ . By construction,  $\mu^*$  is strictly larger than  $\lambda$ .

---

### Algorithm 3 Schedule as You Learn (SYL)

---

**Input:**  $f$  that ensures  $\lambda + \eta \mathbf{1} \preceq \mu^*$  (see Theorem 1)

**while** termination condition is not met **do**

- 1) Observe packet arrivals  $A_k$
- 2) Compute  $\bar{\mu}_k = \frac{\sum_{i=1}^k \alpha_i \mu_i}{\sum_{i=1}^k \alpha_i}$  with Algorithm 2
- 3) Decompose  $\bar{\mu}_k$ , i.e.,

$$\bar{\mu}_k = \sum_{j=1}^{|\mathcal{S}|} \theta_j s_j, \quad s_j \in \mathcal{S},$$

where  $\theta_j \geq 0, \sum_{j=1}^{|\mathcal{S}|} \theta_j = 1$ .

- 4) Select a  $S_k = s_j$  with probability  $\theta_j$ .

**end while**

---

**Lemma 3.** *Suppose that  $\lambda + \eta \mathbf{1} \preceq \mu^*$  for some constant  $\eta > 0$ . Then, Algorithm 2 ensures that*

$$\lambda - \mathbf{E}[\bar{\mu}_k] \preceq -\frac{\eta}{2} \mathbf{1} \quad \forall k \geq \left( \frac{4B}{\eta^2 \sigma m} \right)^2.$$

Vector  $\bar{\mu}_k$  will ensure queue stability for  $k \geq \tau$ . However, the queue updates require that the service vectors are schedules so that they can be implemented in practice. For that, we can decompose  $\bar{\mu}_k$  as the convex combination of schedules, and select one randomly—as when  $\lambda$  is known (see Sec. IV-A). We have arrived at the following theorem.

**Theorem 1.** *Consider an arrival process  $\{A_i\}_{i=1}^k$  with  $\mathbf{E}[A_k] = \lambda$  for all  $k \geq 1$ . Select  $S_k$  as indicated in Algorithm 3, and suppose that  $\lambda + \eta \mathbf{1} \preceq \mu^*$  for some constant  $\eta > 0$ , where  $\mu^*$  is the solution to the optimization in (7). Then, the queuing system is strongly stable.*

Algorithm 3 can be seen as an “online” version of a randomized algorithm when  $\lambda$  is known (Sec. IV-A). It generates randomized schedules in each time slot  $k$  where the expected throughput will, eventually, be strictly larger than the mean arrival rate. Thus, the queuing system will be strongly stable (Proposition 1).

There are multiple ways to ensure  $\lambda + \eta \mathbf{1} \preceq \mu^*$ . One of them is to let  $\mu := \hat{\mu} - \gamma \mathbf{1}$  where  $\hat{\mu} \in C$  and  $\gamma \geq 0$  a slack variable, and use  $f(\gamma) = \gamma^2 - \gamma$ . The objective enforces  $\gamma$  to

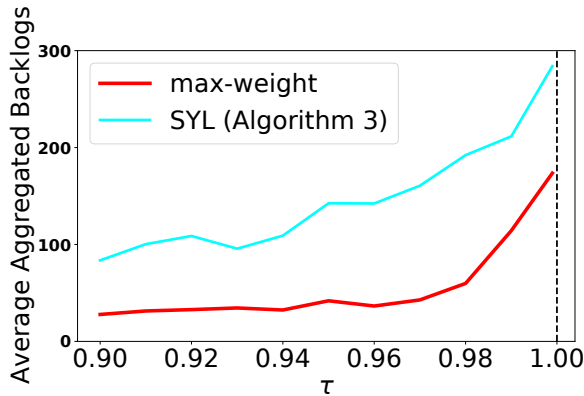


Fig. 4. Aggregated backlogs in the  $3 \times 3$  cross-bar switch example in Sec. VI-A for different arrival rates  $\tau(\lambda/0.9)$ , where  $\lambda$  is given in Eq. (6).

be strictly positive when there exists a  $\hat{\mu} \in C$  that is strictly larger than  $\lambda$ , which is the case by Assumption 1. Thus, we will have  $\mu^* = \hat{\mu}^* - \gamma^* \mathbf{1}$  where  $\gamma^* > 0$  and therefore  $\lambda + \gamma^* \mathbf{1} \preceq \hat{\mu}^*$ . The randomization of the schedules should be over the running average  $\frac{1}{\sum_{i=1}^k \alpha_i} \sum_{i=1}^k \alpha_i \hat{\mu}_i$ .

To conclude, we note some limitations of the approach compared to max-weight approaches:

- (i) Finding a decomposition of  $\bar{\mu}_k$  is computationally more expensive than finding a maximum-weighted matching in a graph. For instance, in cross-bar switches, decomposing  $\bar{\mu}_k$  can require computing up to  $(n-1)^2 + 1$  perfect matchings/linear programs [16], [18]. This has an impact on the frequency in which we can select schedules.
- (ii) To obtain  $\bar{\mu}_k$ , we need to solve a convex optimization problem at each time slot (step 2 in Algorithm 3 and implicit steps in Algorithm 2). While this can be carried out efficiently with off-the-shelf solvers, this is a step that max-weight approaches do not need to carry out.
- (iii) The approach presented assumes that the network connectivity is static, whereas in max-weight approaches, the network connectivity can vary over time. Or, equivalently, when the network connectivity changes, only a fraction of the schedules in  $\mathcal{S}$  are available at a given time. In our approach, the main technical difficulty of having a time varying connectivity is that we may not be able to express  $\bar{\mu}_k$  as the convex combination of the available schedules.

## VI. NUMERICAL EXPERIMENTS

This section presents two numerical experiments for a  $3 \times 3$  cross-bar switch. Experiment A compares the performance of SYL (Algorithm 3) in terms of queue sizes. Experiment B shows an example of how we can improve the latency of a flow with a heuristic strategy based on SYL that prioritizes a flow while aiming to maintain the expected rate  $\bar{\mu}_k$ .

### A. Experiment A

Consider a  $3 \times 3$  cross-bar with arrivals matrix  $\tau(\lambda/0.9)$ , where  $\lambda$  is as in Eq. (6), and  $\tau \in [0, 1)$  a parameter that controls the intensity of the arrivals while being within the network capacity region. The arrivals are assumed to be

Bernoulli distributed with probabilities given by the matrix  $\tau(\lambda/0.9)$ . We run max-weight and SYL for 100k iterations and different values of  $\tau$  in the range  $[0.9, 1)$ . The simulation results are shown in Fig. 4. Observe from the figure that in both policies, the average queue backlog increases as  $\tau \rightarrow 1$  and it explodes as we surpass the capacity boundary. This is the typical behaviour in max-weight approaches when the arrival rate is close the boundary of the capacity region (see, for example, [3, Fig. 4.2]). However, note that max-weight has smaller backlogs sizes compared to SYL. This is because, unlike SYL, max-weight aim to minimize the *total* number of packets in the queues. Minimizing the total number of packets in the queues is, however, not the best strategy if flows have different requirements in terms of latency, as we show in the following example.

### B. Experiment B

In this experiment, we consider the setting in Sec. VI-A for a fixed  $\tau = 0.98$ , and suppose that the flow 1-2, going from input port to 1 to output 2, has low latency requirements.<sup>5</sup> We compare four policies: max-weight [2], delay max-weight [19], SYL (Algorithm 3), and a heuristic variant of SYL. In particular, the SYL variant reserves a fixed number of “tokens” for schedules serving flow 1-2. In each time slot, it selects a schedule using a randomized policy as in SYL, but when sensitive flow packets arrive and tokens are available, it selects a schedule serving these packets and allocates a token to the unused schedule. If tokens are not available, it selects to the randomized schedule. In the absence of packets from the sensitive flow 1-2, if the random selection would serve it, the algorithm instead chooses the schedule with the most tokens and subtract one.

We run the four policies for 100k iterations and show the delay distribution in Figure 5. We only show in the figure flows 1-2, 1-1, and 2-3 due to space constraints, but the behavior of these flows is representative. First, observe that max-weight incurs high delays to flow 1-2 since max-weight prioritizes serving queues with larger backlogs. In other words, max-weight does not prioritize queues that grow “slower”, which is the case of flow 1-2 since this is the flow with lowest intensity. This “unfair” behavior is fixed by the delay max-weight policy, which provides similar delay distributions to all flows. However, making all delay distribution similar does not guarantee low latency to the flow of interest. Finally, observe that the SYL variant provides low latency to flow 1-2 at the price of increasing the delay of the other flows. Importantly, the waiting times of the packets in the queues are bounded, which means that all the packets that get into the queues will eventually get out.

## VII. CONCLUSIONS

This paper has presented a novel approach to designing throughput optimal policies for queueing systems. The approach consists of learning an average rate that ensures queue stability, and then selecting schedules in a “randomized”

<sup>5</sup>i.e., flow 1-2 is the entry with value 0.3 in matrix  $\lambda$  in Eq. (6)






Fig. 5. Comparative histograms of delay distributions across the four scheduling policies (max-weight, delay max-weight, SYL, and SYL-variant) and three different flows, showcasing the probability density of delays for pairs of nodes under three different traffic flows. For the SYL-variant we used 100 tokens for the schedules that serve the sensitive flow 1-2 (see Sec. VI-B).

manner that generate such a rate in expectation. An appealing characteristic of our approach is that the learning process can be offloaded to a subroutine, which allows us to focus on designing a scheduling policy for a particular system. A natural follow-up of this work is to design scheduling policies that are “non-randomized”, and that their expected throughput ensures queue stability.

#### ACKNOWLEDGEMENT

The research work was supported by the Army Research Office MURI under the project number W911NF2110325 and by the National Science Foundation under the project number 2402862.

#### REFERENCES

- [1] A. Wong, C. Chiu, G. Hains, J. Behnke, Y. Khmelevsky, and T. Sutherland, “Network latency classification for computer games,” in *IEEE International Conference on Recent Advances in Systems Science and Engineering (RASSE)*, 2021, pp. 1–6.
- [2] L. Tassiulas and A. Ephremides, “Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks,” *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, 1992.
- [3] L. Georgiadis, M. J. Neely, and L. Tassiulas, “Resource allocation and cross-layer control in wireless networks,” *Foundations and Trends® in Networking*, vol. 1, no. 1, pp. 1–144, 2006.
- [4] M. J. Neely, *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan and Claypool Publishers, 2010.
- [5] A. Shoemaker and S. Vare, “Edmonds’ blossom algorithm,” *CME*, 2016.
- [6] M. J. Neely, “Energy-aware wireless scheduling with near-optimal backlog and convergence time tradeoffs,” *IEEE/ACM Transactions on Networking*, vol. 24, no. 4, pp. 2223–2236, 2015.
- [7] Y. Nesterov, “Primal-dual subgradient methods for convex problems,” *Mathematical programming*, vol. 120, no. 1, pp. 221–259, 2009.
- [8] P. Promponas, V. Valls, and L. Tassiulas, “Full exploitation of limited memory in quantum entanglement switching,” in *IFIP Networking*, 2023, pp. 1–9.
- [9] G. D. Celik and E. Modiano, “Scheduling in networks with time-varying channels and reconfiguration delay,” *IEEE/ACM Transactions on Networking*, vol. 23, no. 1, pp. 99–113, 2013.
- [10] M. G. Markakis, E. Modiano, and J. N. Tsitsiklis, “Delay stability of back-pressure policies in the presence of heavy-tailed traffic,” *IEEE/ACM Transactions on Networking*, vol. 24, no. 4, pp. 2046–2059, 2016.
- [11] M. J. Neely and H. Yu, “Online convex optimization with time-varying constraints,” *arXiv preprint arXiv:1702.04783*, 2017.
- [12] L. Huang, “Deterministic mathematical optimization in stochastic network control,” Ph.D. dissertation, University of Southern California, 2011.

- [13] V. Valls and D. J. Leith, “A convex optimization approach to discrete optimal control,” *IEEE Transactions on Automatic Control*, vol. 64, no. 1, pp. 35–50, 2019.
- [14] T. Chen, A. Mokhtari, X. Wang, A. Ribeiro, and G. B. Giannakis, “Stochastic averaging for constrained optimization with application to online resource allocation,” *IEEE Transactions on Signal Processing*, vol. 65, no. 12, pp. 3078–3093, 2017.
- [15] S. Boyd and A. Mutapcic, “Stochastic subgradient methods,” *Lecture Notes for EE364b, Stanford University*, vol. 9, 2008.
- [16] S. Bojja Venkatakrishnan, M. Alizadeh, and P. Viswanath, “Costly circuits, submodular schedules and approximate carathéodory theorems,” in *SIGMETRICS*, 2016, pp. 75–88.
- [17] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [18] V. Valls, G. Iosifidis, and L. Tassiulas, “Birkhoff’s decomposition revisited: Sparse scheduling for high-speed circuit switches,” *IEEE/ACM Transactions on Networking*, vol. 29, no. 6, pp. 2399–2412, 2021.
- [19] A. Mekittikul and N. McKeown, “A starvation-free algorithm for achieving 100% throughput in an input-queued switch,” in *Proc. ICCCN*, vol. 96, 1996, pp. 226–231.
- [20] S. Kakade, S. Shalev-Shwartz, and A. Tewari, “On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization,” 2009, <http://ttic.uchicago.edu/shai/papers/KakadeShalevTewari09.pdf>.
- [21] R. T. Rockafellar and R. J.-B. Wets, *Variational analysis*. Springer Science & Business Media, 2009, vol. 317.
- [22] Y. Chen, “Lecture 14: Constrained optimization over closed convex sets,” university Of Wisconsin-Madison.

#### APPENDIX

##### A. Proof of Proposition 1

The proof is inspired by the method of proof in max-weight policies [3], but we replace max-weight by a randomized policy. To start, observe

$$\begin{aligned}
 \|Q_{k+1}\|^2 &= \|[Q_k + Z_k]^+\|^2 \\
 &\leq \|Q_k + Z_k\|^2 \\
 &= \|Q_k\|^2 + \|Z_k\|^2 + 2\langle Q_k, Z_k \rangle \\
 &\leq \|Q_k\|^2 + \sigma^2 + 2\langle Q_k, Z_k \rangle
 \end{aligned}$$

Rearrange terms and sum from  $i = 1, \dots, k$

$$\|Q_{k+1}\|^2 \leq \|Q_1\|^2 + k\sigma^2 + 2 \sum_{i=1}^k \langle Q_i, Z_i \rangle.$$

Take expectations w.r.t.  $Z_k$ , and since  $Q_k$  and  $Z_k$  are independent,

$$\mathbf{E}[\|Q_{k+1}\|^2] \leq \|Q_1\|^2 + k\sigma^2 - 2\eta \sum_{i=1}^k \mathbf{E}[\|Q_i\|_1].$$

Rearrange terms and divide by  $2\eta k$  to obtain:

$$\frac{1}{k} \sum_{i=1}^k \mathbf{E}[\|Q_i\|_1] \leq \frac{\|Q_1\|_2^2}{2k\eta} + \frac{\sigma^2}{2\eta}.$$

Finally, since  $\|Q_1\|_2 \leq \|Q_1\|_1$ , take  $k \rightarrow \infty$  to obtain the result.

### B. Proof of Lemma 1

This proof follows the strategy explained in [7, Sec. 2]. To start, let  $l_k(w) := \sum_{i=1}^k \langle \alpha_i \nabla f(x_i), x_i - w \rangle$ . Next, observe that we can write

$$\begin{aligned} l_k(w) &= \sum_{i=1}^k \langle \alpha_i \tilde{\nabla} f(x_i), x_i - w \rangle + N_k \\ &= \sum_{i=1}^k \langle \alpha_i \tilde{\nabla} f(x_i), x_i \rangle + \langle s_k, w \rangle + N_k \end{aligned}$$

where  $N_k := \sum_{i=1}^k \langle \alpha_i \xi_i, x_i - w \rangle$ ,  $\xi_k = \nabla f(x_k) - \tilde{\nabla} f(x_k)$ , and  $s_k = s_0 - \sum_{i=1}^k \alpha_i \tilde{\nabla} f(x_i)$ .

We proceed to upper bound  $\langle s_k, w \rangle$ . Since  $\phi$  is convex, by Fenchel's inequality [17, Sec. 3.3.2],  $\langle s_k, w \rangle \leq \phi(w) + \phi^*(s_k)$ . Thus,

$$l_k(w) \leq \phi(w) + \phi^*(s_k) + \sum_{i=1}^k \langle \alpha_i \tilde{\nabla} f(x_i), x_i \rangle + N_k \quad (11)$$

Now, we upper bound the second and third terms in the right-hand-side of the last equation. First, observe that by the  $\frac{1}{\sigma}$ -smoothness of  $\phi^*$  [20],

$$\begin{aligned} \phi^*(s_k) &\leq \phi^*(s_{k-1}) + \langle \nabla \phi^*(s_{k-1}), s_k - s_{k-1} \rangle + \frac{1}{2\sigma} \|s_k - s_{k-1}\|^2 \\ &= \phi^*(s_{k-1}) - \langle \nabla \phi^*(s_{k-1}), \alpha_k \tilde{\nabla} f(x_k) \rangle + \frac{\alpha_k^2}{2\sigma} \|\tilde{\nabla} f(x_k)\|^2 \end{aligned}$$

where  $s_k - s_{k-1} = -\alpha_k \tilde{\nabla} f(x_k)$ . Apply the argument recursively from  $i = 1, \dots, k$  to obtain

$$\begin{aligned} \phi^*(s_k) &\leq \phi^*(s_0) - \sum_{i=1}^k \langle \nabla \phi^*(s_{i-1}), \alpha_i \tilde{\nabla} f(x_i) \rangle \\ &\quad + \frac{1}{2\sigma} \sum_{i=1}^k \alpha_i^2 \|\tilde{\nabla} f(x_i)\|^2 \end{aligned} \quad (12)$$

where  $\phi^*(s_0) \leq 0$  since  $\phi$  is non-negative and  $s_0 = 0$ . Also, since  $x_k = \nabla \phi^*(s_{k-1}) = \arg \min_{u \in C} \{ \langle s_{k-1}, u \rangle - \phi(u) \}$  for all  $k = 1, 2, \dots$  [21, Proposition 11.3], we can use Eq. (12) in Eq. (11) to obtain

$$l_k(w) \leq \phi(w) + \frac{1}{2\sigma} \sum_{i=1}^k \alpha_i^2 \|\tilde{\nabla} f(x_i)\|^2 + N_k.$$

Finally, note that  $\mathbf{E}[N_k] = 0$  because the noise vector  $\xi_k$  is independent of  $x_k$  by assumption, i.e.,  $\mathbf{E}[\tilde{\nabla} f(x_k) - \nabla f(x_k)] = \mathbf{E}[\xi_k] = 0$ . Hence,

$$\mathbf{E}[l_k(w)] \leq \phi(w) + \frac{1}{2\sigma} \sum_{i=1}^k \alpha_i^2 \|\tilde{\nabla} f(x_i)\|^2.$$

To conclude, note that the left-hand-side of Eq. (4) follows by the convexity of  $f$ .

### C. Proof of Lemma 2

Slater's condition ensures that strong duality holds [17, Chapter 5], i.e.,  $h^* := \sup_{y \geq 0} h(y) = \min_{u \in C, \lambda \preceq u} f(u) =: f^*$ . Next, since  $h(y) \leq h(y^*)$  for every  $y \succeq 0$  and  $h(y^*) = f^*$ , we have

$$\sum_{i=1}^k (h(y_i) - f^*) \leq 0 \quad (13)$$

Now, note that  $h(y_k) = f(\mu_k) + \langle y_k, \lambda - \mu_k \rangle$  where  $\mu_k \in \arg \min_{u \in C} \{ f(u) + \langle y_k, \lambda - u \rangle \}$ . Thus, rearranging terms and multiplying each term in the sum by  $\alpha_k$ , we have

$$\sum_{i=1}^k \alpha_i (f(\mu_i) - f^*) \leq - \sum_{i=1}^k \alpha_i \langle y_i, \lambda - \mu_i \rangle \quad (14)$$

Next, note that  $-h$  is convex, and that the (sub)gradient of  $-h(y_k)$  w.r.t.  $y_k$  is  $-(\lambda - \mu_k)$ . Thus, by Lemma 1 with  $w = 0$  (0 in the domain of  $-h$ ), we have

$$\mathbf{E} \left[ \sum_{i=1}^k \alpha_i (f(\mu_i) - f^*) \right] \leq \phi(0) + \frac{1}{2\sigma} \sum_{i=1}^k \alpha_i^2 \|\lambda - \mu_i\|^2$$

where  $\phi(0) = 0$ . We proceed the lower bound the left-hand-side of the last equation. By the convexity of  $f$ ,

$$\mathbf{E} \left[ \sum_{i=1}^k \alpha_i (f(\mu_i) - f^*) \right] \geq \mathbf{E} \left[ \left( \sum_{i=1}^k \alpha_i \right) (f(\bar{\mu}_k) - f^*) \right].$$

Since  $f(\bar{\mu}_k) - f^* \geq \langle \nabla f(\mu^*), \bar{\mu}_k - \mu^* \rangle + \frac{m}{2} \|\bar{\mu}_k - \mu^*\|^2$  (by the  $m$ -strong convexity of  $f$ ) and  $\langle \nabla f(\mu^*), \bar{\mu}_k - \mu^* \rangle \geq 0$  because  $-\nabla f(\mu^*) \in \mathcal{N}_C(x^*) := \{z \in \mathbf{R}^n \mid \langle z, \mu - \mu^* \rangle \leq 0 \forall \mu \in C\}$  (see [21, Ch. 6b] and [22]), we have

$$\mathbf{E} \left[ \sum_{i=1}^k \alpha_i (f(\mu_i) - f^*) \right] \geq \frac{m}{2} \left( \sum_{i=1}^k \alpha_i \right) \mathbf{E} \left[ \|\bar{\mu}_k - \mu^*\|^2 \right].$$

Finally, divide across by  $\frac{m}{2} \sum_{i=1}^k \alpha_i$  to obtain a bound on  $\mathbf{E}[\|\bar{\mu}_k - \mu^*\|^2]$ , where  $\sum_{i=1}^k \alpha_i \geq \sqrt{k}$  and  $\sum_{i=1}^k \alpha_i^2 \leq \log(k) + 1$  since  $\alpha_k = \frac{1}{\sqrt{k}}$ .

### D. Proof of Lemma 3

We will show that  $\bar{\mu}_k$  generated by Algorithm 2 satisfies  $\lambda + \frac{\eta}{2} \mathbf{1} \preceq \mathbf{E}[\bar{\mu}_k]$  for all  $k \geq \tau$  for some  $\tau \in \{1, 2, \dots\}$ . Since  $\sqrt{x}$  is concave for  $x \geq 0$ , we have that

$$\mathbf{E}[\|\bar{\mu}_k - \mu^*\|] = \mathbf{E} \left[ \sqrt{\|\bar{\mu}_k - \mu^*\|^2} \right] \leq \sqrt{\mathbf{E}[\|\bar{\mu}_k - \mu^*\|^2]}.$$

Hence, from Lemma 2,  $\mathbf{E}[\|\bar{\mu}_k - \mu^*\|] \leq \sqrt{\frac{B}{\sigma m} \frac{\log(k)+1}{\sqrt{k}}}$ , and therefore  $\mathbf{E}[\|\bar{\mu}_k - \mu^*\|] \leq \frac{\eta}{2}$  for  $k \geq \left( \frac{4B}{\eta^2 \sigma m} \right)^2$ .

Finally, since  $\|\cdot\|_\infty \leq \|\cdot\|_2$ , we have that  $\mu^* - \mathbf{E}[\bar{\mu}_k] \preceq \frac{\eta}{2} \mathbf{1}$ . And since  $\lambda + \eta \mathbf{1} \preceq \mu^*$  (by assumption in the lemma), it follows that

$$\lambda + \eta \mathbf{1} - \mathbf{E}[\bar{\mu}_k] \preceq \frac{\eta}{2} \mathbf{1},$$

and, therefore,  $\lambda - \mathbf{E}[\bar{\mu}_k] \preceq -\frac{\eta}{2} \mathbf{1}$  for all  $k \geq \left( \frac{4B}{\eta^2 \sigma m} \right)^2$ .